
Optimising Quantum Networks

Cillian Harney

Doctor of Philosophy.

University of York,
Computer Science

December 2022

Abstract

Quantum communication offers a solution to the threat of quantum computers on cryptographic security and the ability to share quantum information and entanglement on a global scale. The construction of quantum networks and an overarching *quantum internet* will provide the infrastructure needed to facilitate worldwide quantum communication and enable ground breaking advances in science, technology and beyond.

However, quantum networks face challenges that classical networks conveniently avoid. The laws of quantum mechanics impose a fundamental rate-loss tradeoff which critically limits the ability to achieve high rates over long distances. This introduces an intrinsic difference between classical and quantum communication networks with effects that ripple throughout every facet of network design, implementation and utilisation.

In this thesis, we devise new ways to characterise the performance of large-scale quantum networks. Combining expertise from quantum information theory, graph theory and network theory we derive analytical methods with which to inspect the fundamental limits of large-scale quantum networks. We do this through the introduction of an analytic network architecture which exhibits desirable features within high-rate, well connected topologies. These techniques (and variants thereof) are then used in a multitude of contexts: to compare the limits and resource demands of quantum fibre networks and satellite-based quantum repeaters, to benchmark end-to-end network capacity bounds, and to access useful benchmarks for free-space quantum networking.

Motivated by insight from analytical architectures, we inspect the practicality and criticality of realistic quantum networking. We investigate random quantum network architectures and practical end-to-end routing protocols in order to understand the trade-off between network connectivity, resource consumption and performance guarantees. In doing so, we build new and efficient multi-path routing strategies. These analyses are then extended into the multi-user setting, enlightening properties of a reliable quantum internet that can support many communicators.

For Peggy and Joe.
I love and miss you both.

Contents

Abstract	i
List of Tables	vi
List of Figures	vii
Acknowledgments	xvii
Declaration	xix
Chapter 1 Introduction and Motivation	1
1.1 Quantum Information and Communication	1
1.2 Challenging Classical Cryptography	2
1.3 A Quantum Internet	3
1.4 Quantum Networking	4
1.5 Motivation and Thesis Structure	5
Chapter 2 Basics of Quantum Communication Networks	9
2.1 Quantum Systems, Information and Channels	9
2.2 Quantum Communications and Channel Capacities	15
2.3 Quantum Communication Networks	23
2.4 End-to-End Quantum Network Rates and Capacities	32
Chapter 3 Analytical Methods for High-Rate Global Quantum Networks	39
3.1 Introduction	39
3.2 Weakly-Regular Networks (WRNs)	41
3.3 Optimal Performance of Weakly-Regular Networks	48
3.4 Nodal Densities and Bosonic Lossy Weakly-Regular Networks	61
3.5 Comparison with Satellite Quantum Communications	69
3.6 Conclusion	75

Chapter 4 End-to-End Capacities of Imperfect Repeater Quantum Networks . . .	77
4.1 Introduction	77
4.2 Bounds for Realistic Quantum Networks	80
4.3 Network Parameter Benchmarking with Weakly-Regular Networks	83
4.4 Amplitude Damping Networks	88
4.5 Bosonic Thermal-Loss Networks	93
4.6 Conclusion	99
Chapter 5 Free-Space and Hybrid Quantum Network Capacities	101
5.1 Introduction	101
5.2 Quantum Networking over Fading Channels	103
5.3 Free-Space Quantum Communication	105
5.4 Modular Quantum Networks	116
5.5 Fibre/Satellite Configuration	125
5.6 Ground-Based Free-Space/Fibre Configuration	132
5.7 Conclusion	135
Chapter 6 Practical Routing and Criticality in Complex Quantum Networks . . .	137
6.1 Introduction	137
6.2 Random Quantum Networks	139
6.3 Practical Routing in Quantum Networks	144
6.4 Benchmarking Quantum Networks	152
6.5 Numerical Results	156
6.6 Conclusion	163
Chapter 7 Multi-User Limitations of Quantum Communication Networks	165
7.1 Introduction	165
7.2 Multi-User Quantum Networking	166
7.3 Traffic Management Protocols and Network Filters	170
7.4 Numerical Results	178
7.5 Conclusion	184
Chapter 8 Conclusions	187
8.1 Summary of Presented Work	187
8.2 Outlook and Future Research	189
Appendix A Appendices for Chapter 3	191

A.1	Minimum Node Numbers for Internal Weak Regularity	191
Appendix B	Appendices for Chapter 4	193
B.1	Network Parameter Threshold Theorems	193
B.2	Compound Thermal-Loss Channels	194
B.3	Compound Amplitude Damping Channels	196
Appendix C	Appendices for Chapter 5	199
C.1	General Aspects of Quantum Networks with Community Structure	199
C.2	Application to Hybrid Quantum Networks	212
C.3	Collective Node Isolation	215
Appendix D	Appendices for Chapter 7	221
D.1	Extension to Quantum Multicasting	221
Abbreviations	225
Bibliography	227

List of Tables

Table 4.1	Parameter table for realistic CV-QKD protocols through optical-fibre channels.	95
Table 5.1	Parameter table for the fibre/satellite modular network configuration. Here we consider two similar setups using a collimated Gaussian beam at 800 nm wavelength, but differ in initial spot-size w_0 , receiver aperture a_R and frequency filter $\Delta\lambda$	125
Table 5.2	Parameter table for the free-space/fibre modular network configuration. .	132

List of Figures

- Fig. 2.1 (a) A teleportation protocol aims to simulate an identity channel using a maximally entangled state as the shared resource between Alice and Bob. (b) The protocol can be generalised to an arbitrary LOCC \mathcal{T} in which the users share an arbitrary resource state σ , and Alice performs a general quantum operation \mathbb{A}_k on her systems, communicates the classical outcome k to Bob who then applies his own quantum operation \mathbb{B}_k on his systems, leading to a channel simulation $\mathcal{E}(\rho) = \mathcal{T}(\rho \otimes \sigma)$. (c) Teleportation covariant channels are simulated using their Choi matrix $\rho_{\mathcal{E}} = \mathcal{I} \otimes \mathcal{E}(\Phi_{AB})$ 18
- Fig. 2.2 (a) A quantum communication network can be described as a finite, undirected, weighted graph $\mathcal{N} = (P, E)$. Each node represents (b) collection of local quantum registers which can be used to transmit, receive and operate on quantum systems. Network nodes are connected by quantum channels, typically described by (c) optical fibre or (d) free-space. 24
- Fig. 2.3 (a) Point-to-point quantum communications between an end-user pair of nodes Alice \mathbf{a} and Bob \mathbf{b} . (b) A single end-user pair engage in single-unicast quantum networking, utilising end-to-end routes such as ω to establish connections. (c) Multiple-unicast quantum networking concerns multiple end-user pairs (e.g. $\{\mathbf{a}_1, \mathbf{b}_1\}$ and $\{\mathbf{a}_2, \mathbf{b}_2\}$) who may wish to utilise the network simultaneously. 29
- Fig. 3.1 (a) A sub-graph from a $(k, \lambda, \boldsymbol{\mu}) = (6, 2, \{0, 1, 2\})$ -weakly regular network. Considering the yellow node as an end-user, the blue nodes thus represent the user neighbourhood, with a uniform adjacent commonality of $\lambda = 2$. The non-adjacent commonality decreases as nodes increase in distance from the end-user. (b) A $k = 16$ weakly-regular network with inconsistent adjacent commonality properties. This network is scalable so that a single network cell can be concatenated to construct a larger $k = 16$ internally-WR network. For any node in the network, its $\boldsymbol{\lambda}$ will be one of those from the set $\boldsymbol{\Lambda} = \{\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \boldsymbol{\lambda}_3, \boldsymbol{\lambda}_4\}$. Each adjacent commonality multiset is colour

- coded to its corresponding node on the graph. Note that throughout this work we employ a superscript union notation to describe the repeated union of a single set, e.g. $\{x\}^{\cup 3} = \{x\} \cup \{x\} \cup \{x\} = \{x, x, x\}$, etc. 43
- Fig. 3.2 Distinction between (a) genuine weak-regularity and (b) internal weak-regularity. Genuinely WR networks can be embedded on a closed, three-dimensional surface such as a sphere in order to maintain regularity and avoid boundary effects. Internally WR networks satisfy weak-regularity within some nodal boundary P_{bound} , allowing us to investigate open networks which are defined within some two-dimensional area. 47
- Fig. 3.3 Cut-set cardinality with respect to increasing distance from user-node on a honeycomb lattice. We show some example cuts on a honeycomb network of increasing cut-set dimension. The further one moves from a user-node \mathbf{a} , the more edges that must be cut due to k -regularity. \tilde{C}_1 gives the neighbourhood cut-set $E_{\mathbf{a}}$, \tilde{C}_2 gives the smallest cut-set when limited to network-body edges E' , and \tilde{C}_3 gives a wider cut example. 57
- Fig. 3.4 Examples of network cells that can be used to construct quantum WRNs, where k denotes regularity and λ^* is the adjacent commonality multi-set which minimises the quantity in Eq. (3.20). These quantities characterise the network cell and larger WRNs that they can construct. 67
- Fig. 3.5 All plots are colour coded with respect to architectures derived from network cells in Fig 3.4. (a) Relationship between the optimal end-to-end flooding capacity (equal to the min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$) and the maximum link-length $d_{\mathcal{N}}^{\text{max}}$ required to guarantee it for bosonic lossy quantum networks according to Eq. (3.45). Greater network regularity leads to larger permissible ranges of channel lengths. Panel (b) depicts this relationship between minimum nodal density $\rho_{\mathcal{N}}^{\text{min}}$ with respect to optimal performance for bosonic lossy quantum networks, colour coordinated with the network cells. The grey dotted line relates the nodal density to the average flooding capacity between any pair of nodes in a Waxman Network, as in Eq. (3.80) [65]. 68
- Fig. 3.6 The daily secret-key-rate advantage ΔK_{daily} in Eq. (3.87) achieved by fibre-based quantum WRNs and repeater-chains with capacity achieving links over a single, sun-synchronous satellite-based repeater operating at practical, achievable rates from Eqs. (3.83) and (3.84). The architecture of each WRN is shown inside of Panels (d) to (f), such that vertically aligned panels use the same architecture. Panels (a) to (c) plot ΔK_{daily} with respect to

maximum fibre-length permitted within each structure, $d_{\mathcal{N}}^{\max}$. Panels (d) to (f) plot the relationship between the daily rate advantage ΔK_{daily} and minimum nodal density required in each WRN to achieve it. Satellite-based advantage can be achieved when $\Delta K_{\text{daily}} \leq 0$. All considered satellite setup parameters are shown in the table describing the operation time, direction, altitude h , spot-size ω_0 , receiver aperture a_R 74

Fig. 4.1 Optimising the physical orientation of a quantum network. The physical flow of quantum systems is independent of the logical flow of information, and thus physical directions can always be optimised to utilise the directed channels which possess the greatest capacity. In this way, a quantum network described by a directed graph can always be reduced to an undirected form. From a pair of physically-directed channels on a single edge $(\mathcal{E}_{\mathbf{a} \rightarrow \mathbf{x}}, \mathcal{E}_{\mathbf{x} \rightarrow \mathbf{a}})$, we can choose that which has the greatest capacity according to Eq. (4.1), e.g. $\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{a}}$ in the example above. It is always possible to use this superior edge for logical communication in either direction, reducing it to an optimal undirected edge describing the channel $\mathcal{E}_{\mathbf{ax}}^* = \mathcal{E}_{\mathbf{x} \rightarrow \mathbf{a}}$ 80

Fig. 4.2 (a) Node-splitting procedure for quantum amplitude-damping networks. We consider quantum WRNs within some nodal boundary which satisfy regularity $k = 6$ and adjacent commonality $\lambda_{\mathbf{x}} = \{2\}^{\cup 6}$ for any node $\mathbf{x} \in P$. (b) Node-splitting procedure for bosonic thermal-loss quantum networks. We consider quantum WRNs within some nodal boundary which satisfy regularity $k = 8$ and adjacent commonality $\lambda_{\mathbf{x}} = \{2, 4\}^{\cup 4}$ for any node $\mathbf{x} \in P$ 89

Fig. 4.3 Throughout all plots, dashed lines represent bounds obtained using the squashed entanglement as a single-edge capacity upper-bound. Meanwhile, all solid lines plot bounds derived using the RCI as a single-edge capacity lower-bound. Panel (a) plots bounds on the maximum fibre-length permitted within the network $d_{\mathcal{N}}^{\max}$ and a corresponding minimum nodal density $\rho_{\mathcal{N}}^{\min}$ which ensure that an end-user pair can obtain optimal performance $\mathcal{C}(\mathbf{i}, \mathcal{N})$. Panel (b) depicts bounds on the maximum tolerable internal loss $p_{\mathcal{N}}^{\max}$ permitted within the network with respect to maximum fibre-length and nodal density so to guarantee an optimal flooding capacity. 91

Fig. 4.4 Throughout all plots, dashed lines represent bounds obtained using the REE as a single-edge capacity upper-bound. Meanwhile, all solid lines plot bounds derived using the RCI as a single-edge capacity lower-bound. Panel

(a) plots the maximum inter-nodal separation $d_{\mathcal{N}}^{\max}$ permitted within the network, and its associated minimum nodal density, such that optimal performance $\mathcal{C}(\mathbf{i}, \mathcal{N})$ can be obtained. Here, we consider network setups consisting of ideal repeaters or imperfect repeaters using LLOs and TLOs and heterodyne detection. Panel (b) displays the relationship between maximum fibre-length, minimum nodal density and tolerable thermal noise at the receiver $\bar{n}_{\mathcal{N}}^{r, \max}$ throughout the network. The grey areas of each plot illustrate regions of network parameter space for which we are unable to guarantee *any* end-to-end capacity whatsoever; identifying essential properties for realistic thermal-loss networks. 98

Fig. 5.1 Free-space transmission loss associated with (a) ground-based, (b) ground-satellite and (c) intersatellite communication links. In each plot, solid lines depict the average transmissivity (attenuation averaged over fading dynamics), while dotted lines describe the best-case transmissivity (absence of fading). The dashed lines in Panel (b) describe a ground-satellite free-space link with zenith angle $\theta = 1$ radian, while the others consider $\theta = 0$. The operational setup in (a) is consistent with the parameters in Table 5.2 while (b) and (c) are consistent with Setup (#1) in Table 5.1. 114

Fig. 5.2 (a) A modular quantum network architecture built from community sub-networks $\mathcal{N}_{c_{\alpha}}$, $\mathcal{N}_{c_{\beta}}$ and a backbone network \mathcal{N}_b . Each community is connected to the backbone via the sub-networks $\mathcal{N}_{c_{\alpha}:b}$ and $\mathcal{N}_{c_{\beta}:b}$. Nodes from the community c_j which are directly connected to the backbone are contained in $P_{c_j|b}$, while the nodes in the backbone which are connected to the community are contained in $P_{b|c_j}$. (b) We may idealise this modular structure by placing ideal connectivity constraints on the each of the sub-networks. 118

Fig. 5.3 Examples of minimum cardinality intercommunity cut-sets for connections from an arbitrary community to a Manhattan backbone network ($k_b = 4$). These are valid cuts which isolate remote communities (only one community is illustrated here), and are performed exclusively on the backbone. Panel (a) captures the best-case spatial distribution of the largest potential cut-set when no target-nodes share any edges or neighbours, (b) illustrates an example in which neighbour sharing can diminish the overall cut-set size, and (c) describes the worst-case spatial distribution that minimises the cut-set size. 124

-
- Fig. 5.4 The longest possible intersatellite quantum channel is limited by the line-of-sight separation of two satellites. 128
- Fig. 5.5 Optimal end-to-end performance for an ideal modular network consisting of fibre communities interconnected to a satellite-based backbone. In order to guarantee an optimal flooding rate along the x -axis then the maximum internodal separations in each sub-network on the y -axis must be less than or equal to the plotted bounds. We consider operational settings in Setup (#1) for (a), (c) and Setup (#2) for (b), (d) which are described in Table 5.1. The weather/time conditions are those experienced by the worst-case end-user community. Given an optimal flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N})$, we plot the maximum intersatellite separation z_b^{\max} for different backbone connectivity parameters, and the maximum fibre-length in each community $d_{c_j}^{\max}$ for different community connectivity parameters. The dashed lines in Figs (a) and (c) plot an upper bound the maximum intersatellite separation based on the optimal spatial distribution of (a finite number of) community connected satellite nodes $P_{b|c_j}$ at a maximum altitude $h^{\max} = 1500$ km, while the solid lines plot the lower bound based on the worst spatial distribution (for any altitude). The red line indicates the maximum achievable channel length that can be achieved for two satellites at altitude 1500 km, such that $z_{\text{sight}}^{\max} \approx 5428$ km. 130
- Fig. 5.6 Optimal end-to-end performance for an ideal modular network consisting of free-space communities interconnected to a fibre-based backbone. In order to guarantee an optimal flooding rate along the x -axis then the maximum internodal separations in each sub-network depicted on the y -axis must be less than or equal to the plotted bounds. We use the operational settings in Table 5.2 during clear day-time. Given an optimal flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N})$, we plot the maximum fibre-length d_b^{\max} for different backbone connectivity parameters, and the maximum free-space link-length in each community $z_{c_j}^{\max}$ for different community connectivity parameters. The dashed lines plot an upper bound on the maximum fibre-length based on the optimal spatial distribution of community connected backbone nodes $P_{b|c_j}$, while the solid lines plot a lower bound based on the worst-case spatial distribution. 135
- Fig. 6.1 Connectivity properties of bosonic thermal-loss Waxman networks. Panel (a) displays random networks generated under the parameters listed above and using single-edge capacity upper-bounds. The colour intensity of each

- edge is proportional to its capacity. Panel (b) plots the average main component fraction, in which the critical densities necessary for a connectivity phase transition are identified by the black lines. Panel (c) shows the average degree of networks with a fixed number of nodes (given in legend) with respect to variable network density and area respectively. The legend identifies the fixed node number and indicates whether the network uses upper or lower bounds on the rate distributions. 140
- Fig. 6.2 Connectivity properties of bosonic thermal-loss scale-free networks. Panel (a) displays random networks generated under the parameters listed above and using single-edge capacity upper-bounds. The opacity of each edge is proportional to its capacity. Panel (b) plots the average giant component fraction and (c) the average degree of networks with a fixed number of nodes with respect to variable network density and area respectively. The legend identifies the fixed node number and σ exponent. 143
- Fig. 6.3 Generalised Dijkstra's algorithm for end-to-end route optimisation with respect to a cost function \mathcal{F}_ω . Any cost function has a tentative counterpart $F_\omega^{s,x \rightarrow y}$ which is used to evaluate movement throughout the network. The above pseudocode describes the network exploration phase which is followed by a CONSTRUCTPATH subroutine which simply back tracks from the target node \mathbf{t} to \mathbf{s} using the constructed tentative cost and parent node sets. . . 146
- Fig. 6.4 MDPAlg for end-to-end route optimisation with respect to a cost function \mathcal{F}_ω (with corresponding tentative cost function $F_\omega^{s,x \rightarrow y}$). This pseudocode describes the network exploration phase which is followed by a CONSTRUCTMDP subroutine which uses the tentative cost matrix to identify many disjoint paths from the target node \mathbf{t} to \mathbf{s} . Concerning the protocol types considered in this chapter, we can use CONSTRUCTMDP to explicitly identify M paths (fixed route number protocol) or identify as many paths as necessary to guarantee a rate of K^* (rate requirement protocol). 149
- Fig. 6.5 Relationships between performance, routing consumption and nodal density in bosonic thermal-loss Waxman networks. The network link layer in all cases is described by the upper-bound capacity distribution \mathcal{K}_u from Eq. (6.2) and model parameters $(r_0, \beta) = (100, 1)$. Panels (a)-(b) plot the ensemble average end-to-end capacities achieved by different routing protocols and a number of network radii, R (each protocol and network radius is colour coded with the routing protocol and R legends). Panels (c)-(d) show the ensemble average routing consumption of (c) single path routing

-
- and (d) $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing via the MDPAlg. 157
- Fig. 6.6 Illustration of achievable end-to-end routes on an example network (according to the parameters listed) using \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ for users separated by $r_i \approx 800$ km. Black edges in the networks identify unused edges, while red edges are those which engage in the routing protocol. 158
- Fig. 6.7 Waxman quantum network phase characterisations with respect to link layer descriptions and nodal density. Panel (a) outlines connectivity, consumption and performance based critical densities with respect to networks composed of different link layers. These transitions give rise to network phases in which we can expect particular properties. These phases are labelled on the density diagram, while Panel (b) summarises and describes their implications for quantum networking. Phases III₁ and III₂ describe similar network properties, but differ as to whether the maximum routing consumption is surpassed before or after the flooding-based performance transition. 160
- Fig. 6.8 Relationships between performance, routing consumption and nodal density in bosonic thermal-loss scale-free networks. Throughout all plots, the network link layer is described by the upper-bound capacity distribution \mathcal{K}_u from Eq. (6.2) and we use the model parameters $(n_0, m, \sigma_{\text{deg}}) = (10, 5, 1)$. We also consider unique values for the scale-free model parameter $\sigma_r \in \{1, 2\}$ which are distinguished in the legend. Panels (a)-(c) depict the ensemble average capacity with respect to nodal density for \mathcal{P}_{fl} , \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing respectively. Panels (d)-(e) plot the ensemble average routing consumption with respect to nodal density for \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing respectively. Each analysis is performed for a number of network radii R listed in the legend. 162
- Fig. 6.9 Illustration of achievable end-to-end routes on an example network (under the parameters shown) using \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ for users separated by $r_i \approx 200$ km. Black edges in the networks identify unused edges, while red edges are those which engage in the routing protocol. 163
- Fig. 7.1 Quantum register disjoint (QRD) routing: A repeater node \mathbf{x} with finite number of $k_{\mathbf{x}}$ quantum registers and each register can support a single route (two edges). For example, panels (a)-(c) illustrate a node \mathbf{x} with degree $k_{\mathbf{x}} = 6$ and $k_{\mathbf{x}} \in \{1, 2, 3\}$ registers so that a total of $2k_{\mathbf{x}}$ edges can be supported throughout routing. QRD routing generalises edge/node disjoint routing, as Panel (a) recovers the node-disjoint case and (c) captures the edge-disjoint case. 175

- Fig. 7.2 Edge-splitting via multiplexing: Multiple EUPs can exploit the same edge in a network granted that it can be reliably multiplexed a sufficient number of times, at the expense of splitting the point-to-point rate along each channel. Every edge $(\mathbf{x}, \mathbf{y}) \in E$ has a pre-defined multiplexability such that it can be split it into at most m channels each with reduced rates. 176
- Fig. 7.3 Behaviour of the ensemble average end-to-end capacity for a bosonic thermal-loss quantum Waxman networks under simultaneous multiple-unicast communications. Panel (a) displays results for the $\mathcal{P}_{\text{mdp}}^{M=2}$ multi-path routing strategy, while (b) considers the rate-requirement strategy, $\mathcal{P}_{\text{mdp}}^{K^*=1}$, both of which are mediated by an edge-disjoint TM protocol \mathcal{T}_{ED} . Panel (c) illustrates a collection of colour-coded end-to-end routes between $N_{\text{u}} = 5$ simultaneous EUPs on an example $N = 400$ node Waxman network. 181
- Fig. 7.4 Multi-user critical densities for $R = 250$ km quantum Waxman networks for routing protocols (a) $\mathcal{P}_{\text{mdp}}^{K^*=1}$ and (b) $\mathcal{P}_{\text{mdp}}^{M=2}$, under different TM protocols ranging from edge-disjoint ($k_{\mathbf{x}} \rightarrow \infty$) and QRD ($k_{\mathbf{x}} < \infty$). 182
- Fig. 7.5 Waxman quantum network phases: (I) Pre-percolation transition, (II) post-percolation transition, pre-performance transition, (III) single-unicast reliable rates guaranteed by multi-path routing, (IV) multiple-unicast reliable rates up to $N_{\text{u}} = 100$ guaranteed by multi-path routing and edge disjoint TM, and (V) reliable single-unicast rates guaranteed by single-path routing. 184
- Fig. A.1 Minimum node WRNs for a strict satisfaction of internal regularity for (a) honeycomb network, (b) hexagonal network, (c) Manhattan $k = 8$ network and (d) Manhattan $k = 16$ network. Each case resembles the smallest WRN for which there exist a pair of end-user nodes which do not share an edge or a neighbour, and possess minimum cardinality network-bulk cuts which are unaffected by the open boundary. 192
- Fig. C.1 (a) Quotient graph of a modular backbone network under the community equivalence relation for m -communities, resulting in a star network. (b) When considering the multi-path capacity between end-users α and β , the quotient graph can be simplified to a linear chain. 206
- Fig. C.2 Minimum-cut set cardinalities for collective node isolation. Panels (a), (b) and (c) depict different possible cuts that may emerge via the procedure of Section C.3.2 in which subsets of target nodes can be connected via paths to minimise their total cardinality. Panel (a) displays the natural case of

target-node isolation, (b) shows how one can use closed paths, while (c) illustrates how multiple paths can be formed. In this instance, all cut-sets collect an equivalent number of edges. 218

Acknowledgments

To my parents, Colin and Carmel: Any success I have had in life I owe completely to you both, and everything you have done for me, Sé and Aidan. Your ever-present love and support has always been there, and without it I could never have done a thing. Please always know how grateful I am. Thank you, thank you, thank you. I must also give special thanks to my brothers Sé and Aidan for their love, encouragement and ability to drag me out of stress mode through “Harneypalooza’s” and fantasy football.

Nat: Your unwavering love, support, encouragement, patience and all around wonderful existence has made all of this possible. You have kept me sane and smiling throughout the entire PhD, have listened to my stresses, and even joined me in York to save me from the doom of “The Cube”. Without you none of this makes sense. I love you, and thank you. I must also extend a special thank you to Jan and Peter for always opening up your home to me, and always being a wonderful support.

This thesis would not have been remotely possible without the continued guidance of my supervisor Prof. Stefano Pirandola. Every step of the way I have felt supported, trusted, and like part of a team. I could not have dreamed of a better environment in which to carry out my PhD, and am incredibly lucky. Thank you, we have plenty of work left to do.

I also give thanks to all of the members of the quantum information group in York who I have interacted with throughout my PhD: Panos, Jason, Alex, Athena, Kieran, and Prof. Sam Braunstein. In particular, thank you to Alasdair Fletcher who has been a wonderful PhD teammate, collaborator, and great friend throughout this journey.

I am grateful for all of my Irish friends (Michael, Ittai, Ronan, Darragh, Peter, Jack, Ruairi, James and beyond) for your encouragement and general concern for my sanity. With this thesis written hopefully I can reply to texts a little quicker, and see you all more often.

And finally, thank you to my grandparents: Barney, Evelyn, Peggy and Joe. I know that you are always with me.

Declaration

The work in this thesis is based on research carried out at the Department of Computer Science, University of York, United Kingdom during 2019-2022. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Some parts of this thesis have been published as journals articles in all of which I am the first author. The publications which are relevant to this thesis are listed below:

- [1] Cillian Harney and Stefano Pirandola, *Analytical Methods for High-Rate Global Quantum Networks*, PRX Quantum, 3, 010349 (2022)
- [2] Cillian Harney and Stefano Pirandola, *End-to-End Capacities of Imperfect-Repeater Quantum Networks*, Quantum Science and Technology, 7, 045009 (2022)
- [3] Cillian Harney, Alasdair I. Fletcher and Stefano Pirandola, *End-to-End Capacities of Hybrid Quantum Networks*, Physical Review Applied, 18, 014012 (2022)

Beyond these papers, I have also authored a number of other articles which are not directly related the content of this thesis, but marked formative contributions to my PhD:

- Cillian Harney and Stefano Pirandola, *Secure Quantum Pattern Communication*, PRX Quantum, 3, 010311 (2022)
- Cillian Harney and Stefano Pirandola, *Idler-free multi-channel discrimination via multipartite probe states*, npj Quantum Information, 7, 153 (2021)
- Cillian Harney and Stefano Pirandola, *Analytical bounds for dynamic multi-channel discrimination*, Physical Review A, 104, 032402 (2021)

- Cillian Harney, Mauro Paternostro and Stefano Pirandola, *Mixed state entanglement classification using artificial neural networks*, New Journal of Physics, 23, 063033 (2021)
- Cillian Harney, Leonardo Banchi and Stefano Pirandola, *Ultimate Limits of Thermal Pattern Recognition*, Physical Review A, 103, 052406 (2021)
- Cillian Harney, Stefano Pirandola, Alessandro Ferraro and Mauro Paternostro, *Entanglement classification via neural network quantum states*, New Journal of Physics, 22, 045001 (2020)

Copyright © Cillian Harney 2022

The copyright of this thesis rests with the author. Any quotations from it should be acknowledged appropriately

Chapter 1

Introduction and Motivation

1.1 Quantum Information and Communication

Over the past century, our ability to communicate has undergone rapid and dramatic change. From the invention of the telephone, the deployment of the first classical internet [4], to their ubiquity in modern society, communication has become a foundation on which we build our everyday lives. For this reason, a worldwide communication infrastructure spanning land, sea and space has been ever-growing and expanding in an effort to promise three crucial features: (1) *immediacy* (the ability to transfer information quickly), (2) *reliability* (to communicate despite adversarial conditions) and (3) *security* (the guarantee of private, secure communication).

Yet, as science and technology has evolved, so too has the types of information we wish to communicate. Most prominently, the theoretical and experimental development of quantum information science has brought with it new demands and challenges to the landscape of communication technologies. *Quantum technologies* describe a category of technological devices and procedures which exploit fundamental features of quantum mechanics to perform classically impossible tasks, utilising its unique properties such as quantum superposition and entanglement. The modern field of quantum information theory has grown and thrived in the past fifty years to understand how such technologies can be developed and deployed to unlock new, revolutionary abilities: the classically intractable simulation of complex systems for chemistry, material science and drug discovery [5], advanced techniques for biology and medicine [6, 7], astronomy [8], enhanced gravitational wave detection [9], and much more.

1.2 Challenging Classical Cryptography

The research and development of quantum computers are a particularly pertinent focus of quantum information science. A quantum computer is a different form of computing device which is based upon the more general processing of *qubits*, rather than the exclusive processing of classical bits. Quantum computers will help to drive progress in each of the different technologies mentioned before, but perhaps most importantly within communication and cryptography. In Shor’s seminal paper [10], he discovered that quantum computers are capable of factoring large integers efficiently; a task that is not known to be efficiently solvable on classical devices upon which a vast number of classical cryptographic protocols are based. The security of RSA encryption [11] (and other standardised cryptographic methods) relies upon the concept of one-way functions; functions which are easy to compute one-way using some inputs, but are extremely difficult to invert, e.g. integer factorisation, the discrete logarithm problem [12, 13, 14], etc. Consequently, the future arrival of quantum computers which can efficiently defeat one-way functions pose a critical challenge to the current world of classical communications.

Hence, this new wave of technologies which promise a suite of revolutionary abilities simultaneously poses a threat to the way in which we currently guarantee secure/private communications. Different research directions have emerged to address the issue of security post-quantum computers. The aptly named field of *post-quantum cryptography*¹ aims to determine methods (typically classical) to ward off Shor’s factoring algorithm which are usually based upon public-key cryptography and maintain the strategy of one-way functions. In the near-term, while quantum computers remain small and noisy, such approaches may well be sufficient. However, in the long-term vision of fault-tolerant quantum technologies, unless a classical scheme can be devised that is provably intractable on quantum computers they may simply be postponing the inevitable.

Perhaps attacks posed by quantum computers shouldn’t be resisted with classical means, but instead with quantum means? Enter *quantum cryptography*, a rich field of study which fights fire-with-fire, imposing post-quantum security through the manipulation of quantum mechanical systems. Indeed, the protocol of quantum key distribution (QKD) [15, 16, 17] stands as perhaps the most mature and readily deployed quantum technology to date. It is a quantum cryptographic protocol which promises *provably secure* classical communication based upon the fundamental laws of quantum mechanics; going beyond the difficulty of one-way functions, and providing mathematically guaranteed privacy. QKD presents an

¹Commonly referred to in many different ways, e.g. “quantum-safe”, “quantum-proof”, “quantum-resistant”, etc.

exciting solution to navigating Shor's algorithm in a post-quantum world, one that could form the basis of secure communications of the future.

1.3 A Quantum Internet

Quantum communication theory extends into a broader domain, beyond that of securing classical communication against quantum attacks. The more general goal of immediately, reliably and securely communicating quantum information via the exchange of quantum systems opens remarkable new avenues of exploration. A pair of end-users may want to share a *specific* target quantum state, a goal we label *quantum state transfer*. One may consider the alternative task of *entanglement distribution*, in which the goal of communicators is to share a maximally entangled state. As we will review and explore in this thesis, quantum entanglement is pivotal in the theory of quantum communication and the characterisation of channel capacities, thanks to its role in quantum teleportation.

These quantum communication protocols admit no analogy in the classical world, marking a clear departure from classical theory and present novel challenges for theoretical and experimental research. Furthermore, they are necessary to facilitate *distributed quantum information processing* on a large-scale. Much like classical computing, there are many instances in which quantum computing will require the co-operation of many (potentially remote) devices. The true exploitation of quantum information processing will thus demand distributed methods, permitting the execution of quantum algorithms on remote quantum processors [18] and supporting applications such as a fundamentally precise and stable world clock [19], quantum-enhanced telescope arrays [20] and much more.

The core trio of QKD, entanglement distribution, quantum state transfer embody the primary challenges of quantum communications. In tandem, they contribute to the envisioned utility of a greater target: A *quantum internet* [21, 22, 23, 24, 25]. This refers to a large-scale quantum communication network which enables users across the globe to securely communicate, distribute entanglement, perform distributed quantum computing and much more. The quantum internet *does not* stand as a replacement to its classical counterpart, but rather as a partner. In a future world vision, classical and quantum internets will forge a symbiotic relationship in order to enable classical and quantum communication on a grand, worldwide scale; an aspiration that sits as the Holy Grail of quantum communication research.

1.4 Quantum Networking

Since the inception of the BB84 protocol [15] quantum teleportation [26, 27] and their many descendants, the primary theoretical problem of quantum communication has been studied in the *point-to-point* setting. Research has majorly focussed on the ability for a pair of communicators (Alice and Bob) to exchange secret-keys, quantum states or entanglement over an environmental quantum channel which may be subject to attack by an eavesdropper (Eve). Clearly, the study of point-to-point quantum communications serves as a primitive for more complex settings related to quantum networking. Unless we can immediately, reliably and securely communicate over a quantum channel, we cannot do so across a quantum internet.

Fortunately, dramatic progress has been accomplished in understanding the ultimate operational limits of quantum communications in this setting. The year 2009 marked the beginning of the end of a critical and long-standing open question in quantum information/Shannon theory; Given a bosonic lossy quantum channel (the quantum mechanical description of an optical-fibre link, and the basis of free-space media) characterised by transmissivity η (which describes the fraction of photons that survive the link from transmission to detection), what is the *ultimate rate* (channel capacity) at which secret-key bits, entanglement bits or qubits can be transmitted? Indeed, Ref. [28] presented a lower-bound on the quantum channel capacity of bosonic lossy links, and other channel models. Six years later, Pirandola *et al.* corroborated this lower-bound with an equivalent upper-bound, resolving the maximum rate of quantum communications over optical-fibre [29]. This resulted in the Pirandola-Laurenza-Ottaviani-Banchi (PLOB) bound, which shows that the fundamental laws of nature place a limit on the quantum communication rate over a lossy bosonic channel at $-\log_2(1 - \eta)$ bits/channel-use. The PLOB bound would also later be extended to placing limits on the capacities of terrestrial free-space and satellite-based quantum communications [30, 31].

The ability to access the ultimate limits of quantum (or classical) protocols (in communications, sensing or otherwise) is invaluable. Understanding “where the ceiling is” within scientific research is vital, and presents a guiding light for theoretical and experimental investigators alike; if we do not know how nature limits our capabilities then it is remarkably difficult to know (1) if we are doing something well, (2) if we could be doing so much better, or (3) if we are doing something wrong! Within quantum communications, the PLOB bound serves this purpose.

Consequently, it is intuitive to want to extend these limits into the more complex domain of quantum networks. Now, the task of quantum communication is not simply defined over

a single quantum channel but over a general architecture of interconnected channels, each of which may have unique properties of loss and noise. Progress in this direction has been made, in which the point-to-point limits of quantum channels were extended to *end-to-end* limits of quantum communication networks [32]. From these results, some clear insight was uncovered: since quantum mechanics imposes fundamental limits on the attainment of high rates over long distance point-to-point channels, the infrastructure developed to facilitate quantum networking will *need* to take these limitations into account and use them to motivate high-rate architectures. Understanding the trade-offs between end-to-end performance, distance independence and cost efficiency will be crucial for a successful quantum internet.

1.5 Motivation and Thesis Structure

Many glaring gaps still stand in our understanding of quantum communication networks: how should quantum network architectures be designed in a way that provides our key promises from before? What architectural features are necessary to provide high rates over long-distances *while remaining cost effective*? How will quantum systems be exchanged and routed across quantum networks to guarantee efficacy and efficiency? What role will free-space and satellite-based quantum communications play in the future quantum internet? What resource demands will *multi-user* situations place on quantum networks? These questions are in their infancy. But as the practicality of quantum technologies continues to advance, it is essential that we fill these gaps so that our theoretical understanding stays a few steps ahead. The goal of this thesis is to address some of these key open questions.

This thesis is structured as follows. In Chapter 2 we introduce the reader to critical concepts in quantum information, communication and network theory. Information in this chapter should provide useful context and preliminary results pertaining to theory developed in later chapters.

In Chapter 3 we address the challenge of constructing quantum network architectures which are analytically treatable, while retaining realistic physical and topological freedom. Typically, the theoretical capabilities of quantum communication protocols are benchmarked in a network setting via simple linear repeater chains. Repeater chains offer some insight into the ability to extend quantum links over longer distances, but do little to capture more complex features of network connectivity or the utility of multi-path routing. Utilising ideas from quantum information theory, classical networks and graph theory we investigate ideal architectures based on the property of weak-regularity. Weakly-Regular Networks (WRNs) simultaneously (*i*) idealise network connectivity, (*ii*) provide sufficient

freedom to capture a broad class of spatial topologies and (iii) remain analytically treatable so that critical network properties can be rigorously studied. This results in a design with desirable qualities which can efficiently and effectively provide insight for realistic structures. We show that quantum WRNs employing multi-path routing admit remarkably accessible (and achievable) upper-bounds on the end-to-end capacity. This allows for a characterisation of the ideal performance of a fibre-based quantum internet with respect to essential properties such as maximum channel length and nodal density.

In Chapter 4 we investigate a recently developed node-splitting technique which introduces internal losses and noise into repeater devices. This allows us to present achievable end-to-end rates for noisy-repeater quantum networks, obtained by extending the coherent and reverse coherent information (single channel capacity lower bounds) into end-to-end capacity lower bounds. Through this general formalism we show how tight upper-bounds can also be derived by supplementing appropriate single-edge capacity bounds. As a result, we develop tools which provide tight performance bounds for quantum networks constituted of channels whose capacities are not exactly known, and reveal critical network properties which are necessary for high-rate quantum communications. This permits the investigation of pertinent classes of quantum networks with realistic technologies such as qubit amplitude damping networks and bosonic thermal-loss networks.

Chapter 5 combines recent advances in the theory of point-to-point free-space channel capacities and end-to-end network capacities in order to develop crucial tools for the study of hybrid, free-space quantum networks. We present a general formalism for studying the capacities of arbitrary, hybrid quantum networks, before focussing on the regime of atmospheric and space-based quantum channels. We then gather a class of modular quantum network architectures which offer a realistic and readily analysable framework for hybrid quantum networks. By considering a physically motivated, highly connected modular structure we are able to idealise network performance and derive channel conditions for which optimal performance is guaranteed. This allows us to reveal vital properties for which distance-independent rates are achieved, so that the end-to-end capacity has no dependence on the physical separation between users. Our analytical method elucidates key infrastructure demands for a future satellite-based global quantum internet, and for hybrid wired/wireless metropolitan quantum networks.

Having gathered considerable analytical insight from the previous chapters and methods, Chapter 6 moves into the domain of complex, random quantum networks. Indeed, random network theory is enormously useful for the study of realistic features and behaviours of communication networks, and can be used to practically benchmark routing protocols. Using realistic descriptions of quantum networks via random network models and practi-

cal end-to-end routing protocols, we reveal critical phenomena associated with large-scale, optical-fibre quantum networks. Our work reveals the weaknesses of applying single-path routing protocols within quantum networks, observing an inability to achieve reliable rates over long distances. Adapting novel algorithms for multi-path routing, we devise an efficient and practical multi-path routing algorithm capable of boosting performance while minimising costly quantum resources.

Finally, Chapter 7 takes this investigation a step further. End-to-end rates in a single-unicast setting have been used to benchmark performance and qualify critical regimes of network resources required to guarantee reliable rates. While single-unicast studies are insightful, they fail to capture the true demands placed on a quantum network when multiple users are involved. We investigate the implications of routing competition in a multiple-unicast scenario for performance and network resource requirements. We identify a more realistic picture of quantum networking, revealing achievable average end-to-end rates and capacities of realistic, multi-user quantum networks. These results are applied to random, optical-fibre based quantum network models to provide realistic benchmarks for the development of future quantum networks.

Chapter 2

Basics of Quantum Communication Networks

In this chapter we review some of the fundamentals of quantum information, communication and network theory which will be utilised throughout this thesis. These topics include an overview of discretely and continuously encoded quantum information, taking a moment to focus on the pivotal category of Gaussian quantum information. We then describe the role of quantum channels, which leads us into the field of quantum communications. While there, we introduce the concept of quantum channel capacities and their derivation through the tools of general adaptive quantum protocols, channel simulation and protocol stretching. With this knowledge in hand, we follow our review into the realm of quantum network theory and show how quantum channel capacities are extended into network capacities via routing protocols.

2.1 Quantum Systems, Information and Channels

2.1.1 Discrete and Continuous Variable Quantum Systems

The manner in which quantum information is encoded into quantum states is dependent upon the physical system at hand. Indeed, the types of quantum systems utilised for computation, communication, sensing and beyond can vary significantly due to the different properties native to each physical manifestation. In general terms, one can split this landscape into two main categories: discrete-variable (DV) and continuous-variable (CV) quantum systems.

DV quantum systems are those whose quantum states are defined within finite, $d \geq 2$ dimensional Hilbert spaces, such that information is encoded into *discrete* degrees of freedom. Such systems are commonplace in quantum computation and communication. Indeed, the

majority of theoretical and experimental developments in quantum computation are based upon the use of $d = 2$ dimensional quantum systems, giving rise to qubit states with a two-level basis (e.g. $\{|0\rangle, |1\rangle\}$). DV systems are also vital to quantum communications, where photonic states can be used to encode qubits (and beyond). For instance, one can model photonic DV systems by considering discrete bases of polarisation states and low energy Fock states (those with few photons). Such implementations are vital for DV-QKD strategies, such as the classic BB84 protocol.

A CV quantum system is that which has an infinite-dimensional Hilbert space described by observables with continuous eigenspectra [33]. That is, CV quantum information is encoded into continuous degrees of freedom, rather than a finite set of variables as in DV systems. CV quantum information theory admits a rich mathematical structure which has been developed over the past 30 years to forge new pathways for quantum technologies beyond the qubit. An example of a CV quantum system is a *bosonic mode*, corresponding to a quantised radiation mode of the electromagnetic field (a quantum harmonic oscillator). Continuous degrees of freedom emerge from the energy states of a bosonic modes along with its position and momentum quadratures.

CV quantum systems are particularly important for quantum communications. Indeed, a number of CV-QKD protocols are experimentally verified, enjoy advanced information-theoretic security proofs, practically operate at room temperatures and utilise cheap off-the-shelf components. Many mature CV-QKD protocols are based upon the use of coherent states, which are easy to generate and detect via heterodyne/homodyne detectors. The practical relevance of CV systems in quantum communications means that they play a prominent role throughout this thesis, and will be considered frequently.

2.1.2 Gaussian Quantum Information

Here we will review some basic theoretical notions, but refer the reader to Refs. [33, 34, 35, 36] for more comprehensive literature on the topic. Bosonic quantum systems correspond to a collection of quantised modes of a field [33]. A mode of the electromagnetic field corresponds to a quantum harmonic oscillator whose state can be described by the number of photons that occur within the mode, i.e. the energetic properties of the mode. An alternative scenario can be studied by looking at the quantised vibrational modes of quantum systems whose state can be described by the number of phonons that occur within the mode (can be thought of as quantised sound waves, rather than light waves [37, 38]).

In either case, the state basis built from the number of excitations of the quantised field is known as the Fock basis (or number basis), and is infinite-dimensional. From here, we

focus on the theory of electromagnetic bosonic modes. Given a collection of N bosonic modes, each can be associated with an annihilation operator \hat{a}_i and a creation operator \hat{a}_i^\dagger . Collecting these operators into the ordered vector $\hat{\mathbf{b}} := [\hat{a}_1, \hat{a}_1^\dagger, \dots, \hat{a}_N, \hat{a}_N^\dagger]^T$, we then demand that $\hat{\mathbf{b}}$ must satisfy the following commutation relation,

$$[\hat{b}_i, \hat{b}_j] = \Omega_{ij}, \quad \forall i, j \in \{1, \dots, 2N\}, \quad (2.1)$$

where we define Ω_{ij} as the standard N mode symplectic matrix

$$\mathbf{\Omega} := \bigoplus_{i=1}^N \boldsymbol{\omega} = \text{diag}(\boldsymbol{\omega}, \dots, \boldsymbol{\omega}), \quad \boldsymbol{\omega} := \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (2.2)$$

Here, the \oplus operator implies a direct matrix sum. We can also define a number operator $\hat{n} := \hat{a}^\dagger \hat{a}$, which allows us to define the Fock states themselves. Indeed, the Fock basis is defined as the infinite set of eigenstates of the number operator $\hat{n} |n\rangle = n |n\rangle$ such that $|n\rangle$ is the energy eigenstate of a quantum harmonic oscillator with n photons. A creation operator has the effect of exciting a particle within a mode

$$\hat{a}^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle, \quad (2.3)$$

and an annihilation operator removes a particle from a mode

$$\hat{a} |n\rangle = \sqrt{n} |n-1\rangle. \quad (2.4)$$

It is also important to introduce the quadrature field operators, which are derived from the bosonic field operators. Decomposing the bosonic field operators we can write,

$$\hat{q}_j := \frac{1}{\sqrt{2\kappa}} (\hat{a}_j + \hat{a}_j^\dagger), \quad \hat{p}_j := \frac{i}{\sqrt{2\kappa}} (\hat{a}_j^\dagger - \hat{a}_j), \quad (2.5)$$

where κ is a dimensionless constant that controls the variance of the vacuum noise of the system, and is typically set to $\kappa = \frac{1}{2}$ (corresponding to vacuum noise of 1) or $\kappa = 1$ (corresponding to a vacuum noise of $\frac{1}{2}$). Throughout this thesis, let us assume that $\kappa = 1$ (in alignment with important referenced material) and appropriately remark upon it if necessary for the reader. The quadrature field operators are Hermitian, represent dimensionless canonical observables of the system and behave like the position and momentum operators of the quantum harmonic oscillator. For an N mode system, we collect these operators into the $2N$ -element ordered vector $\hat{\mathbf{x}} := [\hat{q}_1, \hat{p}_1, \dots, \hat{q}_N, \hat{p}_N]^T$ which satisfies the commutation relation

$$[\hat{\mathbf{x}}, \hat{\mathbf{x}}^T] = i\mathbf{\Omega}. \quad (2.6)$$

The quadrature operators have continuous eigenspectra $\hat{q}|q\rangle = q|q\rangle$, $\hat{p}|p\rangle = p|p\rangle$ such that $p, q \in \mathbb{R}$, forming two eigensets $\{|q\rangle\}_{q \in \mathbb{R}}$ $\{|p\rangle\}_{p \in \mathbb{R}}$ which can be transformed between one another via a Fourier transform [33]. In general, for multimode states we can write $\hat{\mathbf{x}}^T |\mathbf{x}\rangle = \mathbf{x}^T |\mathbf{x}\rangle$ such that $\mathbf{x} \in \mathbb{R}^{2N}$.

Access to the field quadrature eigenvalues provides a continuous variable representation of a bosonic quantum system. Given an Hilbert space $\mathcal{H}^{\otimes N}$, a quantum state encapsulates all the information about the collective quantum system, characterised by a density operator ρ (a trace-one semi-definite positive matrix). Let us denote the complete state space of density operators as $\mathcal{D}(\mathcal{H})$, to which any ρ belongs. The state space of an infinite-dimensional quantum system is difficult to work with (directly due to its dimensionality). Utilising the field quadratures, it is possible to invoke an equivalent representation defined over the real symplectic space. Consider an N -mode system and let us define the Weyl operator,

$$D(\boldsymbol{\xi}) := \exp(i\hat{\mathbf{x}}^T \boldsymbol{\Omega} \boldsymbol{\xi}), \quad \boldsymbol{\xi} \in \mathbb{R}^{2N}. \quad (2.7)$$

Any density matrix admits a characteristic form, given by

$$\chi(\boldsymbol{\xi}) := \text{Tr}[\rho D(\boldsymbol{\xi})]. \quad (2.8)$$

Applying a Fourier transform to $\chi(\boldsymbol{\xi})$ results in the Wigner function of the state ρ , a quasi-probability distribution which defines the quantum state in real symplectic space

$$W(\mathbf{x}) := \int_{\mathbb{R}^{2N}} \frac{d^{2N} \boldsymbol{\xi}}{(2\pi)^{2N}} \exp(-\mathbf{x}^T \boldsymbol{\Omega} \boldsymbol{\xi}) \chi(\boldsymbol{\xi}). \quad (2.9)$$

where $\mathbf{x} \in \mathbb{R}^{2N}$ are the continuous variables which span the real symplectic space: phase space. Note that this is a quasi-probability distribution as it normalises to one, but is not necessarily always positive. Hence, any N -mode quantum state $\rho \in \mathcal{D}(\mathcal{H}^{\otimes N})$ belonging to an infinite dimensional Hilbert space can be completely defined within a finite, $2N$ -dimensional phase space, a dimensional reduction which proves to be remarkably useful.

Consequently, the concept of Gaussian information and quantum states becomes clear. Gaussian states are those which have a Gaussian Wigner function,

$$W(\mathbf{x}) = \frac{1}{(2\pi)^N \sqrt{\det(\boldsymbol{\sigma})}} \exp\left[-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}})^T \boldsymbol{\sigma}^{-1}(\mathbf{x} - \bar{\mathbf{x}})\right], \quad (2.10)$$

where $\bar{\mathbf{x}}$ and $\boldsymbol{\sigma}$ are the first and second statistical moments of the quantum state. The first moment $\bar{\mathbf{x}} \in \mathbb{R}^{2N}$ (also known as the displacement vector) is computed via

$$\bar{\mathbf{x}} := \text{Tr}[\hat{\mathbf{x}}\rho], \quad (2.11)$$

while the second moment $\sigma \in \mathbb{R}^{2N \times 2N}$ is called the *covariance matrix* of the quantum state such that each element is given by

$$\sigma_{jk} := \frac{1}{2} \text{Tr} [\{\hat{x}_j - x_j, \hat{x}_k - x_k\} \rho], \quad (2.12)$$

where $\{\hat{X}, \hat{Y}\} = \hat{X}\hat{Y} + \hat{Y}\hat{X}$ is the anticommutator. Covariance matrices are $2N \times 2N$, real, positive-definite, symmetric matrices which satisfy the uncertainty principle via $\sigma + i\Omega \geq 0$ [33].

Therefore a Gaussian quantum state can be completely characterised by means of its displacement operator \bar{x} and covariance matrix σ . They are infinite-dimensional, bosonic quantum states which admit a finite-dimensional representation and thus invite a plethora of mathematical simplifications for their study. Not only are they easy to deal with, but Gaussian states and transformations are ubiquitous in real world settings for quantum information processing and are of great practical importance.

2.1.3 Quantum Channels

Given a quantum system in a particular state ρ , a quantum channel describes the transformation of that state as a consequence of interaction with an environment, the action of quantum information processing, or any physical process that may be applied to the system. More precisely, a quantum channel $\mathcal{E} : \mathcal{D}(\mathcal{H}^{\otimes N}) \rightarrow \mathcal{D}(\mathcal{H}^{\otimes N})$ describes a mapping from one quantum state ρ to another valid state $\rho' = \mathcal{E}(\rho)$ within the same state space, i.e. quantum channel is a completely positive trace preserving (CPTP) map [39]. A quantum channel can be equivalently represented using a Stinespring dilation, such that the mapping corresponds to the application of a unitary interaction U upon a joint system of the input state ρ and a pure, environmental state ρ_E ,

$$\mathcal{E}(\rho) := \text{Tr}_E [U(\rho \otimes \rho_E)U^\dagger]. \quad (2.13)$$

There exists a wide variety of important quantum channels which are used to describe essential processes in quantum computation, communication and sensing. Furthermore, quantum channels are dramatically more general than classical channels, due to the fact that quantum states inhabit an infinitely larger state-space than states of classical information.

Quantum channels are fundamental mathematical objects within the study of quantum communications. Quantum systems can never be exchanged perfectly, and will always undergo some interaction with an environment. Furthermore, quantum communications requires the ability to exchange quantum systems resiliently over long-distances. This necessitates the use of *flying quanta*, a carrier of quantum information that can reliably

connected remote parties. To this end, photons are the most promising information carrier for communications, rendering bosonic quantum channels as a crucial category.

Just like bosonic quantum states, bosonic channels can be categorised according to their influence on Gaussianity: a Gaussian quantum channel is that which preserves the Gaussianity of its input state. Otherwise, the channel is non-Gaussian. A Gaussian channel can be formally defined by a Gaussian Stinespring dilation which is a version of Eq. (2.13) with additional constraints. Firstly, let us remark that a Gaussian unitary U_g which transforms a quantum state from $\rho \rightarrow U\rho U^\dagger$, is a unitary transformation whose action can be completely characterised in phase space by means of a displacement vector \mathbf{d} and a symplectic matrix \mathbf{S} that satisfies the identity

$$\mathbf{S}\mathbf{\Omega}\mathbf{S}^T = \mathbf{\Omega}. \quad (2.14)$$

That is, the transformation of a quantum state ρ with first moments $(\bar{\mathbf{x}}, \boldsymbol{\sigma})$ can be described by

$$\bar{\mathbf{x}} \rightarrow \mathbf{S}\bar{\mathbf{x}} + \mathbf{d}, \quad \boldsymbol{\sigma} \rightarrow \mathbf{S}\boldsymbol{\sigma}\mathbf{S}^T. \quad (2.15)$$

It follows that a Gaussian channel \mathcal{E}_g admits a Stinespring dilation in which both the unitary interaction, and the environment state are Gaussian,

$$\mathcal{E}_g(\rho) := \text{Tr}_E \left[U_g(\rho \otimes \rho_{gE})U_g^\dagger \right]. \quad (2.16)$$

One finds that the action of a Gaussian channel on a Gaussian quantum state can be simplified in terms of transformations on its first and second moments,

$$\bar{\mathbf{x}} \rightarrow \mathbf{T}\bar{\mathbf{x}} + \mathbf{d}, \quad \boldsymbol{\sigma} \rightarrow \mathbf{T}\boldsymbol{\sigma}\mathbf{T}^T + \mathbf{N}. \quad (2.17)$$

such that $\mathbf{d} \in \mathbb{R}^{2N}$, $\mathbf{T}, \mathbf{N} \in \mathbb{R}^{2N \times 2N}$, and they satisfy $\mathbf{N} + i\mathbf{\Omega} - i\mathbf{T}\mathbf{\Omega}\mathbf{T}^T \geq 0$. When \mathbf{N} is a zero matrix, and \mathbf{T} is symplectic, then the channel is equal to a Gaussian unitary.

The prevalence of Gaussian states and operations in practice means that Gaussian channels are the most important category of bosonic channel models. Of note, the class of single-mode *phase-insensitive* channels serve critical roles in the characterisations of bosonic quantum communication models. Gaussian phase-insensitive channels are typified by the transformation matrices,

$$\mathbf{T} := \sqrt{\eta}\mathbb{I}, \quad \mathbf{N} := \nu\mathbb{I}, \quad (2.18)$$

where \mathbb{I} is the 2×2 identity matrix and the parameters η and ν are the transmissivity and the induced noise respectively. Their values give rise to a prominent taxonomy of channels:

- A pure-lossy channel \mathcal{E}_η has transmissivity $\eta \in (0, 1)$ and induced noise $\nu = (1 - \eta)$ which transforms input quadratures $\hat{\mathbf{x}} = (\hat{q}, \hat{p})^T$ according to

$$\hat{\mathbf{x}} \mapsto \sqrt{\eta}\hat{\mathbf{x}} + \sqrt{1 - \eta}\hat{\mathbf{x}}_{\text{env}}, \quad (2.19)$$

where the environment is in a vacuum state describing the interaction of bosonic mode with a zero-temperature bath.

- A thermal-loss channel $\mathcal{E}_{\eta, \bar{n}}$ has transmissivity $\eta \in (0, 1)$. This channel can be described by the action of a beam-splitter of transmissivity η which mixes the input mode with an environmental thermal mode with mean photon number

$$\bar{n}_{\text{env}} := \bar{n}/(1 - \eta). \quad (2.20)$$

The thermal-loss channel effectively adds \bar{n} thermal photons to the output state. This transforms the input quadratures of a single-mode input Gaussian state according to Eq. (2.19).

- An amplifier channel has gain $\eta > 1$ corresponds to a two-mode squeezing operation applied to the input mode joint with an environmental mode in thermal or vacuum state.

Of these, the pure-loss and thermal-loss channels are the most important as they represent the underlying physical models for communication lines composed of optical-fibre or established through free-space. Throughout this thesis, we will frequently utilise lossy channels to study the efficacy of bosonic quantum communications in a variety of different contexts.

2.2 Quantum Communications and Channel Capacities

2.2.1 General Adaptive Protocols and Channel Capacities

As discussed, there are many different flavours of quantum communications. For the purposes of studying channel capacities, it is incredibly useful to possess an overarching description of a quantum communication protocol between two parties (Alice and Bob) connected by a generic quantum channel \mathcal{E} . The most general protocol we can consider is a *general adaptive protocol*, within which users can transfer quantum systems and are assisted by adaptive local operations and two-way classical communications (LOCCs). Adaptive operations can be applied to Alice and Bob's local registers of quantum systems, denoted by $\mathbf{a} := \{a_1, a_2, \dots, a_{m_a}\}$ and $\mathbf{b} := \{b_1, b_2, \dots, b_{m_b}\}$, respectively, where each a_i, b_j label individual quantum systems for $i \in \{1, \dots, m_a\}$ and $j \in \{1, \dots, m_b\}$.

The generalised strategy proceeds as follows: Alice and Bob begin with their local quantum registers, and by means of an initialisation LOCC Λ_0 they prepare an initial, separable quantum state ρ_{ab}^0 . Alice then transmits a quantum system a_1 through \mathcal{E} , along

which it is transformed according to the channel description. Bob then adds the received quantum system to his register, and both parties apply another round of LOCCs, Λ_1 . The initial state is therefore transformed into a new, modified state ρ_{ab}^1 . This procedure is repeated, in which quantum systems exchanges are iteratively interleaved with rounds of adaptive LOCCs, Λ_i . Over the course of n transmissions, the initial state is transformed into a final state of the protocol, ρ_{ab}^n . The protocol itself can then be summarised using the collection of LOCCs, $\mathcal{A} := \{\Lambda_0, \Lambda_1, \dots, \Lambda_n\}$.

Ultimately, the objective of quantum communication is to transform the initially separable state ρ_{ab}^0 into a global *target state* ρ_{ab}^* . The rate of a general adaptive protocol is given by R_n^ε if after n uses of the channel, the final shared state ρ_{ab}^n is ε close to the target state ρ_{ab}^* with nR_n^ε bits. That is, the states satisfy $\|\rho_{ab}^n - \rho_{ab}^*\| \leq \varepsilon$ where $\|\cdot\|$ denotes the trace norm¹. Consequently, we can define the capacity of a quantum channel as the rate that can be achieved when using the optimal sequence of LOCCs \mathcal{A} in the asymptotic limit of channel uses $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$,

$$\mathcal{C}(\mathcal{E}) := \sup_{\mathcal{A}} \lim_{n, \varepsilon} R_n^\varepsilon. \quad (2.21)$$

The quantity $\mathcal{C}(\mathcal{E})$ refers to the generic two-way assisted quantum and private capacity of \mathcal{E} . By specifying to particular target state of the protocol we can define more specific capacities. If the target state ρ_{ab}^* is:

- A maximally entangled state, the communication task is *entanglement distribution* and $\mathcal{C} \rightarrow D_2$ is the two-way entanglement distribution capacity (measure in entanglement-bits, i.e. ebits).
- A private state, the communication task is *quantum key distribution*, and $\mathcal{C} \rightarrow K$ is the secret-key capacity (measured in secret-key bits). Note, that this is equivalent to the two-way private capacity P_2 which is the maximum rate at which classical messages can be securely communicated.
- An arbitrary quantum state, the communication task is *quantum state transfer* and $\mathcal{C} \rightarrow Q_2$ is the two-way quantum capacity (measured in qubits). Since an ebit can be used to teleport a qubit, it follows that $Q_2 = D_2$.

¹The trace norm is used as a convenient measure of distinguishability between density matrices. It has many useful properties that can be utilised within information theoretic arguments [39] and is crucial for the concept of composable security [17]. Later, it permits the inequality in Eq. (2.34) which is pivotal in bounding quantum channel capacities.

A hierarchy between these quantities therefore emerges, such that the two-way quantum and entanglement distribution capacity lower bound the secret-key capacity (and therefore the generic two-way capacity), $Q_2 = D_2 \leq K = P_2 = \mathcal{C}$. General adaptive protocols supply an overarching schematic for quantum communication protocols and prove to be remarkably useful in the context of quantum channel (and eventually, network) capacities.

2.2.2 Channel Simulation and Teleportation Stretching

Quantum teleportation [27, 40, 41] is a fundamental quantum protocol that allows users (Alice, Bob) to “teleport” a state from one user to another by means of an entangled resource state and LOCC. Each user in the protocol possess one subsystem of a shared, two-party, maximally entangled state Φ_{AB} , while the sender also prepares some state $\rho_{A'}$ that they wish to teleport. Alice performs a Bell measurement on her subsystems A, A' with a resulting variable k which Alice then classically communicates to Bob. By means of the shared correlations of the resource state, this then informs Bob as to what unitary operation V_k^{-1} he should apply to his subsystem to reconstruct $\rho_{A'}$. See Fig. 2.1(a) for a visual depiction of this protocol. In this sense, the optimal teleportation protocol (where $\rho_{A'}$ is reconstructed at B with maximum fidelity) aims to “simulate” an identity channel.

It is possible to generalise this description, as shown in [29]. Firstly, let us replace the Bell detection and unitary correction with an arbitrary LOCC, \mathcal{T} , composed of quantum operations $\mathbb{A}_k, \mathbb{B}_k$ applied by Alice and Bob on their quantum systems A, A' and B respectively, corresponding to Alice’s classical variable k . Secondly, let us replace the maximally entangled resource state with an arbitrary quantum state σ . Averaging over the classical outcome k then \mathcal{T} is trace-preserving and we can write [29],

$$\mathcal{E}(\rho) = \mathcal{T}(\rho \otimes \sigma). \quad (2.22)$$

A channel \mathcal{E} is “ σ -stretchable” if it can be simulated using a resource state σ and an LOCC \mathcal{T} (depicted in Fig. 2.1(b)). For arbitrary quantum channels, it is not immediately clear what form of resource state is necessary for its simulation.

A quantum channel is defined as *teleportation covariant* if there exist some collection of unitary operators $\{U_\alpha, V_\alpha\}_\alpha$ such that,

$$\mathcal{E}(U_\alpha \rho U_\alpha^\dagger) = V_\alpha \mathcal{E}(\rho) V_\alpha^\dagger. \quad (2.23)$$

Channels that are teleportation covariant promise that there always exists a unitary operation V_α that allows one to undo unitary transformations applied to a state pre-transmission, and therefore construct a reliable teleportation protocol. For Pauli channels $U_\alpha = V_\alpha$ is

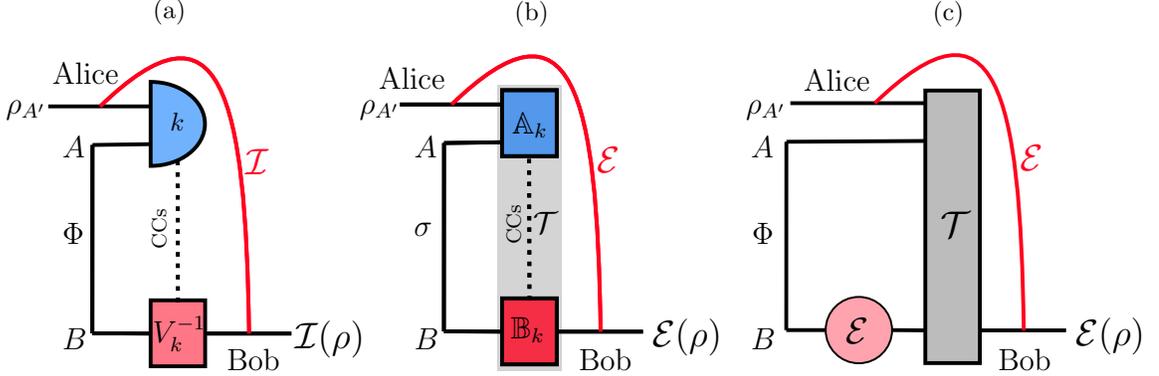


Figure 2.1: (a) A teleportation protocol aims to simulate an identity channel using a maximally entangled state as the shared resource between Alice and Bob. (b) The protocol can be generalised to an arbitrary LOCC \mathcal{T} in which the users share an arbitrary resource state σ , and Alice performs a general quantum operation \mathbb{A}_k on her systems, communicates the classical outcome k to Bob who then applies his own quantum operation \mathbb{B}_k on his systems, leading to a channel simulation $\mathcal{E}(\rho) = \mathcal{T}(\rho \otimes \sigma)$. (c) Teleportation covariant channels are simulated using their Choi matrix $\rho_{\mathcal{E}} = \mathcal{I} \otimes \mathcal{E}(\Phi_{AB})$.

simply the Pauli operators for $d = 2$ and more generally the Weyl-Heisenberg operator basis for finite dimensions $d > 2$. In infinite dimensions all Gaussian channels are teleportation covariant by means of displacement operators $D(\beta)$ [41]. Crucially, teleportation covariant channels can be optimally simulated via their Choi matrices, and are known as Choi-stretchable. The Choi matrix of \mathcal{E} is the result of passing the sub-systems B through the quantum channel, while keeping A preserved [39],

$$\rho_{\mathcal{E}} := \mathcal{I} \otimes \mathcal{E}(\Phi_{AB}), \quad (2.24)$$

where \mathcal{I} denotes the identity channel. Consequently, teleportation covariant channels admit the simulations [29],

$$\mathcal{E}(\rho) = \mathcal{T}(\rho \otimes \rho_{\mathcal{E}}), \quad (2.25)$$

as depicted in Fig. 2.1(c).

It is possible to apply channel simulation to general, two-way adaptive communication protocols in order to simplify the computation of fundamental upper bounds on channel capacities. A two-way adaptive protocol engages in n -rounds of quantum communications followed by two-way classical communication and adaptive LOCC operations. The goal of said protocol is bring an initial, shared, separable state ρ_{ab}^0 (between the registers of two users $\{\mathbf{a}, \mathbf{b}\}$) ε close to some target state ρ_{ab}^* over n uses of the protocol. Let us apply

technique of teleportation stretching in which each use of a channel \mathcal{E} in each adaptive LOCC block (Λ_i) is replaced with a channel simulation using some program state θ . Since the channel is consistent throughout the protocol, the program state can be stretched outside the adaptive LOCC blocks, and the teleportation protocols can be inherited at each stage. By then averaging over all local measurements to produce a trace-preserving LOCC $\bar{\Lambda}(\cdot)$, the n^{th} output state can then be expressed as [29],

$$\rho_{ab}^n = \bar{\Lambda}(\theta^{\otimes n}). \quad (2.26)$$

As discussed in the previous section, we can define the channel capacity as the optimal communications rate R_ε^n in the limit of large n and small ε , optimised over all possible adaptive key generation protocols \mathcal{A} ,

$$\mathcal{C}(\mathcal{E}) := \sup_{\mathcal{A}} \lim_{n, \varepsilon} R_\varepsilon^n. \quad (2.27)$$

We may investigate this quantity by considering the distance between the n^{th} output state of some adaptive protocol and the target state,

$$\|\rho_{ab}^n - \rho_{ab}^*\| \leq \varepsilon. \quad (2.28)$$

The trace distance can be very difficult to work with, so we may alternatively deploy the relative entropy of entanglement (REE) [42, 43, 44]. The REE is based on the quantum relative entropy (QRE), a kind of distance measure between two quantum states where

$$S(\rho\|\sigma) := \text{Tr}[\rho(\log \rho - \log \sigma)], \quad (2.29)$$

such that $S(\rho\|\sigma) \in [0, +\infty)$. Due to its asymmetry and the fact that it is infinite on pure states, it is not a true metric. However, the QRE is an important distinguishability measure between quantum states which provides access to important entropic quantities and inequalities. Minimising the relative entropy with respect to the set of all separable quantum states \mathcal{D}_{Sep} results in the REE,

$$E_R(\rho) := \min_{\sigma \in \mathcal{D}_{\text{Sep}}} S(\rho\|\sigma). \quad (2.30)$$

Interestingly, the REE is sub-additive and in general

$$E_R(\rho \otimes \sigma) \leq E_R(\rho) + E_R(\sigma). \quad (2.31)$$

This lets us define an n -shot REE,

$$E_R^n(\rho) := \frac{1}{n} \min_{\sigma \in \mathcal{D}_{\text{Sep}}} S(\rho^{\otimes n}\|\sigma) \leq E_R(\rho), \quad (2.32)$$

and by taking the limit of $n \rightarrow \infty$ we arrive at the regularised REE [45],

$$E_R^\infty(\rho) := \lim_{n \rightarrow \infty} E_R^n(\rho). \quad (2.33)$$

For d -dimensional states ρ and σ for which $\|\rho - \sigma\| \leq \varepsilon$, one can use the Fannes-type inequality and the REE to write [46],

$$|E_R(\rho) - E_R(\sigma)| \leq 4\varepsilon \log_2 d + 2H_2(\varepsilon), \quad (2.34)$$

where H_2 is the binary Shannon entropy,

$$H_2(p) := -p \log_2 p - (1-p) \log_2 (1-p). \quad (2.35)$$

Now exploiting the monotonicity of the REE under trace-preserving LOCC operations,

$$E_R(\rho_{\mathbf{ab}}^n) \leq E_R(\bar{\Lambda}(\theta^{\otimes n})) \leq E_R(\theta^{\otimes n}). \quad (2.36)$$

Using this bound on the REE, we can consequently place a tight bound on the channel capacity, $\mathcal{C}(\mathcal{E})$. Here we summarise elements of a proof from Ref. [29] that focus on finite dimensional systems with the knowledge that they are readily extended to infinite-dimensional systems. For more information on the complete proof, its extension to bosonic systems and more details on shield systems, please see Supplementary Note 3 of Ref. [29].

Supplementing $\rho_{\mathbf{ab}}^*$ and $\rho_{\mathbf{ab}}^n$ into Eq. (2.34) we can write [46],

$$E_R(\rho_{\mathbf{ab}}^*) \leq E_R(\rho_{\mathbf{ab}}^n) + 4\varepsilon \log_2 d_{\mathbf{ab}} + 2H_2(\varepsilon), \quad (2.37)$$

where $d_{\mathbf{ab}}$ is the total finite dimension of Alice and Bob's registers. In the context of distributing secret keys, $\rho_{\mathbf{ab}}^*$ is a private state with dimension $d_{\mathbf{ab}}$ that can be split into two key systems totalling a dimension d_K^2 , and a shield system d_S which describes the extra systems required to shield the key. Hence $d_{\mathbf{ab}} = d_K^2 d_S$. Following [45], the logarithm of the dimension of the key system defines the rate, so that

$$R_n^\varepsilon := \log_2 d_K \leq K. \quad (2.38)$$

As a consequence, we can connect R_n^ε with Eq. (2.37) to express the bound [29],

$$R_n^\varepsilon \leq \frac{E_R(\rho_{\mathbf{ab}}^n) + 4\varepsilon \log_2 d_{\mathbf{ab}} + 2H_2(\varepsilon)}{n}. \quad (2.39)$$

Considering a quantity $\alpha \geq 2$ which is sufficiently large so that $\alpha n R_n^\varepsilon \geq \log_2 d_{\mathbf{ab}}$, then the previous expression can be simplified,

$$R_n^\varepsilon \leq \frac{E_R(\rho_{\mathbf{ab}}^n) + 2H_2(\varepsilon)}{n(1 - 4\varepsilon\alpha)}. \quad (2.40)$$

Taking the limit of large n , the weak converse limit ($\varepsilon \rightarrow 0$) and an optimization over all general adaptive protocols we finally arrive at the upper bound on the key generation capacity [29],

$$\mathcal{C}(\mathcal{E}) = K(\mathcal{E}) \leq \sup_{\mathcal{A}} \lim_n \frac{1}{n} E_R(\rho_{\mathbf{ab}}^n) \leq E_R^\infty(\theta), \quad (2.41)$$

where we have used the regularised REE from Eq. (2.33). Therefore the fundamental upper bound for a channel capacity of \mathcal{E} for such a protocol can then be hugely simplified provided that the resource state for said channel is known. For teleportation-covariant channels, the optimal resource state is its Choi matrix $\rho_{\mathcal{E}}$, hence

$$\mathcal{C}(\mathcal{E}) \leq E_R^\infty(\rho_{\mathcal{E}}) \leq E_R^n(\rho_{\mathcal{E}}) \leq E_R(\rho_{\mathcal{E}}), \quad (2.42)$$

using the subadditivity of the REE.

2.2.3 Channel Capacity Lower-Bounds

It is always possible to express lower-bounds on the two-way assisted quantum and private capacity of a quantum channel using the coherent information (CI) [47] and the reverse coherent information (RCI) [28]. The CI and RCI of the channel read,

$$I_C(\mathcal{E}) := S[\text{Tr}_A(\rho_{\mathcal{E}})] - S(\rho_{\mathcal{E}}), \quad (2.43)$$

$$I_{RC}(\mathcal{E}) := S[\text{Tr}_B(\rho_{\mathcal{E}})] - S(\rho_{\mathcal{E}}), \quad (2.44)$$

where $S(\sigma) := \text{Tr}[\sigma \log(\sigma)]$ is the von Neumann entropy. These are valuable quantum information theoretic quantities which suggest achievable rates for forward (CI) and backward (RCI) one-way entanglement distillation. Indeed, for any quantum channel \mathcal{E} , the generic two-way assisted capacity $\mathcal{C}(\mathcal{E})$ can always be lower-bounded by the hashing inequality [47],

$$\mathcal{C}(\mathcal{E}) \geq I(\mathcal{E}) := \max\{I_C(\mathcal{E}), I_{RC}(\mathcal{E})\}, \quad (2.45)$$

where the CI lower bound was discovered in Ref. [47] and extended to the RCI in [28], and we have introduced the quantity $I(\mathcal{E})$ which implicitly maximises over them. This lower-bound is true for any quantum channel.

2.2.4 The Case of Bosonic Systems

For bosonic quantum systems the Choi matrix is energy unbounded, since the maximally entangled state is an infinitely squeezed Two-Mode Squeezed Vacuum (TMSV) state. Thus,

additional care must be taken when computing these quantities. Let us denote an energy constrained TMSV state

$$\Phi_{AB}^\mu, \text{ where } \mu := \bar{n} + 1/2, \quad (2.46)$$

and \bar{n} is the mean photon number per mode. Then the maximally entangled state takes the form $\Phi_{AB} := \lim_{\mu \rightarrow \infty} \Phi_{AB}^\mu$. As a result, we can define an asymptotic Choi matrix as the sequence of finite-energy Choi approximation in the limit of infinite squeezing,

$$\rho_{\mathcal{E}} = \lim_{\mu \rightarrow \infty} \rho_{\mathcal{E}}^\mu = \lim_{\mu \rightarrow \infty} \mathcal{I} \otimes \mathcal{E}(\Phi_{AB}^\mu), \quad (2.47)$$

where $\rho_{\mathcal{E}}^\mu := \mathcal{I} \otimes \mathcal{E}(\Phi_\mu)$ is a finite energy quasi-Choi matrix. This treatment must be considered for any functional f of asymptotic Choi matrices, such that it must be computed as the limit $f(\rho_{\mathcal{E}}) = \lim_{\mu \rightarrow \infty} f(\rho_{\mathcal{E}}^\mu)$. This is true for the CI and RCI, and the REE. It can be shown that both hashing inequality in Eq. (2.45) and the teleportation stretching upper-bound in Eq. (2.42) extend to general bosonic systems with constrained energy, and can extend to bosonic Gaussian channels in the limit of infinite energy (see Supplementary Notes 2, 4 of Ref. [29]).

2.2.5 Capacities of Bosonic-Lossy Channels

As discussed, bosonic lossy channels form the most important channel model for quantum communications over optical-fibre and free-space. Here, we briefly summarise the precise capacity statements that can be gleaned from Section 2.2.2. The bosonic pure-loss channel is both teleportation-covariant and distillable, meaning that its upper and lower-bounds coincide,

$$\mathcal{C}(\mathcal{E}_\eta) = I(\mathcal{E}_\eta) = E_R(\mathcal{E}_\eta), \quad (2.48)$$

so that computing the REE of the Choi state of a bosonic pure-loss channel \mathcal{E}_η of transmissivity η reveals its exact capacity. This quantity is known as the Pirandola-Laurenza-Ottaviani-Banchi (PLOB) bound [29], and describes the absolute maximum rate that two parties may transfer quantum states, distribute entanglement, or establish secret-keys over a bosonic lossy channel, equal to

$$\mathcal{C}(\mathcal{E}_\eta) = -\log_2(1 - \eta). \quad (2.49)$$

When considering optical-fibre, assuming a loss rate of γ dB/km (where $\gamma \approx 0.2$ is a typical value), then we can compute the transmissivity with respect to channel length, $\eta(d) = 10^{-\gamma d/10}$. One can then easily directly compute the PLOB bounds with respect to fibre length.

For teleportation-covariant channels, even when they are not distillable it is possible to write an upper-bound on their capacity using the REE of their Choi matrix. Bosonic thermal-loss channels fall into this category of teleportation-covariant which are not distillable. Hence the exact capacity of bosonic thermal-loss channels is not known but can instead be tightly bounded,

$$\mathcal{T}_{\eta, \bar{n}}^l \leq \mathcal{C}(\mathcal{E}_{\eta, \bar{n}}) \leq \mathcal{T}_{\eta, \bar{n}}^u, \quad (2.50)$$

where we define the bounding functions [29],

$$\mathcal{T}_{\eta, \bar{n}}^l := -\log_2(1 - \eta) - h(\bar{n}_{\text{env}}), \quad (2.51)$$

$$\mathcal{T}_{\eta, \bar{n}}^u := \mathcal{T}_{\eta, \bar{n}}^l - \bar{n}_{\text{env}} \log_2(\eta), \quad (2.52)$$

and $h(x) := (x + 1) \log_2(x + 1) - x \log_2(x)$ is an entropic function. Through appropriate modelling of the transmissivity and thermal-noise sources, these quantities can be used to place tight performance bounds on point-to-point capacities over optical-fibre and free-space quantum communication channels.

2.3 Quantum Communication Networks

2.3.1 Networks: From Classical to Quantum

So far we have focussed on the setting of *point-to-point* quantum communications. That is, we have studied the scenario by which a pair of communicators (Alice and Bob) are connected by a quantum channel and wish to exchange quantum/classical information. The study of point-to-point quantum communications has been the primary regime of interest for many years, due to the fact that it forms a primitive for any other possible setting. Nonetheless, the realistic utility of quantum communications on a global scale demands that we go beyond the point-to-point regime and understand the capabilities of *quantum networks*.

The theory of classical communication networks is incredibly well established, borrowing tools from information theory, graph theory, complex network theory and beyond [48, 49, 50, 51]. The sophistication of classical networks such as the internet means that the challenges which face network engineers and theorists today are many layers of abstraction above what we are facing in quantum network theory. The most blatant reason for this is that *the classical internet already exists*; many core mechanisms are well understood, well integrated and operate with state-of-the-art classical performance. Hence research and development aims to advance and build upon technological infrastructure (both hardware and software) that is already in place.

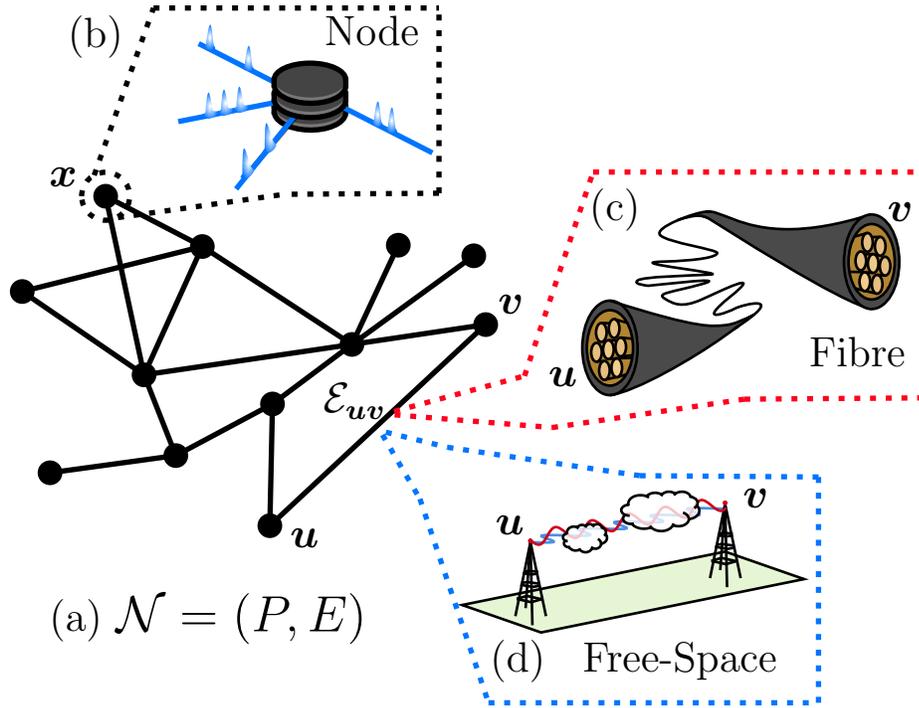


Figure 2.2: (a) A quantum communication network can be described as a finite, undirected, weighted graph $\mathcal{N} = (P, E)$. Each node represents (b) collection of local quantum registers which can be used to transmit, receive and operate on quantum systems. Network nodes are connected by quantum channels, typically described by (c) optical fibre or (d) free-space.

This is not the case for quantum network theory. With quantum networks still in their nascence, the current role of quantum network theory is to motivate future infrastructure and guide best practices through theoretical and experimental benchmarking and investigation; imbuing the field with a broader and more fundamental scope.

2.3.2 Mathematical Modelling of Quantum Communication Networks

General descriptions of a communication networks are based on *graphs*; mathematical structures which are used to describe a set of objects that are connected/related in some discrete manner. More specifically, a graph is composed of nodes (points, or vertices) which are connected to one another by means of edges (links, or arcs). Nodes can technically refer to anything, such as abstract mathematical quantities or concrete physical objects. Similarly, edges can be abstract or physical and have many different properties. They may possess a weight which assigns a specific attribute to the connection between some nodes, or are unweighted and identify a discrete notion of connection. Furthermore, edges can infer

direction, i.e. two nodes are connected in one direction but not another; or are undirected.

To enable an information-theoretic investigation of quantum communication networks, we require a mathematical model which: (1) Can remain agnostic to the precise quantum technologies being deployed, but (2) presents sufficient scope to be specified to particular technologies and protocols if required. Striking a careful balance between these requirements allows for the development of tools that can investigate quantum network capacities *and* specific performance benchmarks for particular protocols.

Consequently, it has been shown that quantum communication networks can be best described by means of finite, weighted, undirected graphs, denoted by

$$\mathcal{N} := (P, E), \quad \text{where } P := \{\mathbf{x}_i\}_i, \quad E := \{(\mathbf{x}, \mathbf{y})_i\}_i. \quad (2.53)$$

The tuple (P, E) collects the two key sets of quantities of network nodes P and network edges E where we use boldface notation to represent nodes $\mathbf{x} \in P$, and denote $(\mathbf{x}, \mathbf{y}) \in E$ as an undirected nodal pair which are connected by a communication channel. In this context, nodes represents either a user of the network (such as a potential communicator Alice or Bob) or a repeater/relay/server-station which is not user-controlled, but facilitates communication across the network nonetheless. The inner workings of quantum network nodes will be be technologically very complex, and may vary in range and utility. Clearly, the precise technologies will be different dependent on the types of quantum systems being deployed (e.g. DV or CV systems), and the communication task being undertaken (entanglement distribution, QKD, etc.). Throughout this thesis, we may focus on a number of different nodal descriptions but they can be generally summarised as an abstract object $\mathbf{x} \in P$ which possesses a number of local quantum registers, from which quantum systems which can be transmitted and collected. We also consider a general information-theoretic definition of a quantum repeater as a middle third-party helping the quantum communication between a sender and a receiver (therefore not connected by a direct link). In practice, there are many possible physical realisations, e.g., see Refs. [52, 53, 54] among others.

Meanwhile, edges within a quantum network $(\mathbf{x}, \mathbf{y}) \in E$ are used to represent a quantum channel $\mathcal{E}_{\mathbf{x}\mathbf{y}}$ that connected nodes \mathbf{x} and \mathbf{y} . The modelling of network edges as quantum channels presents a major generalisation from the classical setting, and broadens the scope of protocols that can be studied across a communication network. The precise description of a channel depends on the medium through which the nodes are connected and the type of quantum systems that are being exchanged. See Fig. 2.2 for a visual depiction of this quantum network description.

It is useful to introduce the graph theoretic property of network nodal degree, which describes the number of nodes to which a given node is connected. Defining the neighbourhood

of a node $\mathbf{x} \in P$ as

$$N_{\mathbf{x}} := \{\mathbf{y} \in P \mid \{\mathbf{x}, \mathbf{y}\} \in E\}, \quad (2.54)$$

then the degree of the node \mathbf{x} is equal to the cardinality of its neighbourhood

$$\deg(\mathbf{x}) := |N_{\mathbf{x}}|. \quad (2.55)$$

Hence, the node \mathbf{x} has exactly $\deg(\mathbf{x})$ neighbours. Similar logic can be followed to define an edge-neighbourhood of \mathbf{x} as all the edges which connect \mathbf{x} to its neighbours,

$$E_{\mathbf{x}} := \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in N_{\mathbf{x}}\}. \quad (2.56)$$

2.3.3 Link Layers

Importantly, we describe quantum networks as *weighted* undirected graphs. That is, each graphical node and each may be endowed with some abstract weight which relates to an important property of the network. For instance, there may exist a nodewise distribution of values $\{f_{\mathbf{x}}\}_{\mathbf{x} \in P}$ such that $f_{\mathbf{x}}$ describes a localised property of each node, e.g. detection efficiency, excess noise etc. Edgewise property distributions can also be constructed, $\{g_{\mathbf{xy}}\}_{(\mathbf{x}, \mathbf{y}) \in E}$ where $g_{\mathbf{xy}}$ is a potentially unique property to each quantum channel in the network.

Throughout the study of quantum network rates and capacities, the most important weight distribution across the network is given by the distribution of point-to-point rates across the architecture. In order to capture the effectiveness of quantum communications across a network topology it is vital to understand the rates at which each node can communicate with one another. For an arbitrary quantum network $\mathcal{N} = (P, E)$ we can define a rate distribution \mathcal{K} as the collection of point-to-point rates

$$\mathcal{K} := \{K_{\mathbf{xy}}\}_{(\mathbf{x}, \mathbf{y}) \in E}, \quad (2.57)$$

where $K_{\mathbf{xy}}$ denotes the rate at which a network edge is able to operate. The rate distribution of any network depends on the end-to-end task at hand (key distribution, entanglement distribution, quantum state transfer), the channel quality and the point-to-point protocols being used.

Crucially, the network rate distribution is completely dependent upon the physical descriptions of network nodes and channels, as well as the point-to-point communication protocols assumed across the architecture. Indeed, \mathcal{K} is a direct consequence of the technological assumptions made when modelling the quantum network. We refer to the global description of the point-to-point network model as its *link layer*, which can be investigated

in a number of different forms. For instance, the \mathcal{K} can analogously become a capacity distribution if we consider that each point-to-point connection operates at its capacity, $K_{xy} = C_{xy}$ and allow us to place performance bounds on quantum networks. Meanwhile, the assumption of achievable link rates allow for the investigation of near-term quantum networking capabilities.

2.3.4 Physical vs. Logical Flow in Quantum Networks

In the context of quantum networks, it is important to make a distinction between *physical flow* and *logical flow*. The logical flow of a quantum communication channel describes the direction in which entanglement, secret-keys, or quantum states are distributed from a node \mathbf{x} to node \mathbf{y} (or vice versa). The physical flow of quantum communication refers to the actual direction of quantum system exchange, i.e. if quantum systems are physically sent in the direction $\mathbf{x} \rightarrow \mathbf{y}$ or $\mathbf{y} \rightarrow \mathbf{x}$. In a quantum network, these concepts can be completely decoupled. This may be due to the fact that the communication task has a symmetric objective i.e. if Alice and Bob wish to share a secret-key, they do not care *who* initiates the exchange of quantum systems. However, it may also be thanks to quantum teleportation; it is always possible to “reverse” the logical direction of communication by means of a teleportation protocol between Alice and Bob.

The independence of physical and logical flow helps us to reliably describe a quantum network as an undirected graph. Any pair of connected network nodes can choose the physical direction in which they wish to exchange quantum systems and may always choose that which has the largest capacity. As a result, we never need to distinguish between forward or backward channels and represent each edge $(\mathbf{x}, \mathbf{y}) \in E$ by the best choice of quantum channel [32].

2.3.5 Communication across Quantum Networks

With the appropriate modelling in place, it is vital to understand *how effectively* quantum communications can be carried out across network architectures, i.e. what effective rates/capacities can be obtained across quantum networks. However, now that we have departed from the point-to-point setting there is no longer a single communication task to be considered. Instead, given a network of many nodes and many potential end-users, there are unique situations that can be considered:

- *Unicast*: During a single-unicast session, a network \mathcal{N} is exclusively used to perform communication between a sender and receiver (also known as an end-to-end protocol).

During multiple-unicast, a network \mathcal{N} is used to transmit unique messages between many dedicated pairs of sender and receiver nodes.

- *Multicast*: A network \mathcal{N} is used to transmit (unique or equivalent) messages between a sender and a pool of receiver nodes. During multiple-multicast, the network is used to transmit unique messages between M senders and M pools of receiver nodes, where the number of receivers in each pool can vary. In the scenario where a receiver pool contains *all* the network nodes, this becomes a broadcast.

In reality, a successful quantum network architecture will enable reliable and strong performance across all communication tasks (or across an appropriate subset for which it was designed). However, the complexity of investigating each type varies significantly. Developing accurate assessments of how many users can simultaneously multicast across a quantum network is much more complicated than that of a single end-user pair. Indeed, there are pieces missing in our theoretical knowledge of *how* to perform these assessments efficiently.

Yet, it is imperative that we start from the simplest communication task and grow our toolset from there. In this way, significant progress has been and is being made. Ultimately, the scenario which acts as a primitive for all others: unicast. Consider a quantum network $\mathcal{N} = (P, E)$ during single-unicast quantum communications, such that a single pair of end-users (Alice and Bob) wish to utilise a quantum network in order to distribute secret-keys, entanglement or transfer a precise quantum state. This stands as the simplest generalising step from point-to-point into end-to-end communications, where Alice and Bob now possess an entire network dedicated to enabling their round of communication.

Hence, the benchmarking of single-unicast quantum networking is vital. If an architecture is incapable of achieving reliable single-unicast rates, then it is doomed to achieve reliable rates for *any other communication protocol*. The work and results presented in this thesis aim to theoretically advance our understanding of unicast quantum networking, firstly from the perspective of single-unicast end-to-end capacities and rates, but then extending these investigations to the domain of multiple-unicast.

2.3.6 End-to-End Routes

Let us mathematically define the idea of an end-to-end route. The ability for information to be exchanged on a network relates to the ability for network users to utilise intermediary nodes to mediate communication. Consider an end-user pair Alice \mathbf{a} and Bob \mathbf{b} . For convenience, let us introduce the notation that collects an end-user pair into a single object,

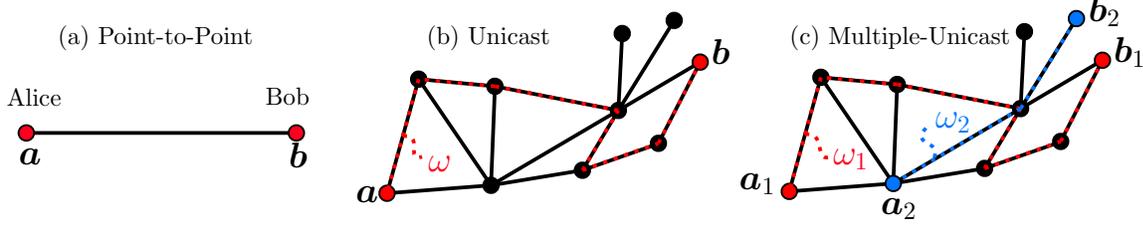


Figure 2.3: (a) Point-to-point quantum communications between an end-user pair of nodes Alice \mathbf{a} and Bob \mathbf{b} . (b) A single end-user pair engage in single-unicast quantum networking, utilising end-to-end routes such as ω to establish connections. (c) Multiple-unicast quantum networking concerns multiple end-user pairs (e.g. $\{\mathbf{a}_1, \mathbf{b}_1\}$ and $\{\mathbf{a}_2, \mathbf{b}_2\}$) who may wish to utilise the network simultaneously.

$$\mathbf{i} := \{\mathbf{a}, \mathbf{b}\}, \quad \mathbf{a} \in P, \quad \mathbf{b} \in P \setminus \{\mathbf{a}\}. \quad (2.58)$$

which will be useful throughout our investigations. For a pair \mathbf{i} to perform quantum communications, they must become connected by a common collection of nodes that have successfully exchanged quantum systems. The simplest way to do this is to identify an end-to-end route (or single-path) a collection of network edges which connect one end-user to the other,

$$\omega := \{(\mathbf{a}, \mathbf{x}_1), (\mathbf{x}_1, \mathbf{x}_2), \dots, (\mathbf{x}_N, \mathbf{b})\}. \quad (2.59)$$

We define an end-to-end route as free of cycles (without loss of generality). Any route ω refers to an associated collection of quantum channels

$$\omega \mapsto \{\mathcal{E}_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x}, \mathbf{y}) \in \omega} = \{\mathcal{E}_{\mathbf{a}\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_1\mathbf{x}_2}, \dots, \mathcal{E}_{\mathbf{x}_N\mathbf{b}}\}, \quad (2.60)$$

through which quantum systems must be exchanged to establish end-to-end quantum communications. The identification of end-to-end routes is a fundamental building block for successful networking in both classical and quantum architectures.

2.3.7 Protocols for End-to-End Routing of Quantum Information

The ability to perform quantum communications across a network depends upon the ability of end-users to construct end-to-end routes. In turn, this requires an effective *end-to-end routing protocol*. Routing protocols are responsible for dictating how information flow takes place across a network in such a way that connects network users. Routing strategies should be designed in an effort to optimise some important features of communication, such

as the end-to-end rate. Indeed, a single-unicast, end-to-end routing protocol can be deemed practical and useful if it can do the following:

1. Identify an end-to-end path (or paths) which maximise the rate between users.
2. Minimise the network resources needed to achieve said end-to-end rate.
3. Efficiently execute and facilitate routing without impeding the end-to-end rate.

In classical networks, these features are not so difficult to achieve, thanks to the robustness of classical information. Such networks can exploit techniques of network coding to satisfy each of these key points, while the fragility of quantum information essentially prohibits the same convenience.

There exist two main classes of end-to-end routing strategy: *single-path* and *multi-path* routing. Single-path routing describes a network protocol, \mathcal{P}_{sp} , in which quantum systems are exchanged from node-to-node throughout the network in a sequential manner. This forges a unique path of interactions through the network, and continues until quantum communication has been established between the end-users. This is the standard strategy used for classical networking and is extremely effective thanks to the robustness of classical information. The vulnerability of quantum networks to decoherence means that it is extremely valuable to explore more general multi-path protocols, \mathcal{P}_{mp} in order to enhance performance. Network nodes may exchange multiple quantum systems with many receiver nodes, repeating until communication is established between the end-users. End-users may explore a variety of end-to-end routes simultaneously in a way which enhances their performance and reliability.

A useful way to intuit end-to-end routing requires the introduction of an edgewise distribution of probabilities which we call the *forwarding probability distribution*, which we denote by

$$\mathcal{Q}_{\mathcal{P}} := \{q_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x},\mathbf{y}) \in E}, \quad 0 \leq q_{\mathbf{x}\mathbf{y}} \leq 1, \forall \mathbf{x}, \mathbf{y} \in P. \quad (2.61)$$

The forwarding probability distribution $\mathcal{Q}_{\mathcal{P}}$ collects the probability that a network edge $(\mathbf{x}, \mathbf{y}) \in E$ is utilised at any point throughout the end-to-end routing process. That is, given many executions of the edge $(\mathbf{x}, \mathbf{y}) \in E$, on average it will be utilised $q_{\mathbf{x}\mathbf{y}}$ times. It follows that $\mathcal{Q}_{\mathcal{P}}$ is conditioned upon the routing protocol, since different protocols will place different restrictions and demands on the use of network nodes and edges. For instance, using single-path routing, then \mathcal{P}_{sp} restricts the forwarding probability distribution from routing along multiple edges from the same node. More complicated routing protocols will invite more complicated relationships with $\mathcal{Q}_{\mathcal{P}}$.

With these ideas in mind, let us provide a generalised description of end-to-end routing protocols on quantum networks, starting from the concept of general adaptive protocols (as discussed in Section 2.2 the context of point-to-point quantum communications). The description we present is similar to those discussed in Ref. [32], to which we direct the reader for more information (or to compare).

Consider a quantum network $\mathcal{N} = (P, E)$ which contains a single pair of end-users (Alice and Bob) $\mathbf{i} := \{\mathbf{a}, \mathbf{b}\}$ who wish to distribute entanglement, secret-keys or perform quantum state transfer. The end-users will utilise a routing protocol \mathcal{P} , a strategy for exploring the network in such a way that facilitates communication. Prior to communication, we may consider that the network is initialised with a random, sub-optimal forwarding probability distribution $\mathcal{Q}_{\mathcal{P}}^1$ which can be adapted and optimised throughout the routing process. Now, communication proceeds: Alice prepares a quantum state within her set of local registers, which is generally multipartite. She transmits quantum systems to all of the nodes $\mathbf{a} \rightarrow N_{\mathbf{a}}$, where $N_{\mathbf{a}}$ is her nodal neighbourhood to which she is directly connected as defined in Eq. (2.54). Each transmission from \mathbf{a} occurs with probability $q_{\mathbf{a}\mathbf{x}} \in \mathcal{Q}_{\mathcal{P}}^1$ for $\mathbf{x} \in N_{\mathbf{a}}$. In general, the set of transmissions is followed by a round of LOCCs at every node (a network LOCC), in which classical information can be exchanged across the network to optimise the routing protocol, and nodes can perform conditional LOs on their local registers. Each of the receiving nodes $\mathbf{x} \in N_{\mathbf{a}}$ will then perform their own point-to-multipoint transmissions with each of their neighbours along any available channels that have not been used (each according to their own forwarding probabilities once again). Each set of transmissions is interleaved with network LOCCs, and the process of multipoint exchanges will continue throughout the network.

The use of adaptive network LOCCs can be equivalently considered as an optimisation of the forwarding probability distribution on the fly. Each round of system exchange between nodes is followed by a network LOCC, in which feedback from the transmission is used to improve future exchanges between all nodes on the network. The manifestation of this feedback is an update of the likelihood of using each edge, i.e. its forwarding probability. Hence, every transmission is followed by an update of the form $\mathcal{Q}_{\mathcal{P}}^1 \rightarrow \mathcal{Q}_{\mathcal{P}}^2 \rightarrow \dots$, and so on. Eventually, after some M steps, a final multipoint-to-point exchange will occur with Bob's node and his neighbourhood $N_{\mathbf{b}} \rightarrow \mathbf{b}$. At this point, communication has been established between the end-users who have successfully exchanged quantum systems from end-to-end.

This entire procedure corresponds to a single-round of end-to-end communication, in which quantum systems have been exchanged across network. Once this first round is complete, the network will have accessed a sequence of M forwarding probability distributions $\{\mathcal{Q}_{\mathcal{P}}^1, \dots, \mathcal{Q}_{\mathcal{P}}^M\}$ which have evolved during the round. Repeating this procedure n times, the

forwarding probability distribution can continue to evolve $\mathcal{Q}_{\mathcal{P}}^1 \rightarrow \dots \rightarrow \mathcal{Q}_{\mathcal{P}}^{nM}$ in which it is consistently being optimised, and forms an increasingly accurate representation of the optimal routing strategy $\mathcal{Q}_{\mathcal{P}} := \lim_n \mathcal{Q}_{\mathcal{P}}^n$. In this limit, the forwarding probability distribution should tend towards some stable, optimal state such that further attempts to improve it are redundant². The end-to-end network protocol is therefore completely characterised by the sequence of forwarding probability distributions, and the collection of LOCCs operated throughout communication.

It is important that we clarify what it means to optimise the forwarding probability distribution. End-users may have different performance measures that quantify the efficacy of end-to-end routing. Clearly the primary goal is to maximise the end-to-end rate of the protocol, just as in the point-to-point setting. However, quantum networking invites the consideration of other factors. End-users may wish to optimise other properties related to routing, and there may exist constraints on the routing protocol \mathcal{P} being deployed. For instance, the end-users may wish to simultaneously *minimise* the number of edges needed to maximise the end-to-end rate; or the routing protocol may be exclusively constrained to using single-paths. For this reason, the optimisation of $\mathcal{Q}_{\mathcal{P}}$ can vary dramatically, dependent on the nature of \mathcal{P} .

Describing end-to-end quantum communication protocols in the language of $\mathcal{Q}_{\mathcal{P}}$ proves to be remarkably useful when we wish to analyse the end-to-end rates and capacities of generic routing strategies, which we explore in the following section.

2.4 End-to-End Quantum Network Rates and Capacities

2.4.1 Network Cuts

An important graph-theoretic concept for quantifying network performance is that of *cuts* and *cut-sets*. Consider a network $\mathcal{N} = (P, E)$ with two remote end-users $\mathbf{a}, \mathbf{b} \in P$. Given an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, we can define a cut C as a bipartition of all network nodes P into two disjoint subsets of nodes $(P_{\mathbf{a}}, P_{\mathbf{b}})$ such that the end-users become completely disconnected, $\mathbf{a} \in P_{\mathbf{b}}$ and $\mathbf{b} \in P_{\mathbf{a}}$, where $P_{\mathbf{a}} \cap P_{\mathbf{b}} = \emptyset$. A cut C generates an associated cut-set; a collection of network edges \tilde{C} which enforce the partitioning when removed,

$$\tilde{C} := \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{a} \in P_{\mathbf{a}}, \mathbf{b} \in P_{\mathbf{b}}\}. \quad (2.62)$$

Under the action of a cut, a network is successfully partitioned

$$\mathcal{N} = (P, E) \xrightarrow{\text{Cut: } C} (P, E \setminus \tilde{C}) = (P_{\mathbf{a}} \cup P_{\mathbf{b}}, E \setminus \tilde{C}), \quad (2.63)$$

²It is important to note that there may be an abundance of optimal states such that the optimal end-to-end routing strategy is degenerate

so that there no longer exists a path between \mathbf{a} and \mathbf{b} . Network cuts play a key role in the derivation of end-to-end network rates and many network optimisation tasks can be reduced to an optimisation over all cuts with respect to single-edge/multi-edge properties.

It is useful to define a multi-edge rate quantity associated with a network cut and the single-edge rates of the edges collected within the cut-set. More precisely, given a network cut C and its associated cut-set \tilde{C} , we have

$$K_{\text{cut}}(C) := \sum_{\mathbf{x}, \mathbf{y} \in \tilde{C}} K_{\mathbf{x}\mathbf{y}}. \quad (2.64)$$

When the rate distribution consists of single-edge capacities, then $K_{\text{cut}}(C) \rightarrow \mathcal{C}_{\text{cut}}(C)$ becomes a multi-edge capacity quantity.

2.4.2 Single-Path Routing

Single-path routing is the simplest network communication method, which utilises point-to-point communications in a sequential manner. Quantum systems can be exchanged from node-to-node along this route, followed by LOCC operations after each transmission until eventually communication is established between the end-users. This kind of strategy is analogous to the use of a repeater-chain and network performance is determined by the strength of each link along an optimal end-to-end route.

Single-path routing is the principal mechanism for classical communications, fundamentally achieved by Dijkstra's algorithm (DA) [55, 56, 57, 58]. Generally, this is a greedy algorithm for single-path route optimisation according to a path defined *cost function*, which measures a property of end-to-end routes that one wishes to optimise, e.g. minimisation of path length or maximisation of bottleneck rate. The challenges which extends from the point-to-point rate limitations of quantum communications means that rate optimisation is the salient task. This is also known solving the *widest path problem*. For a network $\mathcal{N} = (P, E)$, DA locates the widest path efficiently in run-time $\mathcal{O}(|E| + |P| \log_2 |P|)$. Given a single-path protocol \mathcal{P}_{sp} and an end-user pair, the optimal end-to-end rate is given by

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) := \max_{\omega} \min_{(\mathbf{x}, \mathbf{y}) \in \omega} K_{\mathbf{x}\mathbf{y}}, \quad (2.65)$$

where the maximisation is performed over all possible end-to-end routes.

If we are investigating the exact single-path capacities of quantum networks, then we assume a rate distribution of $\mathcal{K} = \{\mathcal{C}_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x}, \mathbf{y}) \in E}$, such that $\mathcal{C}_{\mathbf{x}\mathbf{y}} = \mathcal{C}(\mathcal{E}_{\mathbf{x}\mathbf{y}})$ is the single-edge capacity of each channel in the network. At times in this thesis, we modify our notation to

better identify single-path network capacities according to

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) \leq \mathcal{C}(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) = \max_{\omega} \min_{(\mathbf{x}, \mathbf{y}) \in \omega} C_{\mathbf{x}\mathbf{y}}. \quad (2.66)$$

2.4.3 Flooding

A more powerful strategy is multi-path routing, which properly exploits the multitude of end-to-end routes available in a quantum network. In multi-path protocols, users may simultaneously utilise a number of unique routes $\{\omega_1, \omega_2, \dots, \omega_M\}$ in an effort to enhance their end-to-end rate. A user may exchange an initially multi-partite quantum state with a number of neighbouring receiver nodes, who may each then perform their own point-to-multi-point exchanges along its unused edges. The exchange of quantum systems can be interleaved with adaptive network LOCCs in order to distribute secret correlations and this process continues until a multi-point interaction is carried out with the end-user.

From the perspective of end-to-end rate performance, the optimal multi-path routing strategy operates in such a way that all channels in the network are used once per end-to-end transmission. This is known as a *flooding protocol*, \mathcal{P}_{fl} [50, 51, 32]; each node in the network performs quantum systems exchanges along all its available edges, resulting in non-overlapping point-to-multipoint exchanges between all network nodes. The ability to flood an entire network means that every possible end-to-end route between the end-users are fully explored, allowing them to achieve the optimal end-to-end rate. This greatly enhances end-to-end performance. Given the rate distribution $\mathcal{K} = \{K_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x}, \mathbf{y}) \in E}$ of a network, its flooding rate can be derived by solving the classical maximum-flow minimum-cut problem. It is found by locating the minimum-cut C_{min} which minimises the sum of the single-edge rates in a cut-set, over all cut-sets [32]. More precisely,

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{fl}}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} K_{\mathbf{x}\mathbf{y}}. \quad (2.67)$$

For general quantum networks with arbitrary capacity distributions and network structures, this problem requires a numerical treatment by solving the well known max-flow min-cut problem [59, 60, 61, 62] to find C_{min} .

If we are investigating the flooding *capacity* of a quantum network, then we assume a rate distribution of $\mathcal{K} = \{C_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x}, \mathbf{y}) \in E}$ as we did for single-path capacities. It is important to recognise that since flooding is the best possible routing strategy for maximising end-to-end performance, it follows that the flooding capacity is the ultimate limit of end-to-end quantum communication rates,

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}) \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) := \mathcal{C}(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{fl}}), \quad \forall \mathcal{P}. \quad (2.68)$$

Here, we have introduced $\mathcal{C}(\mathbf{i}, \mathcal{N})$ as the *protocol independent* end-to-end network capacity, equal to the flooding capacity which is a useful quantity throughout this thesis.

Flooding represents a remarkably useful paradigm for benchmarking optimal network performance, but it is not practical in many real world settings. Utilising an entire network to facilitate one end-user pair renders the network useless for any other set of users to communicate simultaneously. In reality, we want to manage resources in a way that enables the concurrent use of a network for many users.

2.4.4 Practical Routing Protocols

It is highly desirable to access a more general platform for the investigation of end-to-end network rates and capacities, broader than that of single-path routing and flooding. In large networks, the use of flooding is not practical, since this requires the use of potentially huge number of edges for a single-pair of uses, while single-path routing is typically unable to guarantee high rates. This motivates the need for broader classes of protocols, and the tools reviewed thus far are insufficient to capture the performance of more general multi-path protocols. Utilising the language of forwarding probability distributions, we can present a unified picture for the computation of end-to-end network rates and capacities for more general routing protocols.

Consider a network $\mathcal{N} = (P, E)$, an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, a rate distribution \mathcal{K} and a generic routing protocol \mathcal{P} . As discussed in Section 2.2, this routing protocol can be mapped onto a forwarding probability distribution $\mathcal{Q}_{\mathcal{P}}$ which infers the probability of any given edge being utilised in an end-to-end execution of \mathcal{P} . Consequently, it's straightforward to compute the end-to-end rate, and once more equates to solving the max-flow min-cut theorem [32]. Using $\mathcal{Q}_{\mathcal{P}}$, we can state that the execution of communication according to \mathcal{P} is equivalent to a flooding protocol in which every network edge operates at an effective rate given by $q_{\mathbf{x}\mathbf{y}}K_{\mathbf{x}\mathbf{y}}$. By modifying the network weights according to the probability of edge usage, the end-to-end rate between \mathbf{i} on the network \mathcal{N} using \mathcal{P} can be computed by

$$K(\mathbf{i}, \mathcal{N}|\mathcal{P}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} q_{\mathbf{x}\mathbf{y}}K_{\mathbf{x}\mathbf{y}}. \quad (2.69)$$

It now becomes intuitive to see how this result generalises the previous end-to-end rate quantities. For instance, one can easily utilise $\mathcal{Q}_{\mathcal{P}}$ to describe *deterministic* routing strategies where a specific subset of the network is utilised for end-to-end routing. A deterministic multi-path protocol can be designed which exclusively performs end-to-end communication along a selection of $M \geq 1$ paths which we collect into a route set,

$$\Omega_{\mathcal{P}} := \{\omega_1, \omega_2, \dots, \omega_M\}. \quad (2.70)$$

By unpacking this route set into all of its constituent edges, we may define an analogous routing edge set $E_{\mathcal{P}}$, given by

$$E_{\mathcal{P}} := \omega_1 \cup \dots \cup \omega_M = \bigcup_{i=1}^M \omega_i. \quad (2.71)$$

More precisely, a protocol \mathcal{P} which deterministically uses a finite set of end-to-end routes $\omega \in \Omega_{\mathcal{P}}$ generates a forwarding probability distribution of the form

$$\mathcal{Q}_{\mathcal{P}} = \{q_{\mathbf{x}\mathbf{y}}\}, \quad \text{such that } q_{\mathbf{x}\mathbf{y}} = \begin{cases} 1 & \text{iff } (\mathbf{x}, \mathbf{y}) \in E_{\mathcal{P}}, \\ 0 & \text{otherwise,} \end{cases} \quad (2.72)$$

for all $(\mathbf{x}, \mathbf{y}) \in E$. Using Eq. (2.69) we can readily compute the end-to-end rates and capacities of more general multi-path routing protocols.

Probabilistic routing strategies can be similarly studied, with the difference that $q_{\mathbf{x}\mathbf{y}} \in [0, 1]$ rather than $q_{\mathbf{x}\mathbf{y}} \in \{0, 1\}$. However, throughout our investigations deterministic routing protocols will be the most commonly utilised approach to quantum networking. This approach readily specifies to the previous results of single-path routing and flooding. It is clear that if we reduce the set of potential paths to only one route $\Omega_{\mathcal{P}} = \{\omega\}$ then we reclaim a single-path protocol, \mathcal{P}_{sp} . Otherwise, we engage in a multi-path protocol, \mathcal{P}_{mp} . Furthermore, the flooding can be reclaimed via a forwarding distribution of $q_{\mathbf{x}\mathbf{y}} = 1$ for all $(\mathbf{x}, \mathbf{y}) \in E$ such that all network edges are used deterministically.

2.4.5 Benchmarking Quantum Networks: A Beginning

A hierarchy of end-to-end network rates and capacities emerges,

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) \leq K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{mp}}) \leq K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{fl}}). \quad (2.73)$$

The optimal single-path routing strategy solved by DA can always be presented as a lower-bound on end-to-end quantum networking. Hence, the deployment of any multi-path protocol can only be deemed useful if it exceeds the single-path rate. On the other side of Eq. (2.73), we see that flooding always presents an optimal upper-bound for end-to-end quantum networking, i.e. one can never do better than exploiting the entire network to establish quantum communications between a single end-user pair. With this information in hand, one may embark upon the task of benchmarking quantum network architectures.

Nonetheless, there are a great number of theoretical open questions regarding quantum communications that the tools we have developed thus far have missed or of which they have made us aware: can we access analytical classes of network architectures to provide

expedient performance benchmarks? Do these architectures remain useful in the context of free-space, hybrid quantum designs? Or do we need to explore numerical, random networks? How about computing the capacities of quantum networks for which the capacity is not exactly known? How do we design practical multi-path routing algorithms that do not resort to flooding? How do we extend our theoretical advancements into the multiple user domain?

In the remainder of this thesis, we embark upon a journey that addresses these questions (and more) in an effort to advance our field's mathematical toolbox and motivate the future development of a quantum internet.

Chapter 3

Analytical Methods for High-Rate Global Quantum Networks

The work in this chapter forms the basis of a paper published in *PRX Quantum*, whose authors are (in order) Cillian Harney and Stefano Pirandola [1].

In this chapter, we tackle the challenge of constructing quantum network architectures which are analytically treatable but have realistic physical and topological freedom. For this purpose, the general class of weakly-regular quantum networks (WRNs) are characterised. Section 3.1 motivates the challenge of analysing large-scale quantum networks, summarising recent research directions in the literature. Section 3.2 introduces the family of WRNs including their mathematical characteristics and a convenient language for their use. Sections 3.3 and 3.4 lay out the utility of WRNs in the context of quantum networking and derive physical conditions that promise optimal end-to-end performance. Finally, Section 3.5 uses the WRN architecture to compare the performance of global fibre-networks with rates that are achievable by satellite quantum communications, assessing the advantages associated with each infrastructure.

3.1 Introduction

Overcoming the point-to-point limitation of quantum communication governed by the PLOB bound requires the construction of quantum networks. Reviewed in Section 2.3, we can combine tools from classical network theory with quantum information theory to derive the end-to-end capacities of quantum networks [32]. These results confirm that the PLOB bound can be beaten via quantum networking, facilitating high rate communication at longer ranges.

Yet, challenges remain. While such bounds are easily expressed in generality for arbi-

trary network topologies, their practical assessment requires the specification of a routing protocol and network architecture. Questions of network topology have been recently considered via the statistical study of complex, random quantum networks [63, 64, 65, 66], which reveal insightful phenomena associated with large-scale network properties. Studies of this kind are extremely valuable and help to unveil important guidelines for the development of future quantum networks. Nonetheless, such analyses are not easy and require significant numerical effort in order to compute important network properties, e.g. critical network densities or maximum fibre-lengths. There is a demand for versatile, *analytical* tools which allow for efficient benchmarking and can motivate the construction of large-scale topologies.

Meanwhile, ground-based quantum channels are not the only conduits available for global quantum communications. A rival infrastructure that may prove superior to fibre-based networks at global distances is Satellite Quantum Communication (SQC) [30, 25, 31, 67, 68, 69, 70, 71]. SQC exploits ground-to-satellite communication channels to overcome the fundamental distance limitations offered by fibre/ground-based mechanisms. A satellite in orbit around the Earth may act as a *dynamic* repeater that physically passes over ground-based users and distributes very long-range entanglement/secret-keys. The ability to exploit a free-space connection with a satellite also carries the possibility of substantially more transmissive channels than optical-fibre, making it ideal for global communication protocols.

The following critical questions emerge:

- Can we develop analytical tools which allows us to place limits on the end-to-end performance of large-scale quantum networks?
- Are fibre-based networks truly the best way to achieve long-distance quantum communication?

The goal of this chapter is to make progress with these challenges. Utilising ideas from quantum information theory, classical networks and graph theory we investigate ideal architectures based on the property of weak-regularity. Weakly-Regular Networks (WRNs) simultaneously (*i*) idealise network connectivity, (*ii*) provide sufficient freedom to capture a broad class of spatial topologies and (*iii*) remain analytically treatable so that critical network properties can be rigorously studied. This results in a design with desirable qualities which can efficiently and effectively provide insight for realistic structures. We show that quantum WRNs employing multi-path routing admit remarkably accessible and achievable upper-bounds on the end-to-end network capacity. This allows for a characterisation of the

ideal performance of a fibre-based quantum internet with respect to essential properties such as maximum channel length and nodal density.

The exact, analytical results derived in this chapter provide an immediate pathway to compare SQC with global ground-based quantum communications. We study the average number of secret bits per day that can be distributed between two remote stations, using large-scale quantum fibre-networks or a single satellite repeater station in orbit. Our findings rigorously prove the superiority of satellite-based quantum repeaters for global quantum communications, derive constraints associated with fibre-based networks and the enormous resource demands required to overcome achievable rates offered by a single satellite.

3.2 Weakly-Regular Networks (WRNs)

In this section we explicitly introduce the concept of weak-regularity and weakly-regular networks (WRNs) using graph theoretic concepts.

3.2.1 Graphs, Neighbour Sharing and Commonality

Consider an undirected, finite graph $\mathcal{N} = (P, E)$ consisting of n nodes in the node set P , and interconnected by edges in the edge set E . The ability to perform communication on a network is characterised by (i) the communication channels which compose the network, and (ii) the distribution of network nodes and edges resulting in a topology. Nodal degree, and its distribution across a network, is hugely influential on the communication capabilities of an architecture. However, the degree alone does not give an indication of *how* a node \mathbf{x} is connected to all of its neighbours. One may ask; are the neighbours also highly connected to one another, or are each of the neighbours distant and disconnected? Answering these questions can be very informative, and provide significant insight into the connectivity and robustness of a network. For this reason, we define useful parameters that contribute to these features. Namely, we use the concept of *commonality*.

Commonality is a pairwise nodal property which describes neighbour sharing between nodes. Given a pair of nodes $\mathbf{x}, \mathbf{y} \in P$ the commonality defines how many neighbours that \mathbf{x} and \mathbf{y} *have in common*. Neighbour sharing behaviour may vary significantly depending on whether \mathbf{x} and \mathbf{y} are already connected (adjacent) and are perhaps close by; or are disconnected (non-adjacent) and perhaps distant. Therefore, we provide the following pair of definitions of commonality:

Definition 3.1 (Adjacent Commonality): *The number of common neighbours shared by adjacent (connected) nodes. Precisely, given that $(\mathbf{x}, \mathbf{y}) \in E$, the adjacent commonality*

between this pair of nodes is $\lambda(\mathbf{x}, \mathbf{y}) := |N_{\mathbf{x}} \cap N_{\mathbf{y}}|$, so that $\lambda(\mathbf{x}, \mathbf{y})$ counts the number of common neighbours shared between the nodes \mathbf{x} and \mathbf{y} .

Definition 3.2 (Non-Adjacent Commonality): *The number of common neighbours shared by non-adjacent (non-directly-connected) nodes. Precisely, given that $(\mathbf{x}, \mathbf{y}) \notin E$, the non-adjacent commonality is computed by $\mu(\mathbf{x}, \mathbf{y}) := |N_{\mathbf{x}} \cap N_{\mathbf{y}}|$, so that $\mu(\mathbf{x}, \mathbf{y})$ counts the number of common neighbours shared between the nodes \mathbf{x} and \mathbf{y} .*

3.2.2 Regular Graphs

Consider an undirected, finite graph $\mathcal{N} = (P, E)$ of n -nodes. A graph is defined as k -regular if all nodes in the graph possess exactly the same degree k , i.e. the neighbourhood of any node consists of strictly k nodes,

$$\deg(\mathbf{x}) = |N_{\mathbf{x}}| = k, \quad \forall \mathbf{x} \in P. \quad (3.1)$$

Regularity significantly simplifies the connective properties of network by assuming a consistency of nodal degree. Clearly, in realistic communication networks there are disparities of nodal degree throughout the network, as some nodes will be highly connected and others less so. Nonetheless, understanding the ability to communicate on a regular graph can help provide important information for more realistic structures. The class of k -regular graphs is very broad, and more detailed classes can be defined.

The class of *strongly regular* (SR) graphs satisfy strict connective properties. A graph $\mathcal{N} = (P, E)$ is SR if it has n -nodes which are k -regular, its commonality properties are constant

$$\lambda(\mathbf{x}, \mathbf{y}) = \lambda, \quad \forall \mathbf{x}, \mathbf{y} \in P \text{ s.t. } (\mathbf{x}, \mathbf{y}) \in E, \quad (3.2)$$

$$\mu(\mathbf{x}, \mathbf{y}) = \mu, \quad \forall \mathbf{x}, \mathbf{y} \in P \text{ s.t. } (\mathbf{x}, \mathbf{y}) \notin E, \quad (3.3)$$

and these parameters follow the relation

$$\mu(n - k - 1) = k(k - \lambda - 1). \quad (3.4)$$

SR graphs may be well connected, but their architectures are very strict; satisfying all of these constraints will typically result in a network with a small number of nodes. Indeed, the parameters k, μ, λ inhibit the ability to use a large number of nodes rendering them impractical for network design.

A more general class is that of *weakly regular* (WR) graphs. Any regular graph that is not SR is technically WR, and can be characterised by a more general set of connectivity properties. We may invite greater generality by loosening the strict values of the

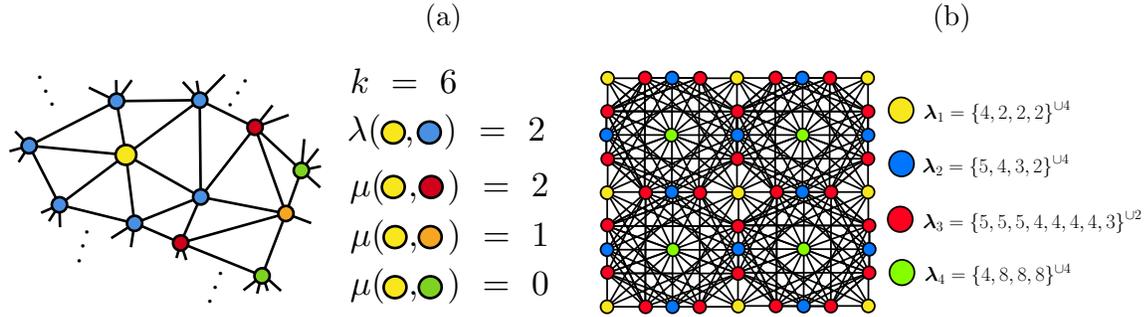


Figure 3.1: (a) A sub-graph from a $(k, \lambda, \mu) = (6, 2, \{0, 1, 2\})$ -weakly regular network. Considering the yellow node as an end-user, the blue nodes thus represent the user neighbourhood, with a uniform adjacent commonality of $\lambda = 2$. The non-adjacent commonality decreases as nodes increase in distance from the end-user. (b) A $k = 16$ weakly-regular network with inconsistent adjacent commonality properties. This network is scalable so that a single network cell can be concatenated to construct a larger $k = 16$ internally-WR network. For any node in the network, its λ will be one of those from the set $\Lambda = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$. Each adjacent commonality multiset is colour coded to its corresponding node on the graph. Note that throughout this work we employ a superscript union notation to describe the repeated union of a single set, e.g. $\{x\}^{\cup 3} = \{x\} \cup \{x\} \cup \{x\} = \{x, x, x\}$, etc.

adjacent/non-adjacent commonalities λ , μ for all nodes. Instead, we may permit nodes within the network to possess different commonality values for different pairs of nodes. To this end, we define sets which contain all the potential values for the commonality properties; an adjacent commonality set and a non-adjacent commonality set respectively,

$$\lambda := \{\lambda_1, \dots, \lambda_l\}, \quad \mu := \{\mu_1, \dots, \mu_m\}. \quad (3.5)$$

These sets summarise *all the non-degenerate values of $\lambda(\mathbf{x}, \mathbf{y})$ or $\mu(\mathbf{x}, \mathbf{y})$ that are possible on a network*. That is for any two nodes $\mathbf{x}, \mathbf{y} \in P$,

$$\lambda(\mathbf{x}, \mathbf{y}) \in \lambda, \quad \mu(\mathbf{x}, \mathbf{y}) \in \mu. \quad (3.6)$$

There is no restriction on the number of potential values that can be contained in λ or μ . Indeed, every graph (regular or not) generate their own versions of these sets.

As a simple example, we illustrate a weakly-regular sub-graph from a larger network in Fig. 3.1. Clearly, the degree of the network is $k = 6$, and the commonality properties are also illustrated by considering adjacent and non-adjacent nodes with respect to a root node. Provided that the regularity in Fig. 3.1 is consistent throughout the network, it can be shown that the adjacent commonality is always $\lambda = 2$ for any pair of nodes and the non-adjacent commonality can be $\mu \in \mu = \{0, 1, 2\}$.

3.2.3 Useful Parameterisation of Weakly-Regular Graphs

In complete generality, a graph can possess unique commonality values between any pair of nodes, leading to a vast collection of possible values in λ and μ with little to no structure. However, for various architectures, such as WRNs, this will not be true and there may exist a level of consistency which simplifies their analytical treatments. Here, we introduce a more useful and intuitive representation of WRNs. Consider a node \mathbf{x} on a k -WR graph. We can define a k element *multiset*¹ which contains all the information about neighbour sharing between the node \mathbf{x} and each of its k neighbours,

$$\lambda_{\mathbf{x}} := \{\lambda(\mathbf{x}, \mathbf{y}_1), \dots, \lambda(\mathbf{x}, \mathbf{y}_k)\} = \{\lambda(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in N_{\mathbf{x}}\}. \quad (3.7)$$

Hence, $\lambda_{\mathbf{x}}$ describes a “local” adjacent commonality description, bespoke to the node \mathbf{x} .

All nodes on any graph (regular or not) possess an adjacent commonality multiset $\lambda_{\mathbf{x}}$ which describes neighbour sharing qualities with respect to their neighbourhoods, so there exists a unique $\lambda_{\mathbf{x}}$ for all $\mathbf{x} \in P$. Therefore we can summarise the neighbour sharing properties of an entire network \mathcal{N} by collecting all of the possible adjacent commonality multisets contained within it into a strict superset

$$\Lambda := \{\lambda_{\mathbf{x}_1}, \lambda_{\mathbf{x}_2}, \dots, \lambda_{\mathbf{x}_n}\} = \{\lambda_{\mathbf{x}} \mid \mathbf{x} \in P\}. \quad (3.8)$$

Hence, for any node \mathbf{x} in a network, the adjacent commonality multiset of this node can be found in Λ .² In many cases of interest, such as WRNs with high levels of symmetry, the adjacent commonalities $\lambda_{\mathbf{x}}$ may be highly degenerate across the network, i.e. many nodes possess similar neighbour sharing qualities. Consequently, since Λ is a strict set it contains *only* the non-degenerate $\lambda_{\mathbf{x}}$ multisets. Here we state a formal definition:

Definition 3.3 (Adjacent Commonality Superset): *Any graph $\mathcal{N} = (P, E)$ possesses an adjacent commonality superset $\Lambda := \{\lambda_{\mathbf{x}} \mid \mathbf{x} \in P\}$ containing all the possible, non-degenerate adjacent commonality multisets $\lambda_{\mathbf{x}}$ which describe the neighbour sharing properties of connected nodes.*

For many of the architectures that we consider in this chapter, Λ only contains one acceptable adjacent commonality multiset, $\Lambda = \{\lambda\}$. This is evident in Fig. 3.1(a) where every pair of adjacent nodes always share exactly two neighbours. In general, this may

¹This is a modified set which may contain multiple copies of the same element, i.e. elements are not necessarily unique.

²Note that we now define Λ as a strict set, *not* a multiset so that it contains only one unique version of each $\lambda_{\mathbf{x}}$.

occur when there are high levels of symmetry/small k regularity in the network structure. However, this is not compulsory. Flexibility in $\mathbf{\Lambda}$ can allow us to describe more complex designs with higher nodal degrees. For instance, Fig. 3.1(b) depicts a $k = 16$ regular network which may be a portion of a larger network. All nodes satisfy $k = 16$, but there are four unique adjacent commonality multisets $\mathbf{\Lambda} = \{\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \boldsymbol{\lambda}_3, \boldsymbol{\lambda}_4\}$ for any network node. It is vital to understand the graphical properties of WRNs in order to properly characterise end-to-end performance for embedded end-users.

In contrast, it's not overly useful to construct an analogous language for the non-adjacent commonality properties of a WR graph, $\boldsymbol{\mu}$. A node-specific non-adjacent commonality object $\boldsymbol{\mu}_x$ would collect the number of shared neighbours between x and *all nodes on the network outside of its neighbourhood*. For large-scale networks this is a potentially huge number of nodes and for the most part will not give valuable information. Hence, in this work we will define WR graphs according to the properties $(n, k, \mathbf{\Lambda}, \boldsymbol{\mu})$ along with the definitions and discussions in this section. This provides us with the most effective language to investigate this interesting graphical class.

3.2.4 Genuine vs. Internal Weak-Regularity

In the context of large-scale quantum networks, some aspects of WR architectures still need to be addressed. Namely, there are properties that are strictly defined by the parameters $n, k, \mathbf{\Lambda}$ and $\boldsymbol{\mu}$ which require some further discussion in order to qualify the WRN structures investigated in our work. This leads to a further sub-categorisation of weak-regularity into two key formats; *genuine* and *internal* weak-regularity.

Genuine Weak-Regularity

Definition 3.4 (Genuine Weak-Regularity): *Consider a network $\mathcal{N} = (P, E)$ which is $(n, k, \mathbf{\Lambda}, \boldsymbol{\mu})$ -weakly-regular. This network is *Genuinely-WR* if there are absolutely no violations of these connectivity properties for any node $x \in P$ within the network.*

While this may seem like a trivial definition, it will become apparent in subsequent sections why it is necessary. Genuine weak-regularity can be readily satisfied but is sometimes quite restrictive. Indeed, a WRN defined within a two-dimensional spatial area may lead to some undesirable characteristics, such as extremely long edges required to satisfy regularity for all nodes; ultimately undermining the integrity of the network.

Nonetheless, genuine weak-regularity conditions can be easily satisfied by considering closed networks embedded on a sphere (or other appropriate closed, three-dimensional ob-

jects). Global quantum networks, in which we consider a network that spans the Earth may be appropriately and ideally modelled via genuinely-WR quantum networks. This is illustrated in Fig. 3.2(a) where a network can be defined on the surface of a sphere.

Internal Weak Regularity

As mentioned, defining regularity conditions on a two-dimensional plane can lead to unwanted features, such as extremely long edges used to “close” the network and satisfy all regularity conditions. Genuine weak-regularity avoids these features by considering closed networks embedded on some three dimensional surface. This may make sense for networks which span a planet, but for smaller areas this is not practical.

Hence, we may provide an alternative model of network connectivity. It is possible to define a network that satisfies the WR connectivity properties *within a network boundary*. That is, one can construct a WRN such that there exists a set of network nodes and edges that form a boundary

$$P_{\text{bound}} = \{\mathbf{p}_1, \dots, \mathbf{p}_m, \dots\}, \quad (3.9)$$

$$E_{\text{bound}} = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x}, \mathbf{y} \in P_{\text{bound}}\}, \quad (3.10)$$

within which all other nodes satisfy some form of weak regularity. In this way, there exists a WR sub-network within this boundary $\mathcal{N}_{\text{int}} = (P_{\text{int}}, E_{\text{int}})$ according to the node and edge sets

$$P_{\text{int}} := P \setminus P_{\text{bound}}, \quad E_{\text{int}} := E \setminus E_{\text{bound}}. \quad (3.11)$$

The total network model \mathcal{N} is clearly not genuinely WR since the boundary nodes $\mathbf{x} \in P_{\text{bound}}$ will violate the weak-regularity conditions. Nonetheless, the internal network \mathcal{N}_{int} will satisfy these conditions, providing a useful architecture which can be readily defined over two-dimensional regions.

Definition 3.5 (Internal Weak-Regularity): *Consider a network $\mathcal{N} = (P, E)$. This network is defined as internally-WR if there exists a network boundary $P_{\text{bound}} \subset P$, $E_{\text{bound}} \subset E$ such that the sub-network $\mathcal{N}_{\text{int}} := (P \setminus P_{\text{bound}}, E \setminus E_{\text{bound}})$ satisfies a form of $(n, k, \mathbf{\Lambda}, \boldsymbol{\mu})$ weak-regularity.*

Fig. 3.2 provides a useful illustration of the difference between genuinely-WR and internally-WR networks. One of the most useful features of this network class is that they are easy to construct, and easy to scale. It is straightforward to construct a regular *network cell*, which when concatenated with many other cells results in an internally-WR

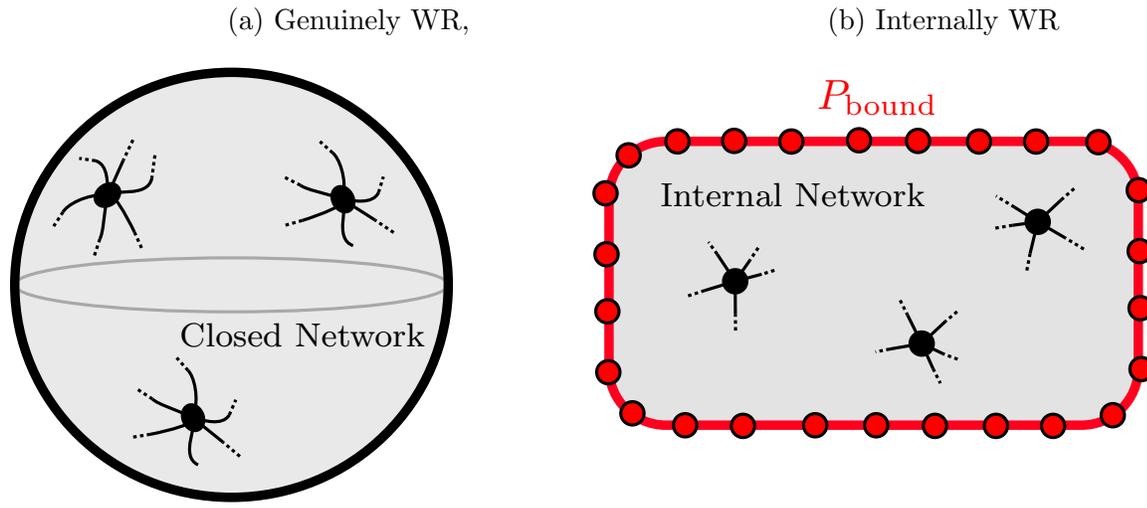


Figure 3.2: Distinction between (a) genuine weak-regularity and (b) internal weak-regularity. Genuinely WR networks can be embedded on a closed, three-dimensional surface such as a sphere in order to maintain regularity and avoid boundary effects. Internally WR networks satisfy weak-regularity within some nodal boundary P_{bound} , allowing us to investigate open networks which are defined within some two-dimensional area.

network. This concatenation process makes it easy to consider large-scale WRNs with open boundaries.

3.2.5 Simplification of Notation

For the purposes of our work, it is possible to simplify the notation we use to describe relevant WR structures. Since we are investigating large-scale network structures, it is not desirable to precisely define the number of nodes n , but allow n to be encoded into other properties, e.g. nodal density, maximum link length, etc. This can be achieved, given the crucial assumption that there are enough network nodes so that our analysis is unaffected by boundary effects or sparsity. The need for this assumption is different depending on whether we consider internal weak-regularity or genuine weak-regularity.

- *Genuine weak-regularity:* By definition, we do not have any issue with boundary effects in this setting since the network is effectively closed, and all nodes are unquestionably k -regular. Consequently, we require a sufficient number of nodes in order to construct the closed network and ensure that there exist two end-user nodes which are not directly connected. This is not a large number of nodes and can be satisfied easily, given a particular architecture.

- *Internal weak-regularity*: In this setting, we assume that end-users nodes that we select always fall within the outer boundary of nodes, and we only consider nodes within this boundary. At the very least, we require that there are enough nodes n within this boundary such that there exist two end-user nodes which are not directly connected (as this would defeat the purpose of investigating end-to-end network protocols). In general, this is a geometric packing problem specific to the weakly-regular architecture we are studying.

Henceforth, these assumptions are implicitly made within each of our WRN models. This allows us to omit the precise number of nodes n from key theory throughout our work and derive general results which apply to a broad range of network structures. The number of nodes and nodal density are revisited later in our studies in order to provide adequate insight to the resource requirements of WRNs.

Finally, we provide one further simplification by removing detailed reference to the possible non-adjacent-commonalities within the network, described by μ . The set μ is important for the characterisation of short-range connective structures, detailing how many shared neighbours two non-connected nodes may possess. For large networks, the vast majority of non-adjacent nodes will simply share no neighbours, $\mu = 0$. This is especially true for networks which obey distance constrained connectivity rules. We find that our subsequent analyses do not require its consistent usage, therefore it can be omitted for the sake of clarity (unless otherwise specified).

Following the implicit assumptions for the necessary number of nodes required to describe a WRN and the ability to ignore the non-adjacent commonalities, we can compactly characterise a class of WR architectures via the parameters: (k, Λ) .

3.3 Optimal Performance of Weakly-Regular Networks

The key mathematical tool we develop in this chapter is the ability to accurately and analytically derive conditions for the optimal performance of quantum WRNs. This requires graph theoretic arguments and a characterisation of minimum network cuts. In this section we elucidate these arguments.

3.3.1 Motivation

As discussed in Section 2.4, computing the flooding capacity for a general network and point-to-point capacity distribution requires a numerical treatment via the max-flow min-cut theorem. However, for any network we can always identify at least one valid cut via

user-node isolation, i.e. given an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ we can partition them by cutting all the edges in the neighbourhood of one of the end-user nodes, $C = E_{\mathbf{a}}$ or $C = E_{\mathbf{b}}$ (where $E_{\mathbf{x}}$ is defined in Eq. (2.56)). This cut totally disconnects an end-user node from the network, resulting in a successful partition. We call the multi-edge capacity associated with this kind of cut as the *min-neighbourhood capacity*. Following the notation from Section 2.4, we can write

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i} := \min_{j \in \{\mathbf{a}, \mathbf{b}\}} \mathcal{C}_{\text{cut}}(E_j), \quad (3.12)$$

$$= \min_{j \in \{\mathbf{a}, \mathbf{b}\}} \sum_{(\mathbf{x}, \mathbf{y}) \in E_j} \mathcal{C}_{\mathbf{x}\mathbf{y}}, \quad (3.13)$$

where we have labelled the min-neighbourhood capacity using $\mathcal{C}_{\mathcal{N}_i}$.

The min-neighbourhood capacity is always at least an upper-bound on the end-to-end flooding capacity. It is a strong indicator of a well connected and thus high-performance network. Networks which are highly connected contain many end-to-end routes between any pair of network nodes. The greater the number of end-to-end routes between an end-user pair, the more difficult it is to partition them via a cut-set, i.e. it requires more and more edges to disconnect them. This initiates a relationship between cut-set cardinality, $|\tilde{C}|$, and distance from an end-user. Performing cuts with edges further away from a user node requires the collection of many more edges to consolidate the partition. The further from the user nodes we begin the cut, the greater the number of potential end-to-end paths we must restrict (since we have permitted a larger flow from the user node) and thus the more edges we must collect. We call this phenomenon *network cut growth*. Once again, for general architectures and topologies it is extremely difficult to investigate the concept of network cut growth and would require numerical treatment. However, WRNs are analytically friendly and an ideal candidate for studying this concept.

Our approach is based on the distinction between two kinds of network cuts; user-node isolation, and network-bulk cuts. Let us formally define the notion of a network-bulk:

Definition 3.6 Consider a network $\mathcal{N} = (P, E)$ and an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ who wish to communicate. We define a network-bulk with respect to this end-user pair as the sub-network $\mathcal{N}' = (P', E')$ which contains all the edges and nodes which are not directly connected to the end-user nodes. That is, the node and edge sets satisfy,

$$P' := \{\mathbf{x} \mid \mathbf{x} \in P \setminus \{\mathbf{a}, \mathbf{b}\}\}, \quad E' := \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{x} \in P \setminus (E_{\mathbf{a}} \cup E_{\mathbf{b}})\}. \quad (3.14)$$

In a large-scale network the network-bulk \mathcal{N}' constitutes the majority of the architecture. A *network-bulk cut* can then be considered as a network-cut C' which is performed by

exclusively collecting edges from the network-bulk rather than the user-neighbourhoods. By our previous arguments, when a network is well connected, cuts performed further away from either end-user refer to collections of edges in the network bulk.

This leads to the primary motivation of our work: via the intuition of network cut growth, we wish to derive a relationship between WR networks, user-node isolation and network-bulk cuts. We wish to show that in highly-connected architectures (such as WRNs), cut growth causes cuts in the network-bulk to be unlikely candidates for the minimum cut. As a result, this allows us to identify conditions for which the upper-bound in Eq. (3.13) is saturated and elucidate network properties for which optimal performance is guaranteed.

3.3.2 Network-Bulk Cuts

In this section, we derive some useful lemmas which help us to understand network cut growth and network-bulk cuts.

Lemma 3.1 *Select two nodes on a genuinely-WR quantum network $\mathcal{N} = (P, E)$ that represent end-users, $\mathbf{a}, \mathbf{b} \in P$, and demand they that they do not share an edge or neighbour. The cut-set \tilde{C} which contains the fewest number of edges collects k edges.*

Proof. Menger's theorem states that for a finite, undirected graph the size of the minimum cut-set is equal to the maximum number of disjoint paths that can be found between any pair of vertices [72, 73]. Here, we are considering a $(k, \mathbf{\Lambda})$ weakly-regular graph with enough nodes to locate a pair of end-users which do not share an edge or neighbour. Every disjoint path will have to use one of the edges from the neighbourhood of an end-user, $N_{\mathbf{a}}$ and $N_{\mathbf{b}}$. After k disjoint paths, all the edges in the neighbourhoods of the end-user nodes will have already been used by one of these paths. Consequently, no more disjoint paths can be found, as the end-users can find no route to the network-bulk. Hence, the smallest cut-set cardinality will always equal k . ■

Lemma 3.2 *Consider a $(k, \mathbf{\Lambda})$ -genuinely-WR quantum network $\mathcal{N} = (P, E)$ such that $\mathbf{\Lambda} = \{\lambda\}$. Select two nodes that represent end-users, $\mathbf{a}, \mathbf{b} \in P$, and demand that they do not share an edge or neighbour. For any cut-set \tilde{C} that is restricted to edges in the network-bulk $e \in E'$,*

$$\text{if } \sum_{j=1}^k \lambda_j \leq k(k-2) \implies |\tilde{C}| \geq k. \quad (3.15)$$

If $\lambda_j = \lambda, \forall j$ then the condition holds if $\lambda \leq k-2$.

Proof. For a genuine $(k, \{\lambda\})$ -regular network there will always exist a cut-set with cardinality $|\tilde{C}_{\text{iso}}| = k$, achieved by isolating the neighbourhoods of either of the end-user nodes. By Lemma 1, we also know that this is the minimum cut-set cardinality. Meanwhile, a network cut which is limited to collecting edges on the network-bulk is unable to directly disconnect the neighbourhoods of \mathbf{a} or \mathbf{b} ($N_{\mathbf{a}}$ and $N_{\mathbf{b}}$ respectively). Hence, any cut which is performed on the network-bulk has to restrict flow from not just the end-users, but each of its neighbours. That is, any alternative cut-set will have to cut the unique edges in the neighbourhoods of the \mathbf{a}/\mathbf{b} 's neighbouring nodes.

Let us consider the cut which restricts flow from all of the neighbouring nodes of either end-user. This cut-set will be either of the following:

$$\tilde{C}_{\mathbf{a}} = \bigcup_{\mathbf{x} \in N_{\mathbf{a}}} \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in N_{\mathbf{x}} \setminus (N_{\mathbf{a}} \cup N_{\mathbf{b}} \cup \{\mathbf{a}, \mathbf{b}\})\}, \quad (3.16)$$

$$\tilde{C}_{\mathbf{b}} = \bigcup_{\mathbf{x} \in N_{\mathbf{b}}} \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in N_{\mathbf{x}} \setminus (N_{\mathbf{a}} \cup N_{\mathbf{b}} \cup \{\mathbf{a}, \mathbf{b}\})\}. \quad (3.17)$$

What are the cardinalities of these cut-sets? Thanks to network regularity this is easy to derive. Our goal is to restrict flow from each of the neighbours of \mathbf{a} or \mathbf{b} . By weak regularity, these neighbours will possess k edges; they will use one edge to connect directly to \mathbf{a} or \mathbf{b} , and λ_j nodes will be connected to *other* neighbours of \mathbf{a} or \mathbf{b} (by definition of adjacent commonality). As a result there will only be $(k - \lambda_j - 1)$ effective edges that permit logical flow *outside* of the end-user neighbour (to the rest of the network). By summing over all of the neighbours in either neighbourhood each of the new cut-set cardinalities are then

$$|\tilde{C}_{\mathbf{a}/\mathbf{b}}| = \sum_{j=1}^k (k - \lambda_j - 1). \quad (3.18)$$

When will this quantity be greater than $|\tilde{C}_{\text{iso}}|$? This is easy to determine and retrieves the condition stated in the lemma,

$$\sum_{j=1}^k (k - \lambda_j - 1) \geq k \implies \sum_{j=1}^k \lambda_j \leq k(k - 2). \quad (3.19)$$

If this condition holds, then we can always write $|\tilde{C}_{\mathbf{a}/\mathbf{b}}| \geq k$ as required. Any cut which is not \tilde{C}_{iso} or $\tilde{C}_{\mathbf{a}/\mathbf{b}}$ will necessarily permit flow into wider parts of the network. This will always increase the number of disjoint paths from \mathbf{a} to \mathbf{b} within the network, since the network is weakly-regular and connectivity properties are consistent throughout the network. It will therefore have a larger cut-set cardinality. ■

Lemma 3.2 serves as a critical tool in our analyses. It relates analytical properties of a WRN with properties of cuts performed on its network-bulk. In abstract terms, it posits weak-regularity conditions for which we can be certain that a WRN undergoes network cut growth.

While Lemma 3.2 has been formalised in the context of a WRN with consistent adjacent commonality properties (i.e. $\mathbf{\Lambda} = \{\boldsymbol{\lambda}\}$ has only one multiset) it can be easily extended to account for multiple possible multisets. For this reason we propose the following definition.

Definition 3.7 (Minimum Adjacent Commonality): *Given a $(k, \mathbf{\Lambda})$ -WR graph, the minimum adjacent commonality multiset $\boldsymbol{\lambda}^* \in \mathbf{\Lambda}$ is that which collects the fewest edges on a network-bulk cut,*

$$\boldsymbol{\lambda}^* = \arg \min_{\boldsymbol{\lambda} \in \mathbf{\Lambda}} \sum_{\lambda \in \boldsymbol{\lambda}} (\lambda - k - 1). \quad (3.20)$$

The minimum adjacent commonality multiset is a characteristic of any WRN. It identifies the network nodes which possess the smallest network-bulk cuts. Ultimately, if Lemma 3.2 is satisfied for the minimum adjacent commonality multiset $\boldsymbol{\lambda}^*$, then it holds for all possible nodes on the network.

Lemma 3.3 *Select two nodes on a genuinely $(k, \mathbf{\Lambda})$ -WR quantum network $\mathcal{N} = (P, E)$ that represent end-users, $\mathbf{a}, \mathbf{b} \in P$, and demand they that they do not share an edge or neighbour. For any cut-set \tilde{C} that is restricted to edges in the network-bulk $e \in E'$, if $\sum_{\lambda \in \boldsymbol{\lambda}^*} \lambda \leq k(k - 2)$ it follows that $|\tilde{C}| \geq k$.*

Proof. Since there is now variation in $\boldsymbol{\lambda}$ from node to node, the network-bulk cut that is associated with $\boldsymbol{\lambda}^*$ collects the least number of edges (by definition). Then any other node with a different $\boldsymbol{\lambda} \in \mathbf{\Lambda}$ must necessarily collect more edges than this. Given this consideration, the proof then follows directly from Lemma 3.2. ■

3.3.3 Network-Bulk Cuts on Internally-WRNs

Proposition 3.1 *Select two nodes on an internally-WR quantum network $\mathcal{N} = (P, E)$ that represent end-users, $\mathbf{a}, \mathbf{b} \in P$, and demand they that they do not share an edge. There exists some minimum number of network nodes n_{\min} for which the the results of Lemma 3.2/3.3 applies to \mathcal{N} .*

This proposition is well motivated, and can be proven for a number of different WR architectures. Open boundary edges add the complication of a potential cut C that utilises

the boundary to find a smaller cut-set than that used in Lemma 3.2. However, it is always possible to construct a sufficiently large network so that a pair of end-user nodes can always be found for which Lemma 3.2 is satisfied. We describe these end-user nodes as *deeply-embedded*. It is possible to provide a general characterisation of n_{\min} by identifying the minimum number of nodes for which there exists a cut-set \tilde{C}'' containing boundary edges $e \in E_{\text{bound}}$ that has a smaller cardinality than the network-bulk cut in Eq. (3.18), i.e. $|\tilde{C}''| < |\tilde{C}_{\mathbf{a}/\mathbf{b}}|$. Indeed this can be achieved, but such generality is not particularly useful in this chapter, and we leave it to future works. For now, we focus on WRN structures for which determining n_{\min} is a basic geometric problem. The quantity n_{\min} is the minimum number of nodes required to locate two end-users which do not share an edge or a neighbour, so that Lemma 3.2 is not compromised by the network boundary.

In Appendix A, Fig. A.1 we provide visual proofs of the minimum number of nodes required to satisfy internal-WR for the structures utilised in this chapter. The WRN properties and minimum number of nodes required in each case is listed below,

$$\begin{aligned}
k = 3, \boldsymbol{\lambda}^* &= \{0\}^{\cup 3} \rightarrow n_{\min} = 54, \\
k = 6, \boldsymbol{\lambda}^* &= \{2\}^{\cup 6} \rightarrow n_{\min} = 89, \\
k = 8, \boldsymbol{\lambda}^* &= \{2, 4\}^{\cup 4} \rightarrow n_{\min} = 120, \\
k = 16, \boldsymbol{\lambda}^* &= \{4, 8, 8, 8\}^{\cup 4} \rightarrow n_{\min} = 197.
\end{aligned} \tag{3.21}$$

In general, we are interested in large-scale networks with many more nodes than any of these values $n \gg n_{\min}$, so clearly the investigation of internally-WR graphs is well justified. When the neighbour sharing condition is relaxed for the end-users, this minimum number of nodes is reduced so that these constructions remain sufficient.

3.3.4 Threshold Capacities

With these lemmas in hand, we can present the key mathematical tools used throughout this chapter and derive the following *threshold theorems*.

Theorem 3.1 *Consider a $(k, \boldsymbol{\Lambda})$ -WR quantum network. Select an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, and demand they are sufficiently distant such that they do not share an edge or neighbour. Then there exists a threshold single-edge capacity \mathcal{C}_{\min} in the network, given by*

$$\mathcal{C}_{\min} := \frac{1}{\delta} \mathcal{C}_{\mathcal{N}_{\mathbf{i}}}, \tag{3.22}$$

where δ is a characteristic property of the network,

$$\delta := \sum_{\lambda \in \boldsymbol{\lambda}^*} k - \lambda - 1, \tag{3.23}$$

such that if all single-edge capacities in the network satisfy this minimum threshold, $\mathcal{C}_{\mathbf{x}\mathbf{y}} \geq \mathcal{C}_{\min}$, $\forall (\mathbf{x}, \mathbf{y}) \in E$ then flooding capacity is guaranteed to satisfy

$$\frac{2(k-1)}{\delta} \mathcal{C}_{\mathcal{N}_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}. \quad (3.24)$$

Proof. Let us denote the $(k, \mathbf{\Lambda})$ -WRN $\mathcal{N} = (P, E)$. The network possesses a large set of valid cuts, $\mathcal{C}_{\mathcal{N}} = \{C_j\}_j$, which collects all of the valid network cuts C_j that can successfully partition the pair of end-users. We can simultaneously define a set of *cut-set cardinalities*, i.e. if there exist M valid cuts, this is a M -element multi-set that counts the number of edges contained in each of the valid network-cuts. More precisely, we can define this multiset

$$c_{\mathcal{N}} = \{|\tilde{C}| \mid \tilde{C} \text{ s.t } C \in \mathcal{C}_{\mathcal{N}}\}. \quad (3.25)$$

By Lemma 1, the minimum-cut-set cardinality for the WRN \mathcal{N} is simply equal to its regularity, i.e. $\min(c_{\mathcal{N}}) = k$, and can be achieved by isolating an end-user (cutting the edges within an end-user neighbourhood). Performing user-node isolation we simply collect the edges from the user-neighbourhood $\tilde{C} = E_i$ to generate the min-neighbourhood capacity,

$$\mathcal{C}_{\mathcal{N}_i} = \min_{j \in \{a, b\}} \sum_{(\mathbf{x}, \mathbf{y}) \in E_j} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (3.26)$$

Now let us consider any network-bulk cut C' and its corresponding cut-set \tilde{C}' which is restricted to collecting edges on the network-bulk \mathcal{N}' . These types of cuts cannot use edges from the end-user-neighbourhoods and will provide a multi-edge capacity,

$$\mathcal{C}_{\text{cut}}(C') = \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}'} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (3.27)$$

In order to ensure $\mathcal{C}_{\mathcal{N}_i}$ is indeed the flooding capacity of the entire network, we must ensure that the minimum network-bulk based cut is *never* a minimum-cut, so that $\mathcal{C}_{\text{cut}}(C') \geq \mathcal{C}_{\mathcal{N}_i}$.

When restricted to performing cuts only on the network-bulk, the set of possible cuts will be different from $\mathcal{C}_{\mathcal{N}}$, since now certain cuts are inaccessible. Instead, we may define a new set of network-cuts $\mathcal{C}_{\mathcal{N}'}$ which are restricted to the network-bulk. This generates an analogous set of cut-set cardinalities

$$c_{\mathcal{N}'} = \{|\tilde{C}'| \mid \tilde{C}' \text{ s.t } C \in \mathcal{C}_{\mathcal{N}'}\}. \quad (3.28)$$

Using Lemma 1 we can determine the smallest network-bulk based cut on a $(k, \mathbf{\Lambda})$ -WR network (with no boundary effects). Since $\mathbf{\Lambda} = \{\lambda_{\mathbf{x}} \mid \mathbf{x} \in P\}$ may contain many different

adjacent commonalities, it is always possible to lower-bound the cardinality of the smallest network-bulk cut-set by using the minimum adjacent commonality multiset λ^* . Then we can write,

$$\min(c_{\mathcal{N}'}) \geq \sum_{\lambda \in \lambda^*} (k - \lambda - 1) = \delta. \quad (3.29)$$

This corresponds to the minimum number of edges that must be cut from the neighbours of *the neighbours* of the minimum end-user (e.g. the green cut-set in Fig. 3.3). This generates \tilde{C}' as the cut-set restricted to the network-bulk with minimum cardinality.

In order to ensure that the flooding capacity is equal to the min-neighbourhood capacity, we want to make sure that this network-bulk cut never generates a multi-edge capacity smaller than $\mathcal{C}_{\mathcal{N}_i}$. That is, we wish to ensure that

$$\min_{C' \in \mathcal{C}_{\mathcal{N}'}} \mathcal{C}_{\text{cut}}(C') \geq \mathcal{C}_{\mathcal{N}_i}. \quad (3.30)$$

Minimising $\mathcal{C}_{\text{cut}}(C')$ is achieved by setting each edge in the network-bulk to its minimum value \mathcal{C}_{\min} and performing the cut which collects the fewest number of edges, such that

$$\min(c_{\mathcal{N}'}) \cdot \min_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}'} \mathcal{C}_{\mathbf{xy}} = \min(c_{\mathcal{N}'}) \cdot \mathcal{C}_{\min} \geq \mathcal{C}_{\mathcal{N}_i}. \quad (3.31)$$

Subsequently we can derive a minimum threshold capacity for any edge in the network,

$$\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}_{\min} = \frac{\mathcal{C}_{\mathcal{N}_i}}{\sum_{\lambda \in \lambda^*} (k - \lambda - 1)} = \frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}, \quad \forall (\mathbf{x}, \mathbf{y}) \in E, \quad (3.32)$$

which ensures that Eq. (3.30) is always upheld. Imposing this threshold constraint ensures that any cut restricted to the network-bulk will generate a multi-edge capacity that is greater than or equal to the min-neighbourhood capacity. As a result, no cut performed exclusively on the network-bulk can ever undermine the flooding capacity.

There is now only one issue; we must identify if there exists any possible *hybrid* cut that might undermine the flooding capacity being equal to the min-neighbourhood capacity. That is, is there a cut that can collect a mixture of edges contained in the user-neighbourhood *and* the network-bulk? Unfortunately there is, and it must be considered. Let us take a worst-case scenario where *all* of the edges in a network-bulk are of minimum threshold capacity \mathcal{C}_{\min} . Furthermore, let's consider that the min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$ is generated by a user-neighbourhood in which has $(k - 1)$ edges of capacity \mathcal{C}_{\min} and one edge with capacity $\mathcal{C}_{\mathcal{N}_i} - (k - 1)\mathcal{C}_{\min}$. That is,

$$\mathcal{C}_{\mathcal{N}_i} = (k - 1)\mathcal{C}_{\min} + [\mathcal{C}_{\mathcal{N}_i} - (k - 1)\mathcal{C}_{\min}]. \quad (3.33)$$

This is a worst-case situation in which one neighbourhood edge contains the majority of the min-neighbourhood capacity. In this scenario, it is possible to cut the $(k - 1)$ edges in the neighbourhood which have the threshold value, and then to cut an additional $(k - 1)$ edges in the network-bulk which are connected to the largest capacity edge in the neighbourhood *instead* of this neighbourhood edge. This results in a hybrid cut C'' which generates a multi-edge capacity

$$\mathcal{C}_{\text{cut}}(C'') \geq 2(k - 1)\mathcal{C}_{\min} = \frac{2(k - 1)}{\delta}\mathcal{C}_{\mathcal{N}_i}. \quad (3.34)$$

This is an absolute worst-case scenario for the network design, placing a lower-bound on the end-to-end flooding capacity.

Consequently, provided that $\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}_{\min}$, $\forall(\mathbf{x}, \mathbf{y}) \in E$ then the flooding capacity always satisfies

$$\frac{2(k - 1)}{\delta}\mathcal{C}_{\mathcal{N}_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}. \quad (3.35)$$

as required. This reveals a single-edge threshold condition for all edges in the network so to ensure that end-to-end performance is guaranteed within tight bounds. ■

Theorem 3.1 therefore allows us to place tight performance bounds on the flooding capacity of a quantum WRN. Using only the connectivity properties of the architecture itself, and a desired end-to-end performance, we can identify a single-edge capacity constraint for all network edges. This is extremely useful, and a key result in this work.

It is also possible to identify what additional constraints are necessary to not just guarantee a tight window of performance, but guarantee exact, optimal performance. This is achieved in the following theorem.

Theorem 3.2 *Consider a (k, Λ) -WR quantum network. Select an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, and demand they are sufficiently distant such that they do not share an edge or neighbour. Then there exists the threshold single-edge capacity \mathcal{C}'_{\min} in the network bulk, and another for the user-connected edges \mathcal{C}^i_{\min} given by*

$$\mathcal{C}'_{\min} := \frac{1}{\delta}\mathcal{C}_{\mathcal{N}_i}, \quad \mathcal{C}^i_{\min} := \left(\frac{1}{k - 1} - \frac{1}{\delta} \right) \mathcal{C}_{\mathcal{N}_i}, \quad (3.36)$$

such that if all single-edge capacities in the network satisfy their minimum thresholds, $\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}'_{\min}$, $\forall(\mathbf{x}, \mathbf{y}) \in E'$ and $\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}^i_{\min}$, $\forall(\mathbf{x}, \mathbf{y}) \in E_{\mathbf{a}} \cup E_{\mathbf{b}}$ then flooding capacity is guaranteed to satisfy

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}. \quad (3.37)$$

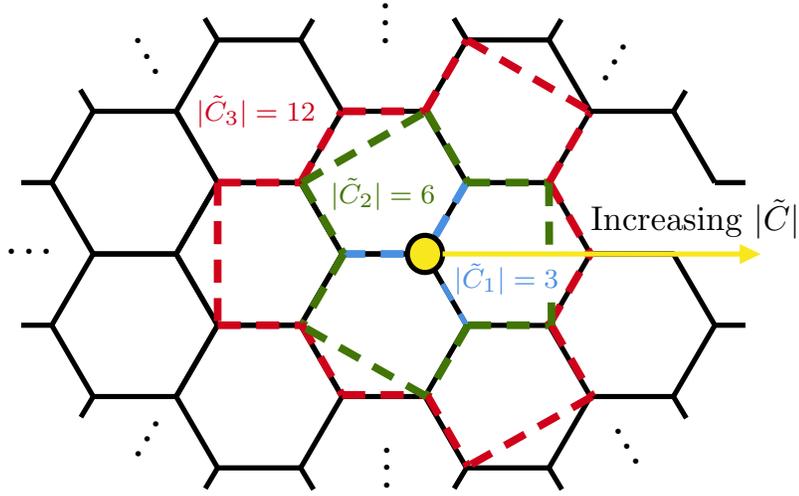


Figure 3.3: Cut-set cardinality with respect to increasing distance from user-node on a honeycomb lattice. We show some example cuts on a honeycomb network of increasing cut-set dimension. The further one moves from a user-node \mathbf{a} , the more edges that must be cut due to k -regularity. \tilde{C}_1 gives the neighbourhood cut-set $E_{\mathbf{a}}$, \tilde{C}_2 gives the smallest cut-set when limited to network-body edges E' , and \tilde{C}_3 gives a wider cut example.

Proof. By Theorem 3.1 we know that if all edges satisfy $\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}'_{\min} = \mathcal{C}_{\mathcal{N}_i}/\delta$, for all $(\mathbf{x}, \mathbf{y}) \in E$, then the flooding capacity satisfies Eq. (3.24). If we want to avoid the worst-case lower-bound it is possible to enforce an additional, slightly stricter constraint on the end-user connected edges, which we label \mathcal{C}_{\min}^i . If we consider the hybrid cut \mathcal{C}'' scenario as in the previous theorem where $(k-1)$ edges from the user-neighbourhood are collected and have minimum capacity \mathcal{C}_{\min}^i along with $(k-1)$ network-bulk edges with capacity \mathcal{C}'_{\min} in order to consolidate the end-user partition. This results in a possible multi-edge capacity

$$\mathcal{C}_{\text{cut}}(\mathcal{C}'') \geq (k-1)\mathcal{C}'_{\min} + (k-1)\mathcal{C}_{\min}^i = \frac{(k-1)}{\delta}\mathcal{C}_{\mathcal{N}_i} + (k-1)\mathcal{C}_{\min}^i. \quad (3.38)$$

To ensure that $\mathcal{C}_{\text{cut}}(\mathcal{C}'') \geq \mathcal{C}_{\mathcal{N}_i}$ we must then demand that

$$\frac{(k-1)}{\delta}\mathcal{C}_{\mathcal{N}_i} + (k-1)\mathcal{C}_{\min}^i \geq \mathcal{C}_{\mathcal{N}_i} \implies \mathcal{C}_{\min}^i \geq \left(\frac{1}{k-1} - \frac{1}{\delta}\right)\mathcal{C}_{\mathcal{N}_i}. \quad (3.39)$$

Therefore, if we demand that all edges in the user-neighbourhoods satisfy $\mathcal{C}_{\mathbf{xy}} \geq \mathcal{C}_{\min}^i$, then the worst-case scenario which generates the lower-bound in Theorem 3.1 disappears and becomes equivalent to the min-neighbourhood capacity. That is, the inequalities become

$$\mathcal{C}_{\mathcal{N}_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i} \implies \mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}, \quad (3.40)$$

as required. We find that $\mathcal{C}_{\min}^i \geq \mathcal{C}'_{\min}$ if $k \leq \frac{\delta}{2} + 1$, which is satisfied in all of our example architectures. ■

3.3.5 Neighbour Sharing End-Users

So far we have considered end-user pairs that are not directly connected and do not share common neighbours. This is appropriate assumption since we are studying global quantum communications over very long distances; it is not interesting to consider short range users separated by single links. Furthermore, it allows for much clearer intuition surrounding increasing cut-set dimension with respect to cuts on the network-bulk as shown in Fig. 3.3. This assumption does not compromise the generality of our arguments, as we show in the following corollary that Theorems 3.1 and 3.2 hold even when end-user nodes share a neighbour.

Corollary 3.1 (Neighbour Sharing): *Consider a (k, Λ) -WR quantum network and an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ within the network that do not share an edge, and possess a min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. Even if the end-user nodes share a neighbour, Theorems 3.1 and 3.2 hold.*

Proof. Consider the (k, Λ) -WR network and assume that the end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ are not directly connected, but share a neighbour. The number of common neighbours that these non-adjacent nodes share is defined by the non-adjacent commonality, $\mu(\mathbf{a}, \mathbf{b}) > 0$. The previous analyses do not directly apply since cuts restricted to the network-bulk will not be able to partition the two users. This is true because there will exist clear paths along the edges connected to the common neighbours of \mathbf{a} and \mathbf{b} . Hence, a valid network-cut of these end-users requires one to collect $\mu(\mathbf{a}, \mathbf{b})$ edges from a user-neighbourhood.

Nonetheless, our results still hold. Let us locate a network-cut that uses the minimum number of user-connected edges possible. This can be considered as a modification to the network-bulk cut which is necessary due to neighbour-sharing. This cut C' still collects at least $\sum_{\lambda \in \lambda^*} (k - \lambda - 1)$ edges, but now $\mu(\mathbf{a}, \mathbf{b})$ of those edges are actually contained in one of the user-neighbourhoods. In a worst-case scenario, one may assume that the user-connected edges which are necessarily cut possess the minimum single-edge capacity in the user-neighbourhoods, defined as

$$\mathcal{C}_{\min}^i = \min_{j \in \{\mathbf{a}, \mathbf{b}\}} \min_{(\mathbf{x}, \mathbf{y}) \in E_j} C_{\mathbf{x}\mathbf{y}}. \quad (3.41)$$

Let all edges in the network obey a threshold capacity $\mathcal{C}_{\min} = \mathcal{C}_{\mathcal{N}_i} / \delta$ as motivated by Theorem 3.1 for non-neighbour sharing end-users. Now, the network-cut C' which collects

the fewest number of edges from the user-neighbourhood will generate a multi-edge capacity,

$$\mathcal{C}_{\text{cut}}(C') \geq \delta \mathcal{C}_{\min} + \mu(\mathbf{a}, \mathbf{b})(\mathcal{C}_{\min}^i - \mathcal{C}_{\min}). \quad (3.42)$$

However, we already know that $\mathcal{C}_{\min}^i = \mathcal{C}_{\min}$ since we stated that all edges in the network obey the same minimum threshold. Therefore,

$$\mathcal{C}_{\text{cut}}(C') \geq \delta \mathcal{C}_{\min} = \mathcal{C}_{\mathcal{N}_i}. \quad (3.43)$$

as required. Therefore, neighbour sharing does undermine the previous threshold theorems. Indeed, introducing a stricter condition on the user-connected edges (as we do in Theorem 3.2) only makes this result stronger, since $\mathcal{C}_{\min}^i \geq \mathcal{C}_{\min}$ will only increase the multi-edge capacity in Eq. (3.42). ■

3.3.6 Bosonic Lossy Weakly-Regular Networks

When considering fibre-based networks, point-to-point links are described by bosonic pure-loss (lossy) channels. For lossy quantum networks, the most important property is channel length, or from a network perspective, *inter-nodal separation*. For a given edge $(\mathbf{x}, \mathbf{y}) \in E$ connecting two users in a network, the inter-nodal separation is simply the distance $d_{\mathbf{x}\mathbf{y}}$ between them. All two-way assisted quantum and private capacities of the lossy channel are precisely known via the PLOB bound from Eq. (2.49) [29],

$$\mathcal{C}(\mathcal{E}_{\eta_{\mathbf{x}\mathbf{y}}}) = \mathcal{C}(d_{\mathbf{x}\mathbf{y}}) = -\log_2 \left(1 - 10^{-\gamma d_{\mathbf{x}\mathbf{y}}} \right), \quad (3.44)$$

where the inter-nodal separation is related to the transmissivity via $\eta_{\mathbf{x}\mathbf{y}} = 10^{-\gamma d_{\mathbf{x}\mathbf{y}}}$. For state-of-the-art fibre-optics the loss rate is $\gamma = 0.02$ per km (which equates to a loss rate of 0.2 dB/km). Since these separations directly dictate the channel quality between nodes they must be precisely engineered and distributed in order to guarantee strong end-to-end performance.

The direct application of Theorem 3.1 to bosonic lossy quantum networks allows us to translate the notion of threshold capacities into something more physical. Indeed, since the capacity of pure-loss channels is known exactly, it is possible to translate the threshold capacity into a *maximum inter-nodal separation*.

Corollary 3.2 *Consider a $(k, \mathbf{\Lambda})$ -WR quantum network which is connected by bosonic lossy channels. Select an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ within the network that do not share an edge, and possess a min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. Then, there exists a maximum inter-nodal separation for all edges within the network,*

$$d_{\mathcal{N}}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}} \right), \quad (3.45)$$

for which the flooding capacity satisfies

$$\frac{2(k-1)}{\delta} \mathcal{C}_{\mathcal{N}_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}. \quad (3.46)$$

If $\exists d_{\mathbf{x}\mathbf{y}} < d_{\mathcal{N}}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E$, this remains an upper-bound on the optimal network performance, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}$.

Proof. Consider a valid pair of end-users $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ embedded within a $(k, \mathbf{\Lambda})$ -WR quantum network. Then as before, there exists a threshold capacity $\mathcal{C}_{\min} = \frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}$ that can be enforced to ensure the flooding capacity between these users is bounded by their min-neighbourhood capacity, $\mathcal{C}_{\mathcal{N}_i}$. Supplanting the PLOB bound into the capacity condition in Theorem 3.1,

$$\mathcal{C}(d_{\mathbf{x}\mathbf{y}}) \geq \mathcal{C}_{\min}, \forall (\mathbf{x}, \mathbf{y}) \in E, \quad (3.47)$$

this readily translates to,

$$d_{\mathbf{x}\mathbf{y}} \leq -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}} \right), \forall (\mathbf{x}, \mathbf{y}) \in E. \quad (3.48)$$

Therefore the threshold capacity becomes an upper-bound on the maximum link-length permitted within the network. We can thus define this maximum length,

$$d_{\mathcal{N}}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}} \right), \quad (3.49)$$

which when satisfied ensures that $2(k-1)\mathcal{C}_{\mathcal{N}_i}/\delta \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}$.

Now suppose that there exists a channel within the network-bulk that violate this max-bulk separation, i.e. $\exists d_{\mathbf{x}\mathbf{y}} > d_{\mathcal{N}}^{\max}$ for $(\mathbf{x}, \mathbf{y}) \in E'$. This violates the threshold capacity condition from Theorem 3.1 meaning that the minimum-cut in the network is not guaranteed to satisfy the performance bounds. However, if the minimum-cut undergoes a transition due to the introduction of poor quality channels in the network-bulk, it cannot improve the network flooding capacity; it can only deteriorate network performance. Therefore the min-neighbourhood capacity remains an upper-bound on the optimal network performance, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}$, as before. ■

In order to achieve a stricter performance guarantee, we can apply Theorem 3.2 to bosonic lossy channels and derive slightly stricter constraints on user-connected channels.

Corollary 3.3 Consider a $(k, \mathbf{\Lambda})$ -WR quantum network which is connected by bosonic lossy channels. Select an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ within the network that do not share an edge, and possess a min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. Then, there exists a maximum link-length in the network-bulk

$$d_{\mathcal{N}}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\frac{1}{\delta} \mathcal{C}_{\mathcal{N}_i}} \right), \quad (3.50)$$

and a maximum link-length for all the user-connected edges,

$$d_{\mathcal{N}_i}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\left(\frac{1}{k-1} - \frac{1}{\delta}\right) \mathcal{C}_{\mathcal{N}_i}} \right) \leq d_{\mathcal{N}}^{\max}, \quad (3.51)$$

which when satisfied guarantee that the flooding capacity is equal to the min-neighbourhood capacity,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}. \quad (3.52)$$

If $\exists d_{\mathcal{N}_i}^{\max} < d_{\mathbf{xy}} \leq d_{\mathcal{N}}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_{\mathbf{a}} \cup E_{\mathbf{b}}$, we regain Theorem 3.2. If $\exists d_{\mathbf{xy}} > d_{\mathcal{N}}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E$, then the performance guarantee is violated, but this remains an upper-bound on the optimal network performance, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}$.

Proof. This is a specification of Theorem 3.2 to bosonic lossy channels where the edges in the neighbourhoods of Alice \mathbf{a} and Bob \mathbf{b} may possess their own, stricter constraint in order to completely guarantee optimal performance. The proof follows directly by supplementing the PLOB bound into the threshold capacity expressions. ■

3.4 Nodal Densities and Bosonic Lossy Weakly-Regular Networks

3.4.1 Sparse Constructions

Nodal density is defined as the number of nodes n per unit area of the network,

$$\rho_{\mathcal{N}} := n/A \quad (3.53)$$

where A is some area in which the network is defined. This is a crucial measure of network resources, especially for quantum networks where there is a very high cost of constructing quantum devices at every node. In many network settings, it is desirable to minimise the nodal density necessary to promise strong end-to-end performance. For this reason, it is also useful to define a *minimum nodal density*. For a class of network, $\mathbf{N} = \{\mathcal{N}_j\}_j$, such that all instances $\mathcal{N} \in \mathbf{N}$ are constrained to some implicit structure, the minimum nodal density describes how it can be constructed in the sparsest way possible. It refers to a limiting scenario in which the network is least dense, and that all other instances of the network topology will possess more nodes per unit area. This is summarised below in a general definition:

Definition 3.8 (Sparse Construction): *Consider a class of network $\mathbf{N} = \{\mathcal{N}_j\}_j$ which imposes a fixed, single-edge distance constraint on its networks $\mathcal{N} = (P, E) \in \mathbf{N}$ so that*

$$d_{\mathbf{xy}} \leq d_{\mathcal{N}}^{\max} \text{ for all } (\mathbf{x}, \mathbf{y}) \in E. \quad (3.54)$$

The sparse construction is an instance of this class which minimises its network nodal density,

$$\rho_{\mathcal{N}}^{\min} = \min_{\mathcal{N} \in \mathbf{N}} \rho_{\mathcal{N}} = \min_{\mathcal{N} \in \mathbf{N}} \frac{n}{A}, \quad (3.55)$$

where $\rho_{\mathcal{N}}^{\min}$ is the minimum permitted nodal density of a network $\mathcal{N} \in \mathbf{N}$.

Clearly, for very general classes of distance-constrained networks this minimisation is extremely difficult. However, for analytical classes such as WRNs, this becomes rather easy and reduces to a geometric packing problem.

To provide simplifications in subsequent arguments, we make the following proposition.

Proposition 3.2 *For a class of single-edge distance constrained networks $\mathcal{N} \in \mathbf{N}$, such that $d_{\mathbf{xy}} \leq d_{\mathcal{N}}^{\max}$ for all $(\mathbf{x}, \mathbf{y}) \in E$, then the minimum nodal density can always be expressed as*

$$\rho_{\mathcal{N}}^{\min} \propto (d_{\mathcal{N}}^{\max})^{-2} = \xi (d_{\mathcal{N}}^{\max})^{-2}, \quad (3.56)$$

such that ξ is a quantity which characterises the network class \mathbf{N} .

It is always possible to express the min-density as proportional to the inverse squared value of the maximum inter-nodal separation in the network. This is obviously true for *any measure of area* since $\rho \propto A^{-1}$ and $A \propto d^2$ where d is some distance measure. Yet, for what follows we find that it is useful to closely relate $d_{\mathcal{N}}^{\max}$ and $\rho_{\mathcal{N}}^{\min}$ in this way. In the follow subsections we endeavour to lower-bound the min-nodal densities for a number of architectures which are classes of WRN.

3.4.2 Honeycomb Network

Our model of a honeycomb network ($k = 3$ and $\boldsymbol{\lambda}^* = \{0\}^{\cup 3}$) is the following: consider a single, initial hexagon consistent of 6-nodes connected by 6 edges. Let us call this the $r = 1$ ring of the network. To construct a larger network, we proceed by adding further hexagons concentrically around the initial shape. Each edge of the $r = 1$ ring is used as an edge of a hexagon in the $r = 2$ ring. We can continue to create a larger and larger network structure by concentrically connecting rings of hexagons to the previous one. As each ring is added, there will be $6r$ hexagons added to the overall structure. See Fig. A.1(a) as an example.

For any fixed number of rings r we can identify the number of nodes within the network. The number of unique nodes added with the addition of each new ring follows a recursive equation

$$\tilde{n}_1 = 6, \tilde{n}_2 = 12, \dots, \tilde{n}_r = 24(r - 1) - \tilde{n}_{r-1}. \quad (3.57)$$

It is simple to solve this set of recursive equations so that \tilde{n}_r takes the form,

$$\tilde{n}_r = 6(2r - 1). \quad (3.58)$$

As a result, given an r -ring honeycomb network structure, the total number of nodes will be

$$n_{\text{hc}}(r) = \sum_{k=1}^r 6(2k - 1) = -6r + 12 \sum_{k=1}^r k = 6r^2. \quad (3.59)$$

We may use this relationship in order to determine the minimum nodal-density ρ_{\min} of a honeycomb network when the maximum permitted fibre-length is $d_{\mathcal{N}}^{\max}$. Since this is the maximum permitted length and a honeycomb lattice can form a regular tiling, then ρ_{\min} is satisfied when every edge in the network is exactly $d_{\mathcal{N}}^{\max}$. Hence, given an r -ring network, the maximum area it will span is

$$A_{\text{hc}}^{\max}(r, d_{\mathcal{N}}^{\max}) = \frac{3\sqrt{3} [1 + 3r(r + 1)] d_{\mathcal{N}}^{\max 2}}{2}. \quad (3.60)$$

Hence, an r -ring minimum nodal density can be computed by

$$\rho_{\min}(r, d_{\mathcal{N}}^{\max}) = \frac{n_{\text{hc}}(r)}{A_{\text{hc}}^{\max}(r, d_{\mathcal{N}}^{\max})}. \quad (3.61)$$

By taking the asymptotic limit of $r \rightarrow \infty$ we can more accurately capture a lower-bound on the nodal density of a honeycomb network which satisfies this fibre-length constraint (as a larger network will permit a more accurate averaging process). As a result, we may compute

$$\rho_{\mathcal{N}}^{\min} \geq \lim_{r \rightarrow \infty} \rho_{\min}(r, d_{\mathcal{N}}^{\max}) = \frac{4}{3\sqrt{3} d_{\mathcal{N}}^{\max 2}} \quad (3.62)$$

as a lower-bound on the nodal-density of a weakly-regular honeycomb network which satisfies a maximum inter-nodal separation. Hence the characteristic quantity of honeycomb networks satisfies $\xi \geq 4/(3\sqrt{3})$.

3.4.3 Hexagonal Network

A class of hexagonal network ($k = 6$ and $\boldsymbol{\lambda}^* = \{2\}^{\cup 6}$) follows the same logic as the honeycomb structure, just with additional nodes located within every hexagon (see Fig. A.1(b)). As a result, we can immediately write

$$\tilde{n}_r = 6(2r - 1) + 6(r - 1) = 6(3r - 2). \quad (3.63)$$

Then, in an r -ring hexagonal structure the total number of nodes is given by,

$$n_{\text{hex}}(r) = 7 + 6 \sum_{k=2}^r (3k - 2) = 1 + 3r(3r - 1). \quad (3.64)$$

Meanwhile, the maximum area spanned by an r -ring hexagonal network is equal to that of the honeycomb network, $A_{\text{hex}}^{\text{max}} = A_{\text{hc}}^{\text{max}}$. Thus, defining

$$\rho_{\text{hex}}(r, d_{\mathcal{N}}^{\text{max}}) := \frac{n_{\text{hex}}(r)}{A_{\text{hex}}^{\text{max}}(r, d_{\mathcal{N}}^{\text{max}})}, \quad (3.65)$$

we can easily compute a lower-bound on the nodal density as before

$$\rho_{\mathcal{N}}^{\text{min}} \geq \lim_{r \rightarrow \infty} \rho_{\text{hex}}(r, d_{\mathcal{N}}^{\text{max}}) = \frac{2}{\sqrt{3} d_{\mathcal{N}}^{\text{max}2}}. \quad (3.66)$$

Hence the characteristic quantity of hexagonal networks satisfies $\xi \geq 2/\sqrt{3}$.

3.4.4 Manhattan-Inspired Networks

Consider a class of WRN such that $k = 8$ and $\lambda^* = \{2, 4\}^{\cup 4}$, as depicted in Fig. A.1(c). To construct this network, we can simply concatenate a cell (consisting of 9 nodes) into an $r \times r$ grid which can be easily evaluated. For a network which is arranged into a r -length square grid there will exist r^2 network cells. In order to maximise the area spanned by each network cell, we assign the longest possible edge in the cell to be of length $d_{\mathcal{N}}^{\text{max}}$. For the $k = 8$ network cell, this means that the diagonal edges in each square must be of length $d_{\mathcal{N}}^{\text{max}}$. Hence, the area of the total 9 node cell will be $\frac{1}{2}d_{\mathcal{N}}^{\text{max}2}$. In an r -ring network this results in a total area of $\frac{1}{2}(rd_{\mathcal{N}}^{\text{max}})^2$. Furthermore, the total number of nodes will be given by,

$$n_{\text{mh:8}}(r) = (r + 1)^2, \quad (3.67)$$

which can be obtained by simply counting the number of nodes on each horizontal/vertical row of the grid. We can thus define the function,

$$\rho_{\text{mh:8}}(r, d_{\mathcal{N}}^{\text{max}}) := \frac{(r + 1)^2}{\frac{1}{2}(rd_{\mathcal{N}}^{\text{max}})^2}. \quad (3.68)$$

As a result, a lower-bound on the minimum nodal density can be readily computed

$$\rho_{\mathcal{N}}^{\text{min}} \geq \lim_{r \rightarrow \infty} \rho_{\text{mh:8}}(r, d_{\mathcal{N}}^{\text{max}}) = \frac{2}{d_{\mathcal{N}}^{\text{max}2}}. \quad (3.69)$$

Therefore the characteristic quantity is lower-bounded by $\xi \geq 2$.

A similar Manhattan-like class can be constructed such that $k = 16$ and $\boldsymbol{\lambda}^* = \{4, 8, 8, 8\}^{\cup 4}$, as depicted in Fig. A.1(d). Using this network cell to construct a larger network, we must constrain the longest edge in the network cell to be of length $d_{\mathcal{N}}^{\max}$. This causes us to constrain the diagonal edge from central nodes on the boundary of the cell to connected nodes at the opposite corner. The maximum area spanned by a network cell is then $\frac{4}{5}d_{\mathcal{N}}^2$. If we again consider an $r \times r$ cell square grid network, then that the total area is $A_{\max}(r) = \frac{4}{5}r^2d_{\mathcal{N}}^{\max 2}$. Via a counting argument, the total number of nodes in an r -radius network will be

$$n(r) = (4r + 1)(r + 1) + 3r(r + 1) + r^2 = (7r + 1)(r + 1) + r^2. \quad (3.70)$$

As a result, we can define the minimum nodal density function,

$$\rho_{\text{mh:16}}(r, d_{\mathcal{N}}^{\max}) := \frac{(7r + 1)(r + 1) + r^2}{\frac{4}{5}(rd_{\mathcal{N}}^{\max})^2}. \quad (3.71)$$

Finally, the lower-bound can be given

$$\rho_{\mathcal{N}}^{\min} \geq \lim_{r \rightarrow \infty} \rho_{\text{mh:16}}(r, d_{\mathcal{N}}^{\max}) = \frac{10}{d_{\mathcal{N}}^{\max 2}}. \quad (3.72)$$

Hence the characteristic quantity can be lower-bounded by $\xi \geq 10$.

3.4.5 Nodal Density and End-to-End Performance

Theorem 3.3 *Consider a $(k, \boldsymbol{\Lambda})$ -WR quantum network $\mathcal{N} = (P, E)$ which is connected by bosonic lossy channels. Select an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ within the network that do not share an edge, and a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. In order to guarantee optimal performance, there exists a minimum nodal density within the network,*

$$\rho_{\mathcal{N}}^{\min} = -\xi\gamma^2 \left[\log_{10} \left(1 - 2^{-\frac{1}{\delta}\mathcal{C}_{\mathcal{N}_i}} \right) \right]^{-2}, \quad (3.73)$$

where ξ is characteristic of the WR architecture being considered.

Proof. In Corollaries 3.2 and 3.3, a global fibre-length constraints are placed on the network in order to guarantee a particular flooding capacity via user-node isolation. Using Corollary 3.2, if all edges $(\mathbf{x}, \mathbf{y}) \in E$ satisfy an maximum link-length constraint,

$$d_{\mathcal{N}}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\frac{1}{\delta}\mathcal{C}_{\mathcal{N}_i}} \right), \quad (3.74)$$

then the flooding capacity is guaranteed to satisfy $2(k - 1)\mathcal{C}_{\mathcal{N}_i}/\delta \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}$. If we apply the additional constraint for user-connected edges such that

$$d_{\mathcal{N}_i}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\left(\frac{1}{k-1} - \frac{1}{\delta}\right)\mathcal{C}_{\mathcal{N}_i}} \right), \quad (3.75)$$

then we can guarantee that $\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}$.

These link-length constraints result in a minimum nodal density for the entire network which is easy to investigate via the appropriate sparse construction. Using Eq. (3.56) we can directly write

$$\rho_{\mathcal{N}}^{\min} = \xi(d_{\mathcal{N}}^{\max})^{-2} = \xi\gamma^2 \left[\log_{10} \left(1 - 2^{-\frac{1}{\delta}\mathcal{C}_{\mathcal{N}_i}} \right) \right]^{-2}, \quad (3.76)$$

where the characteristic quantity, ξ is derived from the sparse construction. This offers a lower-bound on the necessary nodal density required to guarantee a particular flooding capacity.

Summarising, in order for the flooding capacity between \mathbf{a} and \mathbf{b} will be equal to the min-neighbourhood capacity $\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}$, the nodal density must (at least) satisfy the lower-bound $\rho_{\mathcal{N}} \geq \rho_{\mathcal{N}}^{\min}$. ■

The tightness of this lower-bound depends on how ξ is derived. Ideally, one would be able to take into the consideration the stricter constraint $d_{\mathcal{N}_i}^{\max}$ required to guarantee $\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}$ with equality. Solving a sparse construction with multiple edge constraints is not straightforward, hence one may need to use a lower-bound for ξ , as we have in this work. This nonetheless delivers informative bounds on the nodal density required for optimal performance.

3.4.6 Analysis of Bounds

Figure 3.5(c) depicts the connection between flooding capacity and minimum nodal density for a number of types of WRNs. It is clear that there is a trade off between end-to-end performance and regular nodal degree. At low flooding rates ($10^{-2} - 10^{-1}$ bits per network use) the WR structures with lower degrees $k = 3$ and 6 demand fewer resources to achieve the same performance as those with higher degrees $k = 8$ and 16. In this regime, high degrees are not necessary everywhere in the network to achieve the flooding rates; indeed, the consistent connectivity invoked by WR designs help to maintain performance at low densities. Yet, as the flooding capacity transitions towards 1 – 10 bits per network use this behaviour changes; WRNs with low degrees demand shorter and shorter links to achieve the high rates and the inability to involve more connections at each node becomes costly. As can be seen for $k = 3$ the required minimum nodal density rapidly increases, shortly followed by $k = 6$ and 8. Contrarily, the regime of high end-to-end rates is well suited to WRNs with greater regularity, $k = 16$, for which the greater number of connections at each node facilitate a lower overall density.

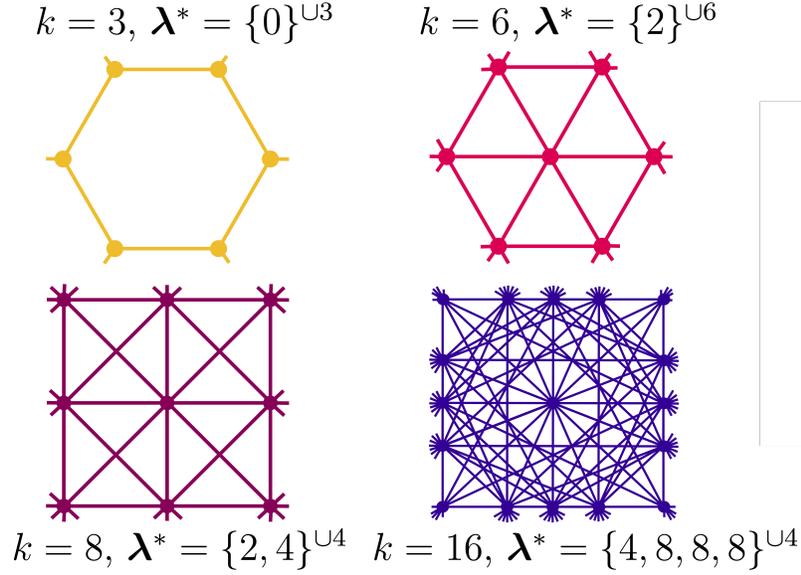


Figure 3.4: Examples of network cells that can be used to construct quantum WRNs, where k denotes regularity and λ^* is the adjacent commonality multi-set which minimises the quantity in Eq. (3.20). These quantities characterise the network cell and larger WRNs that they can construct.

Following recent works which have investigated critical network resources required for effective end-to-end performance on quantum networks, we define a critical nodal density ρ_{crit} as the network density required to achieve an end-to-end rate of 1 bit per network use. For bosonic lossy networks constructed in a weakly-regular structure, we can derive this value analytically.

Corollary 3.4 Consider a $(k, \mathbf{\Lambda})$ -WR quantum network $\mathcal{N} = (P, E)$ which is connected by bosonic lossy channels. The critical nodal-density of the network is lower-bounded by

$$\rho_{\mathcal{N}}^{\text{crit}} \geq -\xi\gamma^2 \left[\log_{10} \left(1 - 2^{-\frac{1}{\delta}} \right) \right]^{-2}, \quad (3.77)$$

where ξ is characteristic of the WR architecture being considered.

In Eq. (3.77), recall that γ is the fibre-loss rate which takes a typical value of $\gamma \approx 0.02$, and δ is defined in Eq. (3.29) as before. For WRNs explored in this chapter we can readily compute their critical nodal densities:

$$\begin{aligned} \rho_{\text{hc}}^{\text{crit}} &\gtrsim \frac{\sqrt{3}}{5625} \left[\log_{10} \left(1 - 2^{-1/6} \right) \right]^{-2} \approx 3.33 \times 10^{-4} \text{ nodes per km}^2, \\ \rho_{\text{hex}}^{\text{crit}} &\gtrsim \frac{\sqrt{3}}{3750} \left[\log_{10} \left(1 - 2^{-1/18} \right) \right]^{-2} \approx 2.28 \times 10^{-4} \text{ nodes per km}^2, \end{aligned} \quad (3.78)$$

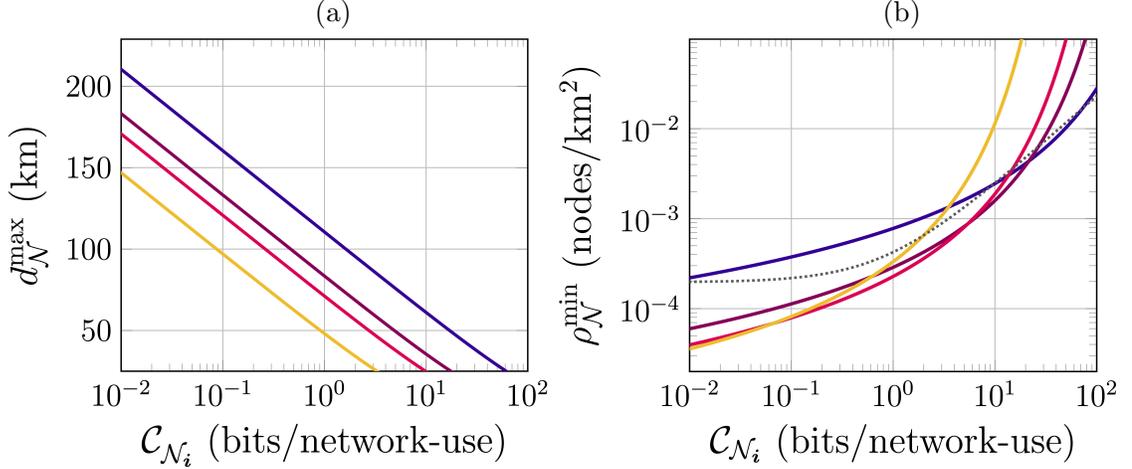


Figure 3.5: All plots are colour coded with respect to architectures derived from network cells in Fig 3.4. (a) Relationship between the optimal end-to-end flooding capacity (equal to the min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$) and the maximum link-length $d_{\mathcal{N}}^{\max}$ required to guarantee it for bosonic lossy quantum networks according to Eq. (3.45). Greater network regularity leads to larger permissible ranges of channel lengths. Panel (b) depicts this relationship between minimum nodal density $\rho_{\mathcal{N}}^{\min}$ with respect to optimal performance for bosonic lossy quantum networks, colour coordinated with the network cells. The grey dotted line relates the nodal density to the average flooding capacity between any pair of nodes in a Waxman Network, as in Eq. (3.80) [65].

$$\begin{aligned} \rho_{\text{mh:8}}^{\text{crit}} &\gtrsim \frac{1}{1250} \left[\log_{10} \left(1 - 2^{-1/32} \right) \right]^{-2} \approx 2.87 \times 10^{-4} \text{ nodes per km}^2, \\ \rho_{\text{mh:16}}^{\text{crit}} &\gtrsim \frac{1}{250} \left[\log_{10} \left(1 - 2^{-1/128} \right) \right]^{-2} \approx 7.78 \times 10^{-4} \text{ nodes per km}^2. \end{aligned} \quad (3.79)$$

In Figure 3.5(c), we plot an approximation of the average flooding capacity between any pair of nodes on a Waxman network with respect to nodal density (dashed grey line) as derived in Ref. [65]. This defines an expected flooding capacity between any pair of users, such that

$$\mathbb{E}_{\mathbf{i}} [\mathcal{C}(\mathbf{i}, \mathcal{N})] \approx \zeta(\rho_{\mathcal{N}} - \rho_{\text{crit}}) - 1, \quad (3.80)$$

where $\rho_{\text{crit}} \approx 4.25 \times 10^{-4}$, $\zeta \approx 4358$ and the average $\mathbb{E}_{\mathbf{i}}[\cdot]$ is taken over all possible end-user pairs in the network. We identify a kinship between the necessary $\rho_{\mathcal{N}}^{\min}$ predicted by WRNs and that derived for Waxman networks. As one may expect, the order and consistency of WRNs is able to promise lower resource demands at lower-rates; resulting in smaller critical nodal density predictions for the necessary density to achieve 1 bit per network use. However, as the flooding performance increases, the flexibility of the Waxman design (its ability to utilise variable nodal degrees) renders it superior to the lower degree

WR structures. In summary, there is good behavioural agreement between these models, corroborating the utility of WR structures as a valuable analytical tool for quantum network design.

3.5 Comparison with Satellite Quantum Communications

3.5.1 Satellite Quantum Communications

Here, we briefly review key results which facilitate a comparison of SQC with idealised, ground-based quantum networks. For more detailed derivations and discussions of these results, please refer to Refs. [31, 30].

Consider two users (Alice and Bob), who choose to communicate by means of an orbiting satellite (a dynamic repeater). Here we consider a ground station G at approximately sea-level, and a satellite S which is in orbit at an altitude $h \geq 100$ km and variable zenith angle θ . Given that the radius of the Earth is $R_E \approx 6371$ km, the slant distance between G and S is $z(h, \theta) = \sqrt{h^2 + 2hR_E + R_E^2 \cos^2(\theta)} - R_E \cos(\theta)$, describing the true distance that an optical beam must travel from G to/from S . We may consider two unique configurations for information transmission; *uplink*, which refers to when G is the transmitter and S is the receiver, and *downlink*, where the converse is true. Both configurations will identically admit the effects of free-space diffraction (beam-spot size widening) and atmospheric extinction (caused by molecular/aerosol absorption as the beam propagates). However additional loss/noise effects emerge with respect to uplink and downlink protocols, which invokes an asymmetry in their communication performance.

The effects of turbulence (caused by fluctuations in the atmospheric refractive index) and pointing errors (alignment of the optical signal with the receiver) are responsible for beam wandering, which instigates a fading process for the communication channels. For uplink protocols, turbulence is a significant factor for loss properties of the ground-satellite channel since it impacts the propagating beam immediately after transmission. However, pointing errors can be reduced thanks to the ability to easily access and optimise adaptive optics at ground level. In downlink these effects are reversed. Turbulence is not a factor until the beam reaches low altitudes, at which point the beam has already spread via diffraction. Hence turbulence can be neglected for downlink, but pointing errors must be considered due to limited onboard access and resources.

Considering each of these physical effects characterising the lossy free-space channel, it is possible to present an ultimate limit on the secret-key capacity K for SQC [31],

$$K \leq -\Delta(\eta, \sigma) \log_2(1 - \eta). \quad (3.81)$$

Here $\Delta(\eta, \sigma)$ is a correction factor to the PLOB bound, where $\eta := \eta(h, \theta)$ is an effective transmissivity which is a function of geometric position, encompassing all the effects of diffraction, extinction, and optical imperfections/inefficiencies. Meanwhile, $\sigma^2 = \sigma_{\text{turb}}^2 + \sigma_{\text{point}}^2$ is the variance of the Gaussian random walk of the beam centroid caused by beam wandering, with contributions from turbulence and/or pointing-errors.

This bound can be further modified to account for the presence of thermal noise, which is highly dependent upon time of day (day or night-time) and weather conditions (cloudy or clear skies). For night-time communications, background noise is practically negligible, and the above bound requires little modification. However, for day-time operations this is generally not the case and the free-space lossy channels must be described as thermal-loss channels which account for additional noise.

3.5.2 Practical Key-Rates for Satellite Quantum Communications

The bound in Eq. (3.81) is an ultimate upper-bound on the capacity of a ground-to-satellite communication channel, it is important to provide an assessment of realistic and practical protocols which embody achievable lower-bounds for SQC. These lower-bounds will facilitate comparisons with global quantum networks, and help deduce the conditions for which we can expect satellite advantage for long-distance quantum communications.

Here we summarise some achievable rates for different satellite configurations. We consider practical, composable-secure secret key-rates achievable from the pilot-guided and post-selected CV-QKD protocol studied in Refs. [30, 31]. The main concept of this protocol is to encode information into Gaussian-modulated coherent states, randomly interleaved by highly energetic pilot pulses used to monitor the transmissivity and fading properties of the free-space channel in real time (facilitating the use of classical post-selection). This protocol has been comprehensively extended to account for the physical scenario of satellite quantum communications, resulting in realistic and practical rates.

We may consider the employment of such a protocol in conjunction with a near-polar sun-synchronous satellite used to communicate between two ground stations. This type of orbit ensures a consistent fly-over time for any point on the Earth's surface, such that the satellite passes over any point at the same local mean solar time each day. This provides the possibility of stable conditions for satellite communications at around the same time each day. Let us assume that the stations lie along the orbital path such that the satellite crosses both of their zenith positions (which happens once per day). We further assume a worst-case scenario such that the stations only interact with the satellite when the zenith positions are crossed, and that both stations assume similar operational conditions.

It is possible to quantify the performance of satellite communications by considering a *daily key rates*, i.e. the number of secret-bits that may be shared per day. This allows us to utilise an average orbital rate R_{orb} associated with up/downlink operations in day/night-time, representing an average secret-key rate per link usage. Thanks to the dynamic nature of SQC, and the fact that we consider communication with both stations only once per day, this daily rate will be constant with respect to ground based end-to-end distances. The number of secret-bits that can be shared in a zenith-crossing passage is then given by the effective transit time for the quantum communications $t_Q(h)$ as a function of the altitude, and a typical clock frequency which we set as $\alpha = 10$ MHz. The average daily-rate in a given configuration is thus

$$R_{\text{daily}}^{\text{sat}} \approx \alpha t_Q(h) R_{\text{orb}}^i, \quad (3.82)$$

for which i labels the up/downlink and day/night-time.

For downlink operations at altitude $h = 530$ km, initial beam-spot-size $\omega_0 = 40$ cm, receiver aperture $a_R = 1$ m, these setup parameters lead to the night-time/day-time rates [31],

$$R_{\text{orb}}^{\text{down}} \approx \begin{cases} 3.066 \times 10^{-2} & \text{bits/use (night),} \\ 3.041 \times 10^{-2} & \text{bits/use (day).} \end{cases} \quad (3.83)$$

For uplink, we consider an altitude $h = 103$ km and similar setups (but now with a spot-size $\omega_0 = 60$ cm and wider aperture $a_R = 2$ m) leading to the rate,

$$R_{\text{orb}}^{\text{up}} \approx \begin{cases} 4.244 \times 10^{-2} & \text{bits/use (night),} \\ 2.737 \times 10^{-2} & \text{bits/use (day).} \end{cases} \quad (3.84)$$

Notice that in both configurations the day and night time rates are very similar. This is thanks to effective noise-filtering that can be performed with this kind of CV-QKD protocol. Such protocols are able to realistically exploit CV quantum systems and interferometric measurements in order to achieve much narrower frequency filters than is possible with DV protocols (see Ref. [31] for more details). As a result, the increased background thermal noise experienced at the receiver in day time does not significantly deteriorate the rate.

3.5.3 Comparison with Ground-Based Networks

As we have established in previous sections, end-to-end distance independence is a critical design feature for the construction of effective quantum networks. It is a feature that can be achieved, provided that one carefully monitors link-length, nodal density and

the limits of quantum communication rates. Yet, as shown in the previous section, it can be very costly to promise strong end-to-end rates between long-distance end-users if we choose to solely utilise ground-based fibre networks. For this reason, it is important to understand the limits of large-scale quantum networks for long-range communication. Moreover, it is invaluable to determine when SQC may be superior and offer a cost-efficient route to global quantum communication.

Determination of *when* SQC is advantageous requires a strict, quantitative comparison with ground-based fibre networks. In this section we aim to benchmark the optimal performance of global quantum fibre networks against practical, near-term SQC capabilities. More precisely, we compare daily secret key-rates obtained between globally distant end-users via:

- (i) A global-scale $(k, \mathbf{\Lambda})$ -WR fibre network with capacity achieving links.
- (ii) A single, sun-synchronous satellite operating at the achievable rates in Eqs. (3.82)-(3.84) using realistic devices and the practical CV-QKD protocol discussed in Section 3.5.2.

Clearly, the resources accessed by an ideal $(k, \mathbf{\Lambda})$ -WR fibre network are significantly greater than the single satellite, and a fairer comparison would be to consider a constellation of satellites; but that is the point. If a single, sun-synchronous satellite, operating at realistic rates is able to outperform a global fibre network within a meaningful resource regime, this offers clear evidence for the superiority (and necessity) of SQC for global quantum communications. Using the tools developed throughout this chapter, our comparison can be carried out analytically.

Assume two globally distant end-users, Alice and Bob. We need not consider a specific end-to-end distance, since the $(k, \mathbf{\Lambda})$ -WRNs are end-to-end distance independent. By considering a daily key rate and the operational setup explained in Section 3.5.2, SQC is also end-to-end distance independent. We are left to compute the daily capacity of the WR fibre network. We consider that the fibre network operates constantly for a day using capacity achieving links with maximum link length $d_{\mathcal{N}}^{\max}$. Given $t_{\text{daily}} = 8.64 \times 10^4$ s as the number of seconds in a day, and again assuming $\alpha = 10$ MHz, it can be shown the average number of secret-key bits per day satisfies

$$R_{\text{daily}}^{(k, \mathbf{\Lambda})}(d_{\mathcal{N}}^{\max}) \lesssim -\frac{\alpha t_{\text{daily}}}{\delta} \log_2(1 - 10^{-\gamma d_{\mathcal{N}}^{\max}}), \quad (3.85)$$

where δ is defined in Eq. (3.20). Repeater-chains can be considered in a similar manner. The repeater-chain capacity is equal to the single-edge capacity associated with the longest

inter-nodal separation in the chain. Hence, the average daily secret-key rate of a repeater-chain is [32]

$$R_{\text{daily}}^{\text{chain}}(d_{\mathcal{N}}^{\text{max}}) \lesssim -\alpha t_{\text{daily}} \log_2(1 - 10^{-\gamma d_{\mathcal{N}}^{\text{max}}}). \quad (3.86)$$

In order to perform a quantitative comparison between satellite and ground-based quantum communications, we can compute the log-ratio between their daily-rates,

$$\Delta K_{\text{daily}} := 10 \log_{10} \left(\frac{R_{\text{daily}}^{(k, \Lambda)}}{R_{\text{daily}}^{\text{sat}}} \right), \quad (3.87)$$

which determines a *daily-rate advantage* in decibels (dB). An analogous quantity can be derived for the repeater chain. By studying the daily-rate advantage as a function of maximum inter-nodal separation and nodal density, we can then determine conditions for which SQC begins to outperform the global, ground-based networks. That is,

$$\begin{aligned} \Delta K_{\text{daily}} > 0 &\implies \text{Fibre-Network Advantage,} \\ \Delta K_{\text{daily}} = 0 &\implies \text{Equal Performance,} \\ \Delta K_{\text{daily}} < 0 &\implies \text{Satellite Advantage.} \end{aligned} \quad (3.88)$$

Hence, there exists a critical inter-nodal separation $d_{\mathcal{N}}^*$ and a critical nodal density $\rho_{\mathcal{N}}^*$ for which $\Delta K_{\text{daily}} = 0$. Beyond $d_{\mathcal{N}}^*$ or below $\rho_{\mathcal{N}}^*$, a single, sun-synchronous satellite quantum repeater is more effective than a global fibre-network.

Fig. 3.6 illustrates results for the daily-rate advantage over SQC for a repeater-chain, and a number of quantum WRNs with various connectivity properties. In particular, we compare the resource demands of SQC with $k = 6, 8$ and 16 WRNs based on the network-cells shown in Fig. 3.5(a). Each architecture will possess its own unique critical values, defining a limiting property of the network. This comparison involves the consideration of a number of SQC operational setups and conditions which are summarised in Table I in Fig. 3.6; regarding the time of operation (night or day), physical direction of communication (uplink or downlink), satellite altitude, initial beam spot-size and receiver aperture radius. It is important to note that we can *always* exploit the superior communication direction (downlink) for the purposes of QKD between end-users, thanks to the independence of physical and logical flow (as discussed in Section 2.3.4). Therefore the critical properties $\rho_{\mathcal{N}}^*$ and $d_{\mathcal{N}}^*$ are computed as the values for which $\Delta K_{\text{daily}} = 0$ with respect to SQC downlink rates.

In Fig. 3.6(a)-(c) we plot the maximum tolerable fibre-length permitted in a repeater chain and each WRN required to guarantee ΔK_{daily} advantage over the single satellite repeater. The critical fibre-length for a quantum repeater chain operating at the ultimate

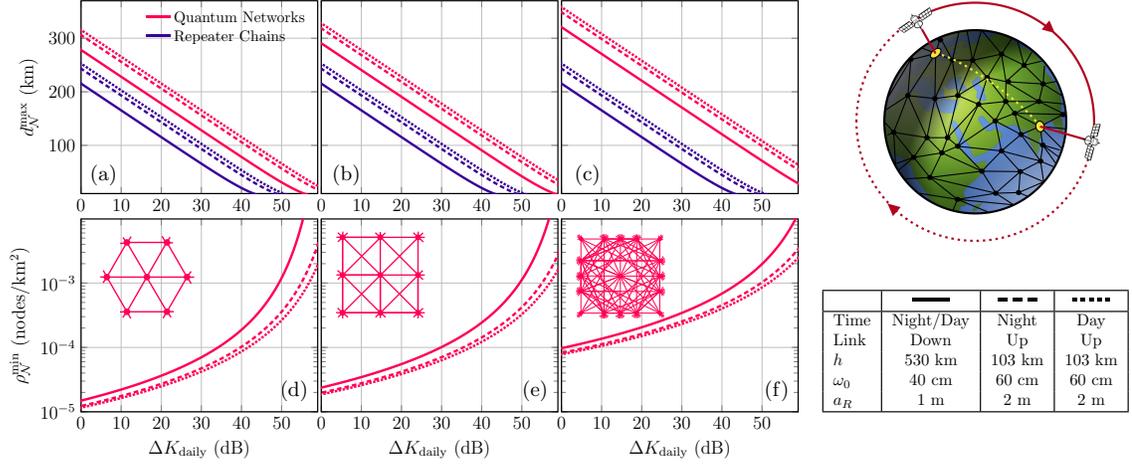


Figure 3.6: The daily secret-key-rate advantage ΔK_{daily} in Eq. (3.87) achieved by fibre-based quantum WRNs and repeater-chains with capacity achieving links over a single, sun-synchronous satellite-based repeater operating at practical, achievable rates from Eqs. (3.83) and (3.84). The architecture of each WRN is shown inside of Panels (d) to (f), such that vertically aligned panels use the same architecture. Panels (a) to (c) plot ΔK_{daily} with respect to maximum fibre-length permitted within each structure, $d_{\mathcal{N}}^{\text{max}}$. Panels (d) to (f) plot the relationship between the daily rate advantage ΔK_{daily} and minimum nodal density required in each WRN to achieve it. Satellite-based advantage can be achieved when $\Delta K_{\text{daily}} \leq 0$. All considered satellite setup parameters are shown in the table describing the operation time, direction, altitude h , spot-size ω_0 , receiver aperture a_R .

limit is $d_{\text{rep}}^* \approx 215$ km, which offers a lower-bound on repeater-assisted, ground-based strategies. This can be extended by quantum networks using multi-path routing strategies, as WRNs are able to tolerate longer lossy channels at the expense of greater resource demands. This is clear from the results in Fig. 3.6, where extending the critical separation by approximately 100 km requires a $k = 16$ regular network, for which $d_{\mathcal{N}}^* \approx 320$ km.

We may also identify the minimum required WRN nodal density, $\rho_{\mathcal{N}}^*$, for obtaining ground-based advantage over a single satellite, plotted in Fig. 3.6(d)-(f). Analysis of this property provides an appreciation of the resources demanded by these fibre-networks. While the WRNs with lower regularity are constrained to shorter link-lengths, the required nodal density at poorer end-to-end rates (low levels of advantage) is smaller than that of better connected designs. We find that the critical nodal densities $\rho_{\mathcal{N}}^*$ are of order 10^{-5} nodes per km^2 , e.g. for $k = 6$ we find that $\rho_{\mathcal{N}}^* \sim 1 \times 10^{-5}$, for $k = 16$ we gather $\rho_{\mathcal{N}}^* \sim 9 \times 10^{-5}$ (nearly a whole order of magnitude larger).

These are expensive values when put into the perspective of a global communication scenario. Let us take a naïve scenario from a practical point of view, but one that is infor-

mative nonetheless. Consider quantum communication between distant end-users located in remote cities across continental Europe (e.g. Paris to Moscow) whose land surface area spans approximately $A \approx 1 \times 10^7 \text{ km}^2$. In terms of truly global communications this is relatively local. We can choose to communicate between remote cities using a satellite in orbit acting as a dynamic quantum repeater. Alternatively, we can construct a quantum fibre-network across the continent. In this scenario, for an ideal $k = 6$ WR quantum fibre-network operating at its ultimate flooding capacity to simply match the already achievable daily-rate of a single, sun-synchronous satellite, would require at least $n \geq A\rho_{\mathcal{N}}^* \approx 150$ repeater stations operating constantly for 24 hours. Clearly, a network of this form operating at realistic rates, under stricter physical conditions (considering thermal noise) would demand even greater resources.

While the classical internet can exploit fibre-optic links which are thousands of kilometres long, a fibre-based quantum internet is severely limited by short link-lengths, resulting in remarkably costly resources for tasks that are already within reach of SQC. These results strongly suggest that a future quantum internet will significantly benefit from the use of SQC, and will be integral to the construction of global quantum communication networks.

3.6 Conclusion

In this work, we have investigated the optimal performance of global, quantum communication networks to characterise the ultimate limits of a fibre-based quantum internet. This analysis is based on an underlying network architecture that exploits weak-regularity to construct powerful, highly-connected networks. Crucially, these bounds allow us to benchmark the performance of a global quantum network versus that of a single sun-synchronous satellite acting as a dynamic repeater. The result of this comparison emphasises the power of SQC, and vast network resources that are required to outperform a single satellite in orbit at global distances. These findings strongly motivate the utilisation of ground-satellite connections within large-scale quantum networks. It is clear that free-space ground-satellite links will be integral to long-range quantum communications, as their co-operation with ground-based infrastructure as dynamic repeaters will be invaluable.

This work introduces useful, analytical techniques for the study of ideal quantum networks which can be readily employed for future investigative paths. Indeed, the study of hybrid fibre/satellite networks is a topic of immediate interest; exploiting the power of SQC to enhance (rather than compete with) ground-based networks. Furthermore, the expansion of these methods to incorporate multiple satellites introduces the possibility of highly transmissive satellite-satellite channels at high altitudes.

Chapter 4

End-to-End Capacities of Imperfect Repeater Quantum Networks

The work in this chapter forms the basis of a paper of the same name which has been published in *Quantum Science and Technology*, whose authors are (in order) Cillian Harney and Stefano Pirandola [2]. This chapter proceeds as follows: Section 4.1 details some preliminary notions and motivations. In Section 4.2 we present general end-to-end capacity bounds for quantum networks with arbitrary topologies. In particular, we present achievable end-to-end rates which are universal for any quantum network, and detail how upper-bounds can be found when single-edge capacity bounds are known. Furthermore we describe a node-splitting technique and discuss how these tools can be used to capture critical properties of quantum network architectures. In Section 4.3 we employ these new theoretical tools in the context of WRNs, elucidating a network benchmarking procedure via end-to-end capacity bounds and WR architectures. These methods and results are then applied in the context of qubit amplitude-damping and bosonic thermal-loss WRNs in Sections 4.4 and 4.5 respectively. Finally, Section 4.6 provides concluding remarks and future directions of study.

4.1 Introduction

Many of the underlying challenges associated with quantum communications and networking propagate from a salient fact: the ability to transmit, detect and preserve quantum information is remarkably more difficult than that of classical information. Classical information is robust, can be copied and stored reliably. On the other hand, quantum information is inherently fragile and the laws of quantum mechanics prohibit its cloning [39]. Hence, quantum networks face many formidable obstacles from which their classical

counterparts are spared. These can be attributed to two key regimes: inevitable *external* decoherence experienced by quantum systems along communication channels and *internal* decoherence due to imperfect devices attempting to preserve or operate on quantum systems. The impact of external decoherence on the performance of quantum communication networks is understood via channel capacities, and through the appropriate modelling of communication conduits between users.

As explored in the previous section, one can accurately benchmark the optimal performance of optical fibre networks via a rate distribution defined according to the PLOB bound, i.e. given a network $\mathcal{N} = (P, E)$ each point-to-point link of length $d_{\mathbf{x}\mathbf{y}}$ has capacity $\mathcal{C}_{\mathbf{x}\mathbf{y}} = -\log_2(1 - 10^{-\gamma d_{\mathbf{x}\mathbf{y}}})$ where γ is the fibre loss rate. This is a sufficient description of the external decoherence experienced in an optical fibre network. But unfortunately, there will always exist unavoidable internal decoherence that quantum systems will experience at each network node. This may be due to sub-optimal detection or transmission, imperfect quantum memories, electronic and environmental noise, or any other inevitable practical imperfection. The nature of this decoherence depends on the type of quantum system being used (CV or DV), the protocol being employed and the realistic technologies available to the network. To properly assess the end-to-end performance limits of quantum networks it is necessary to integrate internal decoherence into the network description.

In this chapter, we provide a framework for *bounding* the end-to-end flooding capacities of noisy-repeater networks. Recently, the performance limits of quantum networks with imperfect repeaters have been studied using a *node-splitting* technique; each channel in the network is split into compound channels which incorporate repeater imperfections via additional internal channels [74]. This work focussed on the impact of internal loss and the class of distillable channels [29], but it is possible to extend these results into a more general setting. We begin by translating the coherent information (CI) and reverse coherent information (RCI) from point-to-point achievable quantum communication rates into lower-bounds on end-to-end network capacities. The resulting bounds are universal, regardless of network topology or channel composition. Combined with relevant upper-bounds, we are able to apply a node-splitting technique which accounts for both internal loss and noise. In particular, we unveil realistic performance limits and infrastructure requirements of networks composed of channels whose capacities are not exactly known.

In order to evaluate end-to-end capacity bounds and utilise the node splitting procedure, we need to investigate suitable network architectures. It is an open question as to how quantum networks should be best constructed on mid-to-large scales in order to balance high rates with cost-efficient resources. Random network architectures (such as Waxman, Erdős-Rényi and scale-free networks) are useful but rely heavily upon numerical

treatments. On the other hand, analytical treatments of quantum repeater networks have been mostly limited to linear networks. Linear networks (or repeater chains) are effective for studying extended point-to-point communications, but are too simplistic to model large, inter-connected structures. To address a common ground between complex architectures and simplified repeater chains we employ WRNs as described in Chapter 3. Crucially, WRNs admit an analytical form that allow us to derive *threshold theorems*; theorems which reveal threshold single-edge capacities \mathcal{C}_{\min} that guarantee specific end-to-end performance bounds. Chapter 3 introduced these threshold theorems and applied them to bosonic lossy WRNs. In this work, we generalise threshold theorems to study the optimal performance of WRNs constituted of any quantum channel, even if their exact capacity is not known. Using the end-to-end capacity bounds and the technique of node splitting, we are able to derive bounds on critical parameters such as maximum channel length, maximum internal loss or maximum thermal noise which guarantee optimal performance.

4.1.1 Optimising the Physical Orientation of Realistic Quantum Networks

In Section 2.3.4 we discussed the nuance concerning with how quantum networks can be treated as undirected graphs. Let us solidify this concept, as it is of importance throughout the arguments in this chapter. Under the assistance of two-way classical communications (CCs), the optimal transmission of quantum information is connected with optimal entanglement distribution, not on the direction of physical system exchange but on the LOs that are applied at each point. The physical orientation of any channel $\mathcal{E}_{\mathbf{x}\mathbf{y}}$ in the network can be forwards or backwards and still facilitate communication in either logical direction. If a channel is physically symmetric then the physical forward and backward channels between two network nodes $\mathbf{x}, \mathbf{y} \in P$ are identical $\mathcal{E}_{\mathbf{x}\mathbf{y}} = \mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}} = \mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}}$. Networks built up of physically symmetric channels are easy to represent on an undirected graph as the physical channel direction is then completely irrelevant.

But we must take particular care when there exist channels that are *physically asymmetric*, i.e. the physical forward and backward channels are not identical, $\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}} \neq \mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}}$. In this case, a quantum network can be described as a collection of forward and backward channels associated with each undirected edge, $\{(\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}}, \mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}})\}_{(\mathbf{x}, \mathbf{y}) \in E}$. One may assume that it is sensible to convert the undirected network graph into a directed graph to account for the unique physically directed channels. Yet, the logical flow of information remains independent from physical flow of quantum systems, as it is always possible to invoke a teleportation protocol in the opposite direction to physical quantum system exchange. But if a channel is physically asymmetric there will be an imbalance between the capacity of

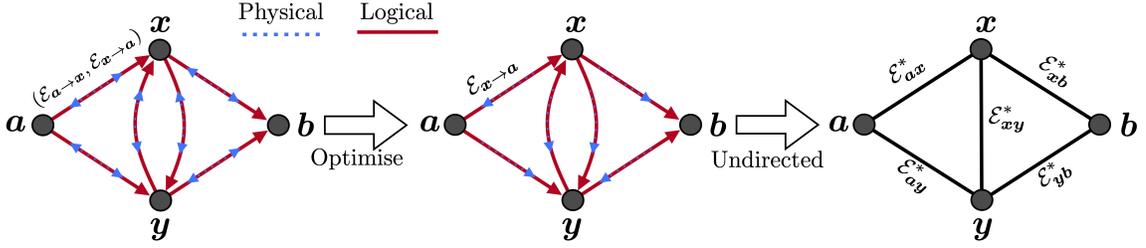


Figure 4.1: Optimising the physical orientation of a quantum network. The physical flow of quantum systems is independent of the logical flow of information, and thus physical directions can always be optimised to utilise the directed channels which possess the greatest capacity. In this way, a quantum network described by a directed graph can always be reduced to an undirected form. From a pair of physically-directed channels on a single edge $(\mathcal{E}_{a \rightarrow x}, \mathcal{E}_{x \rightarrow a})$, we can choose that which has the greatest capacity according to Eq. (4.1), e.g. $\mathcal{E}_{x \rightarrow a}$ in the example above. It is always possible to use this superior edge for logical communication in either direction, reducing it to an optimal undirected edge describing the channel $\mathcal{E}_{ax}^* = \mathcal{E}_{x \rightarrow a}$.

the forward and backward channels, since they are generally unique. As a result, there will exist an optimal physical direction in which quantum systems should be exchanged.

This identifies a crucial optimisation step of the physical orientation of a quantum network, and corresponds to fixing an optimal physical direction to all channels within the network. Consider an arbitrary edge $(x, y) \in E$ and its associated pair of physical forward and backward channels $(\mathcal{E}_{x \rightarrow y}, \mathcal{E}_{y \rightarrow x})$. The directed channel pair should be mapped to the channel which maximises the point-to-point capacity,

$$(\mathcal{E}_{x \rightarrow y}, \mathcal{E}_{y \rightarrow x}) \mapsto \mathcal{E}_{xy}^* := \arg \max_{\mathcal{E} \in \{\mathcal{E}_{x \rightarrow y}, \mathcal{E}_{y \rightarrow x}\}} \mathcal{C}(\mathcal{E}), \quad (4.1)$$

where $\mathcal{C}(\mathcal{E})$ is the capacity of the channel \mathcal{E} . By performing this optimisation for all $(x, y) \in E$, the network can be represented as an undirected graph $\mathcal{N} = (P, E)$ interconnected by the optimal physically directed set of channels $\{\mathcal{E}_{xy}^*\}_{(x,y) \in E}$ for quantum communication throughout the network. Figure. 4.1 illustrates this optimisation process. Throughout this chapter, we will denote channels using this notation, implying the application of a physical orientation pre-optimisation step.

4.2 Bounds for Realistic Quantum Networks

Computing exact network capacities requires knowledge of the exact single-edge capacities of all channels in the network, $\mathcal{C}(\mathcal{E}_{xy}^*)$. For some important classes of quantum channels, exact capacities are known; such as the class of distillable quantum channels,

which include the bosonic lossy channel, quantum-limited amplifiers and dephasing channel [29]. As a consequence, their end-to-end network capacities have been fully characterised for arbitrary topologies [32]. In general, this is not the case and the capacities of many quantum channels remain undetermined. Hence, we rely on upper and lower-bounds in order to understand their efficacy for quantum communication. In this section we extend point-to-point channel capacity bounds into end-to-end capacity bounds to characterise the performance of general quantum networks.

4.2.1 Bounding End-to-End Network Capacities

Consider a quantum network $\mathcal{N} = (P, E)$ with an optimal physical orientation of physical channels $\{\mathcal{E}_{\mathbf{x}\mathbf{y}}^*\}_{(\mathbf{x}, \mathbf{y}) \in E}$. As per Section 2.2.3, we know that the maximised RCI/CI $I(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*)$ is an achievable rate for any channel in the network, it is possible to compute end-to-end capacity lower-bounds by supplementing it into Eq. (2.69). For a routing protocol \mathcal{P} we can define a network generalisation of the CI, such that

$$I(\mathbf{i}, \mathcal{N} | \mathcal{P}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} q_{\mathbf{x}\mathbf{y}} I(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*), \quad (4.2)$$

where $q_{\mathbf{x}\mathbf{y}} \in \mathcal{Q}_{\mathcal{P}}$ is an element of the forwarding probability distribution invoked by the routing protocol. This represents an achievable end-to-end rate of a quantum network with arbitrary topology and channel composition. One can then lower-bound the optimal performance of a quantum network (flooding capacity) via a flooding-defined network CI,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \geq I(\mathbf{i}, \mathcal{N}) = I(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{fl}}). \quad (4.3)$$

When all the channels $\mathcal{E}_{\mathbf{x}\mathbf{y}}^*$ considered within the network structure are distillable then these achievable network capacities become the exact capacities.

Analogously, it is always possible to derive end-to-end capacity upper-bounds by supplementing upper-bounding quantities into the rate distribution. Let $F(\mathcal{E})$ be a function which computes an upper-bound on the capacity of a channel \mathcal{E} , i.e. $\mathcal{C}(\mathcal{E}) \leq F(\mathcal{E})$. Then in the network setting we can write

$$F(\mathbf{i}, \mathcal{N} | \mathcal{P}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} q_{\mathbf{x}\mathbf{y}} F(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*). \quad (4.4)$$

Once more, we can then present an upper-bound on the flooding capacity, using the notation

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq F(\mathbf{i}, \mathcal{N}) = F(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{fl}}). \quad (4.5)$$

Indeed, for teleportation-covariant channels, even when they are not distillable, it is possible to write an upper-bound on their capacity using the REE of their Choi matrix. Hence, one could write the following upper-bound for networks composed of non-distillable teleportation-covariant channels,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}|\mathcal{P}) \leq E_R(\mathbf{i}, \mathcal{N}|\mathcal{P}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} q_{\mathbf{x}\mathbf{y}} E_R(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*). \quad (4.6)$$

The REE can then be used to write end-to-end capacity bounds for networks consistent of bosonic thermal-loss channels, Pauli channels, and more.

4.2.2 Node Splitting

Ideally, quantum repeaters are completely lossless, noiseless, and fully error-corrected. Indeed, Ref. [32] derives end-to-end quantum network capacities under the assumption of perfect repeaters, only considering the unavoidable decoherence due to quantum channels which are external to the repeater devices. In reality, there exist a number of internal loss/noise contributions that should be considered. A first step in this direction was to consider the presence of loss within quantum repeaters due to sub-optimal detection efficiency, channel-memory coupling losses, memory loading and readout [74]. In this chapter, we present the general scenario where repeaters are affected by both loss and noise, incorporating electronic and environmental noise affects that may occur.

Consider a single quantum repeater contained within a quantum network, $\mathbf{x} \in P$. To account for internal imperfections, we may perform *node splitting* [74]. A repeater node $\mathbf{x} \in P$ can be split into a trio of internal nodes $\mathbf{x} \rightarrow \{\mathbf{x}^r, \mathbf{x}^u, \mathbf{x}^s\}$: a receiver node \mathbf{x}^r , a user node \mathbf{x}^u , and a sender node \mathbf{x}^s . The user node \mathbf{x}^u represents the only valid node from which communication can originate or end, or where a user can actually be situated. Through node splitting, we can represent decoherence effects due to imperfect transmission and reception via additional quantum channels between the internal nodes. Internal noise and loss due to imperfect reception and storage of quantum information at the node \mathbf{x} can be described by a quantum channel physically directed from the receiver node to the user node $\mathcal{E}_{\mathbf{x}^r \rightarrow \mathbf{x}^u}$. Similarly, imperfections in the memory loading and transmission process can be captured via an internal channel between the user node and the sender node, $\mathcal{E}_{\mathbf{x}^u \rightarrow \mathbf{x}^s}$. Each internal channel will possess unique properties according to the technologies used throughout the network.

We may consider communication between two repeater nodes $\mathbf{x} = \{\mathbf{x}^r, \mathbf{x}^u, \mathbf{x}^s\}$ and $\mathbf{y} = \{\mathbf{y}^r, \mathbf{y}^u, \mathbf{y}^s\}$ within a quantum network \mathcal{N} . We do not assume the precise nature of the external or internal channels, focussing first on a general picture. Since the internal

imperfections at each node in general are unique, then the physically directed channels between \mathbf{x} and \mathbf{y} will be asymmetric. For quantum communication in the physical direction $\mathbf{x} \rightarrow \mathbf{y}$ the complete channel is a compound channel given by

$$\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}} = \mathcal{E}_{\mathbf{y}^r \rightarrow \mathbf{y}^u} \circ \mathcal{E}_{\mathbf{x}^s \rightarrow \mathbf{y}^r} \circ \mathcal{E}_{\mathbf{x}^u \rightarrow \mathbf{x}^s}, \quad (4.7)$$

accounting for each channel between the user node in $\mathbf{x}^u \in \mathbf{x}$ and the user node in $\mathbf{y}^u \in \mathbf{y}$. For the physical exchange of quantum systems in the opposite direction, the compound channel reads

$$\mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}} = \mathcal{E}_{\mathbf{x}^r \rightarrow \mathbf{x}^u} \circ \mathcal{E}_{\mathbf{y}^s \rightarrow \mathbf{x}^r} \circ \mathcal{E}_{\mathbf{y}^u \rightarrow \mathbf{y}^s}. \quad (4.8)$$

Applying node splitting to every edge in the network $(\mathbf{x}, \mathbf{y}) \in E$, and selecting the best physically directed channel for each edge in the network, we can then retrieve the optimal physical orientation of a realistic quantum network composed of imperfect repeaters, $\{\mathcal{E}_{\mathbf{x}\mathbf{y}}^*\}_{(\mathbf{x}, \mathbf{y}) \in E}$.

4.3 Network Parameter Benchmarking with Weakly-Regular Networks

The invaluable mathematical tool within this chapter is the theory of threshold theorems for WRNs. These are theorems which utilise the connectivity properties of $(k, \mathbf{\Lambda})$ -WRNs (as introduced and defined in Chapter 3, Section 3.2) in order to derive *single-edge threshold conditions* necessary to guarantee strong end-to-end communication performance. These single-edge threshold conditions define upper or lower bounds on physical properties of any single node or channels within the network, e.g. channel length, thermal noise at a receiver node, internal loss, and more. Such threshold values help us to understand the resilience of quantum networks given some desirable level of performance, helping to benchmark the requirements of realistic quantum technologies necessary to perform at high rates.

4.3.1 Threshold Theorems for Network Parameters

In Chapter 3 we state and prove threshold theorems with respect to exact single-edge channel capacities. This is the same as answering the following question: given an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, what is the minimum single-edge capacity that is allowed in a quantum WRN to ensure that the end-to-end flooding capacity is optimally achieved by the min-neighbourhood capacity, $\mathcal{C}_{\mathcal{N}_i}$? Theorem 3.1 solved this question, and helps us to derive a global single-edge capacity constraint in order to guarantee the performance bounds in Eq. (3.24). If we wish to guarantee exactly optimal performance, we must place an additional constraint on user-connected edges, which is solved in Theorem 3.2.

These capacity-defined threshold theorems provide valuable tools for understanding single-edge capacity requirements for WRNs so that performance bounds or optimal performance can be guaranteed. The threshold capacities derived are useful, but it is even more useful to identify a relationship between end-to-end network performance and *physical properties* of the network nodes and channels. Hence, our goal is to translate these abstract threshold theorems into tangible relationships between physical channel parameters and end-to-end performance, which we address with the following corollary:

Corollary 4.1 *Consider a (k, Λ) -WR quantum network $\mathcal{N} = (P, E)$, an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ and a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. Consider a single-edge channel property $\xi_{\mathbf{x}\mathbf{y}}$ for which the point-to-point capacity $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ is monotonic. Then, if $\mathcal{C}_{\mathcal{N}_i}$ is attainable, there exists a threshold parameter $\xi_{\mathcal{N}}^*$ such that*

$$\xi_{\mathcal{N}}^* := \arg \min_{\xi} \left| \mathcal{C}(\xi) - \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta} \right|, \quad (4.9)$$

which represents a maximum or minimum tolerable value of $\xi_{\mathbf{x}\mathbf{y}}$ for any channel in the network:

$$\xi_{\mathcal{N}}^* = \begin{cases} \xi_{\mathcal{N}}^{\max}, & \mathcal{C}(\xi) \text{ is decreasing,} \\ \xi_{\mathcal{N}}^{\min}, & \mathcal{C}(\xi) \text{ is increasing.} \end{cases} \quad (4.10)$$

If $\xi_{\mathcal{N}}^*$ is obeyed for all $(\mathbf{x}, \mathbf{y}) \in E$ then the flooding capacity is guaranteed to satisfy

$$\frac{2(k-1)}{\delta} \mathcal{C}_{\mathcal{N}_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{\mathcal{N}_i}. \quad (4.11)$$

Proof. Via Theorem 3.1, we know that there exists a threshold capacity $\mathcal{C}_{\min} = \mathcal{C}_{\mathcal{N}_i}/\delta$ which when respected throughout the network ensures that the performance bounds in Eq. (3.24) hold. Now, consider a physical property of a single network edge $(\mathbf{x}, \mathbf{y}) \in E$ denoted by $\xi \rightarrow \xi_{\mathbf{x}\mathbf{y}}$, e.g. channel length. Suppose that the single-edge capacity is a monotonic function of the single-edge property $\xi_{\mathbf{x}\mathbf{y}}$. Then the threshold capacity \mathcal{C}_{\min} can be translated into a threshold condition on $\xi_{\mathbf{x}\mathbf{y}}$, since we can write

$$\mathcal{C}_{\mathbf{x}\mathbf{y}} = \mathcal{C}(\xi_{\mathbf{x}\mathbf{y}}) \geq \mathcal{C}_{\min}. \quad (4.12)$$

Therefore, there must exist a critical threshold parameter $\xi = \xi_{\mathcal{N}}^*$ for which the single-edge threshold capacity is exactly satisfied,

$$\mathcal{C}(\xi_{\mathcal{N}}^*) = \mathcal{C}_{\min}. \quad (4.13)$$

The quantity $\xi_{\mathcal{N}}^*$ thus represents some limiting feature of each network edge necessary to uphold the optimal performance bounds.

While we may not know exactly what form the single-edge capacity function takes, we can still determine the threshold value $\xi_{\mathcal{N}}^*$ as the value of ξ which satisfies Eq. (4.13). Yet, we must be slightly careful here. Let us define our capacity function more formally. The single-edge capacity $\mathcal{C}(\xi)$ is a function which maps single-edge network parameters ξ from a domain \mathcal{X} of possible values $\xi \in \mathcal{X}$, to a codomain \mathcal{Y} of potential values $\mathcal{C}(\xi) \in \mathcal{Y}$. This codomain is necessarily a subset of the set of non-negative real numbers \mathbb{R}_0^+ so that the outputs of \mathcal{C} represent meaningful capacity values. More precisely,

$$\mathcal{C} : \mathcal{X} \rightarrow \mathcal{Y} \subseteq \mathbb{R}_0^+. \quad (4.14)$$

We are only interested in capacity functions \mathcal{C} which are monotonic with respect to ξ so this is necessarily a one-to-one correspondence.

We may then state the following: consider a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$ and a single-edge capacity function $\mathcal{C}(\xi)$ which is monotonic with respect to the network parameter ξ . A threshold parameter value $\xi_{\mathcal{N}}^*$ will exist if and only if $\mathcal{C}_{\mathcal{N}_i}/\delta$ falls within the codomain \mathcal{Y} . Otherwise, there will not exist a physical value of ξ for which $\mathcal{C}_{\mathcal{N}_i}$ is attainable. With these considerations in mind, we can state that a network threshold parameter $\xi_{\mathcal{N}}^*$ can be appropriately computed such that

$$\xi_{\mathcal{N}}^* = \arg \min_{\xi} \left| \mathcal{C}(\xi) - \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta} \right| \iff \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta} \in \mathcal{Y}. \quad (4.15)$$

If $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ is a monotonically decreasing function, then $\xi_{\mathcal{N}}^*$ must be a *maximum threshold value* $\xi_{\mathcal{N}}^{\max}$ because increasing it further would decrease the capacity below the threshold, i.e. $\mathcal{C}(\xi_{\mathcal{N}}^* + \varepsilon) < \mathcal{C}_{\min}$, where $\varepsilon > 0$. If $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ is a monotonically increasing function, the opposite is true and $\xi_{\mathcal{N}}^*$ must be a *minimum threshold value* $\xi_{\mathcal{N}}^{\min}$ because decreasing it further would reduce the capacity below the threshold, i.e. $\mathcal{C}(\xi_{\mathcal{N}}^* - \varepsilon) < \mathcal{C}_{\min}$ where $\varepsilon > 0$. This completes the result. ■

Analogously, we can translate Theorem 3.2 with respect to a tangible, threshold network parameter. The translation is immediate, such that the single threshold parameter in Corollary 4.1 is simply extended to a pair of parameters for each condition (see Appendix B for a full description). Corollary 4.1 therefore identifies maximum or minimum threshold parameters which when respected throughout the network are able to guarantee tight performance bounds (or extended to optimal performance). An interesting threshold parameter might be the maximum link-length, the maximum thermal noise that can be

tolerated at a receiver, etc. These theorems reveal extremely useful relationships between the end-to-end rate and physical properties of interest; providing invaluable guidance for future quantum network design.

4.3.2 Threshold Theorems with Capacity Bounds

Up to this point we have been deriving exact threshold quantities under the assumption that an expression is known for the point-to-point capacity $\mathcal{C}(\xi)$ with respect to some physical property ξ . Yet, when the exact nature of a quantum channel capacity is not known it is necessary to make use of capacity bounds, as have been explored in this chapter. Consider a single-edge channel property ξ for which the point-to-point capacity $\mathcal{C}(\xi)$ is monotonically bounded by lower and upper bounding functions $F_j \in \{F_l, F_u\}$ respectively. We introduce the following “cost function” to evaluate the difference between a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$ and a capacity bound,

$$\mathcal{B}_{F_j}(\xi, x) := \left| F_j(\xi) - \frac{\mathcal{C}_{\mathcal{N}_i}}{x} \right|. \quad (4.16)$$

Here x is a free parameter used to scale the capacity bound in accordance with the network connectivity properties. The goal of constructing this cost function is to ensure that $\mathcal{B}_{F_j}(\xi, x)$ is minimised (tends to zero) when $\xi \rightarrow \xi_{\mathcal{N}}^*$. Clearly, we may define this function in many different ways provided that this behaviour is maintained. In our numerical studies, we choose to minimise the log-ratio instead,

$$\mathcal{B}_{F_j}(\xi, x) = |\log(xF_j(\xi)) - \log(\mathcal{C}_{\mathcal{N}_i})|. \quad (4.17)$$

Given a cost function, the goal is to determine bounds on the threshold value $\xi_{\mathcal{N}}^*$ which helps to guarantee optimal performance bounds. Hence, we define the quantity

$$\xi_{F_j}^*(x) := \arg \min_{\xi} \mathcal{B}_{F_j}(\xi, x), \quad (4.18)$$

where $j \in \{l, u\}$, which presents a lower or upper an upper or lower bound on $\xi_{\mathcal{N}}^*$.

Thanks to single-edge capacity bounds, we can generalise Corollary 4.1 to derive conditions for end-to-end performance on networks consisting of quantum channels whose exact capacities are not known. This results in the following corollary.

Corollary 4.2 *Consider Corollary 4.1 and a single-edge channel property $\xi_{\mathbf{x}\mathbf{y}}$ for which the point-to-point capacity $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ is monotonically bounded by lower and upper bounding functions $F \in \{F_l, F_u\}$ respectively. The threshold parameter $\xi_{\mathcal{N}}^*$ satisfies*

$$\xi_{F_j}^*(\delta) \leq \xi_{\mathcal{N}}^* \leq \xi_{F_k}^*(\delta), \quad (4.19)$$

where $j \neq k \in \{l, u\}$ such that

$$\begin{cases} \xi_{F_u}^{\max}(\delta) \leq \xi_{\mathcal{N}}^{\max} \leq \xi_{F_l}^{\max}(\delta), & \mathcal{C}(\xi) \text{ is decreasing,} \\ \xi_{F_l}^{\min}(\delta) \leq \xi_{\mathcal{N}}^{\min} \leq \xi_{F_u}^{\min}(\delta), & \mathcal{C}(\xi) \text{ is increasing.} \end{cases} \quad (4.20)$$

Proof. We may not know $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ exactly, but instead have a pair of single-edge bounding functions F_u and F_l which similarly depend on the same single-edge parameter

$$F_l(\xi_{\mathbf{x}\mathbf{y}}) \leq \mathcal{C}(\xi_{\mathbf{x}\mathbf{y}}) \leq F_u(\xi_{\mathbf{x}\mathbf{y}}). \quad (4.21)$$

When this is the case, it's not possible to determine the network threshold parameter $\xi_{\mathcal{N}}^*$ exactly. Instead, we can provide upper and lower bounds on its value. We can always demand that

$$F_u(\xi_{\mathbf{x}\mathbf{y}}) \geq \mathcal{C}(\xi_{\mathbf{x}\mathbf{y}}) \geq F_l(\xi_{\mathbf{x}\mathbf{y}}) \geq \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta}, \quad (4.22)$$

be satisfied for all $(\mathbf{x}, \mathbf{y}) \in E$. Given that F_u and F_l are both monotonic with respect to $\xi_{\mathbf{x}\mathbf{y}}$, by independently solving the following pair of equations

$$F_l(\xi) = \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta}, \text{ and } F_u(\xi) = \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta}, \quad (4.23)$$

there will exist a pair of unique values $\xi_{F_u}^*$ and $\xi_{F_l}^*$ which appropriately bound the threshold parameter.

Once again it is important to be careful, and so let us define our capacity bounding functions more formally. The function F_k (for $k \in \{u, l\}$) is that which maps single-edge network parameters ξ from a domain \mathcal{X} of possible values $\xi \in \mathcal{X}$, to a codomain \mathcal{Y}_{F_k} of potential bounding values $F_k(\xi) \in \mathcal{Y}_{F_k}$. The meaningful part of this codomain is necessarily a subset of the set of non-negative real numbers \mathbb{R}_0^+ so that the outputs of F_k represent valid capacity values (bounding functions may have less well behaved codomains, but this is always rectifiable by truncating such regions). More precisely,

$$F_k : \mathcal{X} \rightarrow \mathcal{Y}_{F_k} \subseteq \mathbb{R}_0^+. \quad (4.24)$$

Since we are only interested in capacity bounding functions $F_k(\xi)$ which are monotonic with respect to ξ , then this is necessarily a one-to-one correspondence. Consider a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$ and an appropriate single-edge capacity bounding function $F_k(\xi)$ which is monotonic with respect to the network parameter ξ . A network threshold parameter bounding value $\xi_{F_k}^*(\delta)$ for $k \in \{u, l\}$ will exist if and only if $\mathcal{C}_{\mathcal{N}_i}/\delta$ falls within

in the codomain \mathcal{Y}_{F_k} . Otherwise, there will not exist a physical value of ξ for which $\mathcal{C}_{\mathcal{N}_i}$ is attainable. Therefore, a network threshold parameter $\xi_{\mathcal{N}}^*$ is bounded according to $\xi_{F_u}^*$ and $\xi_{F_l}^*$ such that

$$\xi_{F_k}^* := \arg \min_{\xi} \mathcal{B}_{F_k}(\xi, \delta) \iff \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta} \in \mathcal{Y}_{F_k}, \quad (4.25)$$

where $k \in \{u, l\}$ is used to indicate whether it is derived using the upper or lower bounding function F_k respectively.

As a result, for physical single-edge properties ξ for which the capacity function is monotonic, there exists a critical threshold parameter in the network $\xi_{\mathcal{N}}^*$ for which the capacity conditions in Theorem 3.1 are satisfied, and thus the end-to-end capacity is guaranteed to satisfy Eq. (4.11). Even when the capacity function is not exactly known, we can instead use upper and lower capacity bounding functions F_u and F_l to appropriately bound $\xi_{\mathcal{N}}^*$. The nature of these bounds depends on the increasing or decreasing nature of the monotonic bounding functions as shown in the corollary. ■

Of course, an analogous result can be presented for exactly guaranteed optimal performance (see Appendix B). For the networks, capacity bounds and critical parameters studied in this Chapter, monotonicity was satisfied; allowing us to exploit these results. However, it is still possible to glean valuable information about threshold parameters when the capacity function (or its bounding functions) is not monotonic with respect to a single-edge network property ξ_{xy} . When this is the case, Eq. (4.23) will not have unique solutions, but there will exist a number of suitable values of ξ for each bounding function. This gives rise to *threshold ranges* of values of ξ for which the flooding capacity will have performance guarantees. This extension is intuitive, and it is left for future studies wherever it is physically relevant.

Using the tools of the previous sections we can now benchmark the optimal performance of quantum networks with imperfect repeaters. In the following sections, we incorporate realistic, noisy channels throughout the network for which the exact capacities are not known. As example applications, it allows us to investigate the end-to-end rates of lesser studied network models such as amplitude-damping networks and bosonic thermal-loss networks.

4.4 Amplitude Damping Networks

4.4.1 Network Model

We begin with networks consistent of amplitude damping (AD) quantum channels. AD channels are qubit channels that describe the process of energy dissipation through

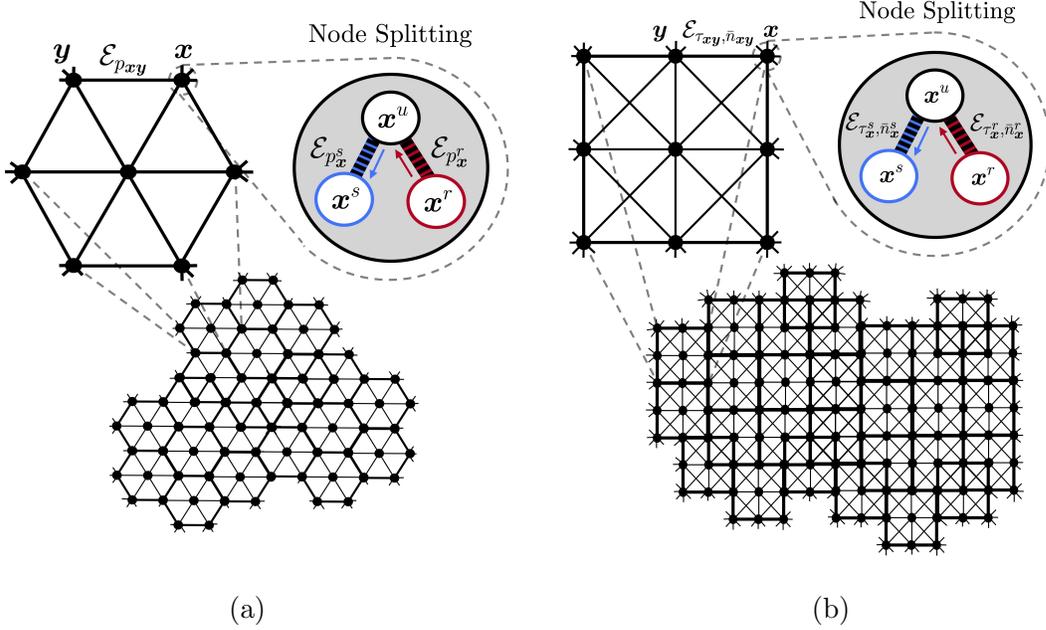


Figure 4.2: (a) Node-splitting procedure for quantum amplitude-damping networks. We consider quantum WRNs within some nodal boundary which satisfy regularity $k = 6$ and adjacent commonality $\lambda_x = \{2\}^{\cup 6}$ for any node $x \in P$. (b) Node-splitting procedure for bosonic thermal-loss quantum networks. We consider quantum WRNs within some nodal boundary which satisfy regularity $k = 8$ and adjacent commonality $\lambda_x = \{2, 4\}^{\cup 4}$ for any node $x \in P$.

spontaneous emission, representing the analogue of a bosonic lossy channel constrained to a two-level system. They are ubiquitous in the modelling of many important physical processes in communications and computation. For damping-probability $p \in (0, 1)$, the AD channel can be defined via the Kraus decomposition,

$$\mathcal{E}_p(\rho) = \sum_{i=0}^1 K_i \rho K_i^\dagger, \quad (4.26)$$

$$K_0 := |0\rangle\langle 0| + \sqrt{1-p}|1\rangle\langle 1|, \quad K_1 := \sqrt{p}|0\rangle\langle 1|. \quad (4.27)$$

Interestingly, the AD channel is not distillable and so its exact capacity is unknown.

Let us consider an arbitrary quantum network composed of AD channels, \mathcal{N}_{AD} . For communication between two network nodes $\mathbf{x} = \{x^r, x^u, x^s\}$ and $\mathbf{y} = \{y^r, y^u, y^s\}$, we can describe each of the internal and external channels as AD channels with a unique damping probability. More precisely, let any external channel between two nodes \mathbf{x} and \mathbf{y} be AD channels with the damping probability $p_{\mathbf{x}\mathbf{y}} = p_{x^s \rightarrow y^r} = p_{y^s \rightarrow x^r}$. The internal channels

can also be considered AD channels with fixed physical directions,

$$\mathcal{E}_{i^r \rightarrow i^u} = \mathcal{E}_{p_i^r}, \quad \mathcal{E}_{i^u \rightarrow i^s} = \mathcal{E}_{p_i^s}, \quad \mathbf{i} \in \{\mathbf{x}, \mathbf{y}\}. \quad (4.28)$$

For the physical exchange of quantum systems in either direction, the complete compound channels read

$$\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}} = \mathcal{E}_{p_{\mathbf{y}}^r} \circ \mathcal{E}_{p_{\mathbf{x}\mathbf{y}}} \circ \mathcal{E}_{p_{\mathbf{x}}^s}, \quad (4.29)$$

$$\mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}} = \mathcal{E}_{p_{\mathbf{x}}^r} \circ \mathcal{E}_{p_{\mathbf{x}\mathbf{y}}} \circ \mathcal{E}_{p_{\mathbf{y}}^s}. \quad (4.30)$$

For any quantum network topology, we can use these compound channels and their capacities to assign an optimal physical orientation and regain an undirected graphical representation. To do this, the RCI places an achievable lower-bound on the single-edge capacity of each compound channel throughout the network. For AD channels, this results in the following achievable rate for each network edge [28],

$$I(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*) = \max_{\mathbf{i} \neq \mathbf{j} \in \{\mathbf{x}, \mathbf{y}\}} \max_u [H_2(u) - H_2(up_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{tot}})]. \quad (4.31)$$

Here, $H_2(u) := -u \log_2(u) - (1-u) \log_2(1-u)$ defines the binary Shannon entropy function, and $p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{tot}}$ is the total damping probability associated with a compound channel directed from node \mathbf{i} to \mathbf{j} ,

$$p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{tot}} := 1 - (1 - p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{int}})(1 - p_{\mathbf{i}\mathbf{j}}), \quad (4.32)$$

$$= 1 - (1 - p_{\mathbf{j}}^r)(1 - p_{\mathbf{i}\mathbf{j}})(1 - p_{\mathbf{i}}^s), \quad (4.33)$$

where we simultaneously define an effective internal loss parameter $p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{int}} := p_{\mathbf{i}}^r(1 - p_{\mathbf{j}}^s) + p_{\mathbf{j}}^s$ which captures both sending and receiving inefficiencies. Eq. (4.31) can be substituted into the single-path and multi-path expressions in Eq. (4.2) to compute achievable network rates.

Upper-bounds on the single-edge capacity of AD channels can also be expanded to provide upper-bounds on the end-to-end capacities. The best known single-edge bound is given by the squashed entanglement [29],

$$E_{\text{sq}}(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*) = \max_{\mathbf{i} \neq \mathbf{j} \in \{\mathbf{x}, \mathbf{y}\}} H_2\left(\frac{1}{2} - \frac{p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{tot}}}{4}\right) - H_2\left(1 - \frac{p_{\mathbf{i} \rightarrow \mathbf{j}}^{\text{tot}}}{4}\right). \quad (4.34)$$

This can then be used within Eq. (4.4) to compute upper-bounds on the ultimate limits of quantum communications over AD networks.

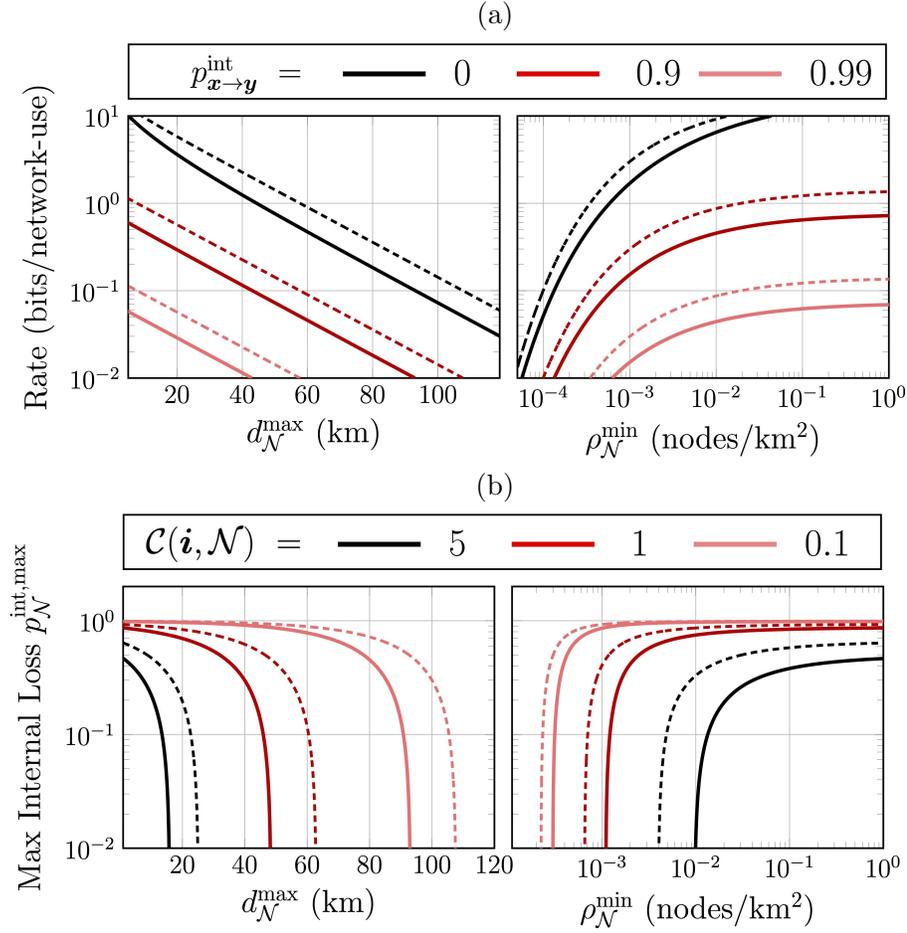


Figure 4.3: Throughout all plots, dashed lines represent bounds obtained using the squashed entanglement as a single-edge capacity upper-bound. Meanwhile, all solid lines plot bounds derived using the RCI as a single-edge capacity lower-bound. Panel (a) plots bounds on the maximum fibre-length permitted within the network $d_{\mathcal{N}}^{\text{max}}$ and a corresponding minimum nodal density $\rho_{\mathcal{N}}^{\text{min}}$ which ensure that an end-user pair can obtain optimal performance $\mathcal{C}(i, \mathcal{N})$. Panel (b) depicts bounds on the maximum tolerable internal loss $p_{\mathcal{N}}^{\text{int,max}}$ permitted within the network with respect to maximum fibre-length and nodal density so to guarantee an optimal flooding capacity.

4.4.2 Benchmarking

To make use of these bounds and the tools of Section 4.3, we can investigate the end-to-end performance of weakly-regular AD networks. In this section, we consider a $k = 6$ weakly-regular network, which adopts equivalent connectivity properties to a triangular lattice. In Fig. 4.2(a) a single cell of this architecture is illustrated, and it is shown how a $k = 8$ WRN can be constructed. We investigate large networks consisting of many of these

cells connected together, and consider end-users which are deeply-embedded in the network so that any network boundary effects can be ignored¹.

All edges $(\mathbf{x}, \mathbf{y}) \in E$ in the network are modelled by AD channels of length $d_{\mathbf{x}\mathbf{y}}$, used to describe optical-fibre. As such, the damping-probability of each edge is given by

$$p_{\mathbf{x}\mathbf{y}} = 1 - 10^{-\gamma d_{\mathbf{x}\mathbf{y}}}, \quad (4.35)$$

where $\gamma = 0.02$ per km is the state-of-the-art loss-rate for fibre (0.2 dB per km). Node-splitting is applied throughout the network in order to incorporate internal loss, as depicted in Fig. 4.2(a). By specifying a threshold theorem to the property of channel length, we are able to place tight bounds on the *maximum tolerable fibre-length* permitted within the network, $d_{\mathcal{N}}^{\max}$, with respect to a desired flooding capacity. That is, for a given flooding capacity between a pair of end-users $\mathcal{C}(\mathbf{i}, \mathcal{N})$ and fixed internal losses $p_{\mathbf{x} \rightarrow \mathbf{y}}^{\text{int}}$, we plot the maximum fibre-length that is allowed within the network-bulk so that we can guarantee the flooding capacity is optimal.

Thanks to the consistent, analytical connectivity properties of WRNs, the maximum link-length can also be used to derive a *minimum nodal density*, $\rho_{\mathcal{N}}^{\min}$. Given a WRN with consistent connectivity rules and a maximum link-length, the minimum nodal density presents a lower-bound on the number of nodes per unit area when defined over a spatial area. For analytical structures such as WRNs, a relationship between $d_{\mathcal{N}}^{\max}$ and $\rho_{\mathcal{N}}^{\min}$ can be identified by determining a least dense configuration of the architecture. Using this relationship, it is possible to extend the link-length threshold theorem to place bounds on the nodal density requirements of a quantum WRN necessary to guarantee high performance. For the $k = 6$ WRN considered here it can be shown that $\rho_{\mathcal{N}} \geq \rho_{\mathcal{N}}^{\min} \geq \frac{2}{\sqrt{3}}(d_{\mathcal{N}}^{\max})^{-2}$ as derived in Section 3.4.

Fig. 4.3(a) emphasises that end-to-end performance has a clear and obvious dependence on internal decoherence; repeater inefficiencies ultimately limit end-to-end performance, and as a result place stricter constraints on the maximum permitted fibre-length and resources required throughout the network. For lossy network nodes with a total internal efficiency of $p_{\mathbf{x} \rightarrow \mathbf{y}}^{\text{int}} = 0.9$ (10% efficient) at any node require that fibre-lengths are limited to approximately 100 km in order to achieve a flooding capacity of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 10^{-2}$ bits per network use. This corresponds to a minimum nodal density of approximately $\rho_{\mathcal{N}}^{\min} \approx 1 \times 10^{-4}$ nodes per km².

¹This is a very reasonable assumption, and provides a better picture of end-to-end performance within a large network setting. Formal requirements for this were established in Chapter 3 with respect to internal weak-regularity.

Fig. 4.3(b) explores this relationship further. Given some end-to-end optimal performance $\mathcal{C}(\mathbf{i}, \mathcal{N})$, we plot the maximum tolerable internal loss $p_{\mathcal{N}}^{\text{int,max}}$ permitted at each node the network, with respect to the maximum fibre-length (and minimum nodal density) necessary to achieve it. This elucidates the *required* efficiency of repeater stations throughout the network, given some maximum fibre-length and desired end-to-end capacity. We see clearly that for very high rates $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 5$, the maximum fibre-length must be limited to approximately 15–25 km, otherwise the tolerable internal loss tends to zero, i.e. it can only be achieved via perfect devices. Similarly, the corresponding minimum nodal density in this setting is very large, requiring on the order of $\sim 10^{-2}$ nodes per km^2 . This is the expected resource requirements for high-rate DV quantum communications within a metropolitan setting. For lower rates, each node can tolerate greater inefficiencies over longer channel lengths. Nonetheless, to achieve a flooding rate of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 0.1$, channel lengths must be limited to below 93 km in the worst-case or 107 km in the best-case ².

4.5 Bosonic Thermal-Loss Networks

4.5.1 Network Model

Let us consider a general thermal-loss quantum network \mathcal{N}_{TL} , such that each channel in the network is a bosonic thermal loss channel $\mathcal{E}_{\eta, \bar{n}}$, as described in Chapter 2. We will assume that two repeater nodes $\mathbf{x} = \{\mathbf{x}^r, \mathbf{x}^u, \mathbf{x}^s\}$ and $\mathbf{y} = \{\mathbf{y}^r, \mathbf{y}^u, \mathbf{y}^s\}$ are connected via a thermal-loss channel with the attenuation and thermal noise properties

$$\eta_{\mathbf{x}\mathbf{y}} = \eta_{\mathbf{x}^s \rightarrow \mathbf{y}^r} = \eta_{\mathbf{y}^s \rightarrow \mathbf{x}^r}, \quad \bar{n}_{\mathbf{x}\mathbf{y}} = \bar{n}_{\mathbf{x}^s \rightarrow \mathbf{y}^r} = \bar{n}_{\mathbf{y}^s \rightarrow \mathbf{x}^r}. \quad (4.36)$$

Furthermore, let us model the internal repeater channels as thermal-loss channels which have fixed physical directions,

$$\mathcal{E}_{\mathbf{i}^r \rightarrow \mathbf{i}^u} = \mathcal{E}_{\tau_{\mathbf{i}}^r, \bar{n}_{\mathbf{i}}^r}, \quad \mathcal{E}_{\mathbf{i}^u \rightarrow \mathbf{i}^s} = \mathcal{E}_{\tau_{\mathbf{i}}^s, \bar{n}_{\mathbf{i}}^s}, \quad \mathbf{i} \in \{\mathbf{x}, \mathbf{y}\}. \quad (4.37)$$

Hence, the complete compound channels in either physical direction are

$$\mathcal{E}_{\mathbf{x} \rightarrow \mathbf{y}} = \mathcal{E}_{\tau_{\mathbf{y}}^r, \bar{n}_{\mathbf{y}}^r} \circ \mathcal{E}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}} \circ \mathcal{E}_{\tau_{\mathbf{x}}^s, \bar{n}_{\mathbf{x}}^s}, \quad (4.38)$$

$$\mathcal{E}_{\mathbf{y} \rightarrow \mathbf{x}} = \mathcal{E}_{\tau_{\mathbf{x}}^r, \bar{n}_{\mathbf{x}}^r} \circ \mathcal{E}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}} \circ \mathcal{E}_{\tau_{\mathbf{y}}^s, \bar{n}_{\mathbf{y}}^s}. \quad (4.39)$$

In general these compound channels are physically asymmetric. Thus for any quantum network topology, we need assign an optimal physical orientation to regain an undirected

²The range of values corresponds to worst-case and best-case bounds on the threshold parameter using the upper or lower-bounds on the end-to-end capacity

graphical representation. Once again, this can be achieved using the RCI to lower-bound the capacity of any compound thermal-loss channel in the network [28]. More precisely, for the forward channel we compute,

$$I(\mathcal{E}_{\mathbf{x}\rightarrow\mathbf{y}}) = -\log_2(1 - \eta_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}}) - h\left(\frac{\bar{n}_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}}}{1 - \eta_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}}}\right), \quad (4.40)$$

where we have made use of the entropic function $h(x) := (x+1)\log_2(x+1) - x\log_2(x)$, and the total point-to-point transmissivity and thermal noise parameters

$$\eta_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}} := \tau_{\mathbf{y}}^r \tau_{\mathbf{x}}^s \eta_{\mathbf{x}\mathbf{y}}, \quad (4.41)$$

$$\bar{n}_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}} := \bar{n}_{\mathbf{y}}^r + \tau_{\mathbf{y}}^r \bar{n}_{\mathbf{x}\mathbf{y}} + \eta_{\mathbf{x}\mathbf{y}} \tau_{\mathbf{y}}^r \bar{n}_{\mathbf{x}}^s. \quad (4.42)$$

These total parameters are derived by finding a single channel representation of compound thermal-loss channels, which can be found in Appendix B. For the backward channel we can simply reverse the order of the nodal directions in Eqs. (4.41) and (4.42), retrieving $I(\mathcal{E}_{\mathbf{y}\rightarrow\mathbf{x}})$. These offer lower-bounds on the capacity of each compound channel. By comparing these quantities, we can then identify the optimal physical channel $\mathcal{E}_{\mathbf{x}\mathbf{y}}^*$ for each edge in the network whose capacity is lower-bounded by,

$$I(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*) = \max_{i \neq j \in \{\mathbf{x}, \mathbf{y}\}} I(\mathcal{E}_{i \rightarrow j}). \quad (4.43)$$

This capacity lower-bound can then be substituted into Eq. (4.2) to compute achievable end-to-end rates for arbitrary topologies and network protocols.

Thermal-loss channels are teleportation-covariant, but are not distillable. Hence, we can compute upper-bounds on the network capacities using the REE single-edge upper-bounds [29]. For either physically directed channel this upper-bound takes the form,

$$E_R(\mathcal{E}_{i \rightarrow j}) = I(\mathcal{E}_{i \rightarrow j}) - \frac{\bar{n}_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}}}{1 - \eta_{\mathbf{x}\rightarrow\mathbf{y}}^{\text{tot}}} \log_2(\eta_{i \rightarrow j}^{\text{tot}}), \quad (4.44)$$

where we have used the RCI from Eq. (4.40). By optimising the physical orientation as before, we can substitute

$$E_R(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*) = \max_{i \neq j \in \{\mathbf{x}, \mathbf{y}\}} E_R(\mathcal{E}_{i \rightarrow j}), \quad (4.45)$$

into Eq. (4.4) to produce network capacity upper-bounds. Since this is an upper-bound, it is not known if they are achievable rates. However, using $I(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*)$ and $E_R(\mathcal{E}_{\mathbf{x}\mathbf{y}}^*)$ it is possible to accurately bound the end-to-end capacities.

<i>Parameter</i>	<i>Symbol</i>	<i>Value</i>
Wavelength	λ	800 nm
Detector Efficiency	τ^r	0.8
Detector Shot-noise	ν^r	1 - Homodyne 2 - Heterodyne
Detector bandwidth	W	100 MHz
Channel noise	\bar{n}_B	0.002
Noise Equivalent Power	NEP	6 pW/ $\sqrt{\text{Hz}}$
LO Power	P_{LO}	100 mW
Line width	l_W	1.6 KHz
Clock	C	5 MHz
Pulse Duration	$\Delta t, \Delta t_{\text{LO}}$	10 ns
Modulation	μ	10

Table 4.1: Parameter table for realistic CV-QKD protocols through optical-fibre channels.

4.5.2 Practical CV-QKD Setups

CV quantum communication protocols which make use of coherent detection (such as homodyne or heterodyne measurements) require the use of a local-oscillator (or phase reference). A phase reference allows the sender and receiver to exploit both quadratures of the mode, and can be established via a transmitted local-oscillator (TLO) or local local-oscillator (LLO) [76, 77]. A TLO is an additional mode which is co-propagated along with signal-mode from the sender to receiver, carrying the relevant phase information. An LLO uses interleaving signal pulses with bright reference pulses which are used to reconstruct the local-oscillator locally at the receiver [17]. The uniqueness of these techniques lead to unique noise/loss sources at the receiver, manifesting in different performance characteristics. For more details, see Appendix B of Ref. [30].

Let us consider CV quantum communication between two arbitrary nodes $\mathbf{x}, \mathbf{y} \in P$ using the physical channel direction $\mathbf{x} \rightarrow \mathbf{y}$. Firstly, we may assume that upon transmission, both LO techniques induce negligible loss or noise, i.e. $\bar{n}_{\mathbf{x}}^{s,\text{TLO}} \approx \bar{n}_{\mathbf{x}}^{s,\text{LLO}} \approx 0$ and $\tau_{\mathbf{x}}^{s,\text{TLO}} \approx \tau_{\mathbf{x}}^{s,\text{LLO}} \approx 0$. But at the receiver, there will be decoherence associated with sub-optimal detection and potential phase errors. The detector will not be perfectly efficient, limited by lossy fibre-couplings and limited quantum efficiency. This is captured by the detection loss, given by $\tau_{\mathbf{y}}^r = \tau_{\text{eff}}$.

There will also be noise induced by the detector and the local-oscillator. The use of an LLO introduces phases errors when it is being reconstructed at the receiver, since this

reconstruction will never be perfect. As such, we conceive a phase noise parameter

$$\Theta_{\text{ph}} := \frac{\pi(\mu - 1)l_{\text{W}}}{C}, \quad (4.46)$$

where μ describes the modulation of the transmitted pulses, C is the operational rate (clock) and l_{W} is the average linewidth of the light source.

Meanwhile, both techniques will be exposed to electronic noise due to sub-optimal detection and imperfect fibre-couplings. Let us define the electronic noise parameter,

$$\Theta_{\text{el}} := \frac{\nu_{\text{det}} \text{NEP}^2 W \Delta t_{\text{LO}}}{2h\nu P_{\text{LO}}^{\text{det}}}. \quad (4.47)$$

Here, ν_{det} is the detector shot-noise where $\nu_{\text{det}} = 1$ for homodyne and $\nu_{\text{det}} = 2$ for heterodyne. This quantity also depends on the noise equivalent power (NEP), the bandwidth W , the duration of the local-oscillator pulse Δt_{LO} , the frequency of the light ν , and the power of the local-oscillator at detection, $P_{\text{LO}}^{\text{det}}$. Since the LLO is reconstructed locally, the power at detection is simply the desired power $P_{\text{LO}}^{\text{det}} = P_{\text{LO}}$. On the other hand, using a TLO diminishes its power at detection due to loss suffered throughout its transmission $P_{\text{LO}}^{\text{det}} = \eta_{\mathbf{xy}} \tau_{\text{eff}} P_{\text{LO}}$.

Collecting these noise sources for each method, we can write

$$\bar{n}_{\mathbf{y}}^{r,\text{LLO}} \approx \eta_{\mathbf{xy}} \tau_{\text{eff}} \Theta_{\text{ph}} + \Theta_{\text{el}}, \quad \bar{n}_{\mathbf{y}}^{r,\text{TLO}} \approx \frac{\Theta_{\text{el}}}{\eta_{\mathbf{xy}} \tau_{\text{eff}}}. \quad (4.48)$$

It is clear that there is a trade-off between the phase-errors induced by a LLO and the electronic noise induced by a TLO. The reciprocal dependence of $\bar{n}_{\mathbf{y}}^{r,\text{TLO}}$ on the channel transmissivity means that as longer distances it will introduce greater levels of noise (which is precisely the behaviour shown in Fig. 4.4). As such, it becomes wise to make use of the LLO method, yet its technical requirements are somewhat more demanding. Table 4.1 collects typical values for the setup parameters which contribute to these noise quantities when considering optical-fibre connections.

4.5.3 Benchmarking

We can now benchmark the end-to-end limits of a bosonic thermal-loss network with imperfect repeaters. Here, we consider a $k = 8$ weakly-regular network, inspired by a Manhattan-like structure. A single network cell is depicted in Fig. 4.2(b) which is used to construct larger designs. Recall that we only demand connectivity constraints, and place no requirements on the spatial or topological properties of the network. The node-splitting procedure is similarly performed, adopting lossy and noisy channels between internal nodes

to capture repeater inefficiencies. Once again, we assume that end-users are located within some nodal boundary where boundary effects are unimportant.

We assume all network edges $(\mathbf{x}, \mathbf{y}) \in E$ are thermal-loss channels used to model optical-fibre so that the transmissivity is given by $\eta_{\mathbf{x}\mathbf{y}} = 10^{-\gamma d_{\mathbf{x}\mathbf{y}}}$ using the fibre-loss rate $\gamma = 0.02$ as before. Furthermore, there is unavoidable background thermal noise \bar{n}_B which is added to the propagating mode through the fibre-channel for which we assume the typical value of $\bar{n}_B \approx 0.002$ in our numerical investigations. As a result, any external fibre-channel of length $d_{\mathbf{x}\mathbf{y}}$ in a bosonic CV quantum network can be modelled as a thermal-loss channel with these parameters. As explored with regards to AD networks, we are able to utilise threshold theorems to derive maximum fibre-lengths $d_{\mathcal{N}}^{\max}$ necessary to guarantee optimal performance bounds. Similarly, this information can be used to identify network nodal density requirements using the minimum nodal density. For the $k = 8$ WR architecture considered here, these quantities are connected via $\rho_{\mathcal{N}}^{\min} \geq 2/(d_{\mathcal{N}}^{\max})^2$ as derived in Section 3.4.

Thanks to the node-splitting technique, we can also incorporate internal decoherence. In particular, we can investigate the end-to-end performance limits of realistic CV-QKD networks by considering setup noise/loss introduced by specific protocols. Typical CV protocols will make use of either homodyne or heterodyne measurements, both of which rely upon the use of a local-oscillator (an LLO or TLO) as reviewed in the previous section. By considering internal noise and loss sources alongside the external contributions from the fibre channel, we can gain insight into the realistic limits of CV-QKD networks which rely upon these techniques.

In Fig. 4.4(a) we plot bounds on the maximum fibre-length, and corresponding minimum nodal densities of bosonic thermal-loss networks necessary to guarantee an optimal flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N})$ between an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$. We do this for a number of setups; a pure-loss network with perfect repeaters, a thermal-loss network with perfect repeaters, and thermal-loss networks with imperfect repeaters using either TLOs/LLOs with heterodyne measurements according to the details and parameter table in the previous section.

For pure-loss networks with ideal repeaters, the maximum inter-nodal separation can become very large as we loosen our demands on the flooding capacity. Indeed, in this ideal scenario $d_{\mathcal{N}}^{\max} \approx 183$ km for end-to-end rates of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 10^{-2}$ bits per network use. However, upon realistic consideration of thermal noise, we realise that this upper-bound is very optimistic. Factoring background noise along each edge, is clear that $d_{\mathcal{N}}^{\max}$ saturates within some limiting range. For rates on the order of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 10^{-2}$ bits per network use, the maximum inter-nodal separation is found in the interval $d_{\mathcal{N}}^{\max} \in [91, 126]$ km. This is clearly much stricter than what is predicted when only external loss is considered.

Consideration of internal imperfections causes $d_{\mathcal{N}}^{\max}$ to become even more strict. In

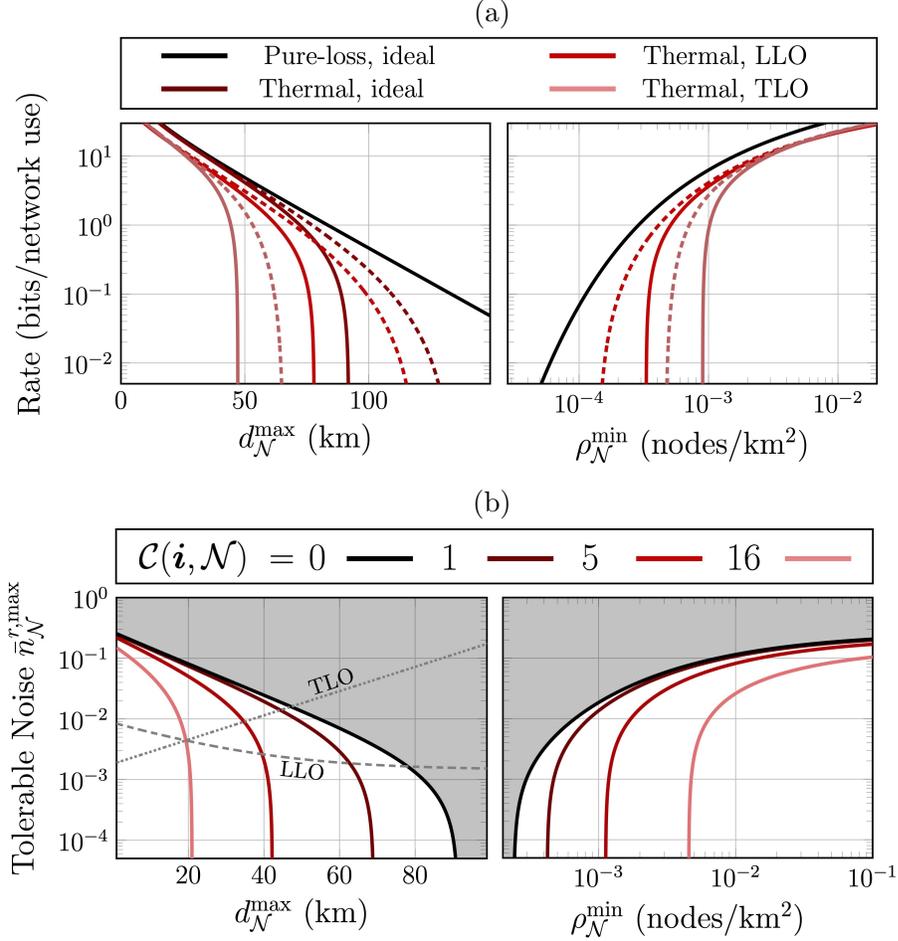


Figure 4.4: Throughout all plots, dashed lines represent bounds obtained using the REE as a single-edge capacity upper-bound. Meanwhile, all solid lines plot bounds derived using the RCI as a single-edge capacity lower-bound. Panel (a) plots the maximum inter-nodal separation $d_{\mathcal{N}}^{\max}$ permitted within the network, and its associated minimum nodal density, such that optimal performance $\mathcal{C}(i, \mathcal{N})$ can be obtained. Here, we consider network setups consisting of ideal repeaters or imperfect repeaters using LLOs and TLOs and heterodyne detection. Panel (b) displays the relationship between maximum fibre-length, minimum nodal density and tolerable thermal noise at the receiver $\bar{n}_{\mathcal{N}}^{r,\max}$ throughout the network. The grey areas of each plot illustrate regions of network parameter space for which we are unable to guarantee *any* end-to-end capacity whatsoever; identifying essential properties for realistic thermal-loss networks.

Fig. 4.4(a) we study the impact that the use of practical CV-QKD setups (using heterodyne detection) have on network resource requirements. It can be seen that LLO based protocols are superior in maintaining a larger tolerable channel length throughout a network, compared to a TLO approach. While LLOs introduce phase errors, the electronic

noise imparted as the receiver is independent from channel transmissivity and thus channel length. As a result, it does not degrade the tolerable channel length, unlike TLO based protocols. Indeed, for end-to-end rates of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 10^{-2}$ bits per network use, CV-QKD protocols which utilise LLOs can tolerate at least ~ 29 km additional channel length for every point to point link (at most ~ 48 km by considering upper-bounds).

Stricter link-length demands have a substantial impact on the nodal density requirements of bosonic thermal-loss networks. This is clear within our results; when the maximum fibre-length saturates, so too does the minimum nodal density. Indeed, the consideration of thermal noise within practical CV-QKD networks demands that the minimum nodal density always be at least of order 10^{-4} nodes per km^2 in order to achieve any non-zero end-to-end rate. This has significant ramifications on the resource requirements of realistic bosonic quantum networks.

A threshold theorem can also be derived with respect to thermal noise at the receiver $\bar{n}_{\mathbf{x}}^r$, so that we can further study the interplay between maximum channel length permitted in the WRN and tolerable receiver noise. Fig. 4.4(c) illustrates this relationship for a number of desired flooding capacities. Here we plot lower-bounds on the maximum tolerable noise $\bar{n}_{\mathcal{N}}^{r, \max}$ for any node in the network, using the RCI based bound. This reveals a permissible region of network parameters for which not only optimal performance is guaranteed, but non-zero rates are guaranteed. In order to guarantee a non-zero end-to-end capacity (even within this highly-connected architecture) channel lengths should be kept below 91 km at worst, and 126 km at best. Fibre-networks in this configuration which have channels longer than this are not guaranteed to have a non-zero rate. Globally, these channel length constraints manifest in the required nodal density so that we can identify a permissible region of nodal densities to guarantee non-zero rates.

Furthermore, we compare the tolerable noise bounds to the actual noise properties of CV-QKD protocols using TLOs and LLOs. It is once more illustrated that LLOs are much more effective within large scale networks, as the associated internal noise scales more favourably with increasing fibre-lengths.

4.6 Conclusion

In this chapter we have presented general bounds for the end-to-end capacities of arbitrary quantum networks. This includes achievable lower-bounds based on the coherent/reverse coherent information which apply to networks composed of any type of quantum channel. We also show how upper-bounds can be obtained using appropriate single-edge capacity bounding functions. Employing these bounds in conjunction with a recently de-

veloped node-splitting technique, we showed ways to bound the end-to-end capacity of quantum networks with lossy and noisy repeaters. As a result, we provide versatile tools to investigate the internal and external decoherence properties of realistic quantum networks.

Making use of these general results, we apply the node-splitting technique to qubit amplitude-damping networks, and bosonic thermal-loss networks; channel models for which the single-edge capacity is not exactly known. Using the class of highly-connected, WR quantum networks we are able to illuminate critical network properties upon which high-rate quantum communications rely. This allows us to identify internal loss and thermal noise thresholds permitted within quantum repeaters which guarantee optimal end-to-end performance.

Our results find valuable insight for the infrastructure requirements of future quantum networks. Most prominently, we emphasise the necessity to consider both internal and external thermal noise when designing quantum architectures, as these noise sources can severely restrict the ability to use long fibre-channels while maintaining high rates. Even when channels are limited in length, quantum repeaters that are insufficiently protected from noise will compromise performance. Future investigative paths will aim to exploit these bounds to both study and motivate realistic and high-rate network designs and unveil the resiliency of quantum networks to unavoidable thermal noise.

Chapter 5

Free-Space and Hybrid Quantum Network Capacities

The work in this chapter forms the basis of a paper published in *Physical Review Applied*, whose authors are (in order) Cillian Harney, Alasdair I. Fletcher and Stefano Pirandola [3]. This chapter is structured as follows: Section 5.1 introduces the context of the chapter, before Section 5.2 overviews the capacities of general fading channels and the end-to-end network capacities of quantum fading networks. Section 5.3 specifies this review to optical free-space quantum communications, summarising recent progress in the determination of ultimate limits in a number of key settings. In Section 5.4, we formalise a network architecture for the study of hybrid, modular quantum networks. We further specify an idealised network architecture which allows us to establish properties that guarantee optimal end-to-end performance and distance-independence. Sections 5.5 and 5.6 applies the machinery from the previous sections to investigate the optimal performance of hybrid quantum architectures. In particular, we establish network constraints for communication between remote fibre-based sub-networks connected to a satellite backbone, and for ground-based free-space sub-networks connected to a fibre backbone. Concluding remarks and future investigative paths are then discussed in Section 5.7.

5.1 Introduction

Quantum networks will not be limited to just optical-fibre but will collaborate with free-space methods of communication. On the ground, the flexibility of free-space links are obviously more suitable for mobile quantum devices and short-range connections. Meanwhile, the ability to establish ground-to-satellite and intersatellite free-space connections offers remarkable short-cuts for global quantum communications [69, 67, 68, 71, 70, 24, 25].

Such connections bypass many decibels of loss that would be otherwise experienced on the ground and utilise the dynamic nature of satellites to achieve high rates over global distances.

Determining the ultimate limits of free-space quantum channels is difficult, requiring tools from quantum information theory [39, 33, 36], optics [78, 79, 80] and turbulence theory [81, 82, 83, 84]. Recent advancements have placed tight upper-bounds on the quantum capacities of point-to-point free-space channels, using a modified PLOB bound that accounts for atmospheric fading processes [30, 31]. With these results in hand, we have the ingredients to go beyond the point-to-point scenario and quantitatively study the ultimate limits of communications in free-space quantum networks.

In this chapter, we combine results from Refs. [32] and [30, 31] in order to place bounds on the end-to-end capacities of generally hybrid quantum networks. In particular, we put forward a formalism for studying the capacities of quantum networks whose channels are described by free-space, fibre, or any medium that can be generalised as a fading channel. This treatment is then specified to fading processes that are experienced by optical transmissions through the atmosphere, or in space. Furthermore, we introduce a framework to investigate the ultimate limits of hybrid, modular quantum networks. We focus on a modular network design which consists of disjoint sub-networks (or communities) connected to a large-scale backbone network used to mediate intercommunity quantum communication. This provides the tools to investigate highly relevant quantum network models, such as (1) globally distant fibre sub-networks connected to a satellite backbone network, offering insight into the resource requirements of a satellite-based quantum internet; and (2) wireless, free-space sub-networks on the ground interconnected via a fibre backbone, presenting a useful model for studying hybrid metropolitan networks.

Extending the techniques of Ref. [1] we employ ideally connected structures within different parts of the modular network. In doing so, we are able to derive simple, yet powerful analytical constraints which promise distance-independent, optimal rates for modular quantum architectures. These results provide valuable insight into the ultimate limits of hybrid networks, can help to motivate future quantum network design and provide a valuable platform upon which to further develop realistic free-space quantum networks.

5.2 Quantum Networking over Fading Channels

5.2.1 Fading Channels

The effect of *fading* refers to the temporal variation of transmissivity along a bosonic lossy channel. The transmissivity along a fading channel is not fixed, but instead follows a probability distribution described by the dynamics of the environment. For example, the propagation of bosonic modes through low-altitude free-space instigates a fading channel thanks to chaotic processes in the atmosphere. The impact of fading on a communications channel is described via its *speed*, i.e. the ability for a receiver to resolve the dynamics of the transmissivity fluctuations. Slow-fading implies that the users can resolve the fading dynamics and accurately perform channel estimation because either the fading process is weak or the users possess sufficiently fast detectors. On the other hand, fast-fading refers to the situation where the users cannot reconcile the dynamics of the channel and can only estimate the statistical distribution of the channel transmissivity [85, 86]. It is clear that fast-fading poses a more formidable task for communicators.

More precisely, a bosonic lossy fading channel is defined as an ensemble of lossy channels in accordance with some probability density function $F(\tau)$ which describes the instantaneous transmissivity along the channel. We denote a lossy fading channel as the ensemble

$$\mathcal{E}_F(\eta) := \{F(\tau); \mathcal{E}_\tau\}, \quad (5.1)$$

where \mathcal{E}_τ is a lossy channel with fixed, instantaneous transmissivity $\tau \in [0, \eta]$ and η is the maximum transmissivity that is attainable along the channel.

5.2.2 Capacities of Fading Channels

While the PLOB bounds assumes a fixed transmissivity η it can be readily employed to study fading channels [29, 87]. Thanks to convexity properties of the relative entropy of entanglement (REE) over ensembles of channels [44, 42], the capacity of a lossy fading channel can be bounded according to

$$\mathcal{C}[\mathcal{E}_F(\eta)] \leq \mathcal{B}_F^\eta := - \int_0^\eta d\tau F(\tau) \log_2(1 - \tau), \quad (5.2)$$

where we have defined \mathcal{B}_F^η as the capacity function for lossy fading channels. This can be interpreted as a generalisation of the PLOB bound, modified to include potential fading processes. Indeed, it is simple to retrieve the standard bound $-\log_2(1 - \eta)$ for fixed lossy channels by considering a trivial probability distribution where only one transmissivity value

is possible, η . Hence, this format is conveniently general and allows one to describe any lossy bosonic channel (with or without fading).

Analogous to the pure-loss setting, a thermal-lossy fading channel can be described by the ensemble

$$\mathcal{E}_F(\eta, \bar{n}) := \{F(\tau, \bar{n}); \mathcal{E}_{\tau, \bar{n}}\}, \quad (5.3)$$

where it is possible that both transmissivity and thermal noise are probabilistic and described within a probability density function $F(\tau, \bar{n})$. Typically, thermal noise can always be considered constant by either assuming stable operational conditions, or by minimising (maximising) its potential value for best-case (worst-case) rates. This allows us to consider the simpler ensemble $\mathcal{E}_F(\eta, \bar{n}) = \{F(\tau); \mathcal{E}_{\tau, \bar{n}}\}$ on which we place the following upper-bound of its capacity [30, 31],

$$\mathcal{C}[\mathcal{E}_F(\eta, \bar{n})] \leq \mathcal{T}_F^{\eta, \bar{n}} := \int_{\bar{n}}^{\eta} d\tau F(\tau) E_R(\mathcal{E}_{\tau, \bar{n}}), \quad (5.4)$$

$$= - \int_{\bar{n}}^{\eta} d\tau F(\tau) \left[\log_2 \left[(1 - \tau) \tau^{\frac{\bar{n}}{1-\tau}} \right] + h \left(\frac{\bar{n}}{1 - \tau} \right) \right]. \quad (5.5)$$

Here we have defined $\mathcal{T}_F^{\eta, \bar{n}}$ as a tight capacity bounding-function for thermal-lossy fading channels¹. Intuitively, one can never outperform the pure-loss PLOB bound in the presence of thermal-noise, hence we can always write

$$\mathcal{C}[\mathcal{E}_F(\eta, \bar{n})] \leq \mathcal{T}_F^{\eta, \bar{n}} \leq \mathcal{B}_F^{\eta}. \quad (5.6)$$

5.2.3 Capacities of Fading Networks

We can combine the theory from these previous sections in order to provide a general model for quantum networks with fading channels. Indeed, we may construct a quantum network $\mathcal{N} = (P, E)$ such that all edges $(\mathbf{x}, \mathbf{y}) \in E$ are generally associated with a unique thermal-lossy fading channel,

$$\mathcal{E}_{\mathbf{x}\mathbf{y}} = \mathcal{E}_{F_{\mathbf{x}\mathbf{y}}}(\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}), \quad \forall (\mathbf{x}, \mathbf{y}) \in E. \quad (5.7)$$

In this way, each network edge not only possesses a unique maximum transmissivity $\eta_{\mathbf{x}\mathbf{y}}$ and thermal-noise properties $\bar{n}_{\mathbf{x}\mathbf{y}}$, but also a unique instantaneous transmissivity probability density function $F_{\mathbf{x}\mathbf{y}}$ through which each edge can adopt its own fading dynamics

¹The word *tight* in this context refers to how close the upper-bound is from its best known lower-bound. Indeed, there exists a lower-bound on the capacity of a point-to-point thermal-loss channel based on its reverse coherent information (RCI) [88]. Hence, throughout our work we implicitly refer to tight upper-bounds on thermal-loss channel capacities (and subsequently, network capacities) as those which in conjunction with the RCI can tightly sandwich the exact capacity.

(or lack thereof). This allows for a description of network channels within different environmental media such as fibre channels, ground-based free-space channels, or free-space channels beyond the atmosphere. Furthermore, we can retrieve pure-loss fading channels via $\bar{n}_{xy} = 0$.

As explored in the previous chapters, we know that the optimal network capacity is associated with a flooding protocol, and is found by locating the entanglement cut C_{\min} which minimises the multi-edge capacity over all cut-sets. For lossy and thermal-lossy fading networks we can utilise the following flooding capacities in order to place ultimate upper-bounds on end-to-end fading network performance,

$$\mathcal{B}(\mathbf{i}, \mathcal{N} | \mathcal{P}) \leq \mathcal{B}(\mathbf{i}, \mathcal{N}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} \mathcal{B}_{F_{\mathbf{x}\mathbf{y}}}^{\eta_{\mathbf{x}\mathbf{y}}}, \quad (5.8)$$

$$\mathcal{T}(\mathbf{i}, \mathcal{N} | \mathcal{P}) \leq \mathcal{T}(\mathbf{i}, \mathcal{N}) := \min_C \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} \mathcal{T}_{F_{\mathbf{x}\mathbf{y}}}^{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}}, \quad (5.9)$$

so that we can write, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{T}(\mathbf{i}, \mathcal{N}) \leq \mathcal{B}(\mathbf{i}, \mathcal{N})$.

5.3 Free-Space Quantum Communication

Consider two remote parties Alice and Bob who are separated by a distance z , and employ quantum communications based upon a quasi-monochromatic optical mode ($\Delta\lambda$ -nm large and Δt -sec long). This may be characterised by a Gaussian beam with wavelength λ , initial field spot-size w_0 and curvature R_0 . Communication consists of transmitting a directed beam towards a receiver with circular aperture of radius a_R . Here we assume that the initial spot-size w_0 is sufficiently small with respect to the transmitter aperture of radius a_T so that there is no relevant diffraction caused by the transmitter.

The atmospheric effects which characterise free-space channels are variable with respect to altitude, due to changes in atmospheric density. Therefore specifying the trajectory of a Gaussian beam through free-space is pivotal in capturing channel quality. To this end, for any point-to-point communications task we may assume a general beam trajectory L and introduce the following altitude/propagation functions respectively, $h_L(z)$ and $z_L(h)$. Using these functions we can retain a geometry independent framework for our study until we wish to specify to a particular setting.

5.3.1 Free-Space Transmissivity

Free-space diffraction is a universal contributor to loss. As a beam propagates in free-space its waist will widen as a function of the distance that it travels,

$$w_d^2(z) = w_0^2 \left[\left(1 - \left(\frac{z}{R_0} \right) \right)^2 + \left(\frac{z}{z_R} \right)^2 \right], \quad (5.10)$$

where $z_R := \pi w_0^2 / \lambda$ is the Rayleigh range. A target receiver with a circular aperture of radius a_R will then only detect a portion of the spread beam since its aperture is finite in size, inducing a diffraction-limited transmissivity [30],

$$\eta_d(z) = 1 - \exp \left[-\frac{2a_R^2}{w_d^2} \right]. \quad (5.11)$$

It is also useful to define a diffraction induced transmissivity in the far-field regime, $z \gg z_R$, making the approximation $\eta_d \approx \eta_d^{\text{far}} := 2a_R^2/w_d^2$. This loss quantity exists regardless of the specific environmental setting considered, from ground-based links to intersatellite connections.

Propagation through the atmosphere incurs further loss due to aerosol absorption and Rayleigh/Mie scattering; an effect known as atmospheric extinction. At a fixed altitude h , this loss can be accurately described via the Beer-Lambert equation [80]. Since beam trajectories may be variable in altitude, we can generally define the extinction-induced transmissivity as [89, 90, 30]

$$\eta_{\text{atm}}(z) = \exp \left[-\int_0^z dz \alpha[h_L(z)] \right], \quad (5.12)$$

where $\alpha(h) = \alpha_0 e^{-h/\tilde{h}}$ is the extinction factor, $\tilde{h} = 6600$ m, and α_0 is the extinction factor at sea-level. For $\lambda = 800$ nm it follows that $\alpha_0 \approx 5 \times 10^{-6} \text{ m}^{-1}$

Finally, there exist inevitable internal losses associated with the detector setup, due to imperfect fibre-couplings, sub-optimal quantum detector efficiency, and more [91, 92]. This inefficiency-induced transmissivity can be as low as $\eta_{\text{eff}} \approx 0.4$ and must be considered to capture realistic performance. All of these effects can be used to describe a fixed, maximum transmissivity of a free-space connection,

$$\eta(z) := \eta_{\text{eff}} \eta_{\text{atm}}(z) \eta_d(z). \quad (5.13)$$

Importantly, η can be readily modified to consider variable altitude beam trajectories and written as a function of a chosen spatial geometry to account for different extinction properties throughout the atmosphere.

5.3.2 Atmospheric Fading

It is remarkably optimistic to assume that a free-space transmission deterministically undergoes a pure-loss channel characterised by Eq. (5.13) only. The chaotic behaviour of air-flow, temperature and pressure throughout the atmosphere invites further complications for free-space transmissions, causing inaccuracies in the point-to-point trajectory known as *beam wandering*. As a result, we must incorporate fading for a more accurate characterisation [93, 94, 95, 30].

Turbulence is used to describe how a free-space propagating beam is perturbed by fluctuations in the atmospheric refractive index, caused by spatial variations in pressure and temperature. Propagating beams interact with small turbulent air-flows on a fast time-scale, too fast for communicators to monitor or resolve. This causes the beam waist to broaden and forces us to define a *short-term spot-size* w_{st} which is larger than the diffraction-induced spot size, $w_{\text{d}} < w_{\text{st}}$. On a slower time-scale, the beam will undergo deflections by significantly larger eddies in the atmosphere. This slower time-scale may be reconcilable by the communicators, and manifests as a wandering of the beam centroid. This wandering can be described by a Gaussian random walk of the centroid with variance σ_{t}^2 which is a functional of the beam trajectory, operational setup, conditions, and more.

Wandering is not exclusively caused by turbulence, and one must also consider pointing errors caused by jitter and imperfect targeting. These effects also occur on a reasonably slow time-scale of order 0.1 – 1 s, and may be resolved by the receiver. This introduces an additional wandering variance σ_{p}^2 , e.g. a 1 μrad pointing error at the transmitter causes a variance $\sigma_{\text{p}}^2 \approx (10^{-6}z)^2$ (where z is in meters). Overall, these effects combine to induce Gaussian centroid wandering with variance $\sigma^2 = \sigma_{\text{t}}^2 + \sigma_{\text{p}}^2$ [30].

The ability for communicators to resolve these wandering dynamics is dependent on their time-scale [96]. The behaviour of turbulence is variable, with regimes ranging from weak to strong turbulence. Increasing turbulent strength can be modelled as an increasingly faster fading process, such that a receiver loses the ability to reconcile the wandering dynamics. For stronger levels of turbulence, it is possible to define a *long-term spot-size* w_{lt} which averages over the wandering caused by both small turbulent eddies and larger eddy deflections, $w_{\text{d}} < w_{\text{st}} < w_{\text{lt}}$. Indeed, the turbulence-induced variance is defined with respect to the long-term and short-term quantities $\sigma_{\text{t}}^2 = w_{\text{lt}}^2 - w_{\text{st}}^2$. However, rigorous studies of strong turbulence will require further considerations which have been addressed in Ref. [97].

Here, we focus on the regime of weak turbulence and the concept of short-term beam spot sizes. These can be used to provide precise descriptions of free-space quantum channels on the ground at short-range, and for ground-to-satellite communication along trajectories

with small zenith angles [30, 96, 98].

5.3.3 Weak Turbulence

For communications undergoing weak turbulence, the beam wandering acts on a time scale of 10 – 100 ms and can be fully resolved with a sufficiently fast detector. In this case, analytical expressions can be found for the short-term spot size w_{st} and the centroid wandering variance σ^2 . Consider a beam with wavenumber $k = 2\pi/\lambda$ following a free-space trajectory L (and its associated altitude function $h_L(z)$). Then the spherical-wave coherence length is given by [30],

$$\rho_0(L) = \left[1.46k^2 \int_0^z d\zeta \left(1 - \frac{\zeta}{z} \right)^{\frac{5}{3}} C_n^2[h_L(\zeta)] \right]^{-\frac{3}{5}}, \quad (5.14)$$

where C_n^2 denotes the refractive index structure constant, used to measure the strength of fluctuations in the atmospheric refractive index. This quantity has an explicit dependence on the beam's trajectory, since this may be variable in altitude, and is typically described via the Hufnagel-Valley model (See Appendix C of Ref. [30]). Provided that Yura's condition is satisfied $\phi := 0.33(\rho_0/w_0)^{\frac{1}{3}} \ll 1$ [99] then we can write [30],

$$w_{\text{st}}^2 \approx w_{\text{d}}^2 + 2 \left(\frac{\lambda z}{\pi \rho_0} \right)^2 (1 - \phi)^2, \quad (5.15)$$

$$\sigma_{\text{t}}^2 \approx 2 \left(\frac{\lambda z}{\pi \rho_0} \right)^2 [1 - (1 - \phi)^2]. \quad (5.16)$$

The short-term spot size can be used to update the diffraction induced transmissivity to account for fast beam interaction with small turbulent eddies in the atmosphere. That is,

$$\eta_{\text{st}} := 1 - \exp \left[-\frac{2a_R^2}{w_{\text{st}}^2} \right] \underset{z \gg z_R}{\approx} \eta_{\text{st}}^{\text{far}} := \frac{2a_R^2}{w_{\text{st}}^2}, \quad (5.17)$$

where we have simultaneously introduced a far-field approximation, $\eta_{\text{st}}^{\text{far}}$ when the propagation distance is very large $z \gg z_R$.

Updating the diffraction-induced transmissivity in Eq. (5.13), we may write a new maximum transmissivity incorporating weak-turbulent effects, $\eta = \eta_{\text{eff}} \eta_{\text{atm}} \eta_{\text{st}}$. This represents the optimal transmissivity parameter that can be achieved when the beam centroid \vec{x}_C is perfectly aligned with the receiver centroid \vec{x}_R , i.e. the centroid deflection is $r := \|\vec{x}_C - \vec{x}_R\| = 0$. However, due to turbulence and pointing errors, the beam centroid now undergoes a Gaussian random walk with variance σ^2 [100] invoking a fading channel.

We can then connect the non-zero centroid deflection $r \geq 0$ to an instantaneous transmissivity $\tau(r)$ to precisely capture the fading process. Gaussian wandering induces a Weibull distribution for the centroid deflection, which results in an instantaneous transmissivity probability density function $F_\sigma[\tau(r)]$ [30]. Defining the functions,

$$f_0(x) := [1 - \exp(-2x)I_0(2x)]^{-1}, \quad (5.18)$$

$$f_1(x) := \exp(-2x)I_1(2x), \quad (5.19)$$

where I_n is the modified Bessel function of the first kind for $n = 0, 1$, we can introduce the following shape and scale parameters,

$$\gamma = \frac{4\eta_{\text{st}}^{\text{far}} f_0(\eta_{\text{st}}^{\text{far}}) f_1(\eta_{\text{st}}^{\text{far}})}{\ln [2\eta_{\text{st}} f_0(\eta_{\text{st}}^{\text{far}})]}, \quad r_0 = \frac{a_R}{\ln [2\eta_{\text{st}} f_0(\eta_{\text{st}}^{\text{far}})]^{\frac{1}{\gamma}}}. \quad (5.20)$$

With these, we can now write the instantaneous transmissivity probability density function [30],

$$F_\sigma(\tau) = \frac{r_0^2}{\gamma\sigma^2\tau} \ln\left(\frac{\eta}{\tau}\right)^{\frac{2}{\gamma}-1} \exp\left[-\frac{r_0^2}{2\sigma^2} \ln\left(\frac{\eta}{\tau}\right)^{\frac{2}{\gamma}}\right]. \quad (5.21)$$

We are left with a free-space, lossy fading channel $\mathcal{E}_{F_\sigma}(\eta) = \{F_\sigma(\tau); \mathcal{E}_\tau\}$. Using the tools from Section 5.2, we can study the capacities of free-space connections.

Hence, the capacities for free-space quantum communications (entanglement distribution or secret-key distribution) are upper bounded according to [30]

$$\mathcal{C} \leq \mathcal{B}_{F_\sigma}^\eta = -\Delta(\eta, \sigma) \log(1 - \eta), \quad (5.22)$$

where Δ represents a correction factor to the PLOB bound due to imperfect alignment,

$$\Delta(\eta, \sigma) = 1 + \frac{\eta}{\ln(1 - \eta)} \int_0^\infty dx \frac{\exp\left[\frac{-r_0^2}{2\sigma^2} x^{\frac{2}{\gamma}}\right]}{e^x - \eta}. \quad (5.23)$$

Through specification to a free-space trajectory, one can easily determine geometry dependent expressions for this ultimate limit. Importantly, for channels which are accurately described as ensembles of pure-loss channels (thermal noise is negligible), then Eq. (5.22) is in fact an achievable and optimal rate, $\mathcal{C} = \mathcal{B}_{F_\sigma}^\eta$. For all other scenarios where thermal noise is non-negligible, it remains an effective upper-bound.

5.3.4 Thermal Noise

As discussed previously, pure-loss based bounds remain ultimate bounds in the presence of thermal noise. Yet, it is still possible to construct tighter performance bounds by considering fading channels which are ensembles of thermal-loss channels. Let \bar{n}_T be the mean

number of input photons transmitted towards a receiver via a single free-space mode. For an instantaneous transmissivity τ the mean photon number collected at the receiver will be $\bar{n}_R = \tau\bar{n}_T + \bar{n}$, where \bar{n} describes the total environmental thermal noise added to the signal. It is useful to define contributions to this environmental noise via

$$\bar{n} := \eta_{\text{eff}} \bar{n}_B + \bar{n}_{\text{ex}}, \quad (5.24)$$

where the receiver collected \bar{n}_B mean background photons with detector efficiency η_{eff} , and \bar{n}_{ex} accounts for excess setup noise. In the study of ultimate limits, \bar{n}_{ex} can be considered to be approximately zero, or can be attributed to trusted noise.

For free-space links, the primary source of thermal noise is attributed to natural brightness within the field of view of the transmission, i.e. the sky, Sun, Moon, etc. Using a receiver of aperture a_R , angular field of view Ω_{fov} , a detector with time window Δt and frequency filter $\Delta\lambda$ around λ , then the number of background thermal photons per mode is [101, 91]

$$\bar{n}_B = H_\lambda \Gamma_R, \text{ where } \Gamma_R := \Delta t \Delta\lambda \Omega_{\text{fov}} a_R^2. \quad (5.25)$$

Here, H_λ describes the spectral irradiance of the environment in units of photons $\text{m}^{-2} \text{s}^{-1} \text{nm}^{-1} \text{sr}^{-1}$, and is unique to the operational setting and trajectory. Using the general bound from Eq. (5.5) and specifying to free-space beam wandering dynamics with variance σ^2 , we can write the free-space thermal upper-bound,

$$\mathcal{C} \leq \mathcal{T}_{F_\sigma}^{\eta, \bar{n}} = \mathcal{B}_{F_\sigma}^\eta - \mathcal{U}_{F_\sigma}^{\eta, \bar{n}}, \quad (5.26)$$

where the thermal correction is given explicitly by [30],

$$\mathcal{U}_{F_\sigma}^{\eta, \bar{n}} := \left[1 - \exp\left(\frac{-r_0^2}{2\sigma^2} \ln\left[\frac{\eta}{\bar{n}}\right]^{\frac{2}{\gamma}}\right) \right] \times \left[\frac{\bar{n} \log_2(\bar{n})}{1 - \bar{n}} + h(\bar{n}) \right] - \mathcal{B}_{F_\sigma}^{\bar{n}}. \quad (5.27)$$

This result applies to settings of weak and intermediate turbulence, such that one can substitute the appropriate reconcilable wandering variance and maximum transmissivity into this result.

5.3.5 Noise Suppression and Frequency Filters

As seen in Eq. (5.25), the number of background thermal photons per mode has a strong dependence on the frequency filter, $\Delta\lambda$. The frequency filter assists in blocking out noise, and thus the use of ultra-narrow filters is highly desirable. In discrete-variable quantum communications, physical frequency filters are typically limited to around $\Delta\lambda = 1 \text{ nm}$. However, using CV quantum systems and appropriate interferometric measurements it is

possible to achieve much narrower effective filters. Discussed in Ref. 4.5.2, CV protocols rely on the use of a local-oscillator in order to perform homodyne or heterodyne measurements at the output. We recall that a local-oscillator can be co-transmitted with signal pulses (TLO) or can be reconstructed locally at the receiver as an LLO. Reconstructing a LLO involves interleaving the signal pulses with strong reference pulses that carry information about the local-oscillator [17]. Since the output of a homodyne measurements is proportional to the mean photon number in the local-oscillator modes, the ability to utilise bright reference pulses over free-space channels introduces an *effective homodyne filter*. Thermal noise mode-matching with the local-oscillator and the signal will be detected, but all other noise will be filtered out. This allows for the implementation of ultra-narrow effective filters on the order of $\Delta\lambda = 0.1$ pm with practical CV protocols, and can dramatically reduce the magnitude of the thermal background noise (see Ref. [30, 76, 77] for more details).

5.3.6 Ground-Based Channels

Wireless classical communication networks are ubiquitous and fundamental to everyday modern life. Thus the desire for a free-space quantum analogue is obvious, enabling access to future wireless quantum technologies. Nonetheless, it is intuitive that such communication will be limited to short-range due to prominent decoherence obtained at ground-level. At a fixed altitude, beam trajectories are horizontal paths with the simple altitude/propagation functions $h_L(z) = h$, $z_L(h) = z$. The absence of a variable altitude in the beam path simplifies a number of key quantities such as the extinction-induced transmissivity,

$$\eta_{\text{atm}}(z) = \exp[-\alpha(h)z], \quad (5.28)$$

and the spherical-wave coherence length

$$\rho_0 = [0.548k^2C_n^2(h)z]^{-\frac{3}{5}}, \quad (5.29)$$

which can be used to accurately describe decoherence and fading dynamics on the ground. Here, turbulence is a major factor and must be stringently considered. A useful parameter for assessing the validity of turbulent regimes on the ground is the Rytov variance,

$$\sigma_{\text{Ry}}^2 = 1.23 k^{\frac{7}{6}} z^{\frac{11}{6}} C_n^2(h). \quad (5.30)$$

Weak turbulence requires that $\sigma_{\text{Ry}}^2 \lesssim 1$. Using a Gaussian beam with $\lambda = 800$ nm and altitudes close to sea-level during typical day-time conditions, weak turbulence is only guaranteed for distances of $z \lesssim 1$ km. Beyond this, as in the intermediate ($\sigma_{\text{Ry}}^2 \gtrsim 1$) and

strong ($\sigma_{\text{Ry}}^2 \gg 1$) turbulent regimes, the long-term spot size must be adopted, leading to weaker channel capacities [97].

Fig. 5.1(a) illustrates the behaviour of transmissivity in ground-based free-space channels with respect to propagation length. Within the weak-turbulence regime the loss properties of free-space channels limited to ~ 4 dB for communications over 1 km, encouraging the utility of short-range, optical free-space quantum communications.

For the assessment of thermal bounds, the primary source of thermal-noise at ground-level is attributed to the brightness of the sky. This provides a spectral irradiance ranging from [101],

$$H_{\lambda}^{\text{sky}} \approx \begin{cases} 1.9 \times 10^{13}, & \text{full-Moon, clear night,} \\ 1.9 \times 10^{18}, & \text{cloudy day time,} \end{cases} \quad (5.31)$$

in units of photons $\text{m}^{-2}\text{s}^{-1}\text{nm}^{-1}\text{sr}^{-1}$. Using this information, the expressions in Eqs. (5.28) and (5.29), and the general capacity bounds developed in the previous sections, we can accurately assess the ultimate limits of free-space quantum communications on the ground (see Ref. [30] for further details and derivations).

5.3.7 Ground-Satellite Channels

For communication between ground/satellite stations, there are two unique configurations that must be considered: Transmissions directed from the ground towards a satellite (uplink) or from a satellite towards the ground (downlink). The quantum channel descriptions of these configurations are very different.

Consider a Gaussian beam propagated in uplink. The beam immediately undergoes turbulence upon generation at low altitude, and thus has a large decohering impact which must be carefully considered. Pointing errors are less critical $\sigma_{\text{p}}^2 \ll \sigma_{\text{t}}^2$ thanks to the availability of adaptive optics to optimise the beam trajectory from the ground station. Therefore we must model uplink as a fading channel predominantly due to turbulent effects. Meanwhile, a Gaussian beam in downlink experiences the opposite; the beam does not undergo serious levels of turbulence until it reaches lower altitudes. But by this point, its spot-size has already been spread by diffraction, hence turbulence does not present a serious factor and $\sigma_{\text{t}}^2 \approx 0$. Yet, in this setting pointing errors become much more relevant due to the lack of onboard access and optimisation ability. Hence, atmospheric decoherence associated with uplink and downlink is physically asymmetric, invoking two unique fading channels.

We specify the trajectory of ground-satellite communication according to a target satellite altitude h and zenith angle θ , which describes the angle formed between the zenith

point at the ground station and the direction of observation towards the satellite. The zenith angle takes values $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, such that when $\theta = 0$ the satellite is at the zenith. The distance that the beam physically travels from its point of generation z (known as its slant distance) can then be expressed with respect to this geometry. Defining the functions,

$$\begin{aligned} h_\theta(z) &= \sqrt{R_E^2 + z^2 + 2zR_E \cos \theta} - R_E, \\ z_\theta(h) &= \sqrt{h^2 + 2hR_E + R_E^2 \cos^2 \theta} - R_E \cos \theta, \end{aligned} \quad (5.32)$$

we may then introduce the altitude/propagation functions with respect to uplink and downlink communications [31],

$$z_\theta^{\text{up}}(h) = z_\theta(h), \quad z_\theta^{\text{down}}(h) = z_\theta(h_{\text{max}}) - z_\theta(h), \quad (5.33)$$

$$h_\theta^{\text{up}}(z) = h_\theta(z), \quad h_\theta^{\text{down}}(z) = h_\theta[z_\theta(h_{\text{max}}) - z]. \quad (5.34)$$

Fig. 5.1(b) illustrates the behaviour of transmissivity in ground-satellite channels with respect to uplink, downlink and satellite altitude. Here we plot both the the expected transmissivity when averaged over the respective fading processes and maximum transmissivity (a best-case loss in the absence of fading). Crucially, it can be shown that for beam trajectories with relatively small zenith angles ($\theta \leq 1$ radian) we can assume the regime of weak turbulence for the ground-satellite fading channel (see Appendix C of [30]). Within this angular window we can accurately resolve the fading dynamics, and by inserting the beam trajectory expressions into the machinery of Sections 5.3.2 - 5.3.5, it is possible to derive loss-based ultimate limits for both uplink and downlink quantum communications using the Eq. (5.22), and thermal-loss-based limits using Eq. (5.26).

The sources of environmental thermal-noise are also unique to both uplink and downlink configurations, and operational settings such as the time of day and weather. In uplink during the day, the primary source of thermal-noise is sunlight being reflected from the Earth to the satellite detector. Meanwhile, at night, this noise is diminished but there still exists sunlight being reflected from the Moon to the Earth and back towards the satellite. For uplink, we may write [31],

$$\bar{n}_B^{\text{up}} = \kappa H_\lambda^{\text{sun}} \Gamma_R. \quad (5.35)$$

Here κ is a parameter that accounts for the Earth/Moon albedos and ranges from $\kappa_{\text{night}} = 7.36 \times 10^{-7}$ for a clear night with a full Moon, to $\kappa_{\text{day}} = 0.3$ during clear day-time. Meanwhile, for the optical wavelength $\lambda = 800$ nm, we can approximate that in uplink the solar spectral irradiance is $H_\lambda^{\text{sun}} = 4.61 \times 10^{18}$ photons $\text{m}^{-2} \text{s}^{-1} \text{nm}^{-1} \text{sr}^{-1}$.

For downlink, the receiver is now a detector on the ground and the main source of noise is more simply attributed to the sky (as it was in the ground-based scenario). In this

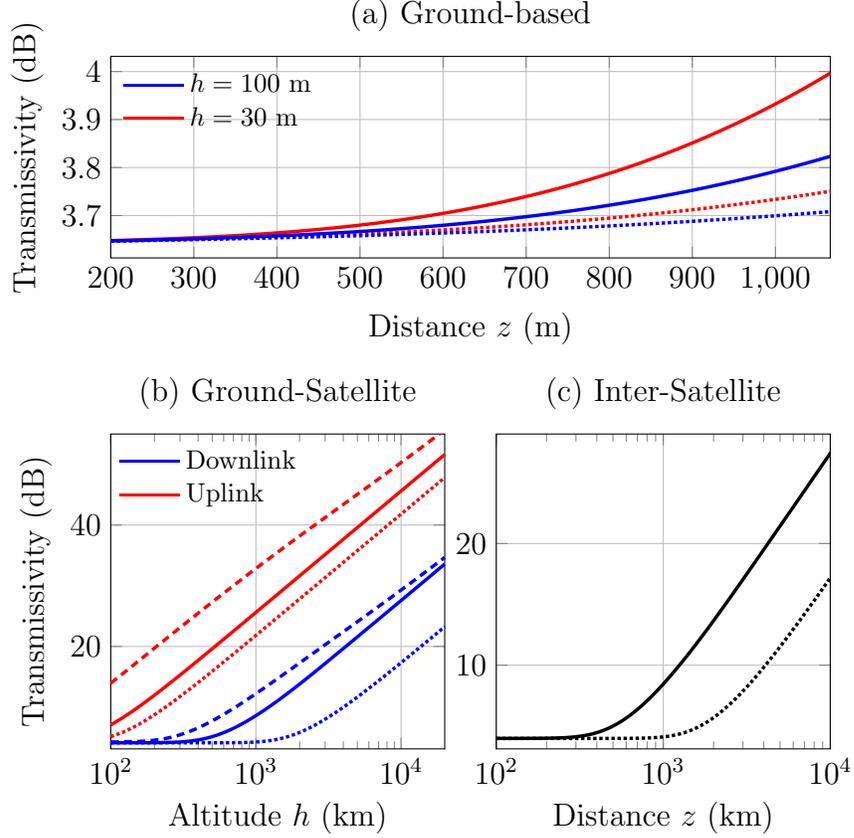


Figure 5.1: Free-space transmission loss associated with (a) ground-based, (b) ground-satellite and (c) intersatellite communication links. In each plot, solid lines depict the average transmissivity (attenuation averaged over fading dynamics), while dotted lines describe the best-case transmissivity (absence of fading). The dashed lines in Panel (b) describe a ground-satellite free-space link with zenith angle $\theta = 1$ radian, while the others consider $\theta = 0$. The operational setup in (a) is consistent with the parameters in Table 5.2 while (b) and (c) are consistent with Setup (#1) in Table 5.1.

setting, and for $\lambda = 800$ nm, the spectral irradiance of the sky follows Eq. (5.31). For a much more detailed analysis, see Appendix D, Ref. [31].

5.3.8 Intersatellite Channels

Finally, we can consider free-space quantum communication between satellites in orbit. This represents a high quality free-space quantum channel which is free from atmospheric decoherence, and thus will not experience losses due to extinction nor undergo turbulence. Indeed, these intersatellite link losses are characterised by free-space diffraction only. Assuming negligible pointing errors, then the intersatellite channel is simply a lossy channel

with transmissivity given by $\eta_d(z)$ as a function of the propagation distance between satellites, z . This lets us write an ideal upper-bound on the intersatellite channel capacity,

$$\mathcal{C} \leq -\log_2(1 - \eta_d) = \frac{2a_R^2}{w_d^2(z) \ln 2}. \quad (5.36)$$

Due to the lack of onboard access and adaptive optics, it is possible that pointing errors become important and must be considered. If pointing errors are non-negligible, $\sigma_p^2 > 0$, then we instead must consider a lossy fading channel $\mathcal{E}_{F\sigma_p} = \{F_{\sigma_p}; \mathcal{E}_\tau\}$ with maximum transmissivity $\eta_d(z)$. As discussed in earlier sections, pointing errors occur on a sufficiently slow time-scale such that they are reconcilable by the receiver. Hence, the capacity for this channel can be accessed via Eq. (5.22), such that

$$\mathcal{C} \leq \mathcal{B}_{F\sigma_p}^{\eta_d} = \frac{2a_R^2}{w_d^2(z) \ln 2} \Delta(\eta_d(z), \sigma_p), \quad (5.37)$$

where Δ acts as a correction factor to the PLOB bound. It is clear that when $\sigma_p^2 = 0$ we retrieve Eq. (5.36). In Fig. 5.1(c) the loss properties of an optical intersatellite channel are illustrated with respect to distance between communicating satellites. This depicts similar transmissivity behaviour to ground-satellite downlink channels with zenith angle $\theta = 0$ without the additional degradation associated with atmospheric interactions.

We have some important considerations to note. First of all, intersatellite channels can only be formed between satellites that fall within each other's line-of-sight. This naturally implies a limit to the maximum distance over which an intersatellite channel can be physically established. For any two satellites in circular orbits, at some point the Earth blocks the free-space between them, prohibiting transmittance. This is derived through basic geometric considerations conveyed in subsequent sections.

Secondly, let us justify the modelling of intersatellite channels as pure-loss channels. The number of thermal photons impinging upon a satellite detector is determined by the orientation and field of view of the detector. For communication between satellites, the transmitters and detectors do not occupy fixed orientations with respect to the main sources of brightness. Indeed, there will exist best case and worst case orientations: in the best-case scenario, the satellite detector will face completely away from the Earth or Moon, so that their albedos are not within the detector's field of view whatsoever. In a worst case scenario, the detector will be oriented directly facing the Earth (as in uplink).

However, point-to-point quantum communication can always be optimised by choosing the physically directed channel which results in less thermal background photons at the detector, irrespective of the logical direction of communication. Each intersatellite channel

can exchange quantum systems in the direction which achieves the best detector orientation with respect to background noise. By optimising the physical orientation of an intersatellite quantum network, each receiver will only ever experience a fraction of the worst background noise experienced by satellite uplink channels for which thermal corrections are minimal for link lengths of $z \lesssim 10000$ km [31]. We leave more formal treatments of these channel properties to future works, with the confidence that pure-loss channels accurately model such free-space links.

Hence, we can reliably model intersatellite free-space links as pure-loss channels. As such, we treat the upper-bound in Eq. (5.37) as an achievable rate so that $\mathcal{C} = \mathcal{B}_{F_{\sigma_p}}^{\eta_d}$ can be accomplished by an optimal point-to-point protocol.

5.4 Modular Quantum Networks

5.4.1 Network Model

In this chapter, we construct a simple model for the study of modular quantum networks. Namely, we consider a global network $\mathcal{N} = (P, E)$ which consists of a collection of sub-networks called *communities*, where the i^{th} community is denoted by the undirected sub-graph

$$\mathcal{N}_{c_i} = (P_{c_i}, E_{c_i}), \quad P_{c_i} \subset P, \quad E_{c_i} \subset E. \quad (5.38)$$

Here, P_{c_i} defines a subset of all network nodes that compose the i^{th} community, while E_{c_i} denotes the subset of all network edges that connect them. For now, we consider each community network to be completely general, and can adopt an arbitrary topology. We focus on quantum networks which observe *spatial-modularity* [102], such that communities are spatially separated. This means that each community is completely disconnected from every other community, i.e. the community node sets are all pairwise disjoint $P_{c_i} \cap P_{c_j} = \emptyset$, for all i, j .

In order to mediate communication between different communities, we introduce a *backbone network* $\mathcal{N}_b = (P_b, E_b)$. This is a large-scale network for which none of its nodes $\mathbf{x} \in P_b$ are user nodes used purely to facilitate end-to-end communications between users contained in different communities. Crucially, we assume that each community possesses a set of undirected edges which connect a set of community nodes to backbone network nodes. We refer to these as *intercommunity edges*, such that the set of intercommunity edges

$$E_{c_i:b} := \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x} \in P_{c_i}, \mathbf{y} \in P_b\}, \quad (5.39)$$

gives each community access to the backbone.

We can define an intercommunity sub-network $\mathcal{N}_{c_i:b} = (P_{c_i:b}, E_{c_i:b})$, which describe the undirected graph that emerges between the i^{th} community and the backbone. The set $P_{c_i:b}$ defines the complete collection of nodes that are interconnected between the community and the backbone. However, the nodes $\mathbf{x} \in P_{c_i:b}$ are already contained within \mathcal{N}_{c_i} or \mathcal{N}_b . Hence, it is important to distinguish between the community nodes and the backbone nodes which comprise this sub-network. For this, we introduce the notation

$$P_{c_i|b} := P_{c_i:b} \cap P_{c_i} \subseteq P_{c_i}, \quad (5.40)$$

$$P_{b|c_i} := P_{c_i:b} \cap P_b \subseteq P_b. \quad (5.41)$$

Intuitively, $P_{c_i|b}$ can be thought of as the subset of nodes from the community P_{c_i} conditioned on being connected to \mathcal{N}_b (and vice versa for $P_{b|c_i}$).

This modular structure takes a very intuitive form and is remarkably useful for modelling realistic, hybrid quantum networks. When an equivalence relation is enforced between nodes in similar communities, the network quotient graph can be viewed as a star-network². It allows us to completely separate communities and the backbone from one another. This makes it easier to compartmentalise different sub-network structures which may operate in completely different physical domains. Furthermore, it helps to derive independent network conditions on each of the sub-networks in accordance with some global objective. We summarise this architecture in the following definition which has also been illustrated in Fig. 5.2(a).

Definition 5.1 (Modular Network): *A modular network $\mathcal{N} = (P, E)$ is a network architecture constituent of n community sub-networks $\{\mathcal{N}_{c_i}\}_{i=1}^n$, and a backbone sub-network \mathcal{N}_b . Each community sub-network is connected to the backbone via a set of edges $E_{c_i:b}$, described by the intercommunity sub-networks $\{\mathcal{N}_{c_i:b}\}_{i=1}^n$, and there are no direct links between communities.*

5.4.2 Modular Network Capacities

As discussed in Section 2.4, the optimal end-to-end performance within a quantum network is quantified by its flooding capacity $\mathcal{C}(i, \mathcal{N})$, which describes the optimal number of target bits that can be transmitted between end-users per use of a flooding protocol. Any quantum network $\mathcal{N} = (P, E)$, including the modular designs introduced, can be represented

²In Appendix C we present more general aspects of networks with community structures from which this modular network emerges as a useful and highly desirable class

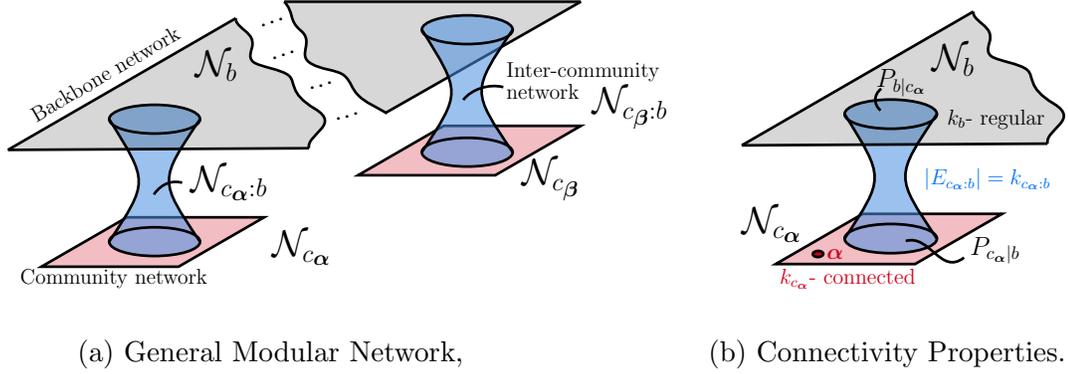


Figure 5.2: (a) A modular quantum network architecture built from community sub-networks \mathcal{N}_{c_α} , \mathcal{N}_{c_β} and a backbone network \mathcal{N}_b . Each community is connected to the backbone via the sub-networks $\mathcal{N}_{c_\alpha:b}$ and $\mathcal{N}_{c_\beta:b}$. Nodes from the community c_j which are directly connected to the backbone are contained in $P_{c_j|b}$, while the nodes in the backbone which are connected to the community are contained in $P_{b|c_j}$. (b) We may idealise this modular structure by placing ideal connectivity constraints on the each of the sub-networks.

by a global distribution of channels $\{\mathcal{E}_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x},\mathbf{y})\in E}$ and a corresponding distribution of single-edge channel capacities. For general fading networks, it is always possible to use these distributions and the general expressions from Eqs. (5.8) and (5.9) in order to determine the flooding capacity.

However, the translation into a modular architecture means that there exist particular classes of network cuts which are performed on different sub-networks. It becomes very useful to formally define a number of the important multi-edge capacities associated with these classes of cuts. In each of the following settings, we consider a pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ contained within remote communities of a generic, global modular network i.e. $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$ such that $c_\alpha \neq c_\beta$ ³. It is now useful to denote community sub-networks with respect to the end-user that they contain, i.e. we may write c_α and c_β respectively. We assume each sub-network adopts arbitrary topologies and capacity distributions.

Local-Community Capacities

We define a *local-community cut* C_{c_j} as that which partitions two end-users within the network by exclusively collecting edges within one of the user communities c_j , for either $j \in \{\alpha, \beta\}$. That is, a local-community cut-set takes the form $\tilde{C}_{c_j} = \{(\mathbf{x}, \mathbf{y}) \in E_{c_j} \mid \mathbf{x} \in$

³Note that in this chapter we choose to label end-user nodes via α and β rather than \mathbf{a} and \mathbf{b} . This is a stylistic choice used to better distinguish between end-user nodes and the backbone network which we is identified via the subscript b .

$\mathbf{A}, \mathbf{y} \in \mathbf{B}\}$. This restricted form of network cut will generate an associated multi-edge capacity, which we label a *local-community capacity*,

$$\mathcal{C}_{c_j} := \mathcal{C}_{\text{cut}}(C_{c_j}) = \min_{C_{c_j}} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}_{c_j}} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (5.42)$$

For end-user nodes \mathbf{j} which do not share a direct connection with the backbone (i.e. $\mathbf{j} \notin P_{c_j|b}$) then this form of restricted cut always exist.

However, if an end-user node does share a direction connection with the backbone, then a valid local-community cut will not exist. In this case it is never sufficient to remove edges solely from the community networks, and one must cut at least one edge from the set of intercommunity edges $E_{c_j:b}$. To this end, we must slightly modify the local-community cut so that it removes any direct connections from the user node to the backbone, and then to identify the optimal set of edges to be removed from the community. Hence, a valid cut-set becomes $\tilde{C}'_{c_j} = \{(\mathbf{j}, \mathbf{y}) \in E \mid \mathbf{y} \in P_b\} \cup \tilde{C}_{c_j}$. We can then define an analogous local-community capacity according to this class of network cut.

Backbone Capacities

A *backbone cut* C_b is a network cut that exclusively collect edges on the backbone network in order to partition the two end-users. This kind of cut-set takes the form $\tilde{C}_b = \{(\mathbf{x}, \mathbf{y}) \in E_b \mid \mathbf{x} \in \mathbf{A}, \mathbf{y} \in \mathbf{B}\}$, which generates an associated multi-edge *backbone capacity*,

$$\mathcal{C}_b := \mathcal{C}_{\text{cut}}(C_b) = \min_{C_b} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}_b} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (5.43)$$

When considering end-users contained in unique communities within the modular network architecture we are investigating, these kinds of cuts always exist. It is always sufficient to perform a cut on the backbone since there does not exist any other collection of edges that can be used to form a valid path between communities.

Global-Community Capacities

Finally, we can formalise a multi-edge capacity associated with exclusively collecting intercommunity edges. The end-user communities \mathcal{N}_{c_α} and \mathcal{N}_{c_β} are connected to the backbone via the sets of intercommunity edges $E_{c_\alpha:b}$ and $E_{c_\beta:b}$ respectively. If we removed either of these sets of edges, then the two remote users would be automatically partitioned. Hence, the edge sets $E_{c_\alpha:b}$ and $E_{c_\beta:b}$ both correspond to valid cuts on the network and each

generate a multi-edge capacity

$$\mathcal{C}_{c_j:b} := \mathcal{C}_{\text{cut}}(E_{c_j:b}) = \sum_{(\mathbf{x}, \mathbf{y}) \in E_{c_j:b}} \mathcal{C}_{\mathbf{x}\mathbf{y}}, \quad (5.44)$$

for $j \in \{\alpha, \beta\}$. We can then minimise over the end-users to define a multi-edge capacity,

$$\mathcal{C}_{c:b} := \min_{j \in \{\alpha, \beta\}} \mathcal{C}_{c_j:b}. \quad (5.45)$$

Clearly, this form of network cut always exists. We refer to this kind of partitioning as community isolation, since it isolates a community sub-network entirely from the rest of the network. Furthermore, we name $\mathcal{C}_{c:b}$ the *global-community capacity*, as it refers to globally isolating the entire community sub-network⁴.

5.4.3 Idealised Modular Networks

Arbitrary architectures can always be treated using the capacity expressions from Section 5.2 for general fading networks. However, the generality of these arguments make it difficult to present rigorous analytical statements about specific features or tangible network properties. In order to understand the ultimate potential of quantum networks, we need to simultaneously optimise the point-to-point channels and the network architecture in which they are arranged. Hence, it is desirable to strike a balance between realism and ideality in such a way that allows us to derive informative results about quantum networks and end-to-end performance. In the following we propose sub-network connectivity constraints that strike this balance.

Backbone Regularity

Firstly, we can impose regularity on the network backbone, demanding that the degree of each node is constant. This leads to a highly-connected network structure which is ideal for multi-path routing strategies. Let the function $\text{deg}(\mathbf{x})$ compute the degree of the node \mathbf{x} . Then we impose

$$\text{deg}(\mathbf{x}) = k_b, \forall \mathbf{x} \in P_b, \quad (5.46)$$

which defines the regularity parameter of the backbone. It is important to make clear that these constraints only apply to intra-network connections. Indeed, a node on the backbone

⁴While it might be more convenient to call this the *intercommunity capacity*, such a name might be confused as a more general term for the capacity when the end-users are located in different communities (which is implied). The global-community capacity is intended to be more distinct than this as it specifies a particular network cut.

can have k_b connections to neighbours on the backbone network, but also possess additional intercommunity connections via the sub-network $\mathcal{N}_{c_j:b}$, without any further constraint. It is useful to quantify the number of intercommunity connections permitted between the backbone and communities using the notation,

$$k_{c_j:b} = |E_{c_j:b}|, \quad \mathbf{j} \in \{\boldsymbol{\alpha}, \boldsymbol{\beta}\}. \quad (5.47)$$

While regularity is an idealised property of realistic networks, in the context of a non-user repeater network such as the backbone, it is feasible and extremely useful in order to understand the limits of quantum networks.

Community Connectivity

Community sub-networks are likely to be smaller scale and less predictable structures than the backbone, partly due to the presence of user-nodes. Flexibility in their design is important. Here, we do not impose regularity but instead define classes of communities in accordance with the smallest local-community cut that they contain.

Definition 5.2 (k_c -connectivity): *Consider a community sub-network \mathcal{N}_c . We say the community is k_c -connected if k_c is the smallest number of edges that must be removed in order to disconnect a pair of community nodes, minimised over all possible node pairs $\mathbf{x} \neq \mathbf{y} \in P_c$. More precisely,*

$$k_c := \min_{\mathbf{x} \neq \mathbf{y} \in P_c} |\tilde{C}_c|, \quad (5.48)$$

where \tilde{C}_c denotes a community cut-set between the nodes \mathbf{x} and \mathbf{y} .

Hence, k_c defines the minimum local community cut-set cardinality, given some network topology and choice of end-users. This is a completely general property which is unique for all community networks, using the most easily disconnected pair of nodes in the network as a metric for how well it is connected. Regular networks are an example of an architecture for which their k_c -connectivity is simply equal to the network regularity. Hence, we can consider community sub-networks to be k_c -connected while encompassing a very large set of architectures.

Idealised Modular Network

Combining the constraints of regularity on the backbone and k_c -connectivity on the community sub-networks, it is possible to define an ideal modular quantum network architecture in terms of these parameters. This generates a structure that can be investigated analytically in the following sections.

Definition 5.3 (Ideal Modular Network): *An ideal modular network $\mathcal{N}^* = (P, E)$ is a network architecture constituent of n -community sub-networks $\{\mathcal{N}_{c_i}\}_{i=1}^n$ each of which are k_{c_i} -connected, and a backbone sub-network \mathcal{N}_b which is k_b -regular. Each community sub-network is connected to the backbone via $k_{c_i:b}$ edges, described by the intercommunity sub-networks $\{\mathcal{N}_{c_i:b}\}_{i=1}^n$, and there are no direct links between communities.*

An illustration of this architecture can be found in Fig. 5.2(b). When focusing on a particular pair of end-user nodes $\mathbf{i} = \{\alpha, \beta\}$ from two remote communities in the global network, we can then specify their k_{c_j} -connectivity properties.

5.4.4 Minimum-Cut as Community Isolation

Care must be taken when constructing this form of modular network so to ensure not only high-rate communication within each community, but also high-rate communication between different communities mediated by the backbone. If the backbone network is poorly connected or possesses weak links, it will not effectively assist long-distance communication. Meanwhile, even if communities are connected to a high quality backbone, insufficiently strong capacities in a local-community can compromise its use. Hence, there exists a careful balance between all of the sub-networks in the modular model, and their connectivity/capacity properties throughout. It is therefore highly desirable to identify a relationship between the quality of channels within the backbone and the quality of channels within the communities.

In order to better grasp these relationships, we can investigate the ideal modular networks \mathcal{N}^* defined in Definition 5.3. Regular networks (such as that on the backbone) possess very convenient qualities which allow for useful insight into minimum network cuts. As such, they can be analytically studied as highly connected, ideal network structures and used to reveal fundamental limitations for end-to-end communication.

Our mission becomes the following: to derive conditions on each of the sub-networks such that the flooding capacity between the remote users is always their global-community capacity. In this way, the minimum cut is always achieved by community isolation on either of the end-user communities. Equivalently, it means that the minimum cut can always be found on a simplified quotient graph of the modular network, vastly simplifying our analyses. When this is the case, the end-to-end capacities between any two unique communities are always *distance-independent*, i.e. the ultimate rate between two end-user communities does not change with respect to the physical separation of those communities. This is an highly desirable property of a quantum network, particularly on large-scales.

If a modular network satisfies this property, it means that (i) the backbone network is of sufficiently high quality that it never impedes the network performance over (potentially very) long distances, and (ii) that the local-communities are of sufficiently high quality that neither compromises local or network-wide communication. By imposing that the minimum cut be the intercommunity edges, we can reveal unique constraints on each sub-network which are summarised in the following threshold theorem.

Theorem 5.1 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3. Select any pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ contained in remote communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$. For all $\mathbf{j} \in \{\alpha, \beta\}$, there exist single-edge threshold capacities on the communities $\mathcal{C}_{c_j}^{\min}$ and backbone \mathcal{C}_b^{\min} sub-networks for which the network flooding capacity is given by the global-community capacity,*

$$\left. \begin{aligned} \mathcal{C}_{\mathbf{x}\mathbf{y}} &\geq \mathcal{C}_{c_j}^{\min}, \forall (\mathbf{x}, \mathbf{y}) \in E_{c_j}, \\ \mathcal{C}_{\mathbf{x}\mathbf{y}} &\geq \mathcal{C}_b^{\min}, \forall (\mathbf{x}, \mathbf{y}) \in E_b, \end{aligned} \right\} \implies \mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{c:b}. \quad (5.49)$$

The threshold capacities are given by,

$$\mathcal{C}_{c_j}^{\min} := \frac{\mathcal{C}_{c:b}}{k_{c_j}}, \quad \mathcal{C}_b^{\min} := \frac{\mathcal{C}_{c:b}}{H_{\min}^*}, \quad (5.50)$$

where H_{\min}^* is the minimum cut-set cardinality on the backbone network. If these threshold capacities are violated, then the global-community capacity becomes an upper-bound on the end-to-end capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c:b}$.

A detailed proof can be found in Section C.1.6. Thanks to backbone regularity and community connectivity, the minimum cut-set cardinalities that occur within each sub-network can be easily identified. Then, it is straightforward to enforce single-edge capacity constraints which ensure that the local-community/backbone capacities are always larger than the global-community capacity.

In this theorem, we have used the fact that the cardinality of the smallest backbone cut-set between two end-users in remote communities can be analytically derived as a consequence of network regularity. This minimum cardinality takes the form

$$H_{\min}^* := \min_{\mathbf{j} \in \{\alpha, \beta\}} H_{\min}(k_b, P_{b|c_j}), \quad (5.51)$$

where $H_{\min}(k_b, P_{b|c_j})$ is a function that computes the minimum number of edges that must be cut to isolate all the nodes $P_{b|c_j}$ on the backbone which are also connected to the community c_j . The explicit form of this expression is found in Section C.3, and depends

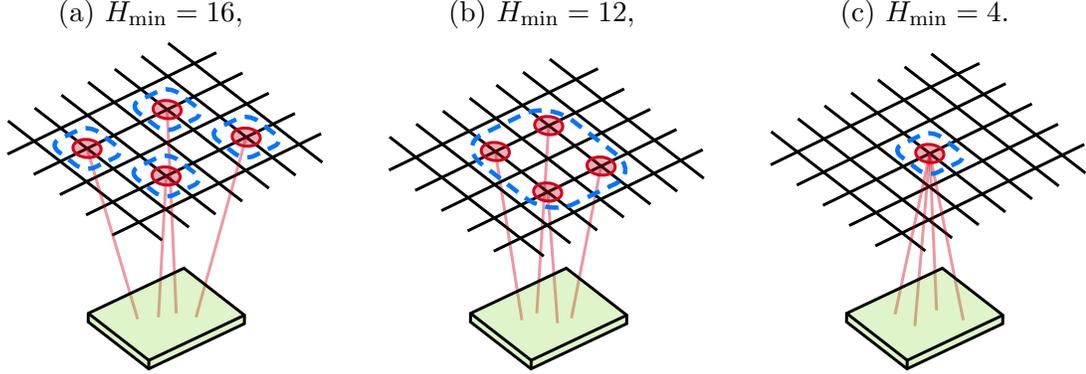


Figure 5.3: Examples of minimum cardinality intercommunity cut-sets for connections from an arbitrary community to a Manhattan backbone network ($k_b = 4$). These are valid cuts which isolate remote communities (only one community is illustrated here), and are performed exclusively on the backbone. Panel (a) captures the best-case spatial distribution of the largest potential cut-set when no target-nodes share any edges or neighbours, (b) illustrates an example in which neighbour sharing can diminish the overall cut-set size, and (c) describes the worst-case spatial distribution that minimises the cut-set size.

on the precise spatial arrangement of connections from the community to the backbone. However, we can generally bound this quantity using

$$k_b \leq H_{\min}(k_b, P_{b|c_j}) \leq k_b |P_{b|c_j}|. \quad (5.52)$$

The lower-bound k_b corresponds to a worst-case spatial distribution of community-to-backbone connections, when all the community nodes are connected to the same node on the backbone, i.e. $|P_{b|c_j}| = 1$. Then it is sufficient to isolate just one backbone node to perform a valid end-user cut, collecting only k_b edges (since the backbone is k_b -regular). The upper-bound corresponds to a best-case scenario; when all the community nodes are connected to backbone nodes which don't share any neighbours or edges. In this case, the smallest cut-set restricted to the backbone is found by isolating all nodes individually. As a result, this cut collects exactly $k_b |P_{b|c_j}|$ edges.

Fig. 5.3 depicts a number of examples of minimum backbone cut-sets for remote communities connected to a Manhattan backbone ($k_b = 4$). In these figures we display only one end-user community and assume that the other end-user community is sufficiently distant that it does not share intercommunity connected nodes on the backbone.

As a result, we can always present a best and worst-case single-edge threshold capacity for the backbone network, C_b^{\min} . We can sandwich the backbone threshold capacity according to

$$\frac{C_{c:b}}{k_b |P_{b|c_j}|} \leq C_b^{\min} \leq \frac{C_{c:b}}{k_b}, \quad \mathbf{j} \in \{\alpha, \beta\}. \quad (5.53)$$

<i>Parameter</i>	<i>Symbol</i>	<i>Value</i>
Beam Curvature	R_0	∞
Wavelength	λ	800 nm
Initial spot-size	w_0	40 cm - Setup (#1) 20 cm - Setup (#2)
Receiver Aperture	a_R	1 m - Setup (#1) 40 cm - Setup (#2)
Detector Efficiency	η_{eff}	0.4
Detector Noise	\bar{n}_{ex}	≈ 0
Pointing error	σ_p^2	$1 \mu\text{rad} \approx (10^{-6}z)^2$
Pulse Duration	Δt	10 ns
Field of View	Ω_{fov}	10^{-10} sr
Frequency Filter	$\Delta\lambda$	0.1 pm - Setup (#1) 1 nm - Setup (#2)
Intercommunity Link	ICL	Downlink
Fibre Loss-Rate	γ	0.02 per km

Table 5.1: Parameter table for the fibre/satellite modular network configuration. Here we consider two similar setups using a collimated Gaussian beam at 800 nm wavelength, but differ in initial spot-size w_0 , receiver aperture a_R and frequency filter $\Delta\lambda$.

The more effectively that the intercommunity connections are dispersed across the backbone, the weaker the single-edge constraint that must be forced upon it.

5.5 Fibre/Satellite Configuration

5.5.1 Motivation

An interesting modular configuration consists of fibre-based community networks which are interconnected via a backbone satellite network. This model captures a realistic satellite-based model of the quantum internet, in which dynamic intersatellite links are used to facilitate long distance quantum communication at high rates. In this scenario, the weakest links are typically the ground-to-satellite free-space connections owing to the impact of atmospheric decoherence and turbulence on a transmitted beam. Therefore, the constraints revealed in Theorem 5.1 are realistic as community isolation is likely to be the minimum cut in many settings.

In Theorem 5.1 we devise single-edge capacity lower bounds on the community networks which guarantee the network flooding capacity is equal to the global-community capacity. For fibre-based networks, these single-edge lower bounds can be used to identify a *maximum*

tolerable fibre-length $d_{c_j}^{\max}$ that is permitted within the fibre-network. In the context of a satellite-based backbone network, the single-edge capacity lower bound can be translated into a *maximum intersatellite separation*, z_b^{\max} which describes the maximum propagation distance that is permitted for free-space channels between satellites in the backbone. These are critical quantities which directly motivate the construction of ground-based and satellite-based networks for global quantum communication.

5.5.2 Optimal Performance

We wish to enforce that the minimum cut is always achieved by community isolation, generating the global-community capacity $\mathcal{C}(i, \mathcal{N}) = \mathcal{C}_{c:b}$. In this physical setting, each intercommunity edge is described by ground-to-satellite channel which may be an uplink or downlink channel. Thanks to teleportation, a network protocol can always choose the physical channel direction that maximises its point-to-point capacity independently from the desired logical direction of community. Downlink channels are always superior to uplink, and therefore we can simply model the global community capacity as the sum of a downlink capacities. This multi-edge capacity will be bounded by

$$\mathcal{C}_{c:b} \leq \min_{j \in \{\alpha, \beta\}} \sum_{(\mathbf{x}, \mathbf{y}) \in E_{c_j:b}} \mathcal{T}_{F_{\mathbf{x}\mathbf{y}}}(\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_j), \quad (5.54)$$

$$\leq \min_{j \in \{\alpha, \beta\}} \sum_{(\mathbf{x}, \mathbf{y}) \in E_{c_j:b}} \mathcal{B}_{F_{\mathbf{x}\mathbf{y}}}(\eta_{\mathbf{x}\mathbf{y}}), \quad (5.55)$$

where $F_{\mathbf{x}\mathbf{y}}$ and $\eta_{\mathbf{x}\mathbf{y}}$ capture the fading dynamics and maximum transmissivity of each downlink channel that connect c_j to the backbone, and depend on beam trajectory. Meanwhile, \bar{n}_j infers community-wide thermal-noise conditions. Since all of the intercommunity edges in $E_{c_j:b}$ are connected to a relatively small area, we can assume identical operational conditions for all downlink edges. However, these operational conditions will not be consistent for both end-users; when communicating on a global scale, one user may be in night-time while the other is in day-time with independent weather conditions.

We can derive single-link distance constraints which guarantee $\mathcal{C}_{c:b}$ to be the optimal network capacity. These conditions follow directly from Theorem 5.1 and are summarised in the following corollary:

Corollary 5.1 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3, and assume optical-fibre communities networks \mathcal{N}_{c_α} , \mathcal{N}_{c_β} and a satellite-based backbone \mathcal{N}_b . Select any pair of end-users $\{\alpha, \beta\}$ located in remote communities $\alpha \in P_{c_\alpha}$*

and $\beta \in P_{c_\beta}$. There exists a maximum fibre-length in each community

$$d_{c_j}^{\max} := -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-C_{c:b}/k_{c_j}} \right), \quad (5.56)$$

and a maximum intersatellite separation in the backbone

$$z_b^{\max} := \arg \min_z \left| \log \left(\frac{H_{\min}^* \mathcal{B}_{F_{\sigma_p}}(\eta)}{C_{c:b}} \right) \right|, \quad (5.57)$$

for which the network flooding capacity is equal to the global-community capacity,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = C_{c:b}. \quad (5.58)$$

Otherwise, if any intersatellite links violate this condition $\exists z_{\mathbf{xy}} > z_b^{\max}, (\mathbf{x}, \mathbf{y}) \in E_b$ or the local community links are in violation, $\exists d_{\mathbf{xy}} > d_{c_j}^{\max}, (\mathbf{x}, \mathbf{y}) \in E_{c_j}$, for either $j \in \{\alpha, \beta\}$, then this becomes an upper-bound on the network flooding capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq C_{c:b}$.

The analytical simplicity of the maximum fibre-length follows from the remarkably compact PLOB bound for bosonic lossy channels. However, the maximum intersatellite separation in Eq. (5.57) must be computed numerically due to the more complex PLOB bound which accounts for fading due to pointing errors. The lack of onboard access makes it difficult to perfectly optimise beam trajectory, and thus pointing errors cannot be ignored. However, it is possible to analytically upper and lower bound the quantity z_b^{\max} .

5.5.3 Analytical Bounds for the Maximum Intersatellite Separation

The maximum intersatellite separation z_b^{\max} describes a maximum *tolerable* channel length permitted within the backbone network. Yet, it is not always true that such a channel length is achievable due to line-of-sight limitations associated with orbital geometry. This is quantified by the maximum line-of-sight distance which is a function of the altitudes of the communicating satellites. Consider two satellites in circular orbits around the Earth at altitudes. First, suppose the two satellites are at positions A and B which have equivalent altitudes h . By drawing a chord AB tangential to the Earth's surface from one satellite to the other, we can identify the maximum line-of-sight separation, z_{sight}^{\max} (see Fig. 5.4 for a geometrical insight). Label the centre of the Earth E , its radius R_E and the point at which the chord touches the Earth's surface S . Denoting the angle $\angle EAS$ as α , this will satisfy

$$\sin \alpha = \frac{R_E}{R_E + h}. \quad (5.59)$$

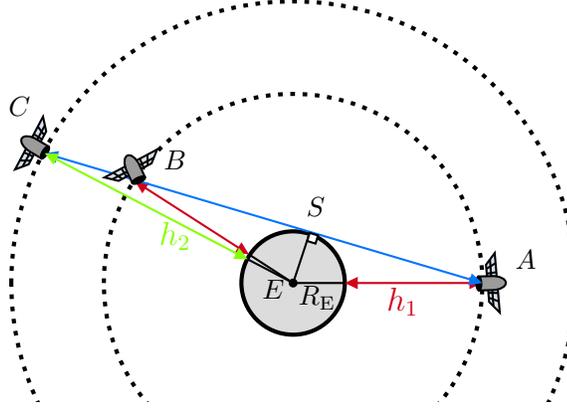


Figure 5.4: The longest possible intersatellite quantum channel is limited by the line-of-sight separation of two satellites.

With this in hand, we find that the distance $AB = 2AS$, since the triangle EBA is clearly isosceles. That is,

$$z_{\text{sight}}^{\max} = 2(h + R_E) \cos \alpha = \frac{2h(h + 2R_E)}{h + R_E}. \quad (5.60)$$

For two satellites which are at different altitudes $h_1 \leq h_2$, this maximum distance is extended to

$$z_{\text{sight}}^{\max} = \frac{h_1(h_1 + 2R_E)}{h_1 + R_E} + \frac{h_2(h_2 + 2R_E)}{h_2 + R_E}, \quad (5.61)$$

which follows intuitively from the previous geometrical considerations. Crucially, if we find that $z_b^{\max} \geq z_{\text{sight}}^{\max}$ for some network configuration and desirable rate, this means that the satellites within the backbone can reliably communicate with *any* other satellite that fall within its line-of-sight, without compromising performance. This is an extremely useful property, providing significant flexibility for satellite backbone networks.

With this limitation in mind, an upper-bound is found by considering a lack of pointing errors, which means the channel is no longer a fading channel but is instead a fixed lossy channel with the maximum possible transmissivity. This idealises the intersatellite channel by removing the potential for beam wandering, resulting in an upper-bound for the maximum separation. Using Corollary 5.1 we can always write the upper-bound,

$$z_b^{\max} \leq z_R \sqrt{\frac{2a_R^2}{w_0^2 \ln \left[\frac{\eta_{\text{eff}}}{\eta_{\text{eff}} - 1 + 2^{-C_{c:b}/H_{\text{min}}^*}} \right]} - 1}, \quad (5.62)$$

derived using the ideal pure-loss single-edge capacity upper-bound from Theorem 1. Meanwhile, we can find a lower-bound on the maximum intersatellite separation by considering the use of slow detectors. A slow detector at the receiver will not be able to resolve pointing errors, resulting in a lossy channel with fixed transmissivity averaged over the entire fading process. For intersatellite channels, this is considered through the long-term spot size, by replacing the ideal diffraction limited spot size w_d^2 with $w_{\text{lt}}^2 = w_d^2 + \sigma_p^2$ into the capacity formula,

$$C \leq \mathcal{B}_{\text{slow}}(\eta_{\text{lt}}) = \frac{2a_R^2}{w_{\text{lt}}^2 \ln 2}. \quad (5.63)$$

Interestingly, the rate in bits *per channel use* via slow detection can be higher than that for fast detectors which actually resolve the fading dynamics. But do not be mistaken; the slower detection time severely limits the operational rate at which the channel can actually be used (or clock rate). The point-to-point communication rate via slow detection will be orders of magnitude smaller than those with fading-resolving setups. Therefore, it is essential to explicitly consider the clock rate α (channel uses/second) when comparing fast and slow detector protocols [30, 31]. In any case, the maximum intersatellite separation will be lower-bounded by

$$z_b^{\text{max}} \geq \sqrt{\frac{2a_R^2}{\left(\frac{w_0^2}{z_R^2} + \epsilon_p^2\right) \ln \left[\frac{\eta_{\text{eff}}}{\eta_{\text{eff}} - 1 + 2^{-\tilde{\alpha} C_{c:b}/H_{\text{min}}^*}} \right]}} - w_0^2, \quad (5.64)$$

where $\epsilon_p = 10^{-6}$ comes from the point error variance $\sigma_p^2 = (\epsilon_p z)^2$, and $\tilde{\alpha} = \alpha_{c:b}/\alpha_b$ is the ratio between the clocks used by the intercommunity sub-network and the backbone network.

5.5.4 Discussion

Fig. 5.5 offers insight into the constraints proposed by Corollary 5.1 for satellite-fibre modular networks corresponding to a number of different physical settings and network properties. Here we consider two free-space communication setups described in Table 5.1: Setup (#1) in Figs. (a) and (c) and Setup (#2) in Figs. (b) and (d).

Consider a flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N})$ that is desired between the two end-users who are located in remote, fibre communities. The actual ground distance between the users or unique communities is irrelevant, and can be arbitrarily situated at any location across the Earth. If that flooding capacity is to be achieved, then for a given modular architecture there exists a maximum fibre-length $d_{c_j}^{\text{max}}$ permitted within the user community c_j , and a maximum intersatellite separation z_b^{max} permitted throughout the backbone network.

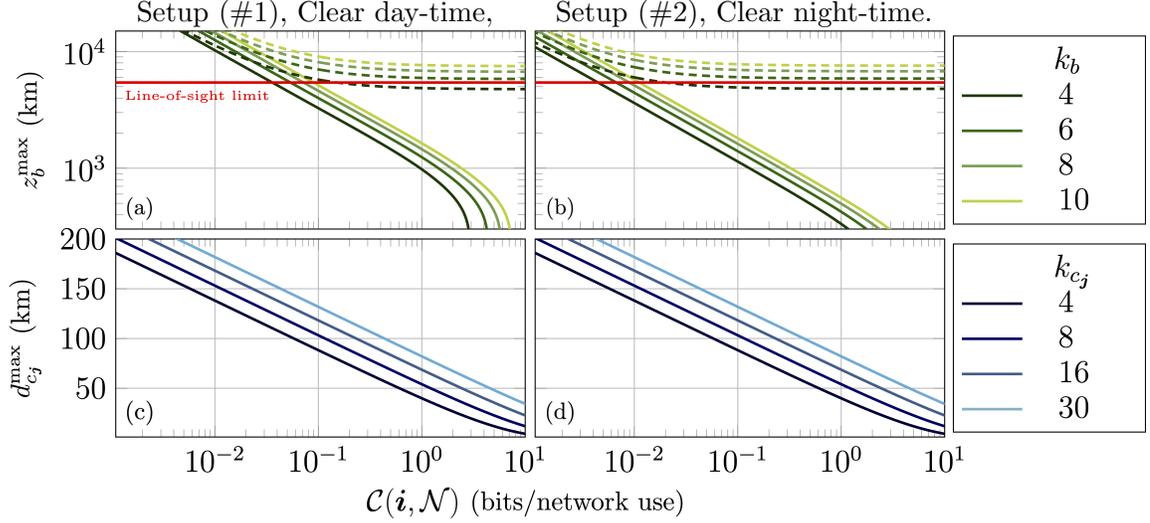


Figure 5.5: Optimal end-to-end performance for an ideal modular network consisting of fibre communities interconnected to a satellite-based backbone. In order to guarantee an optimal flooding rate along the x -axis then the maximum internodal separations in each sub-network on the y -axis must be less than or equal to the plotted bounds. We consider operational settings in Setup (#1) for (a), (c) and Setup (#2) for (b), (d) which are described in Table 5.1. The weather/time conditions are those experienced by the worst-case end-user community. Given an optimal flooding capacity $\mathcal{C}(i, \mathcal{N})$, we plot the maximum intersatellite separation z_b^{\max} for different backbone connectivity parameters, and the maximum fibre-length in each community $d_{c_j}^{\max}$ for different community connectivity parameters. The dashed lines in Figs (a) and (c) plot an upper bound the maximum intersatellite separation based on the optimal spatial distribution of (a finite number of) community connected satellite nodes $P_{b|c_j}$ at a maximum altitude $h^{\max} = 1500$ km, while the solid lines plot the lower bound based on the worst spatial distribution (for any altitude). The red line indicates the maximum achievable channel length that can be achieved for two satellites at altitude 1500 km, such that $z_{\text{sight}}^{\max} \approx 5428$ km.

In Figs. 5.5(a) and (b) we plot the behaviour of the maximum intersatellite separation with respect to desired flooding capacity. In the solid lines, we plot the worst-case z_b^{\max} , which corresponds to the situation where all the downlink channels are connected to the same node on the backbone, allocating a single satellite to connect to a community. This is a worst-case situation because it means that the minimum cut on the backbone is very small, $H_{\min}^* = k_b$. Yet, even in this scenario, the lack of atmospheric decoherence means that very large distances are permitted between satellites, such that $z_b^{\max} \sim 10^3 - 10^4$ km can still ensure high flooding rates between the end-users communities on the Earth.

Meanwhile, the dashed lines plot z_b^{\max} for the best-case spatial distribution of downlink connections on the backbone when the maximum satellite altitude is $h^{\max} = 1500$ km and

all downlink beam trajectories are within a 1 steradian angular window. This means that the smallest backbone cut-set has the total number of edges,

$$H_{\min}^* = k_b |P_{b|c_j}|. \quad (5.65)$$

In this case, the minimum cut-set cardinality on the backbone is very large, as the number of downlink channels must be increased in order to obtain the chosen flooding capacity. In this best-case scenario, as $\mathcal{C}(\mathbf{i}, \mathcal{N})$ increases z_b^{\max} begins to plateau, permitting large intersatellite separations even at large flooding capacities. This confirms a strong dependence between the distribution of intercommunity edges and the single-edge capacity properties of a backbone network. For all other distributions of intercommunity connections $P_{b|c_j}$, the behaviour of the maximum intersatellite separation falls between these bounds.

We also display the maximum line-of-sight distance $z_{\text{sight}}^{\max} \approx 5428$ km between any pair of satellites orbiting at an altitude $h^{\max} = 1500$ km. This is the longest intersatellite channel that can be established due to orbital geometry. Interestingly, even in the worst-case backbone configuration (each community possesses many connections to a single satellite) the line-of-sight limit is exceeded by z_b^{\max} at relatively good rates such that $\mathcal{C}(\mathbf{i}, \mathcal{N}) \in [10^{-2}, 10^{-1}]$ bits per network use. When $z_b^{\max} \geq z_{\text{sight}}^{\max}$ is true, satellites in the backbone may connect to any other satellite within its line-of-sight. This promises achievable and flexible constraints for intersatellite networks.

Figs. 5.5(c) and (d) depict the maximum fibre-lengths permitted within k_{c_j} -connected community networks to ensure a desired end-to-end flooding capacity. Of course, the quality of the bosonic lossy channels do not change with respect to Setups (#1) and (#2) and therefore Figs. (c) and (d) are identical. As one would expect, the permissible channel lengths for strong end-to-end rates depend upon the community channels being $d_{c_j}^{\max} \lesssim 100$ km, even in a highly connected network setting. But thanks to the modular network configuration, this is not problematic. In this configuration, the community fibre-networks are designed to cover small areas relative to the satellite backbone and facilitate local communication. Quantum communication over global distances is then appropriately mediated by the satellite backbone.

As an example, let us focus on Setup (#1) and consider a satellite backbone network with regularity $k_b = 4$ used to mediate long-distance quantum communication between two end-users $\mathbf{i} = \{\alpha, \beta\}$ contained within fibre-networks which are $k_{c_\alpha} = 4$ and $k_{c_\beta} = 8$ connected. What are the network constraints required to ensure that their flooding capacity is $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 1$ bit per network use? Provided that $z_b^{\max} \lesssim 1000$ km, that $d_{c_\alpha}^{\max} \lesssim 30$ km and $d_{c_\beta}^{\max} \lesssim 50$ km, then it is guaranteed that this flooding rate is achievable. This provides extremely valuable information for future quantum network designs; if an ideal modular

<i>Parameter</i>	<i>Symbol</i>	<i>Value</i>
Beam Curvature	R_0	∞
Wavelength	λ	800 nm
Initial spot-size	ω_0	5 cm
Receiver Aperture	a_R	5 cm
Detector Efficiency	η_{eff}	0.5
Detector Noise	\bar{n}_{ex}	0.05
Pointing error	σ_p^2	$1 \mu\text{rad} \approx (10^{-6}z)^2$
Pulse Duration	Δt	10 ns
Field of View	Ω_{fov}	10^{-10} sr
Frequency Filter	$\Delta\lambda$	1 nm
Altitude	h	30 m
Fibre Loss-Rate	γ	0.02 per km
Intercommunity Link	ICL	Free-Space (Clear day-time)

Table 5.2: Parameter table for the free-space/fibre modular network configuration.

network cannot exceed these constraints, then less ideal structures should take even stronger heed of them.

5.6 Ground-Based Free-Space/Fibre Configuration

5.6.1 Motivation

It is also interesting to investigate the limits of ground-based quantum networks which are composed from a mixture of fibre channels and free-space channels. For this purpose, modular network architectures offer an appropriate and physically relevant model. One may consider a metropolitan network area which is spanned by a collection of free-space quantum networks, or “hotspots”. These are short-range communities within which reliable free-space quantum communications can take place. In order to communicate over a larger area and between free-space communities we can use an underlying optical-fibre backbone which mediates longer distance communication.

Utilising the recently derived ultimate limits of ground-based, free-space quantum communication [30] we wish to determine whether free-space links are reliable enough to enable high-rate quantum communication in this setting. Furthermore, it is important to understand the requirements of the optical-fibre backbone required to facilitate wireless quantum networking.

5.6.2 Optimal Performance

It is possible to once more translate Theorem 5.1 to establish conditions for which the flooding capacity is given by the global-community capacity, ensuring optimal end-to-end performance. Now, each community is a ground-based free-space community located at an altitude of $h = 30$ m, and we consider the intercommunity edges connecting each community to the backbone to also be free-space links. Furthermore, since our rigorous free-space capacities are restricted to the regime of weak-turbulence, then we must investigate free-space channels $\mathcal{E}_{\mathbf{x}\mathbf{y}}$ which are no longer than $z_{\mathbf{x}\mathbf{y}} \approx 1066$ m [30].

While this may at first appear restrictive, we remind the reader of the physical context. Free-space communities are inherently designed for short-range networks with mobile users. Indeed, with network nodes that are limited to line-of-sight connections in a potentially urban area, focusing on the weakly turbulent range is natural. This leaves us with the remaining questions: are free-space quantum channels resilient enough within this range to offer high-rate communication and what are the resource requirements of the fibre backbone? We provide insight in the following corollary:

Corollary 5.2 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3, and assume free-space community networks \mathcal{N}_{c_α} , \mathcal{N}_{c_β} and an optical-fibre backbone \mathcal{N}_b . Select any pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ located in unique communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$. There exists a maximum free-space link length in each community*

$$z_{c_j}^{\max} \leq \arg \min_z \left| \log \left(\frac{k_{c_j} \mathcal{T}_{F_\sigma}(\eta, \bar{n}_j)}{\mathcal{C}_{c:b}} \right) \right|, \quad (5.66)$$

and a maximum fibre length in the backbone

$$d_b^{\max} := -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c:b}/H_{\min}^*} \right), \quad (5.67)$$

for which the network flooding capacity is equal to the global-community capacity,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{c:b}. \quad (5.68)$$

Otherwise, if any fibre links violate this condition $\exists d_{\mathbf{x}\mathbf{y}} > d_b^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_b$ or the local community links are in violation, $\exists z_{\mathbf{x}\mathbf{y}} > z_{c_j}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_{c_j}$, for either $\mathbf{j} \in \{\alpha, \beta\}$, then this becomes an upper-bound on the network flooding capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c:b}$.

Notice that we now obtain an upper-bound on the maximum free-space link length, as it is not known whether the single-edge quantity $\mathcal{T}_{F_\sigma}(\eta, \bar{n}_j)$ is achievable or not. However, this bound has been shown to be tight and thus offers an accurate bound on $z_{c_j}^{\max}$ [30].

Furthermore, this maximum free-space link length must be computed numerically due to the complex nature of the free-space PLOB bound which accounts for fading and thermal effects. Yet, the maximum fibre length within the backbone can be readily determined for an arbitrary distribution of intercommunity connections.

5.6.3 Discussion

Fig. 5.6 provides example network constraints using Corollary 5.2 for ideal modular networks and a variety of community and backbone connectivity properties. Operational parameters are found in Table 5.2 for this modular architecture. Given a desired end-to-end flooding capacity, we generate a maximum fibre length in the backbone d_b^{\max} and maximum free-space link length in each community $z_{c_j}^{\max}$ in Figs. (a) and (b) respectively, such that this flooding capacity is achieved by the global-community capacity.

Immediately we notice that the flooding capacities plotted are large. This is because, as seen in Fig. 5.6(b), the free-space links are sufficiently capable in the weakly turbulent regime so that $z_{c_j}^{\max} > 1$ km for flooding capacities as high as $\mathcal{C}(\mathbf{i}, \mathcal{N}) \approx 2$ bits/network use, even when the community connectivity is low e.g. $k_{c_j} = 4$. As the community connectivity gets larger, the free-space capacities become increasingly reliable within this distance range, and do not compromise the minimum cut until the flooding capacity becomes very large.

Yet, these large end-to-end capacities simultaneously place greater demands on the backbone network, demanding shorter links as the global-community capacity increases. The solid lines in Fig. 5.6(a) plot the maximum fibre-length corresponding to the worst-case spatial distribution of free-space connections from the communities to the backbone, i.e. all intercommunity links are focussed on a single backbone node. Meanwhile, the dashed lines consider a best-case scenario in which all the intercommunity links are of maximum length $z_{c:b} = 1$ km, and are oriented such that they maximise the backbone cut-set cardinality $H_{\min}^* = k_b |P_{b|c_j}|$.

We find that this free-space/fibre modular architecture reports very feasible constraints on the free-space hotspots and fibre-backbone in order to guarantee a high end-to-end performance. For a regular fibre-based backbone with $k_b = 4$, and end-user communities which are $k_{c_j} > 4$ connected, then one can guarantee an achievable flooding capacity of $\mathcal{C}(\mathbf{i}, \mathcal{N}) = 2$ bits/network use given that the free-space links all fall within the weakly turbulent range, and at worst $d_b^{\max} \lesssim 25$ km. Within a metropolitan setting, such constraints can be satisfied with realistic resources, supporting the development of wireless quantum networks. Furthermore, confidence in the use of free-space links within this setting reduces the need for wired fibre connections in small areas.

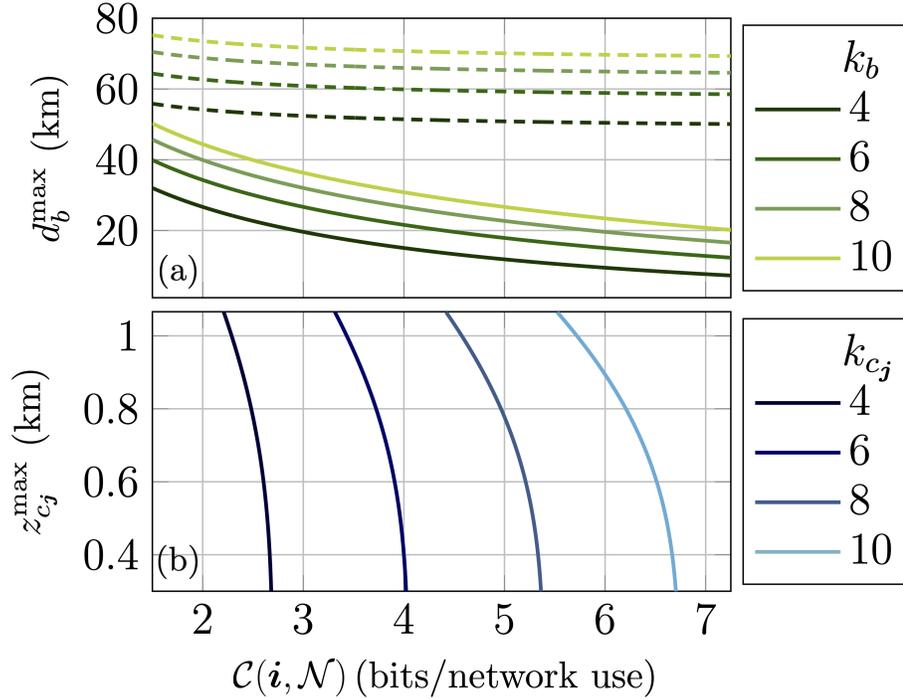


Figure 5.6: Optimal end-to-end performance for an ideal modular network consisting of free-space communities interconnected to a fibre-based backbone. In order to guarantee an optimal flooding rate along the x -axis then the maximum internodal separations in each sub-network depicted on the y -axis must be less than or equal to the plotted bounds. We use the operational settings in Table 5.2 during clear day-time. Given an optimal flooding capacity $\mathcal{C}(i, \mathcal{N})$, we plot the maximum fibre-length d_b^{\max} for different backbone connectivity parameters, and the maximum free-space link-length in each community $z_{c_j}^{\max}$ for different community connectivity parameters. The dashed lines plot an upper bound on the maximum fibre-length based on the optimal spatial distribution of community connected backbone nodes $P_{b|c_j}$, while the solid lines plot a lower bound based on the worst-case spatial distribution.

5.7 Conclusion

In this chapter we have investigated the end-to-end capacities of free-space and hybrid quantum networks, combining recently developed results in quantum information theory and well established theories of free-space optical communication. After collecting and reviewing these recent results, we introduced a modular network architecture for the purposes of constructing hybrid quantum networks using free-space and fibre links. With these tools in hand, we specified our analysis to ideal modular networks which utilise an underlying regular backbone. Through this ideality it was possible to study ultimate limits for highly relevant modular architectures, revealing critical network properties that assure optimal

performance.

We have performed a detailed analysis of the ultimate limits of a satellite-based quantum internet; leveraging the properties of fibre-networks on the ground, ground-satellite connective structures and intersatellite networks in space. This theoretically demonstrates that high-rate global quantum communication can be efficiently mediated by a satellite quantum network with realistic connectivities and tolerable intersatellite separations on the order of $\sim 10^3 - 10^4$ km. Such designs allow for effective quantum communication between arbitrarily distant end-users on the Earth. Our analyses also indicate that careful consideration of the spatial distribution of ground-satellite connections can more effectively alleviate separation constraints, rather than increasing the nodal degree.

Furthermore, we studied the ultimate limits of a free-space/fibre modular network configuration, discussing the efficacy of free-space sub-networks within metropolitan areas. We have shown that within the weakly-turbulent regime (where free-space links are limited to ~ 1 km) high-rate intercommunity communication can be readily achieved, using a fibre-backbone with realistic resources.

The results of this chapter offer promising first steps in the direction of understanding the ultimate limits of free-space and hybrid quantum networks; motivating its future study both theoretically and experimentally. Our analyses offer a rigorous demonstration of the efficacy of free-space quantum links in a network setting, emphasising that the integration of free-space and fibre can be reliably performed within future quantum networks. Hybrid architectures can and should be designed to take advantage of the strengths of different modes of quantum communication. This work may serve as a platform for future investigations that account for full technical details of the nodes; exploiting these tools to study more realistic, random architectures of hybrid networks which can be benchmarked against the ideal designs studied here.

Chapter 6

Practical Routing and Criticality in Complex Quantum Networks

The work in this chapter forms the basis of a paper being prepared for submission, whose authors are (in order) Cillian Harney and Stefano Pirandola. Section 6.1 introduces the goals of this chapter, to carry forward the important sentiments of the previous chapters regarding network capacities and use them to motivate a more realistic assessment of the limitations of quantum networks. In Section 6.2 we begin this journey by reviewing realistic models of network architectures, namely *random quantum networks*. In Section 6.3 we discuss routing schemes and their practicality in a realistic network settings. Section 6.4 then develops quantitative and qualitative notions of criticality in quantum networks. Section 6.5 then applies these notions to random network architectures in order to reveal key insight for future quantum networks, before Section 6.6 concludes and discusses future investigations.

6.1 Introduction

Glaring gaps exist in our understanding of the realistic performance limits of quantum networks and their practical resource demands. The most expensive resource in a quantum internet will be the number of quantum repeaters needed in a physical network area to guarantee effective communications, i.e. the network nodal density. While classical repeaters are cheap and easy to deploy, quantum repeaters are costly and should not be wasted as non-user nodes. Recent studies have been able to identify important nodal density conditions for reliable end-to-end performance on quantum networks [1, 64, 65]. Of these investigations, Ref. [64] principally focuses on network connectivity properties and a specific single-photon transmission protocol through pure-loss channels, while Refs. [1, 65] evaluate network capacities with a protocol agnostic approach, focussing on networks composed of

pure-loss channels and end-to-end communication achieved via optimal flooding protocols.

These works provide meaningful insight to the performance/resource limitations of realistic quantum networks, and some do so via the study of *random quantum networks*. These are network architectures which are generated in a probabilistic manner according to a set of well defined connection rules. The process of building a random network is (in general) evolutionary; network nodes are added to a structure, which are then connected to pre-existing nodes according to some probability. Classes of networks generated via random models are incredibly useful for gathering knowledge on the realistic performance of communication networks. Such models are designed to manifest particular features that are observed in the real world, i.e. clustering within dense node populations, realistic distributions of edge lengths according to channel efficacy, and connection preferences to well connected nodes. The complexity introduced through random network modelling helps to replicate the complexity of real-world quantum networks in the future.

While random quantum network models have been deployed in Refs. [64, 65] there remains room for progress. Firstly, Ref. [65] clearly asserts that network connectivity investigations are useful, but are not completely sufficient to benchmark end-to-end performance (as connectivity analyses formed the basis of Ref. [64]). Quantum networks can be classified as “well connected” from the perspective that each node can identify an end-to-end path between every other node, yet their end-to-end capacities remain incredibly weak. This emphasises the importance of quantum network capacities and their evaluation, a recurring theme within this thesis. However, the capacity evaluations carried out in Ref. [65] focus strictly on flooding protocols. This is completely fitting for studying the ultimate limits of quantum networks, but the employment of flooding protocols on a large-scale is highly impractical. It is not realistic to assume the use of *every* edge in a potentially global network to facilitate communication between a single end-user pair. Finally, both works model quantum networks using pure-loss channels. Beyond pure-loss models, optical fibre is most accurately described using thermal-loss channels which account for environmental and experimental thermal noise. Hence, while some of these assumptions are effective and completely suitable for network capacity assessments, they are over-optimistic in the context of revealing stricter, practical insight for quantum networking.

In this chapter, we tighten the assessment of critical resource requirements in realistic quantum networks with two crucial improvements. Firstly, we consider random network architectures composed of realistic link-layers; using point-to-point channels which account for thermal noise, experimental imperfections and practical point-to-point protocols. This immediately improves our characterisation of realistic quantum networks, and reveals the impact of point-to-point limitations on network connectivity models.

Secondly, we explore the efficacy of *practical end-to-end routing* in quantum networks to identify critical resource requirements of quantum networks. Moving away from impractical flooding schemes, we study the feasibility of single-path routing and multi-path routing, adopting recent results in the efficient generation of end-to-end multi-paths [103]. Our results suggest that single-path routing is an inefficient strategy for reliable rates on large-scale quantum networks. Instead, there exist multi-path routing protocols which can efficiently achieve close to optimal rates without consuming the entire network. Furthermore, different network architectures observe a suitability to quantum multi-path routing which others do not. All of these results identify important design criteria for future, large-scale quantum networks while motivating the further development of practical multi-path protocols.

6.2 Random Quantum Networks

6.2.1 Link Quality and Edge Pruning

The study of random network models is a rich and wide-spanning field, within which many key tools have been developed. Of course, random network models cannot capture *all* of the features of a future quantum internet; these features will emerge as the technology is developed and deployed. Nonetheless, relevant random networks are able to capture and predict important behaviours of realistic complex networks. In this chapter, we consider the classes of Waxman and scale-free networks with two important considerations: realistic link-layer descriptions and practical end-to-end protocols.

In this chapter we consider bosonic thermal-loss networks $\mathcal{N} = (P, E)$ and compare different rate distributions that are either bounds on the capacity distributions, or achievable rate distributions based on practical protocols. Most importantly, we consider the bounding capacity distributions,

$$\mathcal{K}_l = \{\mathcal{T}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}}^l\}_{(\mathbf{x}, \mathbf{y}) \in E}, \quad (6.1)$$

$$\mathcal{K}_u = \{\mathcal{T}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}}^u\}_{(\mathbf{x}, \mathbf{y}) \in E}, \quad (6.2)$$

where $\mathcal{T}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}}^l$ and $\mathcal{T}_{\eta_{\mathbf{x}\mathbf{y}}, \bar{n}_{\mathbf{x}\mathbf{y}}}^u$ are defined in Eqs. (2.51) and (2.52). To address the weak rates achieved by quantum links over long distances, we impose a modification to the standard random network models. Network edges which possess a point-to-point rate below some threshold value ε are *pruned* and removed from the network edge set after graph generation. Under this modification networks can only be considered completely connected if they are *competently connected* by edges with $K_{\mathbf{x}\mathbf{y}} \geq \varepsilon$. Given that desirable clock rates C for

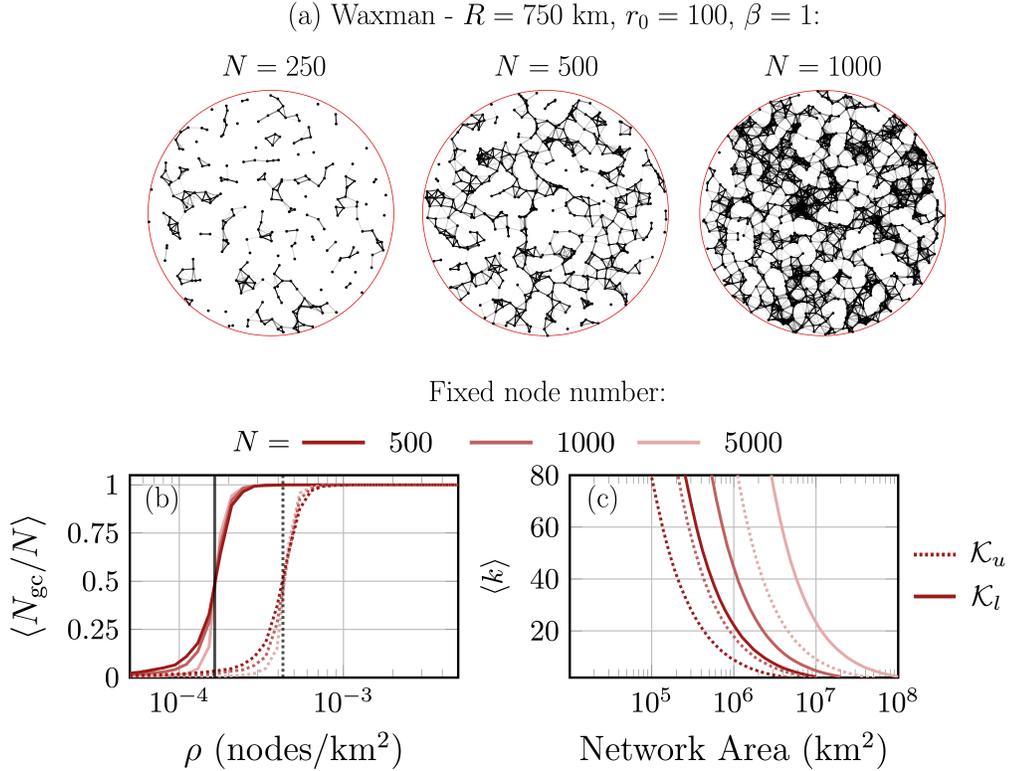


Figure 6.1: Connectivity properties of bosonic thermal-loss Waxman networks. Panel (a) displays random networks generated under the parameters listed above and using single-edge capacity upper-bounds. The colour intensity of each edge is proportional to its capacity. Panel (b) plots the average main component fraction, in which the critical densities necessary for a connectivity phase transition are identified by the black lines. Panel (c) shows the average degree of networks with a fixed number of nodes (given in legend) with respect to variable network density and area respectively. The legend identifies the fixed node number and indicates whether the network uses upper or lower bounds on the rate distributions.

quantum communications are on the order of GHz, we set $\varepsilon \sim 10^{-12}$ so that valid network edges must guarantee at least $CK_{xy} \gtrsim 1$ mbit per second¹. The notion of pruning goes towards preserving network resources that are otherwise wasted, and offers a more accurate representation of network connectivity. Importantly, it affects our random network models in unique ways, reflecting the potential impact of poor link quality in a realistic quantum internet.

¹That is, 1 milli-bit per second. Note that this is already a lenient condition, as clock rates in many quantum communication protocols are thus far restricted to the MHz range.

6.2.2 Waxman Networks

Networks from the Waxman class $\mathcal{N} \in \mathcal{W}$ are constructed as follows: a number of N nodes are generated in a region of area A , and any pair of nodes \mathbf{x}, \mathbf{y} are connected with a probability that decays exponentially with respect to their point-to-point separation, $r_{\mathbf{x}\mathbf{y}}$. More precisely, a channel $\mathcal{E}_{\mathbf{x}\mathbf{y}}$ between the nodes \mathbf{x}, \mathbf{y} is created with probability

$$p_{\mathbf{x}\mathbf{y}}^{\mathcal{W}} := \beta e^{-\frac{r_{\mathbf{x}\mathbf{y}}}{r_0}}. \quad (6.3)$$

The parameters β, r_0 are characteristic of the model and influence generation process; $\beta \in (0, 1]$ defines a maximum probability of connection, while $r_0 \in (0, \infty)$ dictates the speed at which the connection probability exponentially decays. Erdős-Rényi networks are a specification of the Waxman model such that $r_0 \rightarrow \infty$, i.e. there is no decay with respect to channel length but the probability of connection is simply $p_{\mathbf{x}\mathbf{y}}^{\text{ER}} := \beta \leq 1$.

One may think of the generation process for Waxman networks as a *static* generation process. The probability of nodes becoming connected according to Eq. (6.3) is completely independent of every other network node or edge, and only on the spatial separation of those nodes. Consequently, building a Waxman network involves (1) generating N nodes within a defined spatial region, and (2) connecting each pair of nodes in parallel according to the probability in Eq. (6.3).

Examples of N node networks contained in circular areas of radius R can be seen in Fig. 6.1(a). In order to analyse these networks, we define N_{gc} as number of nodes in the giant component of the network (largest subset of nodes between which all nodes possess paths to one another). Furthermore, we denote the nodal degree via $k_{\mathbf{x}} := \text{deg}(\mathbf{x})$ which (as we know) describes the number of nodes to which \mathbf{x} is connected. When N is very small (or R is very large) the network is sparse and poorly connected; the average distance between nodes is large so that links will connect with very low probability under Eq. (6.3). As a result, the network will be clustered and incompletely connected². This can be seen in Fig. 6.1(a) for $N = 250$ nodes. Yet, as N is increased (or R is decreased) then the average nodal separation will shrink, increasing the likelihood of connections. Eventually, the nodal density becomes large enough to promise that the number of nodes in the giant component $N_{\text{gc}} = N$ and there exists end-to-end paths between all nodes.

This behaviour in Waxman networks is well known, and gives rise to a critical connectivity phase transition. As seen in Fig. 6.1(b) There exists a critical nodal density ρ_{G}^* at which the model abruptly transits from poorly connected ($N_{\text{gc}} < N/2$) to well connected ($N_{\text{gc}} \geq N/2$). The phase transition is illustrated in Fig. 6.1(b) in the context of bosonic

²The number of nodes in the giant component is less than the number of network nodes N

thermal-loss networks, where we plot the average fraction of the network contained in the giant component, $\langle N_{\text{gc}}/N \rangle$. One can similarly see in Fig. 6.1(c) how the average nodal degree $\langle k \rangle$ undergoes a collapse as the network area is expanded, leading to the connectivity transition. It can be seen that the critical density is bounded between

$$1.6 \times 10^{-4} \lesssim \rho_{\text{G}}^* \lesssim 4.3 \times 10^{-4} \text{ nodes per km}^2. \quad (6.4)$$

This is nearly two orders of magnitude larger than that predicted by bosonic pure-loss networks, for which $\rho_{\text{G}}^* \sim 7 \times 10^{-6}$ nodes per km^2 [64, 65]. This emphasises thermal decoherence as a vital consideration since the impact of environmental noise alone can significantly degrade connectivity.

6.2.3 Scale-free Networks

Another important random architecture is that of scale-free networks \mathbf{S} . A network is scale-free if their degree distribution (probability distribution of nodes in \mathbf{S} having degree k) follows a power law, i.e. $p_k \propto k^{-\gamma}$ where γ is some real number characterising the distribution. Many real world networks (such as the classical internet) are thought to exhibit scale-free properties, however it is rare that a network is precisely scale-free [104]. Here, we study scale-free networks generated dynamically using Yook's model [105, 106] in which new nodes \mathbf{y} are iteratively added to an initially small, n_0 node connected network. Each new node \mathbf{y} is attached to a collection of m existing nodes $\{\mathbf{x}_i\}_{i=1}^m$ with probability

$$p_{\mathbf{x}\mathbf{y}}^{\mathbf{S}} \propto k_{\mathbf{x}}^{\sigma_{\text{deg}}} / r_{\mathbf{x}\mathbf{y}}^{\sigma_r}, \quad (6.5)$$

where $\sigma_{\text{deg}}, \sigma_r \in \mathbb{R}_0^+$ are model parameters which controls the influence degree and link length have on connection probability.

Some example networks are illustrated in Fig. 6.2(a). In general, the connectivity properties of scale-free networks are very different to that of Waxman networks. Fig. 6.2(b) shows the behaviour of the giant component with respect to fixed node number and nodal density. When there is no connection dependence on link length ($\sigma_r = 0$) there is no critical density. The giant component eventually transits to $\langle N_{\text{gc}}/N \rangle = 1$, but less abruptly, and at a higher density than that of Waxman networks. Increasing σ_r , one can see that the network becomes fully connected more quickly and a critical density begins to emerge as we recover the transition shape from Fig. 6.1(b).

There are also notable differences in the context of average nodal degree, $\langle k \rangle$. The average degree of Waxman networks scales exponentially with respect to nodal density. In the scale-free setting, the relationship between connection probability and nodal degree

gives rise to nodal hubs (i.e. nodes to which most others are connected) but also a significant level of sparsity, so that most nodes have a very low degree. This is due to the mechanism of preferential attachment, as nodes which have high degrees are more likely to gain more connections. As seen in Fig. 6.2(c), at high densities, scale-free networks saturates to a maximum value $\langle k \rangle \rightarrow 2m$.

However, at lower network densities this is not the case and the average degree undergoes a collapse. This can be understood as follows: when new nodes are added, they attempt to connect with m other existing network nodes. At low densities, nodes are typically too far away to forge competent links. When a new node is added, the $\sigma_r = 0$ model has no preference in choosing m nodes within a quality connection range and thus new nodes fail to connect reliably and the average degree collapses. For $\sigma_r = 1$, there exists some preference

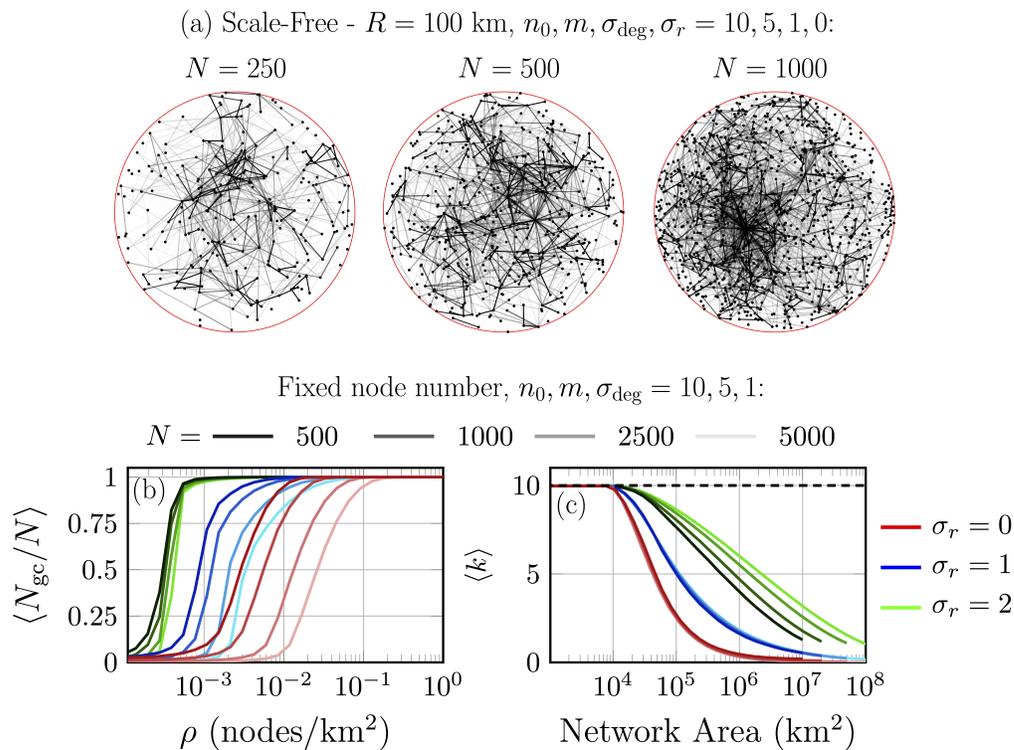


Figure 6.2: Connectivity properties of bosonic thermal-loss scale-free networks. Panel (a) displays random networks generated under the parameters listed above and using single-edge capacity upper-bounds. The opacity of each edge is proportional to its capacity. Panel (b) plots the average giant component fraction and (c) the average degree of networks with a fixed number of nodes with respect to variable network density and area respectively. The legend identifies the fixed node number and σ exponent.

for choosing nearby nodes which is capable of slowing the $\langle k \rangle$ decay. Forcing a stronger link length dependence at $\sigma_r = 2$ slows this rate even further and raises the degeneracy with respect to node number.

Clearly, setting the parameter $\sigma_r = 2$ breaks the “scale-freedom” of this model, since the degree behaviour is no longer independent from N . However, this dependence is an inevitable consequence of utilising quantum communications. As with the classical internet, it is likely that scale-free properties will emerge in quantum networks. However, network engineers will not turn a blind eye to link length and channel quality, making the $\sigma_r = 0$ model unrealistic, rendering larger values of σ_r meaningful and interesting. Throughout this chapter we focus on $\sigma_{\text{deg}} = 1$, $\sigma_r > 0$ models as better reflections of future quantum networks which we refer to as the network class S_{σ_r} .

6.3 Practical Routing in Quantum Networks

In practice, one must be able to deploy a routing algorithm which informs network nodes how to interact and establish communication between the end-users. This algorithm characterises the network protocol, for which there are some intuitive and practical objectives:

1. *Rate Optimisation*: Identify an end-to-end path (or paths) which maximise the rate between users. Alternatively, the protocol should surpass a target rate requirement.
2. *Resource Minimisation*: Minimise the network resources needed to achieve said end-to-end rate, i.e. minimise the necessary network nodal density as well as the number of links/nodes which participate in communication.
3. *Computational Efficiency*: Be efficient enough to execute and facilitate routing without impeding the end-to-end rate.

It is extremely important that each of these points are considered for a routing strategy to be deemed practical. An easily executed protocol is useless if it identifies poor routes, while a protocol which establishes high rate routes very slowly is equally undesirable. Meeting these requirements in quantum networks is more challenging than in classical settings, extending from the point-to-point rate limitations of quantum communications.

As we know, single-path routing is the principal mechanism for classical communications, fundamentally achieved by Dijkstra’s algorithm (DA) [55, 56, 57, 58]. which was reviewed briefly in Section 2.4.2. For the purposes of rate optimisation, DA is used to solve the widest path problem, in which the goal is to locate the end-to-end route which maximises its minimum single-edge rate (has the largest bottleneck rate). As we will explore in this

chapter, DA and its variants can be modified to locate routes which are optimal with respect to different cost functions. Recall that a network $\mathcal{N} = (P, E)$, DA locates the widest path efficiently in run-time $\mathcal{O}(|E| + |P| \log_2 |P|)$, and for end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$, their optimal single-path end-to-end rate is given by

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) := \max_{\omega} \min_{(\mathbf{x}, \mathbf{y}) \in \omega} K_{\mathbf{x}\mathbf{y}}, \quad (6.6)$$

where the maximisation is performed over all possible end-to-end routes. To mitigate single-path performance limitations, multi-path routing emerges as good solution. Flooding represents a useful paradigm for benchmarking optimal network performance, but it is not practical in many real world settings. Utilising an entire network to facilitate one end-user pair renders the network useless for any other set of users to communicate simultaneously. In reality, we want to enable the concurrent use of a network for many users. Therefore, the development of practical multi-path protocols is an area of major interest.

A basic approach to building high-rate multi-paths is to perform an iterative form of DA. This is a modified algorithm which locates multiple end-to-end routes which are either edge-disjoint (no two paths share the same edge) or node-disjoint (no two paths visit the same node). Both of these versions involve executing a Dijkstra search, followed by a network modification where edges included in previous paths (or in the neighbourhoods of nodes in previous paths) are not included in the next search. Locating M end-to-end paths requires M executions of DA, leading to an increased time complexity $\mathcal{O}(M(|E| + |P| \log_2 |P|))$. This is a costly scaling factor which does not lend well to deployability in large-scale networks.

To expand our suite of routing strategies, it is useful to possess a clear picture of the utility of DA. One can generalise DA in order for it to be integrated into other routing strategies, as will be explored within this chapter.

6.3.1 Generalisations of Dijkstra's Algorithm

It is possible to construct general versions of DA that optimise different path-wise properties associated with end-to-end routes. It is well known that DA can be used to minimise path length (shortest path problem) or modified to maximise the bottleneck rate (widest path problem). But one can be more general and can define a global cost function \mathcal{F}_{ω} such that goal of DA is to find

$$\omega^* = \arg X_{\omega}(\mathcal{F}_{\omega}), \quad (6.7)$$

such that $\arg X \in \{\arg \min, \arg \max\}$. In this way, \mathcal{F}_{ω} may admit a more complex characterisation of routing value.

Algorithm 1 Generalised Dijkstra Algorithm

Inputs: Network - $\mathcal{N} = (P, E)$, Source - \mathbf{s} , Target - \mathbf{t} ,
 Min/Max Functions - $(X, f_X) \in \{(\min, <), (\max, >)\}$.
 Tentative Cost Function - F_ω ,

```

1: procedure OPTCOSTPATH( $\mathbf{s}, \mathbf{t}, \mathcal{N}$ )
2:   Priority queue  $Q = \{\mathbf{x}\}_{\mathbf{x} \in P}$ 
3:   Parent node set:  $P = \{\text{undef}\}_{\mathbf{x} \in P}$ 
4:   Tentative cost set:  $T = \{\varepsilon'_{\text{init}}\}_{\mathbf{x} \in P}$ 
5:    $T_{\mathbf{s}} \leftarrow \varepsilon_{\text{init}}$ 

6:   while  $|Q| > 0$  do
7:      $\mathbf{u} \leftarrow \arg \max_{\mathbf{x} \in Q} T_{\mathbf{x}}$  (and remove  $\mathbf{u}$  from  $Q$ )
8:     if  $\mathbf{u} \neq \mathbf{t}$  then
9:       for all neighbours  $\mathbf{v} \in Q$  of  $\mathbf{u}$  do
10:         $a \leftarrow F_\omega(T_{\mathbf{v}}, T_{\mathbf{u}}, \mathbf{c}_{\mathbf{uv}})$ 
11:        if  $f_X(a, T_{\mathbf{v}})$  is true then
12:           $T_{\mathbf{v}} \leftarrow a, P_{\mathbf{v}} \leftarrow \mathbf{u}$ 
13:          Reprioritise  $Q$  wrt  $T$ 

   return CONSTRUCTPATH( $P, T$ )

```

Figure 6.3: Generalised Dijkstra's algorithm for end-to-end route optimisation with respect to a cost function F_ω . Any cost function has a tentative counterpart $F_\omega^{s, \mathbf{x} \rightarrow \mathbf{y}}$ which is used to evaluate movement throughout the network. The above pseudocode describes the network exploration phase which is followed by a CONSTRUCTPATH subroutine which simply back tracks from the target node \mathbf{t} to \mathbf{s} using the constructed tentative cost and parent node sets.

Consider an N -node network $\mathcal{N} = (P, E)$. Let us define an N element *tentative path cost set* $T = \{T_{\mathbf{x}}\}_{\mathbf{x} \in P}$, which will be used to track the cost of routing throughout the search. This set will be initialised with a unique value for the source node $T_{\mathbf{s}} = \varepsilon_{\text{init}}$, while all other nodes are initialised with a different value $T_{\mathbf{x}} = \varepsilon'_{\text{init}}$ for all $\mathbf{x} \in P \setminus \{\mathbf{s}\}$. Thus, T takes the initial form

$$T = \{\varepsilon'_{\text{init}}, \dots, \varepsilon'_{\text{init}}, \underbrace{\varepsilon_{\text{init}}}_{\text{source}}, \varepsilon'_{\text{init}}, \dots, \varepsilon'_{\text{init}}\}. \quad (6.8)$$

Along with this information, we also have access to point-to-point properties of each edge in the network. There may exist m different types of properties for every edge $(\mathbf{x}, \mathbf{y}) \in E$, which we denote via the set $\mathbf{c}_{\mathbf{xy}} := \{c_{\mathbf{xy}}^1, \dots, c_{\mathbf{xy}}^m\}$ ³.

³For instance, one property may be the single-edge rate $c_{\mathbf{xy}}^1 = K_{\mathbf{xy}}$ or the single-edge channel length $c_{\mathbf{xy}}^1 = l_{\mathbf{xy}}$, etc.

The algorithm then operates greedily; starting at an initial source node \mathbf{s} , it traverses throughout the network hopping from node to node while keeping track of a tentative path cost from \mathbf{s} to any other network node $\mathbf{x} \in P$. Suppose that we have traversed from \mathbf{s} to a node \mathbf{x} and must decide whether to move to its neighbour \mathbf{y} or not. Here, we must evaluate the tentative cost to the path if we make this move, since we should only hop to a subsequent node if it optimises the routing cost. The tentative path cost is evaluated via a *tentative cost function* F_ω which admits the form

$$F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}} = F_\omega(T_{\mathbf{x}}, T_{\mathbf{y}}, \mathbf{c}_{\mathbf{x}\mathbf{y}}). \quad (6.9)$$

That is, F_ω is a function of $T_{\mathbf{x}}$ the tentative path cost from $\mathbf{s} \rightarrow \mathbf{x}$, $T_{\mathbf{y}}$ the tentative path cost from $\mathbf{s} \rightarrow \mathbf{y}$ and the known point-to-point properties associated with moving along the edge $(\mathbf{x}, \mathbf{y}) \in E$. Note that F_ω is not equal to the global cost function \mathcal{F}_ω , but they are inexorably tied. The tentative cost function is used to compute the global cost of an end-to-end route *during* the deployment of DA. Consider an end-to-end route

$$\omega = \{(\mathbf{a}, \mathbf{x}_1), (\mathbf{x}_1, \mathbf{x}_2), \dots, (\mathbf{x}_m, \mathbf{b})\}, \quad (6.10)$$

which traverses the network over $m + 2$ nodes (including \mathbf{a} and \mathbf{b}). Then the global cost of the path \mathcal{F}_ω is computed via the nested operation,

$$F_\omega \left(T_{\mathbf{b}}, F_\omega \left(T_{\mathbf{x}_m}, F_\omega \left(\dots F_\omega (T_{\mathbf{x}_2}, F_\omega (T_{\mathbf{x}_1}, T_{\mathbf{a}}, \mathbf{c}_{\mathbf{a}\mathbf{x}_1}), \mathbf{c}_{\mathbf{x}_1\mathbf{x}_2}) \dots \right), \mathbf{c}_{\mathbf{x}_{m-1}\mathbf{x}_m} \right), \mathbf{c}_{\mathbf{x}_m\mathbf{b}} \right). \quad (6.11)$$

Hence, the tentative cost function can be thought of as a *local* translation of the global cost function such that that can be used iteratively within DA.

With the ability to evaluate the impact this hop has on the total path, the value $F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}}$ can be compared with the tentative path cost $T_{\mathbf{y}}$. If the goal is to minimise the path cost, then the search will hop to \mathbf{y} iff $F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}} < T_{\mathbf{y}}$ i.e. moving to this edge better minimises the cost function. Contrarily, if the goal is to maximise the path cost function then the search will hop to \mathbf{y} iff $F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}} > T_{\mathbf{y}}$. The algorithm follows these steps, evaluating, hopping and reevaluating the path cost until eventually the target node \mathbf{t} is reached. Since the algorithm is greedy, it follows an optimal path throughout the network at all times. Once it arrives at a node, one can then be sure that the path followed has been the best one.

Let us specify to a couple of examples. When minimising the path length the global and tentative cost functions take the form

$$\text{Global: } \mathcal{F}_\omega = \sum_{(\mathbf{x}, \mathbf{y}) \in \omega} l_{\mathbf{x}\mathbf{y}}, \quad \text{Tentative: } F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}} = T_{\mathbf{x}} + l_{\mathbf{x}\mathbf{y}}, \quad (6.12)$$

where $l_{\mathbf{x}\mathbf{y}}$ is a measure of point-to-point length of the edge $(\mathbf{x}, \mathbf{y}) \in E$. To minimise this quantity we use the initialisation values are $(\varepsilon_{\text{init}}, \varepsilon'_{\text{init}}) = (0, \infty)$. For maximising the bottleneck rate we define the global cost function

$$\text{Global: } \mathcal{F}_\omega = \min_{(\mathbf{x}, \mathbf{y}) \in \omega} K_{\mathbf{x}\mathbf{y}}, \quad \text{Tentative: } F_\omega^{s, \mathbf{x} \rightarrow \mathbf{y}} = \max(T_{\mathbf{y}}, \min(T_{\mathbf{x}}, K_{\mathbf{x}\mathbf{y}})), \quad (6.13)$$

where $K_{\mathbf{x}\mathbf{y}}$ is the point-to-point rate of the edge $(\mathbf{x}, \mathbf{y}) \in E$. Here, we use $(\varepsilon_{\text{init}}, \varepsilon'_{\text{init}}) = (\infty, -\infty)$.

Fig. 6.3 provides an algorithmic description of generalised DA up to the point of path construction. The CONSTRUCTPATH subroutine is a simple procedure following identically from the original algorithm, which uses the output cost evaluations of the algorithm to traverse back from the target node to the sender node, constructing the optimal path along the way.

6.3.2 Multiple Disjoint Paths Algorithm (MDPAlg)

There are ongoing developments in the study of fast algorithms for multi-path routing [107, 108, 109, 110, 111]. In particular, Lopez-Parajes *et al.* recent devised an efficient, centralised, *one-shot* approach to multi-path routing through their Multiple Disjoint Path Algorithm (MDPAlg) [103]. They recognised that a single execution of DA observes much more information than it actually utilises. Given a source node \mathbf{a} , the standard DA focuses on building a minimum cost tree from which only the optimal end-to-end paths can be located from any other node $\mathbf{x} \in P \setminus \{\mathbf{a}\}$ in the network. It does this by storing a minimum cost set (stores the minimum cost associated with traversing from \mathbf{a} to any node \mathbf{x}) and a parent node set (stores the parent node from which the minimum cost was obtained to \mathbf{x}). With these quantities, DA can then reconstruct a minimum cost path by backtracking from a target node \mathbf{b} along the parent-node tree until \mathbf{a} is reached.

Alternatively, the MDPAlg uses a *cost matrix* to collect additional information on the *aggregated* cost from \mathbf{a} to any other node; not just the minimum cost. Whereas DA would discard information concerning sub-optimal paths, the MDPAlg stores such information in the cost matrix so that they may be used to construct additional end-to-end paths later. Furthermore, it does this through only a single search of the network rather than the many searches that may be required by an iterative Dijkstra approach. The cost matrix can then be used to reconstruct the optimal path, and many other paths between the source \mathbf{a} and $\mathbf{b} \in P \setminus \{\mathbf{a}\}$.

Much like DA, the algorithm is split into two phases, (i) network exploration and (ii) path reconstruction. Throughout network exploration the algorithm proceeds similarly to

Algorithm 2 MDPAlg - Network Exploration

Inputs: Network - $\mathcal{N} = (P, E)$, Source - \mathbf{s} , Target - \mathbf{t} ,
 Min/Max Functions - $(X, f_X) \in \{(\min, <), (\max, >)\}$.
 Number of Routes - M , Rate requirement - K^* .
 Tentative Cost Function - F_ω ,

```

1: procedure MDPALG( $\mathbf{s}, \mathcal{N}$ )
2:   Priority queue  $Q = \{\mathbf{x}\}_{\mathbf{x} \in P}$ 
3:   Tentative cost matrix  $T$  ( $\dim |P| \times |P|$ )
4:    $T_{\mathbf{x}\mathbf{y}} \leftarrow \varepsilon'_{\text{init}}, \forall \mathbf{x}, \mathbf{y} \in P \setminus \{\mathbf{s}\}$ 
5:    $T_{\mathbf{s}\mathbf{s}} \leftarrow \varepsilon_{\text{init}}$ 

6:   while  $|Q| > 0$  do
7:      $\mathbf{u} \leftarrow \arg \min_{\mathbf{x} \in Q} T_{\mathbf{x}}$  (and remove  $\mathbf{u}$  from  $Q$ )
8:     for all neighbours  $\mathbf{v} \in Q$  of  $\mathbf{u}$  do
9:        $a \leftarrow F_\omega(T_{\mathbf{v}\mathbf{v}}, T_{\mathbf{u}\mathbf{u}}, \{c_{\mathbf{u}\mathbf{v}}^i\}_{i=1}^m)$ 
10:      if  $f_X(a, T_{\mathbf{v}\mathbf{v}})$  is true then
11:         $T_{\mathbf{v}\mathbf{v}} \leftarrow a$ 
12:        Reprioritise  $Q$  wrt  $T$ 
13:      else
14:         $T_{\mathbf{u}\mathbf{v}} \leftarrow a$ 

return CONSTRUCTMDP( $T, \mathbf{t}, M, K^*$ )

```

Figure 6.4: MDPAlg for end-to-end route optimisation with respect to a cost function \mathcal{F}_ω (with corresponding tentative cost function $F_\omega^{\mathbf{s}, \mathbf{x} \rightarrow \mathbf{y}}$). This pseudocode describes the network exploration phase which is followed by a CONSTRUCTMDP subroutine which uses the tentative cost matrix to identify many disjoint paths from the target node \mathbf{t} to \mathbf{s} . Concerning the protocol types considered in this chapter, we can use CONSTRUCTMDP to explicitly identify M paths (fixed route number protocol) or identify as many paths as necessary to guarantee a rate of K^* (rate requirement protocol).

DA in which the tentative cost matrix T is constructed. The pseudo-code for this section is outlined in Fig. 6.4, where the primary difference between it and DA is shown in lines 13 and 14: even when the computed tentative cost a is not considered optimal, it is stored as an off diagonal element of $T_{\mathbf{u}\mathbf{v}}$, describing the accumulated cost of travelling from node \mathbf{s} to \mathbf{v} via \mathbf{u} . This cost analysis algorithm is illustrated in Fig. 6.4.

This additional information is then used in the path reconstruction phase to identify many end-to-end routes. The reconstruction is sequential and straightforward: starting at a target node \mathbf{t} a path is built by moving to the neighbour which incurs the lowest cost in the tentative cost matrix, T , until the source node is reached. On the first path

reconstruction, this is simple and the minimum cost path ω_1 is produced, which adds edges to the routing edge set ($E_\omega = \omega_1$). In order to construct subsequent disjoint paths, one must then enforce disjointedness by restricting the use of any edges which were used previously (for link disjointedness). In other words, future reconstructions will ignore the cost matrix elements $T_{\mathbf{x}\mathbf{y}}$ for $(\mathbf{x}, \mathbf{y}) \in E_\omega$. Repeating this process, path reconstruction may occur by moving from the target \mathbf{t} to its *next* minimum cost neighbour, and so on until \mathbf{s} is met again.

Through this sequential process of path building coupled with edge restriction, a number of edge-disjoint paths can be built which have favourable properties. The node-disjoint variant of the MDPAlg can be easily achieved through a small modification; instead of only restricting previously used edges from subsequent path reconstructions, one must restrict the use of any edges connected to nodes used in the previous routes.

6.3.3 Modifying the MDPAlg for Rate Optimisation

The MDPAlg was originally introduced with the intention of identifying multiple disjoint *shortest* paths, that which minimises the cumulative cost of edge-weights along an end-to-end route. Consequently, it is not immediately clear how the MDPAlg can be translated for the purposes of rate maximisation. One might assume that since there exists a variant of DA for this purpose (the widest path algorithm outlined in Section 6.3.1) then it should be able to modify the MDPAlg in an identical way.

Unfortunately, this is not the case. The widest path version of DA locates an end-to-end route which maximises a bottleneck rate between source and target nodes. It does so using via the generalised DA and the cost functions in Eq. (6.13) wherein the tentative cost function takes the form $F_\omega^{s,\mathbf{x} \rightarrow \mathbf{y}} = \max(T_\mathbf{y}, \min(T_\mathbf{x}, K_{\mathbf{x}\mathbf{y}}))$. When this is employed within the MDPAlg, it is capable of identifying a single optimal path *but no more*. It is a consequence of the fact that this cost function is extremely degenerate; once the first optimal path is constructed, aggregated cost information stored within the cost matrix are now filled with highly redundant information linked to the initial optimal route. This invokes a “push and pull” effect in subsequent route construction, in which nascent routes are being “pulled” towards the widest path, and subsequently “pushed” away whenever the building process recognises that these edges have already been used. The result is an unstable, ineffective process for building high-rate multi-paths.

As a result, one should avoid the use of routing cost functions that possess high degeneracies, such as the direct translation of the widest path cost functions. Instead, one should explore cumulative costs, similar to that used in the shortest path formulation. Cumulative

cost functions which capture properties of entire paths are more suited to this algorithm and can more effectively motivate end-to-end routing.

Fortunately, there are other heuristics that take this suitable form and that can be used to identify high-rate and resource efficient paths. In this chapter, we modify the MDPAlg in such a way that approximates rate optimisation by minimising the *inverse accumulated rate* over the course of an end-to-end route. That is, we do not minimise the total route length, but rather the sum of its inverted rates. In this way, the algorithm will (typically) identify routes with large bottleneck rates. To achieve rate maximisation, a potential ansatz for the tentative cost function may take the form,

$$F_{\omega}^{s,x \rightarrow y} = T_{\mathbf{x}} + K_{\mathbf{x}\mathbf{y}}^{-\eta} + \epsilon, \quad (6.14)$$

which corresponds to minimising the following global cost function,

$$\mathcal{F}_{\omega} = \sum_{(\mathbf{x},\mathbf{y}) \in \omega} (K_{\mathbf{x}\mathbf{y}}^{-\eta} + \epsilon). \quad (6.15)$$

This aligns with minimising the sum of the inverse point-to-point rates along a path ω , adding a penalty for channel usage ϵ associated with each link. Minimising this cost function simultaneously locates a path which maximises the sum of the point-to-point rates along the path while minimising its length; *indirectly* identifying a high-rate and edge-efficient route.

In this tentative cost function, $\eta, \epsilon \in \mathbb{R}_0^+$ are hyperparameters used to impose a tradeoff between rate and path length. The parameter η can be thought of as a rate motivator; if η is large, then the inverse term $K_{\mathbf{x}\mathbf{y}}^{-\eta}$ forces a large penalty to the cost function when low rate edges are included in the route. If the single edge rate is large, this incurs a low cost, motivating the use of an edge. Meanwhile, ϵ enforces a constant penalty term for edge usage, thus encouraging efficient routing. The edge-usage penalty can be generalised as an edge-wise property $\epsilon_{\mathbf{x}\mathbf{y}}$ such as spatial length etc., to more accurately manage routing efficiency. The ability to control this tradeoff is extremely useful and is typically ignored in the standard widest path routing algorithm.

This is by no means an optimal approach. Nonetheless, the enormous benefit associated with quickly locating additional routes proves to outweigh weaknesses in the approximation. In future investigations, it may be interesting to explore more sophisticated cost analyses, e.g. a neural network variational ansatz. Such ansätze may be of significant benefit when one wishes to optimise more than just rate and path length, e.g. balancing the routing priorities of multiple user pairs, or considering waiting-times in quantum memories for entanglement distribution.

6.3.4 Efficient Multi-path Protocols

There are two potential versions of multi-path protocol that we may consider. To enhance rates, one may insist upon the use of $M > 1$ end-to-end paths, where M is fixed. This is an intuitive approach which we call fixed route number protocols, denoted by $\mathcal{P}_{\text{mdp}}^M$. However, a pair of end-users may be more interested in preserving resources granted that they possess a particular rate guarantee. It is easy to devise a protocol in which there is a rate requirement; via the MDPAlg, we use the least number of end-to-end routes required to achieve K^* bits per protocol use. Denoted by $\mathcal{P}_{\text{mdp}}^{K^*}$, this is easily implemented strategy and reflects the quality of service principle in classical networks.

It is important to note that the MDPAlg is not optimal and may not always match the performance of iterative Dijkstra (or, of course, flooding). However, we will see that any minor cost in performance is vastly outweighed by gains in efficiency, as the MDPAlg can construct end-to-end multi-paths orders of magnitude faster than iterative Dijkstra approaches (see Ref. [103] for cost analyses). This makes it a much more practical method and allows us to perform numerical assessments that would not be possible otherwise.

6.4 Benchmarking Quantum Networks

6.4.1 Benchmarking Performance

Consider a class of quantum networks $\mathbb{N} = \{\mathcal{N}_i\}_i$ and a specific routing protocol \mathcal{P} . For any network $\mathcal{N} = (P, E) \in \mathbb{N}$ drawn from this class, one can define the average end-to-end rate as that which is averaged over all possible end-user pairs in the network. We denote this set of end-users by

$$\mathcal{I} := \{\mathbf{i}_j\}_j = \{\{\mathbf{a}_j, \mathbf{b}_k\}\}_{j \neq k}. \quad (6.16)$$

Then the average end-to-end rate takes the form,

$$\langle K \rangle_{\mathcal{N}|\mathcal{P}} := \frac{1}{|\mathcal{I}|} \sum_{\mathbf{i} \in \mathcal{I}} K(\mathbf{i}, \mathcal{N}|\mathcal{P}), \quad (6.17)$$

where $K(\mathbf{i}, \mathcal{N}|\mathcal{P})$ is the rate achieved between the end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ given that they engage in the routing protocol \mathcal{P} (given in Eq. (2.69)). In an N node network there exist C_N^2 potential end-user pairs (where C_n^k denotes the binomial coefficient), which may be far too large to consider explicitly. In practice, we will sample L end-user pairs from \mathcal{N} to approximate the average end-to-end rate, $\langle K \rangle_{\mathcal{N}|\mathcal{P}} \approx \frac{1}{L} \sum_{j=1}^L K(\mathbf{i}_j, \mathcal{N}|\mathcal{P})$ which can be computed along with statistical error considerations.

This quantity benchmarks the ability of quantum communication on a specific network. To assess the efficacy of a routing strategy more generally, we study the *ensemble average end-to-end rate* sampled across the entire network class. Given that any network $\mathcal{N} \in \mathbf{N}$ is drawn with equal probability, then the ensemble average rate is equal to

$$\langle K \rangle_{\mathbf{N}|\mathcal{P}} := \frac{1}{|\mathbf{N}|} \sum_{\mathcal{N} \in \mathbf{N}} \langle K \rangle_{\mathcal{N}|\mathcal{P}}. \quad (6.18)$$

For the network classes in which we are interested, \mathbf{N} can be incredibly large (or infinite). Hence, we are not able to exactly compute Eq. (6.18) but instead compute an accurate approximation over a sample space of M' networks $\{\mathcal{N}_i\}_{i=1}^{L'}$ such that $\langle K \rangle_{\mathbf{N}|\mathcal{P}} \approx \frac{1}{L'} \sum_{i=1}^{L'} \langle K \rangle_{\mathcal{N}_i|\mathcal{P}}$. This approximation can be made sufficiently accurate by taking enough end-user samples L and network class samples L' .

These concepts can be immediately translated in the context of end-to-end capacities rather than rates. Indeed, one can readily define a specific end-to-end capacity $\mathcal{C}(\mathbf{i}, \mathcal{N}|\mathcal{P})$, an average $\langle \mathcal{C} \rangle_{\mathcal{N}|\mathcal{P}}$ and ensemble average $\langle \mathcal{C} \rangle_{\mathbf{N}|\mathcal{P}}$. The only difference is the description of the point-to-point rates throughout the network; when network links are considered to operate at their capacity, then we are studying the end-to-end capacities. Otherwise, they are sub-optimal end-to-end rates. The same logic can of course be extended to capacity-bounds, as these quantities can be computed with respect to any link layer description.

6.4.2 Benchmarking Efficiency

Routing protocols should not only be benchmarked with respect to performance, but also with respect to efficiency, i.e. what proportion of network resources is required to guarantee a particular level of performance? Therefore, it is useful to define a measure which we call the *routing consumption*, which quantifies the proportion of network edges used via a given routing protocol. Given a network $\mathcal{N} = (P, E)$, an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ and a routing protocol \mathcal{P} , the routing consumption \tilde{E} is the fraction of network edges required to perform communication,

$$\tilde{E}(\mathbf{i}, \mathcal{N}|\mathcal{P}) := \frac{|E_{\mathcal{P}}(\mathbf{i}, \mathcal{N})|}{|E|}, \quad (6.19)$$

where $|E_{\mathcal{P}}(\mathbf{i}, \mathcal{N})|$ is the number of edges in the routing edge set connecting the end-user pair. We can then define an average end-to-end routing consumption by averaging over end-user pairs

$$\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}} := \frac{1}{|\mathcal{I}|} \sum_{\mathbf{i} \in \mathcal{I}} \tilde{E}(\mathbf{i}, \mathcal{N}|\mathcal{P}), \quad (6.20)$$

which of course extends to an ensemble average over a class of networks

$$\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}} := \frac{1}{|\mathcal{N}|} \sum_{\mathcal{N} \in \mathcal{N}} \langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}}, \quad (6.21)$$

following the notation convention used for end-to-end rates. As before, these quantities are estimated by sampling from the set of end-user pairs and the network class.

The routing consumption is a useful measure of how resource efficient a network protocol is at achieving its rates. A protocol which can attain high rates with a low routing consumption is clearly a desirable strategy. Hence, a tradeoff exists between end-to-end rates and routing consumption. While we know flooding to be rate optimal, it will always satisfy $\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}_f} = 1$ because it utilises all network edges. Hence, it is likely sub-optimal from the perspective of routing consumption. For single-path protocols we can typically expect $\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}_{sp}} \ll 1$, but its rates may be poor. It is most interesting to inspect the behaviour of $\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}_{mp}}$ as understanding the relationship between end-to-end rate, resource consumption and practical multi-path routing is poorly understood.

6.4.3 Criticality of Quantum Networks

It is essential to pursue a rigorous and quantitative assessment of practicality with respect to quantum network routing, and allow us to understand the core network features which guarantee both performance and efficiency. To this end, the concept of *network criticality* is vital. A critical transition is an abrupt regime shift in the behaviour of some complex system, such that the system transits from sub-to-super-critical with respect to a particular property. In the context of complex quantum networks, we have already discussed how critical transitions may occur with regards to robustness (connectivity) but it can be extended to notions of performance [65] or routing efficiency. A major contribution of this chapter is to introduce routing consumption-criticality, identifying network regimes within which end-to-end routing is efficiency via practical routing.

We begin with performance-defined criticality; critical properties within a quantum network model which a guarantee a transition from unreliable rates to consistent super-critical rates. A useful manifestation of performance-based criticality that can arise is end-to-end distance independent rates, i.e. the spatial separation of users is independent of communication quality. This is especially important for quantum networks in order to overcome point-to-point limitations. Furthermore, these measures are more appropriate than purely connectivity defined quantities, since they make meaningful assessments of the efficacy of a network to perform quantum communications.

Definition 6.1 (Performance-defined Critical Density): *For a class of quantum networks \mathbf{N} , and a network routing protocol \mathcal{P} , we define $\rho_{\mathbf{N}|\mathcal{P}}^*$ as the minimum nodal density required to guarantee an ensemble average rate of $\langle K \rangle_{\mathbf{N}|\mathcal{P}} \geq 1$ bit/protocol use.*

The critical nodal density $\rho_{\mathbf{N}|\mathcal{P}}^*$ is an extremely important measure since quantum repeaters are costly and they should be minimised as a resource in a future quantum internet. Investigations of this quantity with respect to flooding $\rho_{\mathbf{N}|\mathcal{P}_{\text{fl}}}^*$ and capacity achieving bosonic lossy networks have been carried out which can be considered a lower-bound for all other critical densities. Indeed, minimising this quantity over all protocols is equivalent to the flooding-defined critical density,

$$\rho_{\mathbf{N}}^* := \min_{\mathcal{P}} \rho_{\mathbf{N}|\mathcal{P}}^* = \rho_{\mathbf{N}|\mathcal{P}_{\text{fl}}}^* \leq \rho_{\mathbf{N}|\mathcal{P}}^*. \quad (6.22)$$

Nonetheless, analogous studies have not been carried out for other routing protocols or networks.

The performance-based critical density is a valuable characteristic of a quantum network model, however it lacks insight to the resources required to achieve critical rates. Networks may be super-critical with respect to performance but may demand impractical resources to do so. To address this we can analyse efficiency regimes with respect to routing, defining a consumption-based critical density.

Definition 6.2 (Consumption-defined Critical Density): *Consider a class of quantum networks \mathbf{N} , and let $\mathbf{N}(\rho) \subset \mathbf{N}$ be a subset of this class with nodal density ρ . For a network routing protocol \mathcal{P} , we define $\tilde{\rho}_{\mathbf{N}|\mathcal{P}}^*$ as the density at which the routing consumption is maximised and after which it undergoes consistent decay. More precisely, $\tilde{\rho}_{\mathbf{N}|\mathcal{P}}^* := \arg \max_{\rho} \langle \tilde{E} \rangle_{\mathbf{N}(\rho)|\mathcal{P}}$, such that $\langle \tilde{E} \rangle_{\mathbf{N}(\rho)|\mathcal{P}} \leq \langle \tilde{E} \rangle_{\mathbf{N}(\tilde{\rho}_{\mathbf{N}}^*)|\mathcal{P}}$, $\forall \rho \geq \tilde{\rho}_{\mathbf{N}}^*$.*

Hence, $\tilde{\rho}_{\mathbf{N}|\mathcal{P}}^*$ represents a secondary critical measure which separates network classes into different efficiency regimes. Executing the same protocol, networks which possess a nodal density $\rho > \tilde{\rho}_{\mathbf{N}|\mathcal{P}}^*$ will consume a smaller fraction of network resources during routing. Furthermore, this efficiency increase as the density continues to grow. We may connect the concepts of performance and routing consumption within the following definition.

Definition 6.3 (δ -Critical routing consumption): *Consider a class of quantum network \mathbf{N} . The critical routing consumption $\langle \tilde{E} \rangle_{\mathbf{N}}^*$ is the minimum ensemble average fraction of network edges that must engage in end-to-end routing granted the ensemble average rate*

is $\langle K \rangle_{\mathcal{N}|\mathcal{P}} \geq \delta$ bits per network use, given that $\delta \leq \langle K \rangle_{\mathcal{N}|\mathcal{P}_a}$. More precisely, the critical routing consumption is

$$\langle \tilde{E} \rangle_{\mathcal{N},\delta}^* := \min_{\mathcal{P}|\langle K \rangle_{\mathcal{N}|\mathcal{P}} \geq \delta} \langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}}. \quad (6.23)$$

The definition in Eq. (6.23) can be made intuitive in the following way: there exists a critical routing consumption achieved by an end-to-end protocol which can promise δ bits per network use on average, such that δ is an attainable network rate. Further to this performance guarantee, the protocol minimises the ensemble average fraction of network edges used, i.e. on average, it is the most efficient routing mechanism which promises the target rate. It is by no means clear how to determine the optimal protocol. However, the tools developed in this chapter can effectively bound the critical routing consumption via practical routing schemes,

$$\langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}_{sp}} \leq \langle \tilde{E} \rangle_{\mathcal{N},\delta}^* \leq \langle \tilde{E} \rangle_{\mathcal{N}|\mathcal{P}_{mp}}. \quad (6.24)$$

The optimal single-path protocol can always be a lower-bound which is saturated iff an ensemble average δ bits per network use is achievable via single-path routing. An appropriate multi-path strategy generates the upper-bound. Indeed, there will always exist a multi-path strategy capable of this since one can employ a flooding protocol in the worst-case (least efficient) scenario.

6.5 Numerical Results

In this section we apply the tools developed and discussed throughout this chapter to provide valuable insight into the realistic requirements of quantum networks. We wish to concretely identify meaningful limits of quantum networks exposed by considering realistic link layers and practical routing protocols.

6.5.1 Benchmarking Waxman Fibre Networks

In Fig. 6.5 we present relationships between nodal density, routing consumption and end-to-end capacities on bosonic thermal loss Waxman networks, with an upper-bounding capacity distribution \mathcal{K}_u (lower-bounding capacity distribution \mathcal{K}_l). Each network edge represents an optical-fibre link of loss rate 0.2 dB/km and environmental thermal noise $\bar{n} = 1/500$ connecting ideal transmitter/detectors. The Waxman parameters are set as $r_0, \beta = 100, 1$ ($r_0, \beta = 63, 1$), where the decay parameter corresponds to the approximate

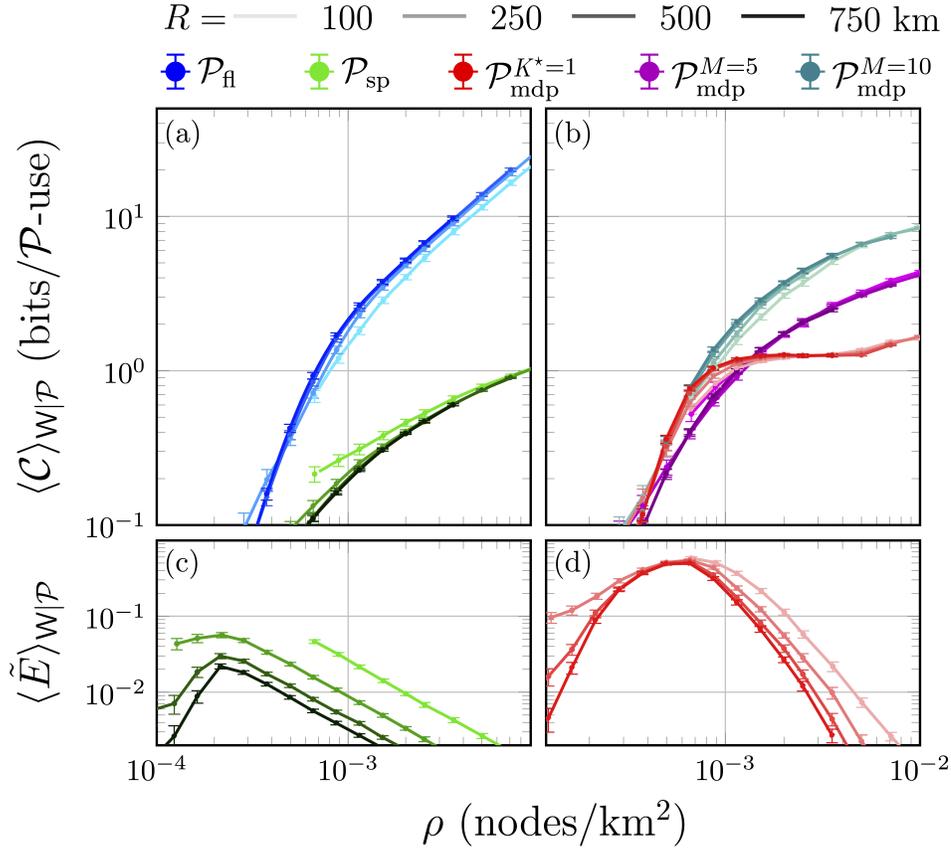


Figure 6.5: Relationships between performance, routing consumption and nodal density in bosonic thermal-loss Waxman networks. The network link layer in all cases is described by the upper-bound capacity distribution \mathcal{K}_u from Eq. (6.2) and model parameters $(r_0, \beta) = (100, 1)$. Panels (a)-(b) plot the ensemble average end-to-end capacities achieved by different routing protocols and a number of network radii, R (each protocol and network radius is colour coded with the routing protocol and R legends). Panels (c)-(d) show the ensemble average routing consumption of (c) single path routing and (d) $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing via the MDPAlg.

distance at which the single-edge capacity upper-bound collapses to zero ~ 100 km (~ 63 km). Since we are considering end-to-end *capacities*, results concerning these networks offer universal benchmarks for any fibre-based Waxman network.

Figs. 6.5(a)-(b) depict the ensemble average end-to-end capacities with respect to nodal density for a number of routing strategies. In Ref. [65], the authors identified a performance-based critical nodal density of $\rho_N^* \approx 4.25 \times 10^{-4}$ nodes/km² associated with flooding and bosonic pure-loss networks. The consideration of environmental thermal-noise naturally increases the critical density such that $\rho_N^* \approx 7.35 \times 10^{-4}$ (1.04×10^{-3}) nodes/km² (the

$$R = 500 \text{ km}, N = 1500, r_i = 800 \text{ km.}$$

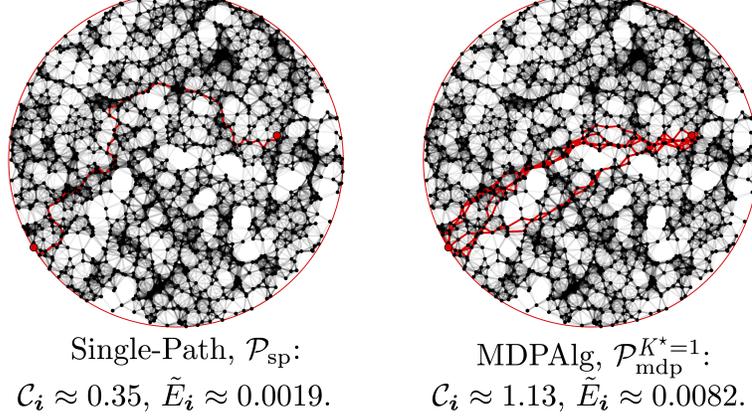


Figure 6.6: Illustration of achievable end-to-end routes on an example network (according to the parameters listed) using \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ for users separated by $r_i \approx 800$ km. Black edges in the networks identify unused edges, while red edges are those which engage in the routing protocol.

density in parentheses refers to the lower-bounding capacity distribution). While significant, this shift remains relatively optimistic since the considered thermal-loss link layer does not consider additional experimental sources of thermal noise in an effort to remain protocol agnostic.

Critical densities with respect to flooding protocols are inherently optimistic, due to the unrealistic resource demand. This optimism is emphasised when one considers the utility of single-path routing. Fig. 6.5(a) clearly illustrates the intense demands on the network density required for single-path routing to reach criticality, and promise effective end-to-end rates. We find the single-path critical nodal density to be approximately 8.52×10^{-3} (8.86×10^{-3}) nodes/km², which is an order of magnitude larger than that predicted by flooding. On large-scales, a stark increase of this magnitude has significant ramifications for the cost and deployability of quantum networks. Therein lies the necessity for efficient multi-path routing protocols. Fig. 6.5(b) illustrates the efficacy of MDPAlg based routing, showing that flooding is not necessary to preserve lower critical densities. Each variant of the multi-path protocol is able to recover a critical density close to that offered by flooding, reinforcing the notion that high rates can be guaranteed with realistic resources.

Further evidence of this is gathered in Figs. 6.5(c)-(d) which plot the routing consumption of the protocols \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ respectively. While single-path protocols are expected to achieve very low routing consumptions, the same expectation is not necessarily held for

multi-path routing. However, the protocol $\mathcal{P}_{\text{mdp}}^{K^*=1}$ can obtain a critical density of approximately 9.26×10^{-4} (1.35×10^{-3}) nodes/km², while maintaining a maximum average routing consumption of $\lesssim 0.5$. That is, no more than half of the network edges are ever required to participate in end-to-end routing. Even before single-path routing is a plausible option, the routing consumption of $\mathcal{P}_{\text{mdp}}^{K^*=1}$ undergoes a rapid decay as the nodal density passes $\sim 6.6 \times 10^{-4}$ (1.15×10^{-3}) nodes/km². For networks with nodal density $\rho \geq 3 \times 10^{-3}$, this protocol can guarantee super-critical rates while consuming $\lesssim 5\%$ of the network edges during routing. Being able to minimise edge usage while guaranteeing high rates is vital to scalability in large quantum networks with many users.

Crucially, in the absence of effective single-path routing, efficient multi-path strategies exist which preserve network resources while obtaining reliable, super-critical end-to-end rates. The importance of this result is emphasised when more realistic link layers are considered, such as those described by practical QKD protocols (and beyond).

6.5.2 Practical Link Layers and Network Phases

Performing the previous analyses of end-to-end rates, routing consumption and nodal density for different link layers, we can build a comprehensive picture of the efficacy of quantum communication networks in different scenarios. As such, we can establish relationships between different critical network properties which give rise to key *network phases*; sub-classes of Waxman networks for which the behaviour of end-to-end quantum communication have similar characteristics. Here, phases are defined through statistical analyses of the properties deemed most important to end-to-end communication; connectivity, performance and routing efficiency. Identifying what these phases are and where they fall within meaningful density ranges can help us to understand the realistic needs of quantum networking.

Fig. 6.7 summarises the Waxman quantum network phases defined in this study for a number of link layer models: bosonic thermal-loss capacity distributions, and asymptotic CV-QKD rate distributions. Here, we have identified six key network phases, explicitly defined in Fig. 6.7(b), ranging from Phase I networks with no connectivity guarantees, to Phase VI networks which promise super-critical performance via single-path routing. In between, there exists a spectrum of phases corresponding to critical changes in the connectivity, routing consumption and performance guarantees.

The most desirable network phase is naturally Phase VI, in which single-path routing is sufficient to perform reliable quantum communication. Unfortunately, our primary takeaway is that the nodal densities necessary to inhabit Phase VI are very large, and become

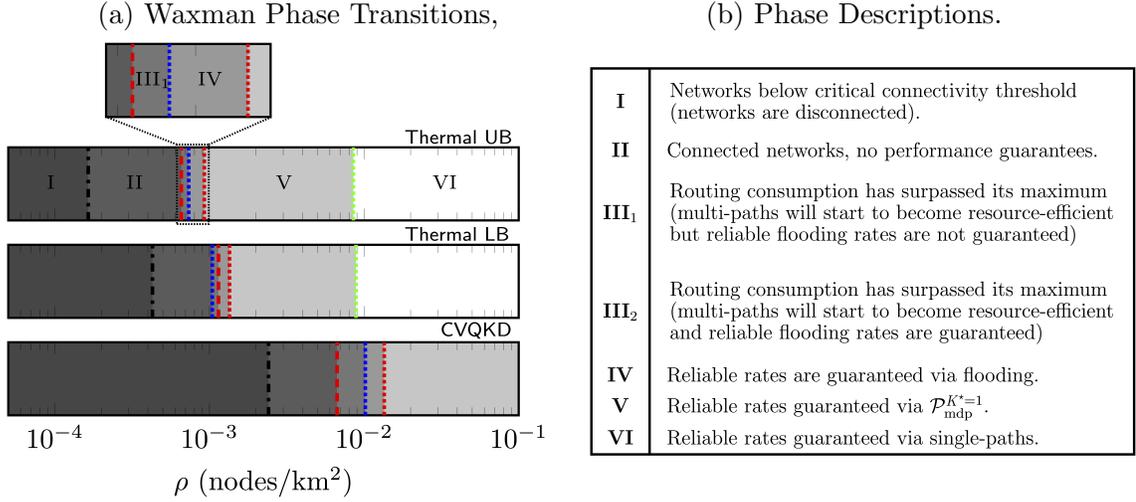


Figure 6.7: Waxman quantum network phase characterisations with respect to link layer descriptions and nodal density. Panel (a) outlines connectivity, consumption and performance based critical densities with respect to networks composed of different link layers. These transitions give rise to network phases in which we can expect particular properties. These phases are labelled on the density diagram, while Panel (b) summarises and describes their implications for quantum networking. Phases III₁ and III₂ describe similar network properties, but differ as to whether the maximum routing consumption is surpassed before or after the flooding-based performance transition.

impractical as point-to-point link layer descriptions become more realistic. Even when considering fully trusted QKD networks described using asymptotic secret key-rates, the nodal densities required to reach Phase VI are more than 10^{-1} nodes/km². For example, the number of nodes needed to deploy a fully trusted QKD network that spans the surface area of Europe ($\sim 10^7$ km²) would be on the order of millions.

Fortunately, efficient multi-path routing strategies can resurrect the utility of lower density networks which achieve reliable rates. Networks which occupy Phase V are supercritical provided that they employ the protocol $\mathcal{P}_{\text{mdp}}^{K^*=1}$, or better. As shown in Fig. 6.7 for each considered link layer, Phase V occupies a more practical nodal density region (around an order of magnitude improvement) for which practical multi-path strategies can promise strong rates using efficient protocols; not just flooding.

An important observation that we wish to reiterate is the existence of Phases II and IV. Phase II portrays the resource gap between connectivity guarantees and the most optimistic performance guarantee (via flooding). The existence of this gap makes it clear that the design and assessment of quantum networks *cannot* solely focus on connectivity analyses, as such promises are not enough to guarantee useful quantum communication. Analogously,

Phase IV identifies a gap between the resources required by flooding and practical multi-path protocols for critical rates. Closing the gap posed between Phase IV and V while minimising routing consumption is the goal of any multi-path strategy.

These insights emphasise the need for explicit analyses of end-to-end performance and reiterate the point that routing in quantum networks cannot naïvely follow its classical counterpart. Multi-path routing techniques do not just represent a means of boosting rates but establish a pertinent method of reducing the resource demands of practical quantum network development.

6.5.3 Scale-Free Properties and Quantum Networks

Scale-free architectures capture important features of real world networks beyond that of the Waxman model. As discussed in Section 6.2.3, scale-free networks obey a power-law cumulative degree distribution. This realises a connectivity structure in which there exist a number of highly connected hubs to which many nodes of low degree are connected, giving rise to a lower average degree than Waxman models. Considering the network class \mathcal{S}_{σ_r} , we investigate the ramifications that such connectivity features have on end-to-end routing and performance

Fig. 6.8 displays analyses of bosonic thermal-loss scale-free networks with respect to nodal density and a number of routing protocols: flooding, single-path, and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing using the edge-disjoint MDPAIlg. We consider architectures generated with the parameter $\sigma_r \in \{1, 2\}$ so to vary the awareness of link-length during network creation and consequently its “strictness” with respect to scale-free behaviour. In Figs. 6.8 (a)-(c) we study the ensemble average end-to-end capacities for each network model and routing strategy. It is clear for $\sigma_r = 1$ (in which scale-free behaviour is followed closely) that performance is restricted by the low average degree. At network scales of $R < 150$ km and increasing density, the flooding capacity is able to reach reliable rates, but only at very high densities. Furthermore, Ref. [65] showed that the optimal flooding capacity of such networks exponentially decays with respect to R (regardless of the number of nodes).

Practical routing strategies cannot do better, as displayed for both single-path routing and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing. Ultimately, the connectivity structure of scale-free networks $\mathcal{N} \in \mathcal{S}_1$ display a poor aptitude for multi-path routing. The utilisation of multiple routes is only helpful if many additional routes can be found and the existence of high-degree hubs limits this possibility. User nodes will typically only connect to a single hub, limiting the ability to reinforce the end-to-end path set. This limitation can be seen in the ensemble average routing consumption of the $\mathcal{P}_{\text{mdp}}^{K^*=1}$ protocol, which remains relatively constant with respect

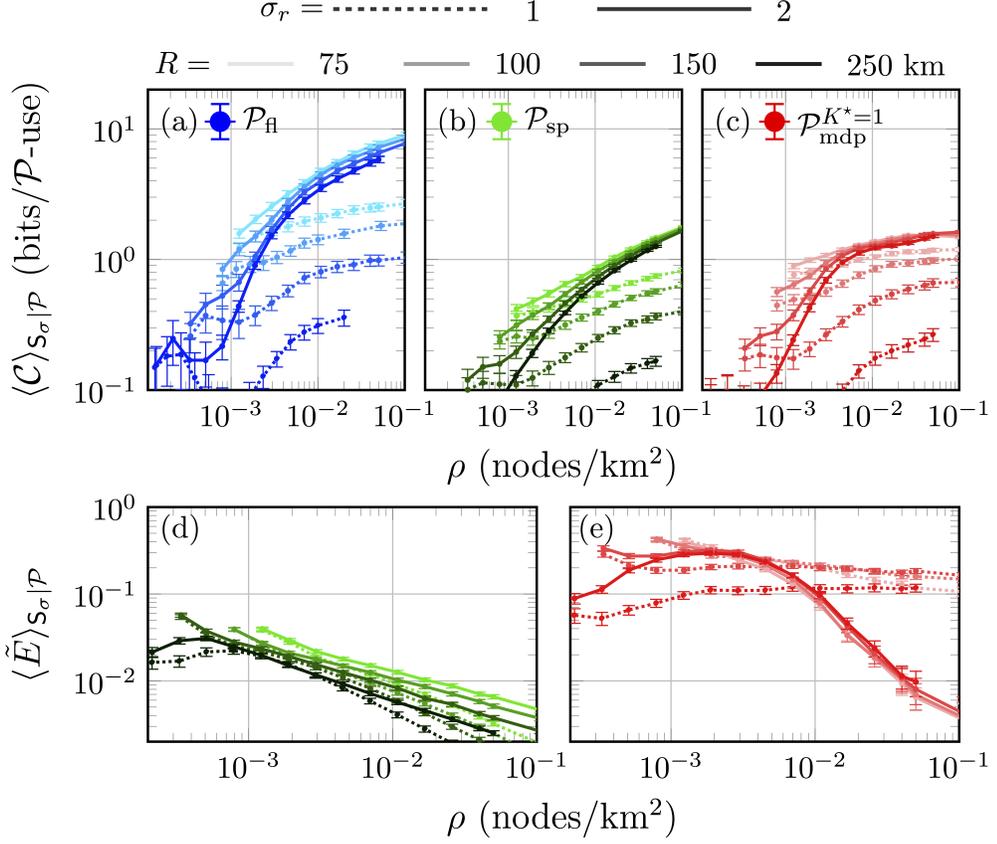


Figure 6.8: Relationships between performance, routing consumption and nodal density in bosonic thermal-loss scale-free networks. Throughout all plots, the network link layer is described by the upper-bound capacity distribution \mathcal{K}_u from Eq. (6.2) and we use the model parameters $(n_0, m, \sigma_{\text{deg}}) = (10, 5, 1)$. We also consider unique values for the scale-free model parameter $\sigma_r \in \{1, 2\}$ which are distinguished in the legend. Panels (a)-(c) depict the ensemble average capacity with respect to nodal density for \mathcal{P}_{fl} , \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing respectively. Panels (d)-(e) plot the ensemble average routing consumption with respect to nodal density for \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ routing respectively. Each analysis is performed for a number of network radii R listed in the legend.

to nodal density. This means that there is little correlation between increasing density and the number of effective end-to-end routes between user nodes.

Breaking from strict scale-freedom, we see that effective performance *can* re-emerge within the network class S_2 . With an increased connection probability dependence on link-length, the ensemble average capacity and routing consumption follow similar trends to the Waxman model: the exponential decay of the flooding capacity with respect to network scale is subdued, and end-to-end capacities of $\langle C \rangle_{S_2 | \mathcal{P}} \geq 1$ bit/protocol-use can be guaran-

$$R = 150 \text{ km}, N = 1000, \sigma_r = 2, r_i = 200 \text{ km}.$$

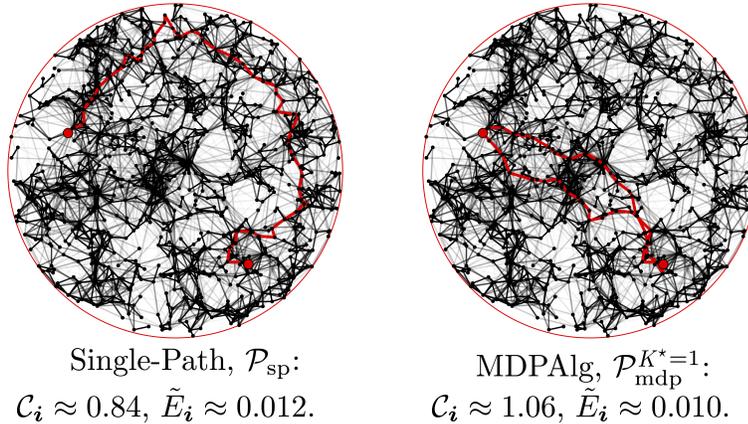


Figure 6.9: Illustration of achievable end-to-end routes on an example network (under the parameters shown) using \mathcal{P}_{sp} and $\mathcal{P}_{\text{mdp}}^{K^*=1}$ for users separated by $r_i \approx 200$ km. Black edges in the networks identify unused edges, while red edges are those which engage in the routing protocol.

teed for each routing protocol (within reasonable nodal density ranges). Furthermore, the multi-path routing consumption reliably decays with increasing density, implying that an applicability for multi-path routing can be reestablished.

6.6 Conclusion

In this chapter, we have investigated the capabilities of quantum communication networks under practical routing strategies, in an effort to gain insight to realistic requirements of future quantum networks. With the goal of providing a comprehensive study of the performance and feasibility of quantum optical fibre-networks, we have combined theory of quantum communication with that from random network models and network routing algorithms. Our work expands upon previous developments in this domain by considering realistic link layer descriptions (capacities and secret-key rates over bosonic thermal-loss channels) alongside practical routing protocols (efficient multi-path and single-path routing methods). These assessments focus on the crucial Waxman and scale-free classes of random quantum networks.

Through our statistical analyses, we reveal vital insights into the practical requirements of critical behaviour in quantum networks, where criticality may be defined with respect to connectivity, performance and routing efficiency. For quantum Waxman networks, a network phase structure is introduced with respect to nodal density; identifying ranges of

nodal densities within which networks can be classified as sub/super-critical with respect to these properties. We show that for increasingly realistic link layer descriptions, the existence of a performance super-critical phase with respect to single-path routing becomes more and more difficult within practical density ranges. Nonetheless, multi-path routing strategies can re-establish practical density ranges and super-critical rates, reiterating the necessity for the development of multi-path routing methods in quantum networks.

It is important to note that the network phases we have identified are classifications which emerge naturally from the heuristics considered in this chapter, but are by no means definitive. With increasingly sophisticated investigations in the future, we expect richer and more detailed phases to be unveiled which provide insight and guidance for quantum network design.

Our findings reiterate the challenging relationship between scale-free networks and quantum communication. While scale-free architectures such as Yook's model [105] offer useful insight into the structure of the classical internet, they do not find analogy with reliable quantum fibre networks. Adherence to scale-freedom in large-scale quantum networks results in a poor communication quality regardless of routing strategy. Adjusting the parameters in Yook's model to generate networks with an increased sensitivity to link length, scale-freedom begins to break while the ability to perform quantum communication strengthens. These results stress the resounding differences between the classical and quantum internet: quantum networks need to be *consistently* well connected to overcome point-to-point limitations, and must be strictly engineered with nodal density and link length in mind.

While our results focus on the efficacy of fibre-networks, it is well motivated that the quantum internet will exploit satellite links and inter-satellite networks to reinforce long-range communication. This is an area of serious interest, and a future investigative path for extending this research; investigating the interplay of realistic, ground-based networks interacting with satellite based infrastructure. Furthermore, the study of ground based free-space networks in both long-range (inter-metropolitan) and mobile (metropolitan) settings are of immediate interest.

Finally, understanding how many end-user pairs may simultaneously communicate across realistic quantum networks is of paramount importance. Indeed, the trade-off between network architecture, end-to-end rates and number of end-users is poorly understood. While the study of end-to-end routing consumption is a useful step in this direction, greater progress must be made in order to develop practical routing strategies for the quantum internet.

Chapter 7

Multi-User Limitations of Quantum Communication Networks

The work in this chapter forms the basis of a paper currently being prepared for submission, whose authors are (in order) Cillian and Stefano Pirandola. Section 7.1 overviews the aims of this chapter, where we wish to expand our theoretical toolbox so to delve into the investigation of multi-user quantum networking. Section 7.2 reviews some previously established theory before asking: how do we benchmark multiple-unicast quantum networking with respect to practical routing protocols? In Section 7.3 we introduce an analytical interpretation of multi-user routing via novel concepts of traffic management and network filters. In this way, we devise naïve strategies for handling simultaneous communicators on a network and provide achievable benchmarks for multiple-unicast quantum communications. Section 7.4 exhibits the utility of these methods by applying them to the class of quantum Waxman networks to provide initial insight to multi-user limitations. Section 7.5 then concludes and discusses future directions.

7.1 Introduction

Throughout this thesis, we have explored challenges that quantum communication networks face. In particular, we have aimed to better understand how the fundamental rate limitations of remotely sharing quantum information manifest within the grander goal of quantum networking. The infrastructure built to facilitate quantum networking will need to take these limitations into account, and use them to motivate high-rate architectures. Understanding the trade-offs between end-to-end performance, distance-independence and cost efficiency will be vital for a successful and practical quantum internet. These trade-offs have been investigated in a number of different ways, via idealised architectures [1, 2, 112],

complex, random models [64, 113] and end-to-end network capacity assessments [65, 66]. Yet, these works have been primarily focussed on the capability of *single-unicast* quantum communication (a single end-user pair communicating freely across a network). For quantum networks to be cost effective, they must enable multi-user functionality and provide practical rates in a multiple-unicast scenario for which the resource/performance relationship is poorly understood.

During multi-user quantum communications the presence of many end-users attempt to route at the same time gives rise to *routing competition*. End-users cannot always build their preferred routes as parts of them may already be in use by other users. Classical networks are capable of overcoming the increased resource demands of multi-user communication through network coding; point-to-point links can be used to simultaneously transmit multiple messages at high rates via pre-transmission encoding and post-transmission decoding. The fragility of quantum information makes network coding significantly less effective, and is only reliable if network nodes possess prior entanglement making it an expensive and impractical approach. Since quantum networks do not have the luxury of expedient network coding, the demand for sophisticated protocols for routing and traffic management are essential.

In this work, we expand the task of quantum network benchmarking into the multi-user domain, devising theoretical tools that allow us to inspect *achievable* multiple-unicast rates on quantum networks. Through the concept of *network filters* we describe a protocol structure by which multi-user objectives during multiple-unicast are offloaded from the end-users who need only execute an independent end-to-end routing strategy. In this way, we are able to benchmark deployable and practical routing protocols in the multiple-unicast communication setting. Consequently, we apply these tools in the context of a random bosonic thermal-loss networks (Waxman model), providing critical benchmarks for realistic, future quantum networks. We identify regimes of network resources for which reliable end-to-end rates can be promised between many users. Our hope is that these results are illuminating for the development of advanced multi-user quantum network protocols, revealing effective standards for quantum networking.

7.2 Multi-User Quantum Networking

7.2.1 Unicast Communications

Consider the single-unicast setting, in which a quantum network is used to facilitate communication between a single end-user pair (EUP) $i = \{\mathbf{a}, \mathbf{b}\}$. Single-unicast is the

scenario that we have investigated throughout this thesis, and serves as a fundamental primitive for all other scenarios in which a single EUP is considered as the only pair of user nodes on the network. As such, they can avail of as much of the network during routing as required/demanded by their routing protocol and can always choose their preferred route. If this cannot be performed efficiently, then no other multi-user protocol will be able to either.

In the multiple-unicast setting, a quantum network is used to facilitate communication between not just one end-user pair (EUP) $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ but *many* EUPs

$$\mathcal{I} := \{\mathbf{i}_1, \dots, \mathbf{i}_{N_u}\}, \quad (7.1)$$

that wish to engage in unique communication sessions [114]. Quantum networks should be designed and operated in such a manner that this is possible, effective and reliable, i.e. high rates should be achievable even under the scrutiny of multiple-unicast communications.

7.2.2 Capacity/Rate Regions

Useful benchmarks for multi-user quantum networking do exist, and have been studied in Refs. [114, 115]. Ultimately, bounding the performance of *many* communicators is more complicated than end-to-end scenarios, as there is no single metric that quantifies the effectiveness of communication for multiple users. Typically, one considers a k -length *capacity/rate region*¹ which is defined as the sum of sender/receiver capacities/rates for every possible combination of k end-user pools. That is for N_u end-user pairs in multiple-unicast such that K_i is the end-to-end rate between the i^{th} EUP then the rate regions are,

$$\underbrace{\{K_i \mid \text{for any } i\}}_{k=1}, \underbrace{\{K_i + K_j \mid \text{for any } i \neq j\}}_{k=2}, \dots, \underbrace{\left\{ \sum_{i=1}^{N_u} K_i \right\}}_{k=N_u}. \quad (7.2)$$

Clearly, when $k = N_u$ then the rate region is simply the sum of all the end-to-end rates within the multi-user scenario. Rate regions are thus able to collectively describe the performance limits of multi-user networking. Relationships forged between particular end-user pairs will manifest within these regions and characterise the network.

It is then possible to provide upper bounds for multi-user quantum networking which are upper bounds on each of the regions with respect to a certain networking scenario, i.e. multiple-unicast, single or multiple multicast. These upper bounds are derived by generalising the methods reviewed in Section 2.4 for bounding the end-to-end rates/capacities

¹We will use capacity regions when referring to networks composed of capacity achieving links, and rate regions when considering networks composed of links operating at achievable rates.

of quantum networks. Given a configuration of multiple senders and receivers, one may describe a general adaptive protocol whose goal is to prepare a multipartite collection of target states (maximally entangled states, private states, etc.) across the network. The exact structure of this target state depends on the communication type, and the number of users that are participating. The tools of channel simulation and teleportation stretching can once more be applied to network such that it can be described via a resource representation, $\sigma(\mathcal{N}) = \{\sigma_{\mathbf{x}\mathbf{y}}\}_{(\mathbf{x},\mathbf{y}) \in E}$. In this way, one can upper bound the capacity of every edge in the network via the REE of its resource state $\mathcal{C}_{\mathbf{x}\mathbf{y}} \leq E_R(\sigma_{\mathbf{x}\mathbf{y}})$. Consequently, the REE of the target network state can be decomposed into the REE of an entanglement cut which can be made across the network. Through this methodology, Pirandola was able to successfully apply upper bounds to all of the capacities of multi-user communication settings of interest (for detailed proofs and discussions, see Refs. [114, 115]).

7.2.3 Multi-User Network Cuts

To produce these bounds, it is necessary to generalise our notion of network cuts. Thus far in this thesis, we have only ever concerned ourselves with network cuts defined with respect to single end-user pairs. Given a network $\mathcal{N} = (P, E)$ we have defined a cut C between end-users Alice \mathbf{a} and Bob \mathbf{b} as that which partitions a network by removing its corresponding cut-set of edges \tilde{C} such that Alice and Bob are separated into two disjoint node-sets, i.e. $\mathbf{a} \in \mathbf{A}$, $\mathbf{b} \in \mathbf{B}$ such that $\mathbf{A}, \mathbf{B} \subseteq P$ and $\mathbf{A} \cap \mathbf{B} = \emptyset$. This is an acceptable definition with respect to a single EUP, but *not* when there are multiple senders and receivers. Hence, let us generalise our definition of a network cut. Let us generally consider a collection end-user pairs $\mathcal{I} = \{\mathbf{i}_j\}_j = \{\mathbf{a}_j, \mathbf{b}_j\}_j$. The object \mathcal{I} can be decomposed into a collection of senders $\{\mathbf{a}_i\}_{i \in \mathcal{I}}$ and a collection of receivers $\{\mathbf{b}_i\}_{i \in \mathcal{I}}$ (where we subscript via \mathbf{a}_i to relate each sender/receiver to their pair-notation). One can specify a network cut

$$C : \{\mathbf{a}_i\} | \{\mathbf{b}_i\}, \forall i \in \mathcal{I}, \quad (7.3)$$

as that which bipartitions the network node set P in such a way that all $\{\mathbf{a}_i\}_{i \in \mathcal{I}} \subseteq \mathbf{A}$ and $\{\mathbf{b}_i\}_{i \in \mathcal{I}} \subseteq \mathbf{B}$, such that $\mathbf{A}, \mathbf{B} \subseteq P$ and $\mathbf{A} \cap \mathbf{B} = \emptyset$.

7.2.4 Upper Bounds for Multiple-Unicast

Through the appropriate resource representation of a quantum network and the expanded definition of a network cut, it is possible to compute upper bounds for multiple-unicast networking. Assuming a collection of EUPs \mathcal{I} and the use of single-path routing,

one can bound the single-path rate regions [114],

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) \leq \min_{C: \mathbf{a}_i | \mathbf{b}_i} \max_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \in \mathcal{I}, \quad (7.4)$$

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) + K(\mathbf{i}', \mathcal{N} | \mathcal{P}_{\text{sp}}) \leq \min_{C: \mathbf{a}_i \mathbf{a}_{i'} | \mathbf{b}_i \mathbf{b}_{i'}} \max_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \neq \mathbf{i}' \in \mathcal{I}, \quad (7.5)$$

⋮

$$\sum_{\mathbf{i} \in \mathcal{I}} K(\mathbf{i}, \mathcal{N} | \mathcal{P}_{\text{sp}}) \leq \min_{C: \{\mathbf{a}_i\} | \{\mathbf{b}_i\}} \max_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \in \mathcal{I}. \quad (7.6)$$

While these single-path benchmarks are useful, they are accompanied with the same caveat we have come across throughout this thesis; single-path routing is weak and typically ineffective for quantum networks (especially on a large-scale). Clearly, it is desirable to extend these bounds to account for multi-path routing strategies.

A multi-path strategy that is easy to inspect is flooding. Flooding allows us to express ultimate upper bounds on the capacity regions of multiple-unicast quantum networking by assuming the use of simultaneous flooding protocols between all end-user pairs. This results in the upper bounds [114],

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}) \leq \min_{C: \mathbf{a}_i | \mathbf{b}_i} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \in \mathcal{I}, \quad (7.7)$$

$$K(\mathbf{i}, \mathcal{N} | \mathcal{P}) + K(\mathbf{i}', \mathcal{N} | \mathcal{P}) \leq \min_{C: \mathbf{a}_i \mathbf{a}_{i'} | \mathbf{b}_i \mathbf{b}_{i'}} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \neq \mathbf{i}' \in \mathcal{I}, \quad (7.8)$$

⋮

$$\sum_{\mathbf{i} \in \mathcal{I}} K(\mathbf{i}, \mathcal{N} | \mathcal{P}) \leq \min_{C: \{\mathbf{a}_i\} | \{\mathbf{b}_i\}} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}} E_R(\sigma_{\mathbf{x}\mathbf{y}}), \text{ for all } \mathbf{i} \in \mathcal{I}. \quad (7.9)$$

While these upper bounds are useful for identifying the ultimate limits of multi-user performance, they are not necessarily realistic. As carried out in the end-to-end setting, they assume all users to indulge in a flooding protocol so that communication between all EUPs necessitates the use of every edge in the network. As before, this is a wasteful procedure that may not be efficient in practice. Furthermore, applying the concept of flooding manages to bypass the key problem of routing within multi-user networking. How do many EUPs know *how* to divide available resources and negotiate reliable routes across the network? During realistic deployment of multiple-unicast sessions, EUPs will need to compete and compose end-to-end strategies in spite of diminishing resources.

7.2.5 The Luxury of Time-Independence

There is a prominent convenience to the study of single-unicast rates in that they are *time-independent*. We focus upon a single EUP as the only network users with total freedom to complete their routing protocol and execute quantum communication without time constraints. The network on which their routing search is performed is fixed at all times during routing and remains static until they communicate. This is highly convenient from the perspective of performance analysis, as we only need to worry about establishing an end-to-end route set and then computing the end-to-end capacity by solving the max-flow min-cut relative to the route set (as introduced in Chapter 2, Section 2.4) without worrying about any time dependencies.

Even in the multiple-unicast domain, the ultimate upper bounds in the previous section are similarly time-independent. This is thanks to the nature of the routing protocol considered, i.e. flooding. The optimal multi-path (flooding) rate/capacity can be computed *without needing to explicit build end-to-end routes*. Flooding implies total use of all network edges, and so the flooding rate regions can be bounded by solving the max-flow min-cut problem with respect to each of the relevant partitions in Eqs. (7.7)-(7.9). The solutions to these cut-set minimisations are global properties of the network, and independent from explicit routing mechanisms. For this reason, they can be computed in a time-independent manner.

Unfortunately, the same approach does not work with respect to realistic, and practical routing protocols. The need to execute many routing protocols between many end-user pairs *potentially simultaneously* means that time is finally a factor. Routing algorithms must execute in a finite period of time, and thus the route building process between different communicators on a network will interact, compete and even collaborate. In order to benchmark and inspect *achievable* multiple-unicast rates using practical routing protocols (such as single-path or realistic multi-path methods) then the luxury of time-independence no longer holds. In the following sections, we introduce an approach which explicitly accounts for practical routing strategies and competitive behaviour in quantum communication networks.

7.3 Traffic Management Protocols and Network Filters

In this section, we focus on the task of multiple-unicast quantum networking as it provides a simpler paradigm in which to introduce the concepts of TM protocols and network filters. The theory developed in this section for multiple-unicast will be directly relevant to

single and multiple multicasting, which will be subsequently discussed.

7.3.1 Managing Multiple-Unicast Traffic

Consider a quantum network $\mathcal{N} = (P, E)$ and a collection of EUPs, $\mathcal{I} := \{i_j\}_{j=1}^{N_u}$. When $|\mathcal{I}| = 1$ this corresponds to single-unicast, and multiple-unicast when $|\mathcal{I}| > 1$. Each EUP employs a preferred end-to-end routing protocol \mathcal{P}_i which describes how they would perform end-to-end routing in the absence of routing competition, i.e. each \mathcal{P}_i is independent from one another and does not take other users into account during routing.

Instead of further complicating the routing strategies directly, we may consider multiple-unicast quantum communications through a higher-level of abstraction; traffic management (TM) protocols. EUPs should be able to offload any responsibility regarding routing competition and multi-user collaboration to a more abstract layer of network protocols. A TM protocol should be used to *organise* end-to-end routing sessions between many EUPs by controlling two degrees of freedom:

- *Space*: The notion of routing competition means that end-to-end protocols \mathcal{P}_i may collide and wish to utilise the same network node/edge. A TM protocol can manipulate where physical nodes and edges are used throughout the network and by whom in accordance with some overarching goal and network constraints. The TM protocol must take into account the capabilities of each network node, e.g. how many quantum registers are available at each node, it may also take the locality of EUPs into account.
- *Time*: The initiation of end-to-end routing searches (routing time) and the physical execution of transmission over point-to-point links (execution time) can occur on unique timescales. A TM protocol is responsible for ensuring the smooth running of end-to-end communications, anticipating the availability of network edges through careful monitoring of time-dependent processes. In this way, EUPs remain responsible for only their own goals while the TM protocol oversees larger, multi-user challenges.

Crucially, end-to-end routing protocols \mathcal{P}_i are unaware of the availability of network edges but a TM protocol can be used to monitor and control the spatial and temporal degrees of freedom during multi-user routing.

In this work, we think of TM strategies as centralised protocols with full knowledge of the network topology and rate distribution, however it is not clear that distributed strategies cannot be constructed. In this way, it is possible to mathematically define a TM protocol using the concept of *network filters*. A network filter provides an EUP (more precisely, their routing protocol \mathcal{P}_i) with a specific “view” of the network topology and rate distribution

at any given time t in the operation cycle. Each filter can be defined as a set of edge-wise parameters conditioned on the TM protocol, such that

$$\mathcal{F}(\mathbf{i}, t | \mathcal{T}) := \{f_{\mathbf{x}\mathbf{y}}^i(t)\}_{(\mathbf{x}, \mathbf{y}) \in E}, \quad f_{\mathbf{x}\mathbf{y}} \in [0, 1], \quad (7.10)$$

such that each parameter $f_{\mathbf{x}\mathbf{y}}^i(t)$ describes how “visible” an edge $(\mathbf{x}, \mathbf{y}) \in E$ is to the routing protocol between the users in \mathbf{i} at time t . When $f_{\mathbf{x}\mathbf{y}}^i(t) = 1$, the edge (\mathbf{x}, \mathbf{y}) is fully available for use by the protocol \mathcal{P}_i ; when $f_{\mathbf{x}\mathbf{y}}^i(t) = 0$ it is completely inaccessible and unavailable for routing at time t ; when $0 < f_{\mathbf{x}\mathbf{y}}^i(t) < 1$ then only a fraction of the point-to-point rate is available along this edge, i.e. the edge has been split via multiplexing and thus its rate is limited for use by any EUP.

A complete TM protocol is thus defined as a collection of network filters over the set of EUPs, $\mathcal{T} \mapsto \{\mathcal{F}(\mathbf{i}, t | \mathcal{T})\}_{\mathbf{i} \in \mathcal{I}}$. To ensure that the maximum load of all network edges is respected, we demand the condition

$$\sum_{\mathbf{i} \in \mathcal{I}} f_{\mathbf{x}\mathbf{y}}^i(t) \leq 1, \quad \forall t \geq 0, \quad (\mathbf{x}, \mathbf{y}) \in E. \quad (7.11)$$

That is, the rate of a point-to-point channel cannot be over-promised to the set of EUPs. The TM protocol is continuously responsible for managing each network filter. Its most basic role is to ensure that the end-to-end routes built by a particular EUP are available at the instant of communication. Sophisticated TM protocols can be used to optimise single or multiple end-user objectives on the fly through modification of each network filter.

Let us consider a single EUP and their process of end-to-end routing. Suppose that they initiate a routing search at time t , and the routing search takes δt_R seconds. During this time interval, their network filter could be modified many times, leading to a continuous ensemble of quantum network filters, topologies and rate distributions that they have been exposed to throughout routing, $\{\mathcal{F}(\mathbf{i}, t' | \mathcal{T})\}_{t'=t}^{t+\delta t_R}$. This ensemble of filters dictates how the routing protocol constructs its end-to-end paths throughout the routing interval. However, the final time-state of each network filter dictates the single-edge rate distribution that is available when communication is finally executed. Hence, the final filter state plays a crucial role in computing the end-to-end rate, which we denote as

$$\tilde{\mathcal{F}}(\mathbf{i} | \mathcal{T}) := \mathcal{F}(\mathbf{i}, t + \delta t_R | \mathcal{T}), \quad \forall \mathbf{i} \in \mathcal{I}. \quad (7.12)$$

Each end-to-end rate in this multiple-unicast setting is explicitly dependent on the evolution of their network filter throughout routing (managed by the TM protocol). Each routing protocol \mathcal{P}_i will identify a set of nodes and edges which can be reliably employed at the time of execution which are necessarily a subset of the filtered network they have

been exposed to during routing. As in the single-unicast case, end-to-end routing generates a forwarding probability distribution which is time-dependent throughout the course of routing itself. The forwarding distribution evolves according to the routing protocol \mathcal{P}_i and the evolution of its network filter $\mathcal{F}(i, t|\mathcal{T})$ during its search, and is thus conditioned upon the routing and TM protocols

$$\mathcal{Q}(i, t|\mathcal{P}, \mathcal{T}) := \{q_{xy}^i(t)\}_{(x,y) \in E}. \quad (7.13)$$

Analogously to the network filter, it will eventually reach a final state $\tilde{\mathcal{Q}}(i|\mathcal{P}, \mathcal{T}) := \mathcal{Q}(i, t + \delta t|\mathcal{P}, \mathcal{T})$ which directly infers the final-state subnetwork over which communication can now be executed. Hence, upon execution of communication each EUP will achieve the rate,

$$K(i, \mathcal{N}|\mathcal{P}, \mathcal{T}) := \min_C \sum_{(x,y) \in \tilde{C}} \tilde{f}_{xy}^i \tilde{q}_{xy}^i K_{xy}, \quad (7.14)$$

where $\tilde{q}_{xy}^i \in \tilde{\mathcal{Q}}(i|\mathcal{P}, \mathcal{T})$, $\tilde{f}_{xy}^i \in \tilde{\mathcal{F}}(i|\mathcal{T})$ are the final-state forwarding/filter parameters of each edge with respect to $i \in \mathcal{I}$. The quantity $K(i, \mathcal{N}|\mathcal{P}, \mathcal{T})$ thus represents an achievable end-to-end rate between the EUP i which has been accomplished via \mathcal{P} and \mathcal{T} . These end-to-end rates can then be used to explicitly construct multiple-unicast rate regions, and thus benchmark the multi-user performance of the network.

It is clear that the consideration of time-dependent routing/execution during multiple-unicast communication deeply complicates the assessment of end-to-end rates and performance. The sheer variability, especially when considering many sets of communicators, leads to an analytically intractable scenario. We could immediately turn to intensive numerical methods such as event-based network simulations, which would provide useful insight. However, our goal is to identify useful lower-bounds on end-to-end performance during multiple unicast. As such, we introduce a range of *naïve* TM protocols which can be used to compute performance benchmarks.

7.3.2 Achievable Benchmarks via Naïve TM Protocols

As discussed, the luxury of time-independence in single-unicast quantum networking is lost in the multiple-unicast domain. The evolutionary behaviour of the TM protocol through network filtering is by no means obvious. One could employ strategic techniques involving time management of edge-usage, locality based routing, etc. To unveil lower-benchmarks on multiple-unicast performance, we need to describe *naïve* strategies which are efficiently implementable and are therefore able to set a baseline for multi-user networking.

The first step we take to identify lower-benchmarks is to model all routing as *simultaneous*. That is, if we are considering N_u EUPs, we demand that they all begin their routing

search at time $t = 0$ and finish at time $t = \delta t_R$. In this way, routing competition is necessarily maximised. With each EUP attempting to utilise the network at the same time, the demand for any given edge will be at its maximum. Any alternative consideration of time, e.g. delayed batches of routing searches, alternating exploration of the network, would only serve to improve performance. More sophisticated considerations of time are only carried out to increase the size of each filtered network. As a result, simultaneous routing identifies a worst-case scenario for any TM protocol.

To complete our naïve lower-bounds, we consider *the absence of strategic traffic management* during simultaneous routing. More precisely, we permit ad-hoc, multi-user routing in the presence of maximum competition. As a result, no EUP is given priority over another, and the only condition for network filtering is to ensure EUPs do not overload the network, i.e. the capacities/rates and topology are respected and unchanged.

With these concepts in mind, it is possible to describe a naïve TM protocol during simultaneous multiple-unicast through the language of network filters. As we have shown, a network filter is an edgewise distribution of real parameters $f_{\mathbf{x}\mathbf{y}}^i \in [0, 1]$ used to restrict information flow between EUPs. The goal of a network filter is two-fold:

1. Preserve the rate/capacity of all point-to-point quantum channels across all EUPs.
2. Provide each EUP with a filtered network which is *sufficiently connected* so that end-to-end routes can be constructed between them.

While the first goal is straightforward, the second goal is less so. A TM protocol is useless if it provides a filtered network for each EUP over which end-to-end paths cannot be located. Nonetheless, naïve TM protocols can be designed through the principles of conservation of information flow and route-disjointness. To achieve this, we can construct network filters based on a *decision function* $D_{\mathcal{T}}$ which is a function of time t , EUP \mathbf{i} , and an arbitrary network edge $(\mathbf{x}, \mathbf{y}) \in E$ and decides what fraction of this channel should be visible to \mathbf{i} at a given time. That is,

$$\mathcal{F}(\mathbf{i}, t|\mathcal{T}) = \{D_{\mathcal{T}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t]\}_{(\mathbf{x}, \mathbf{y}) \in E}. \quad (7.15)$$

In the following sections we design a number of relevant network filter decision functions.

7.3.3 Routing-Disjoint Network Filters

One can design a naïve TM protocol by asking for *routing disjointness*, i.e. we demand that no two end-to-end routes ever share a network edge. We do not apply any strategy to prioritise particular EUPs and simply allow users to compete for effective routes. Let

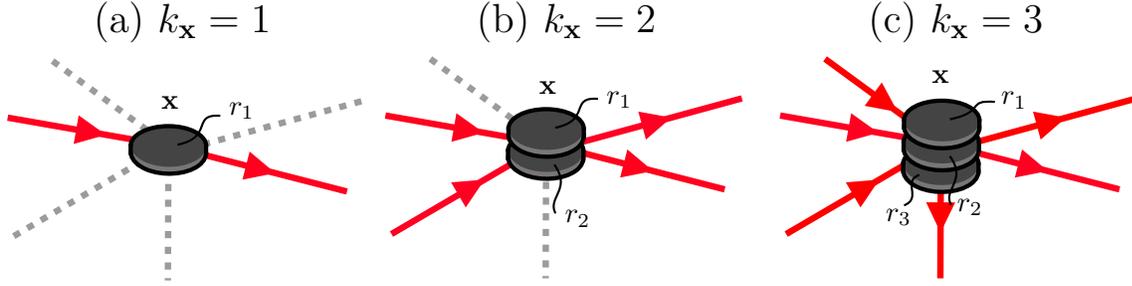


Figure 7.1: Quantum register disjoint (QRD) routing: A repeater node \mathbf{x} with finite number of $k_{\mathbf{x}}$ quantum registers and each register can support a single route (two edges). For example, panels (a)-(c) illustrate a node \mathbf{x} with degree $k_{\mathbf{x}} = 6$ and $k_{\mathbf{x}} \in \{1, 2, 3\}$ registers so that a total of $2k_{\mathbf{x}}$ edges can be supported throughout routing. QRD routing generalises edge/node disjoint routing, as Panel (a) recovers the node-disjoint case and (c) captures the edge-disjoint case.

us first define the following node and edge sets $P_{\mathcal{P}_i}(t)$ and $E_{\mathcal{P}_i}(t)$ respectively as the sets of nodes and edges that are contained within the route-set being built by the EUP \mathbf{i} , at a given time t .

Routing-disjointedness comes in a variety of flavours, depending on how strict we are about node/edge restrictions between routes. We may insist that all end-to-end routes are strictly edge-disjoint (ED), then

$$D_{\text{ED}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t] := \begin{cases} 0, & \exists \mathbf{i}' \in \mathcal{I} \text{ s.t. } (\mathbf{x}, \mathbf{y}) \in E_{\mathcal{P}_{\mathbf{i}'}}(t), \\ 1, & \text{otherwise.} \end{cases} \quad (7.16)$$

That is, an edge is filtered from the EUP \mathbf{i} if it is located in any other EUPs routing subnetwork. We may instead be stricter and ask that all end-to-end routes are node-disjoint (ND), such that

$$D_{\text{ND}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t] := \begin{cases} 0, & \exists \mathbf{i}' \in \mathcal{I} \text{ s.t. } (\mathbf{x} \in P_{\mathcal{P}_{\mathbf{i}'}}(t) \vee \mathbf{y} \in P_{\mathcal{P}_{\mathbf{i}'}}(t)) \\ D_{\text{ED}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t], & \text{otherwise,} \end{cases} \quad (7.17)$$

so that edges are filtered if either of the connected nodes are contained within another routing subnetwork.

More generally, these decision functions can be classed under the umbrella of quantum register disjointedness, such that the number of paths that a single node is able to support is bounded by the number of quantum registers it contains. Let's define the \mathbf{i} -neighbourhood of a node \mathbf{x} as the set of nodes to which it is connected in \mathbf{i} 's routing subnetwork,

$$N_{\mathbf{i}}(\mathbf{x}, t) = \{\mathbf{y} \mid (\mathbf{x}, \mathbf{y}) \in E_{\mathcal{P}_{\mathbf{i}}}(t)\}. \quad (7.18)$$

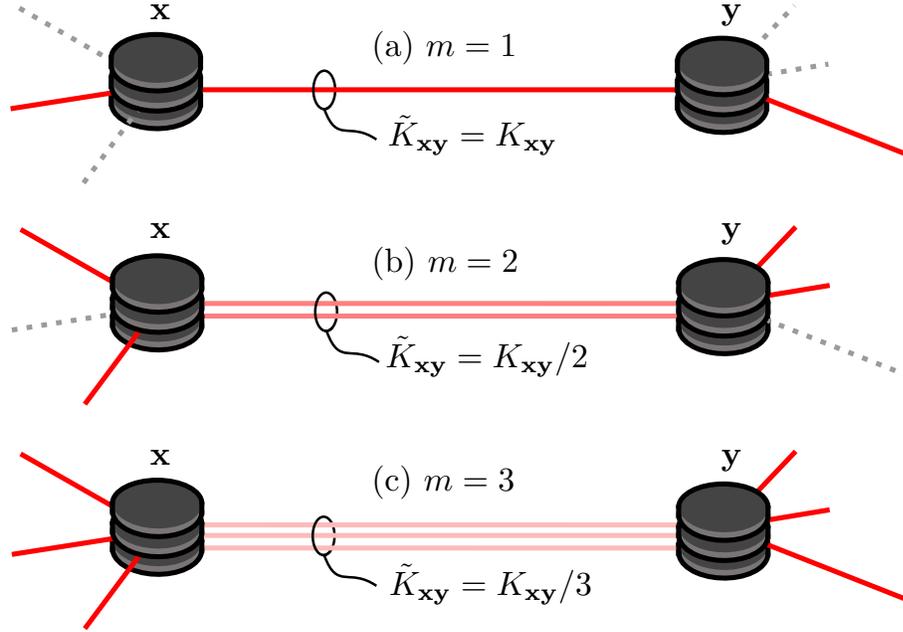


Figure 7.2: Edge-splitting via multiplexing: Multiple EUPs can exploit the same edge in a network granted that it can be reliably multiplexed a sufficient number of times, at the expense of splitting the point-to-point rate along each channel. Every edge $(\mathbf{x}, \mathbf{y}) \in E$ has a pre-defined multiplexability such that it can be split it into at most m channels each with reduced rates.

The cardinality of $N_i(\mathbf{x})$ indicates the number of edge that a node \mathbf{x} is used to support routing for a given EUP. We may use this notation to construct a $k_{\mathbf{x}}$ quantum-register disjoint ($k_{\mathbf{x}}$ -QRD) decision function which generalises the previous functions,

$$D_{k\text{QRD}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t] := \begin{cases} 0, & \left(\sum_{i' \in \mathcal{I} \setminus i} |N_{i'}(\mathbf{x}, t)| \leq 2k \right) \wedge \left(\sum_{i' \in \mathcal{I} \setminus i} |N_{i'}(\mathbf{y}, t)| \leq 2k \right), \\ D_{\text{ED}}[(\mathbf{x}, \mathbf{y}), \mathbf{i}, t], & \text{otherwise.} \end{cases} \quad (7.19)$$

When $k_{\mathbf{x}} = 1$ we regather then ND decision function, while edge-disjointness is recovered by letting $k_{\mathbf{x}}$ tend to infinity².

7.3.4 Multiplexing Network Filters

The previous network filters ensure that no edge is ever shared between end-to-end routes of different end-users. This is not always desirable, as sometimes it can be highly beneficial

²Edge-disjointness does not truly require $k_{\mathbf{x}} \rightarrow \infty$, this is just a concise condition. More precisely, edge-disjoint routing is permitted given that $k_{\mathbf{x}}$ is greater than or equal to half of the maximum degree found within the network $\mathcal{N} = (P, E)$ being considered, i.e. $k_{\mathbf{x}} \geq \lceil \frac{1}{2} \max_{\mathbf{x} \in P} \deg(\mathbf{x}) \rceil$. If this is true, it follows that any node $\mathbf{x} \in P$ can support $2k_{\mathbf{x}} \geq \max_{\mathbf{x} \in P} \deg(\mathbf{x})$ edges so that we recover edge-disjoint routing.

to share point-to-point channels through multiplexing. We need to define a function which, with respect to an EUP i counts how many times an edge is included within other EUPs routing edge-sets. Defining,

$$\mu_{\mathbf{x}\mathbf{y}}^i(t) := \sum_{i' \in \mathcal{I} \setminus i} |\{(\mathbf{x}, \mathbf{y})\} \cap E_{\mathcal{P}_{i'}}|, \quad (7.20)$$

we can use this to define a decision function for a multiplexing-based TM protocol. The following decision function filters any edge in the network by fractioning its rate according to how many edges wish to route through it, i.e. edges are multiplexed rather than being made disjoint. We consider a scenario in which any edge can be multiplexed at most m times. As such, the decision function takes the form

$$D_{m\text{ES}}[(\mathbf{x}, \mathbf{y}), i, t] := \begin{cases} 0, & \mu_{\mathbf{x}\mathbf{y}}^i(t) \geq m, \\ 1/(1 + \mu_{\mathbf{x}\mathbf{y}}^i(t)), & \text{otherwise.} \end{cases} \quad (7.21)$$

That is, edges which have been multiplexed less than m times are visible to EUPs, and but at a fraction of their original, maximum rate. Edges which have been multiplexed more than or equal to m -times are filtered from all other EUPs so to not overload the multiplexability of the channel. In general, the multiplexability of network edges may vary from one to another, but in this work we consider a constant maximum value. If $m = 1$ we recover the routing-disjoint format, since each edge can only ever be used once. When $m \rightarrow \infty$ then edges can be multiplexed as much as necessary. The tradeoff between multiplexing and maintaining disjoint routes is fascinating and may be an adaptive feature of advanced TM protocols.

7.3.5 Evaluation of Achievable Benchmarks

In this section we have developed tools that allow for the performance evaluation of realistic multiple-unicast quantum networking. Clearly, this evaluation is numerical as it depends upon the explicit deployment of end-to-end routing strategies within a numerical framework. Nonetheless, the study of achievable benchmarks via network filters, naïve TM strategies and simultaneous deployment introduces a *significant* simplification of multi-user quantum networking. A typical approach when abandoning the luxury of time-independence (found only within single-unicast) is to move to a completely numerical, event-based simulation framework in which use of the network is truthfully simulated with respect to every communication process, over a period of time.

Fortunately, our approach strikes a middle ground where we rely upon numerical network evaluations, but in such a way that is wildly more efficient than event-based techniques.

Furthermore, these methods can be utilised with respect to any end-to-end routing protocol that possesses a route-building phase during which network filters can be applied (which is effectively every realistic routing protocol). The MDPAIlg is one such end-to-end routing algorithm that offers a perfect platform for the study of multi-user limitations.

7.3.6 End-to-End Routing Protocols during Multiple-Unicast Sessions

The utility of different end-to-end routing protocols may dramatically differ with respect to single or multiple-unicast sessions. The lack of routing competition during single-unicast means that routing protocols which consume large amounts of network resources can promise high rates. While high resource consumption is not desirable, it is not punished. However, this is not true during multiple-unicast, and greedy end-to-end routing strategies \mathcal{P}_i will be punished by poor rates; network nodes and edges are quickly consumed by each EUP in such a way that inhibits their ability to completely construct their desired route sets. Hence, given the available network resources (captured by the nodal density and connectivity properties of the network) and the number of EUPs, there exists a careful balance between maximising end-to-end rates and minimising routing consumption.

For this purpose, the suite of MDPAIlg-defined routing protocols offer effective approaches to multiple-unicast routing. These protocols can be adjusted to locate a fixed number of routes or satisfy some rate requirement. Crucially, utilising an MDPAIlg which optimises the inverse rate-sum cost function in Eq. (6.14) allows EUPs to control their tradeoff between rate and efficiency. Exploring the effectiveness of such flexible routing strategies with naïve TM protocols permits the investigation of realistic multiple-unicast benchmarks and provides insight to the practical capabilities of multi-path routing in busy quantum networks.

7.4 Numerical Results

7.4.1 Average Multiple-Unicast Performance

To make use of the tools developed in the previous section we can provide informative benchmarks for multiple-unicast communications using large-scale, random quantum networks. In this setting, and inspired by Chapter 6, it is useful to inspect the *average end-to-end rate* achieved between end-user pairs within a given architecture. This provides a succinct but powerful metric for characterising the ability for a quantum network to multiple-unicast communications, without becoming overcomplicated by specific rate re-

gions. In the context of multiple-unicast, the average rate equates to a *average simultaneous* end-to-end rate between communicators.

Consider a quantum network \mathcal{N} within which a batch of N_u EUPs $\mathcal{I} = \{\mathbf{i}_j\}_{j=1}^{N_u}$ are simultaneously multiple-unicasting via a routing protocol $\mathcal{P}_i = \mathcal{P}$, for all $\mathbf{i} \in \mathcal{I}$ and a TM protocol \mathcal{T} . Throughout routing, each EUP experiences a different filtered version of \mathcal{N} which evolves over the course of simultaneous routing. To understand the average end-to-end performance of N_u simultaneous EUPs, we must also consider all the possible N_u batches that can be constructed. Let us define the complete collection of EUPs as the multiset,

$$\mathfrak{I}_{N_u} := \bigcup_{\mathcal{I}:|\mathcal{I}|=N_u} \{\mathcal{I}\}, \quad (7.22)$$

which iterates over all possible N_u length batches of unique user pairs. For example, in an $N = 6$ node network in which we are interested the performance of $N_u = 3$ EUPs, one can write $\mathfrak{I}_{N_u} = \left\{ \{ \{1, 2\}, \{3, 4\}, \{5, 6\} \}, \{ \{1, 3\}, \{2, 5\}, \{4, 6\} \}, \dots \right\}$, which already has many possible collections of simultaneous EUPs. The exact average simultaneous N_u -unicast rate on \mathcal{N} is given by

$$\langle K \rangle_{\mathcal{N}|N_u, \mathcal{P}, \mathcal{T}} := \frac{1}{|\mathfrak{I}_{N_u}|} \sum_{\mathcal{I} \in \mathfrak{I}_{N_u}} \frac{1}{N_u} \sum_{\mathbf{i} \in \mathcal{I}} K(\mathbf{i}, \mathcal{N}|\mathcal{P}, \mathcal{T}). \quad (7.23)$$

Clearly, the multiset of end-user batches \mathfrak{I}_{N_u} is enormous and cannot be considered explicitly. Instead, we may approximate this average quantity by sampling a reasonable number of batches B ,

$$\langle K \rangle_{\mathcal{N}|N_u, \mathcal{P}, \mathcal{T}} \approx \frac{1}{BN_u} \sum_{k=1}^B \sum_{j=1}^{N_u} K(\mathbf{i}_j, \mathcal{N}|\mathcal{P}, \mathcal{T}). \quad (7.24)$$

Finally, we can gather an ensemble average simultaneous N_u -unicast rate by studying this quantity across a large number of networks from the network class. As before, this takes the exact form

$$\langle K \rangle_{\mathfrak{N}|N_u, \mathcal{P}, \mathcal{T}} := \frac{1}{|\mathfrak{N}|} \sum_{\mathcal{N} \in \mathfrak{N}} \langle K \rangle_{\mathcal{N}|N_u, \mathcal{P}, \mathcal{T}}, \quad (7.25)$$

which can be readily approximated by taking B' samples from the class,

$$\langle K \rangle_{\mathfrak{N}|N_u, \mathcal{P}, \mathcal{T}} \approx \frac{1}{B'} \sum_{k=1}^{B'} \langle K \rangle_{\mathcal{N}|N_u, \mathcal{P}, \mathcal{T}}. \quad (7.26)$$

7.4.2 Criticality and Performance Benchmarking

To corroborate our theoretical methods, we consider multiple-unicast quantum communications on the classes of Waxman quantum architectures \mathfrak{W} composed of quantum fibre

links modelled via bosonic thermal-loss channels $\mathcal{E}_{\eta_{xy}, \bar{n}_{xy}}$ and assume that they operate at their capacity upper bound $\mathcal{C}_{xy} \leq \mathcal{T}_{\eta, \bar{n}}^u$ where $\mathcal{T}_{\eta, \bar{n}}^u$ is defined in Eq. (2.52), and $\bar{n} = 1/500$ is a typical value for optical-fibre channels. We consider a range of network routing protocols by adapting and modifying the recently developed MDPAlg [111, 103] according to methods described in the previous section and Chapter 6. From the perspective of routing protocols, we exploit the inverse rate-sum ansatz as the MDPAlg cost function, which was shown to be a very useful model particularly when one requires control over routing consumption (a feature that pure rate-optimisation routing strategies do not possess). Hence, for each EUP in the multiple-unicast scenario, we assume that they make use of an MDPAlg routing protocol $\mathcal{P}_{i, \text{mdp}}^M$ which identifies a fixed number of M routes; or $\mathcal{P}_{i, \text{mdp}}^{K^*}$ which builds as many routes as necessary to satisfy some end-to-end rate requirement K^* (if possible).

These routing strategies are then complemented by naïve TM protocols in order to compute achievable benchmarks for multiple-unicast networking. More precisely, we employ the routing disjoint network filters outlined in Section 7.3.2, considering edge-disjoint, node-disjoint, and QR disjoint network filters. In this way, we are able to investigate the balancing act involved with respect to multi-user demand and available routing resources in the network.

As discussed in Chapter 6, quantum Waxman networks undergo a percolation phase transition at a critical connectivity density $\rho_G^* \gtrsim 1.6 \times 10^{-4}$ nodes/km² (for bosonic thermal loss networks) above which end-to-end routes can be found between all network nodes. However, this nodal density is insufficient to guarantee that an EUP can reliably achieve super-critical end-to-end rates such that the ensemble average end-to-end capacity $\langle K \rangle_{W|\mathcal{P}} \geq 1$ bits/ \mathcal{P} -use, regardless of the routing strategy. One can more appropriately identify performance-defined critical densities as in Definition 6.1, i.e. a nodal density which marks the transition from unreliable end-to-end routing into a reliable regime (ensemble average end-to-end rate of at least 1 bits/ \mathcal{P} -use). Previous studies have only assessed single-unicast critical densities for Waxman networks. A fascinating question is to ask how the critical density behaves with respect to the number of simultaneous EUPs. We define the multi-user critical density:

Definition 7.1 (Multi-User Critical Density): *For a class of quantum networks \mathbf{N} , network routing protocol \mathcal{P} and TM protocol \mathcal{T} , we define $\rho_{\mathbf{N}|N_u, \mathcal{P}, \mathcal{T}}^*$ as the minimum nodal density required to guarantee an ensemble average simultaneous N_u -unicast rate of 1 bits/protocol use.*

The multi-user critical density is necessarily larger than the single-user critical density, which stands as a fundamental lower-bound on the critical network density. Understanding

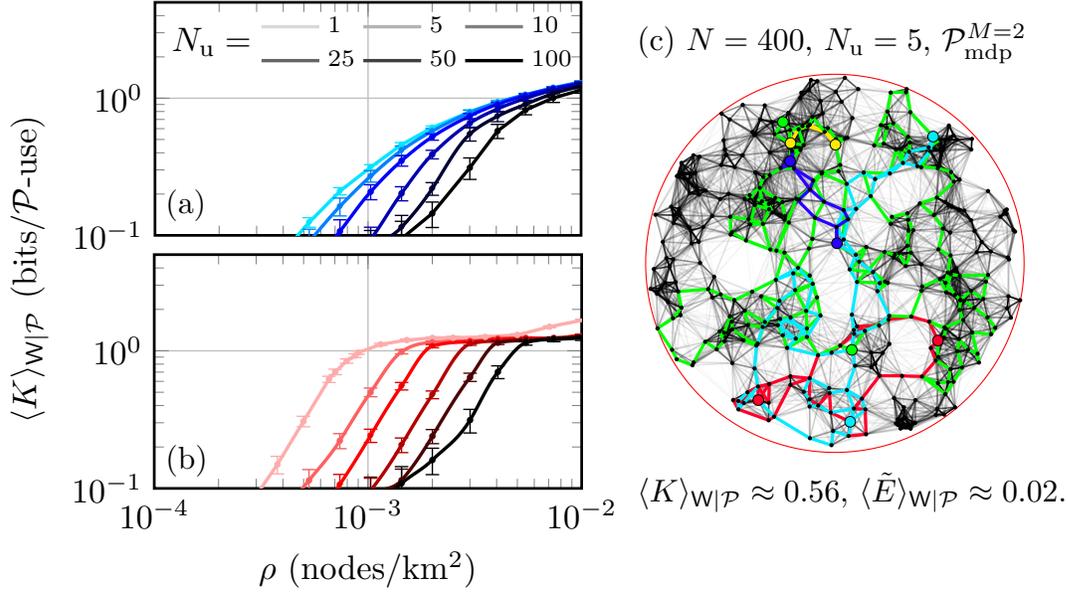


Figure 7.3: Behaviour of the ensemble average end-to-end capacity for a bosonic thermal-loss quantum Waxman networks under simultaneous multiple-unicast communications. Panel (a) displays results for the $\mathcal{P}_{\text{mdp}}^{M=2}$ multi-path routing strategy, while (b) considers the rate-requirement strategy, $\mathcal{P}_{\text{mdp}}^{K^*=1}$, both of which are mediated by an edge-disjoint TM protocol \mathcal{T}_{ED} . Panel (c) illustrates a collection of colour-coded end-to-end routes between $N_u = 5$ simultaneous EUPs on an example $N = 400$ node Waxman network.

the behaviour of the multi-user variant $\rho_{N|N_u, \mathcal{P}, \mathcal{T}}^*$ is incredibly valuable, as it is informative of the *scalability* of this class of networks. If the $\rho_{N|N_u, \mathcal{P}, \mathcal{T}}^*$ explodes with respect to increasing EUPs, clearly the architecture does not efficiently manage high levels of traffic. On the other hand, if an architecture is able to reliably minimise $\rho_{N|N_u, \mathcal{P}, \mathcal{T}}^*$ with respect to many users then it possesses features that are desirable within a large-scale quantum network.

7.4.3 Results and Analysis

Fig. 7.3 reports numerical analyses of multiple-unicast ensemble average end-to-end rates for quantum Waxman networks. Here, we deploy two relevant multi-path routing protocols using the MDPAlg; $\mathcal{P}_{\text{mdp}}^{M=2}$ and $\mathcal{P}_{\text{mdp}}^{K^*=1}$. In each multiple-unicast setting, we consider all EUPs to utilise the same routing protocol while the network is centrally mediated by a naïve edge-disjoint TM protocol, \mathcal{T}_{ED} .

Under the edge-disjoint TM protocol we observe that both routing strategies are able to achieve performance criticality, even under the pressure of many EUPs. Both protocols scale differently with respect to increasing nodal density, an expected feature which emerges from

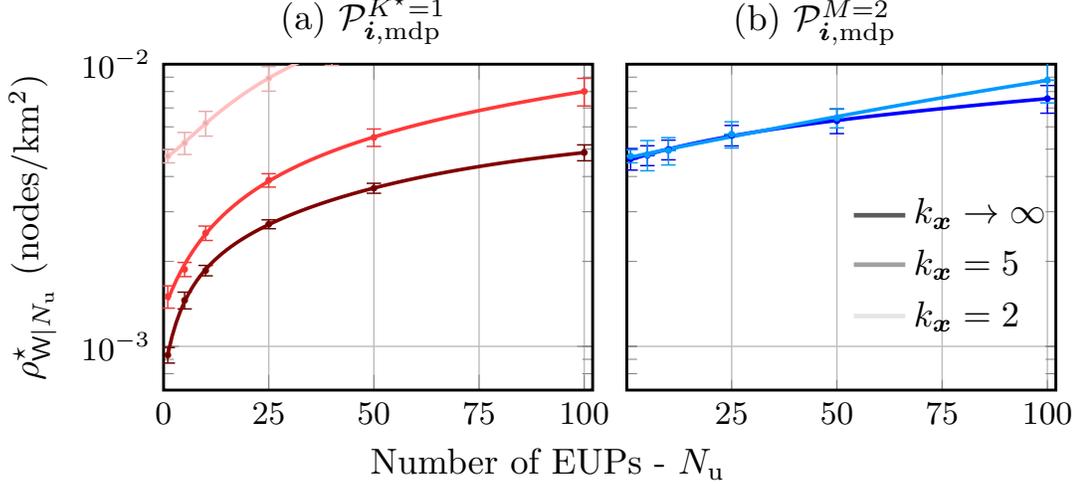


Figure 7.4: Multi-user critical densities for $R = 250$ km quantum Waxman networks for routing protocols (a) $\mathcal{P}_{\text{mdp}}^{K^*=1}$ and (b) $\mathcal{P}_{\text{mdp}}^{M=2}$, under different TM protocols ranging from edge-disjoint ($k_x \rightarrow \infty$) and QRD ($k_x < \infty$).

the tradeoff between routing consumption and routing competition. Indeed, while $\mathcal{P}_{\text{mdp}}^{M=2}$ rigidly commits to precisely two end-to-end routes per EUP, $\mathcal{P}_{\text{mdp}}^{K^*=1}$ is a more flexible protocol so that each EUP will use as much of the network as necessary in order to meet the rate requirement threshold. For this reason, the rate requirement protocol will (in all cases displayed) more rapidly reach performance criticality. Nonetheless, it is clear that for increasing numbers of N_u EUPs, the difference between these protocols diminish. As the number of EUPs increases, there exist less resources with which $\mathcal{P}_{\text{mdp}}^{K^*=1}$ can amplify its end-to-end rates. Consequently, only when the nodal density is sufficiently large will the rate requirement protocol provide an advantage over the fixed route strategy.

We collect a clearer picture of multi-user criticality within Fig. 7.4, where the multi-user critical density is plotted with respect to N_u for both routing protocols, and for a number of routing disjoint TM protocols, \mathcal{T}_{k_x} . Here, k_x is used to define the k_x -QR disjointness, such that each node possesses k_x independent quantum registers and can thus support precisely k_x routes ($2k_x$ edges). When $k_x \rightarrow \infty$ we recover edge-disjointness while for $k_x \rightarrow 1$ we have node-disjointness. All other values of k_x define the QR disjoint regime. When considering an edge-disjoint TM protocol, we are able to fully exploit the connectivity of the quantum Waxman network model. Nodes may have very large degrees, and therefore can help to connect many EUPs and support a vast collection of routes. In this way, edge-disjointness is a valuable asset for a quantum network which is emphasised in Fig. 7.4(a) and (b). We see that using both routing protocols, the multi-user critical

density can be effectively suppressed and kept within an order of magnitude of the single-user critical density. This is particularly true for the routing protocol $\mathcal{P}_{\text{mdp}}^{K^*=1}$ which is effective at maintaining $\rho_{N|N_u, \mathcal{P}, \mathcal{T}}^* \lesssim 5 \times 10^{-3}$ for as large as $N_u = 100$ user pairs. An analogous suppression is shown for $\mathcal{P}_{\text{mdp}}^{M=2}$, yet to a lesser extent due to the fewer resources permitted per EUP.

Unfortunately, as one considers stricter values of k_x the ability to suppress the critical density weakens. Values of $k_x \in \{2, 5\}$ mean that as the nodal density increases, nodes are unable to support as many routes as necessary to boost the end-to-end rate between communicators. Hence, larger nodal densities are required to overcome this limitation and effectively reach performance criticality. This scaling is certainly poor for $k_x = 2$, as we see neither protocol is capable of keeping the critical density below 1×10^{-2} for $N_u \gtrsim 30$ (for the $M = 2$ fixed route protocol, it cannot achieve this at all). However, raising the number of registers to $k_x = 5$ the multi-user critical density scaling is largely reclaimed. This is an important result. It emphasises the importance of considering node-based quantum resources and how they propagate into the properties of quantum networks. Furthermore, we show that even with relatively few nodal resources (quantum registers), practical multi-path routing strategies and a naïve TM protocol multiple-unicast networking can be reliably supported.

Crucially, the employment of (1) practical routing strategies, (2) achievable TM protocols and (3) random network models, allow us to piece together a more realistic image of the multi-user limitations of quantum networks. As we know, quantum Waxman networks offer a strong representation of future quantum networks thanks to their innate aversion to long connections, and willingness to forge strong connectivity which is needed for multi-path routing. Hence, their phase characterisations (as carried out in Chapter 6 for single-unicast) is informative.

In Fig. 7.5 we provide a phase characterisation of the quantum Waxman network model composed of a bosonic thermal-loss link layer and capacity upper bound achieving links. Here we provide a simplified description with respect to Fig. 6.7, but with the additional insight offered by our multi-user analysis. We show that using achievable multi-path routing strategies (via the MDPAlg) and naïve TM protocols, then it is possible to supersede criticality for 100 EUPs *before single-path routing can reach criticality for a single user*. While this pertinent insight has already been made in the previous chapter, our multi-user analyses truly accentuates this point: multi-path routing is incredibly important for efficient quantum networking, for which sophisticated routing and traffic management strategies will be invaluable. Furthermore, by advancing our theoretical toolbox we can help to better develop advanced methodologies for the future quantum internet.

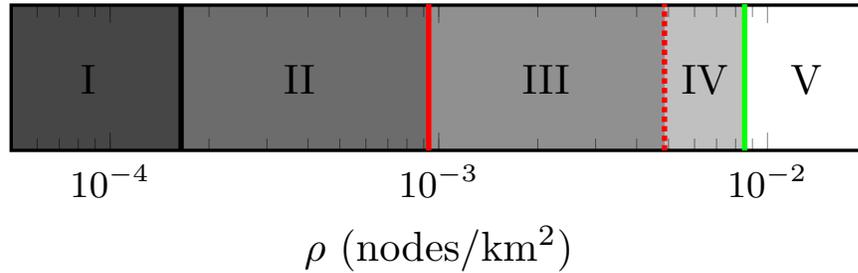


Figure 7.5: Waxman quantum network phases: (I) Pre-percolation transition, (II) post-percolation transition, pre-performance transition, (III) single-unicast reliable rates guaranteed by multi-path routing, (IV) multiple-unicast reliable rates up to $N_u = 100$ guaranteed by multi-path routing and edge disjoint TM, and (V) reliable single-unicast rates guaranteed by single-path routing.

7.5 Conclusion

In this chapter, we have introduced new methods for the performance analysis of multi-user quantum networking. We identified that the analysis of explicit routing protocols within the multi-user network setting introduces novel challenges due to the dependency of routing algorithms and communication execution on time. Multiple users communicating on the same network must compete for nodes and edges to take part in their session, and thus provides a time-dependency from which single-unicast modelling was spared. Through the concept of TM protocols and network filters, we established a useful formalism to describe routing in the presence of competitive users. By exploring naïve and simultaneous TM protocols we showed that explicit routing protocols can then be evaluated in a multiple-unicast setting and used to benchmark achievable end-to-end rates.

These new tools were then applied to the class of bosonic thermal-loss quantum Waxman networks, used to describe realistic optical-fibre networks. In this way, we have provided informative numerical analyses, considering multi-user critical nodal densities associated with realistic quantum Waxman networks. Of most importance, our results reiterate the importance of multi-path routing within quantum architectures and the effectiveness that it offers within highly competitive scenarios, in spite of simplistic TM techniques. To efficiently gather reliable rates through quantum networks, end-users must be able to deploy sophisticated routing strategies beyond single-path routing. Furthermore, network nodes must be built with this concept in mind; nodes *need* to be able to support large numbers of routes for multi-path routing to be effective or else their benefits will be diminished.

The methods and techniques developed in this chapter serve as a platform from which

many different forms of multi-user quantum communications can be investigated. Indeed, we are by no means limited to multiple-unicast scenarios, as one can already extend the network filter concepts of Section 7.3 to single and multiple-multicasting (overviewed in Appendix D). Furthermore, the tools of network filters can be advanced and optimised in a variety of different ways: considering sophisticated TM techniques which can be optimised on the fly, or improved end-to-end multi-path routing protocols beyond the approximate form used here via the inverse rate-sum ansatz.

Chapter 8

Conclusions

In this final chapter, we summarise the work presented within this thesis and discuss future directions for investigation.

8.1 Summary of Presented Work

Chapter 1 and Chapter 2 provided important context and useful reviews for the reader, outlining the overarching theme of research regarding the optimisation and benchmarking of quantum communication networks.

In Chapter 3, we introduced an analytically treatable class of network architectures which can be used as a building block for investigating ideally-connected, high-rate quantum communication networks. Through the use of WR quantum networks we then derived analytical bounds on the performance limits of large-scale quantum networks, connecting the physical limitations of channel length, nodal density with end-to-end network capacities. These developments have provided a useful platform for large-scale quantum network benchmarking without the need for heavy-duty numerics. Our findings emphasised the need for multi-path routing and proved the superiority of satellite-based quantum repeaters for global quantum communications. We revealed constraints associated with fibre-based networks and the enormous resource demands required to overcome achievable rates offered by a single satellite.

In Chapter 4, we expanded our ability to analytically benchmark quantum networks through end-to-end capacity bounds. These network capacity bounds extended from the use of relevant single-edge bounds, e.g. using coherent and reverse coherent information (lower-bounds) and the REE (upper-bounds). Combining this information with the recently developed node-splitting technique and WRN architectures, we were able to analyse the impact of device/channel imperfections upon the ultimate limits of quantum networking.

Our results strongly emphasised the need for close and careful analyses of noisy and realistic quantum devices when investigating the resource requirements of quantum networks. Indeed, the inevitability of noise within quantum repeaters can have a dramatic impact on design considerations for a large-scale structure, as the cumulative cost of imperfections have a ripple effect that influence the global architecture.

Chapter 5 took first steps at investigating large-scale, hybrid quantum network architectures through detailed considerations of free-space quantum channels on the ground and in space. We considered a modular network architecture which connects communities of different network types through a backbone structure, allowing for analytical derivations of critical network properties for end-to-end performance guarantees. Crucially, we illuminated the feasibility of high-rate global quantum communications mediated by a satellite-based backbone network, while theoretically demonstrating the efficacy of free-space quantum communication in hybrid wired/wireless metropolitan networks.

In an effort to better address realistic quantum networking, Chapter 6 moved into the realm of complex network architectures. We considered random networks, realistic link-layer descriptions and practical end-to-end routing protocols. Conducting thorough numerical analyses of Waxman and scale-free quantum networks, we assessed the criticality of these models with respect to connectivity, performance and routing efficiency. Furthermore, by adapting recent advancements in the field of multi-path routing algorithms we were able to develop a modified multiple-disjoint paths algorithm suitable for rate maximisation and the amplification of end-to-end rates. Our results reveal some vital takeaways for quantum network design: Reliable single-path routing places extreme resource demands on quantum networks; practical multi-path protocols exist and re-establish performance using realistic network resources; and adherence to scale-freedom in large-scale quantum networks results in poor end-to-end performance.

Finally, Chapter 7 developed new methods for the investigation of multi-user quantum networking. Via the concepts of routing competition, TM protocols and network filters, we presented a language for describing multi-user end-to-end rates/capacities. These concepts allowed us to adapt known algorithms in the context of single-unicast communications (e.g., MDPAlg) and extend them to relevancy in the multi-user scenario. Focussing on the pivotal task of multiple-unicast quantum communications, we investigated the limitations of multi-user quantum networking on bosonic thermal-loss Waxman networks. Following from the previous chapter, our results further emphasised the importance of multi-path routing for sustaining high rates within quantum architectures. Efficient multi-path routing strategies do exist, and can be used to significantly reduce resource requirements even in the presence of many users, and should be a pertinent point of interest for the coming quantum internet.

8.2 Outlook and Future Research

As discussed in the very beginning of this thesis, there remains a swathe of open questions regarding the implementation of optimisation of quantum communication networks. The fundamental rules and limits imposed by nature upon high-rate quantum communications reverberate throughout every layer of quantum networks. As a result, there are many questions to ask and lessons to learn regarding design, cost, operation and optimisation. Nonetheless, the work presented in this thesis has aimed to provide new tools to address these open questions, and build platforms from which we can investigate them in the future.

The development and theoretical deployment of analytically treatable architectures in Chapters 3-5 allowed for the efficient assessment of end-to-end network rates/capacities by exploiting the concept of network cut growth and flooding protocols. It would be fascinating to extend this tool to multi-user communication scenarios, and to investigate if it can be generalised to a broader range of topologies. Furthermore, comparing the expected performance/properties of WRNs to more realistic (perhaps random) networks which also exhibit cut growth, could be worthwhile. A comparison of this kind could illuminate stricter guidelines for the alignment of analytical expectations with realistic features of less ideal architectures.

Following from the investigation of practical routing, criticality and random networks in Chapters 6 and 7, there are a surplus of essential future research questions. The demonstration of efficient multi-path routing for quantum networks with the modified MDPAlg is highly encouraging, however it is unclear as to how one can most effectively optimise it for rate maximisation and resource minimisation. The current cost analyses in this thesis (while effective) likely leave lots of room for further optimisation. Similar arguments exist with respect to multi-user network benchmarking. Having considered naïve TM strategies, there is much to explore regarding the use of advanced network filters (perhaps boosted by machine learning) and the exploitation of time as a degree of freedom in multi-user routing protocols. Finally, the consideration of these practical analyses with respect to free-space quantum networking is also a crucial domain of future research.

Appendix A

Appendices for Chapter 3

A.1 Minimum Node Numbers for Internal Weak Regularity

In Fig. A.1 we provide visual proofs of the minimum number of nodes required to satisfy internal-WR for the structures utilised in Chapter 3. In each case the red regions of the network describes boundary region, while the white region resembles the internal network which is WR for the end-user nodes (which are coloured yellow). The blue dotted line is the network-bulk cut which collects the number of edges described in Lemma 3.2. The green cut is the smallest cut that exploits the boundary edges in order to reduce the cut-set size. The removal of any node on each network will give rise to a smaller cut than the blue cut by exploiting the boundary edges. We repeat the results of these network structures,

$$\begin{aligned}
 k = 3, \lambda^* &= \{0\}^{\cup 3} \rightarrow n_{\min} = 54, \\
 k = 6, \lambda^* &= \{2\}^{\cup 6} \rightarrow n_{\min} = 89, \\
 k = 8, \lambda^* &= \{2, 4\}^{\cup 4} \rightarrow n_{\min} = 120, \\
 k = 16, \lambda^* &= \{4, 8, 8, 8\}^{\cup 4} \rightarrow n_{\min} = 197.
 \end{aligned}
 \tag{A.1}$$

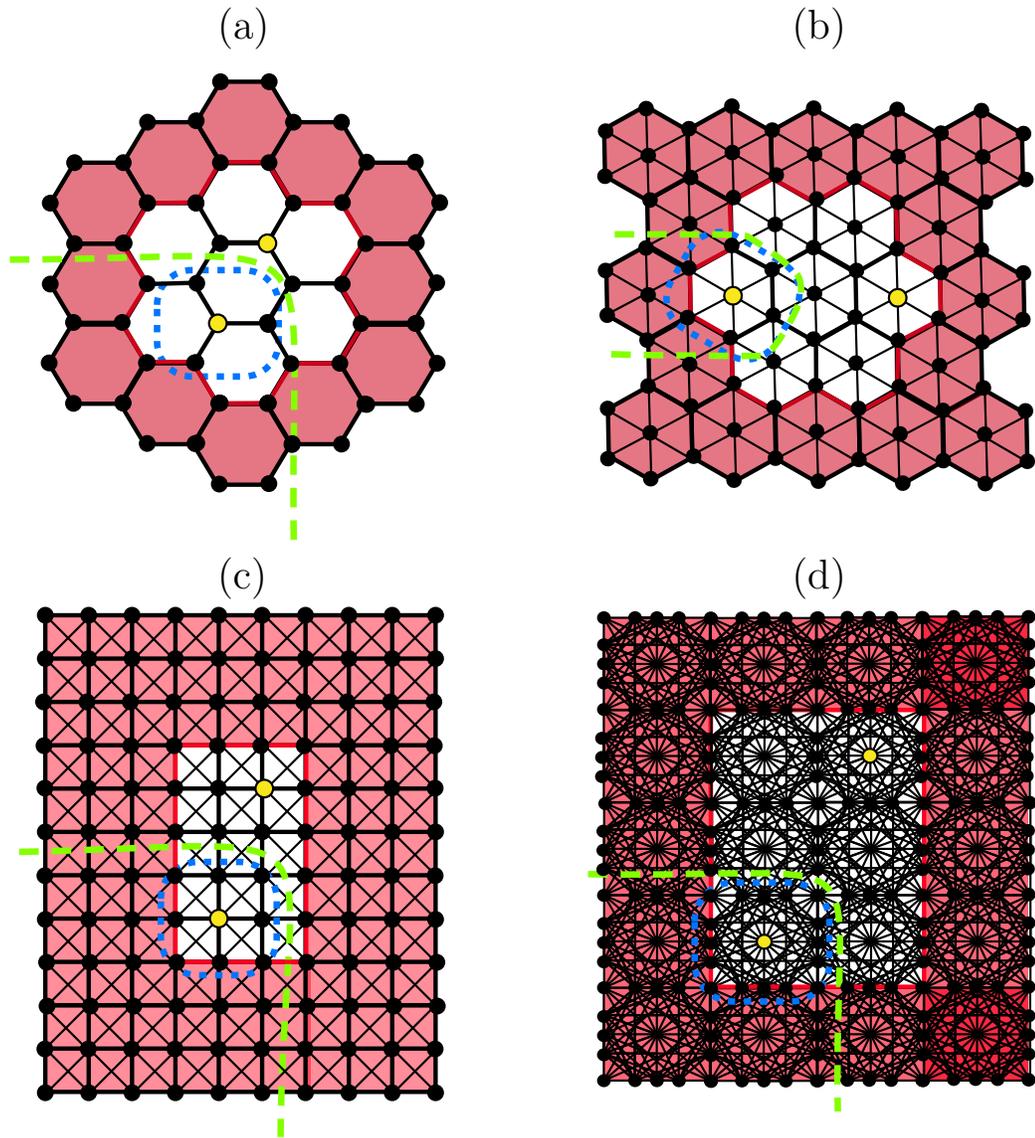


Figure A.1: Minimum node WRNs for a strict satisfaction of internal regularity for (a) honeycomb network, (b) hexagonal network, (c) Manhattan $k = 8$ network and (d) Manhattan $k = 16$ network. Each case resembles the smallest WRN for which there exist a pair of end-user nodes which do not share an edge or a neighbour, and possess minimum cardinality network-bulk cuts which are unaffected by the open boundary.

Appendix B

Appendices for Chapter 4

B.1 Network Parameter Threshold Theorems

Here we present a generalisation of Theorem 3.2 in terms of a generic, single-edge network parameter, ξ .

Corollary B.1 *Consider a $(k, \mathbf{\Lambda})$ -WR quantum network $\mathcal{N} = (P, E)$, an end-user pair $\mathbf{i} = \{\mathbf{a}, \mathbf{b}\}$ and a desired min-neighbourhood capacity $\mathcal{C}_{\mathcal{N}_i}$. Consider a single-edge channel property $\xi_{\mathbf{xy}}$ for which the point-to-point capacity $\mathcal{C}(\xi_{\mathbf{xy}})$ is monotonic. Then, if $\mathcal{C}_{\mathcal{N}_i}$ is attainable, there exist threshold parameters $\xi_{\mathcal{N}}^*$ and $\xi_{\mathcal{N}_i}^*$ which represent maximum or minimum tolerable values of $\xi_{\mathbf{xy}}$ for edge in the network-bulk or user-connected edge respectively:*

$$(\xi_{\mathcal{N}}^*, \xi_{\mathcal{N}_i}^*) := \begin{cases} (\xi_{\mathcal{N}}^{\max}, \xi_{\mathcal{N}_i}^{\max}) & \mathcal{C}(\xi_{\mathbf{xy}}) \text{ is decreasing,} \\ (\xi_{\mathcal{N}}^{\min}, \xi_{\mathcal{N}_i}^{\min}) & \mathcal{C}(\xi_{\mathbf{xy}}) \text{ is increasing.} \end{cases} \quad (\text{B.1})$$

If $\xi_{\mathcal{N}}^*$ and $\xi_{\mathcal{N}_i}^*$ are obeyed in their respective sub-networks, then the flooding capacity is guaranteed to satisfy

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{\mathcal{N}_i}. \quad (\text{B.2})$$

Proof. This corollary is simply a translation of Theorem 3.2 with respect to a single-edge network property $\xi_{\mathbf{xy}}$ for which the capacity function is monotonic. It follows an identical logic to the Corollary 4.1, except now we must identify two different threshold values by solving the equations,

$$\mathcal{C}(\xi_{\mathcal{N}}^*) = \mathcal{C}'_{\min} = \frac{\mathcal{C}_{\mathcal{N}_i}}{\delta}, \quad (\text{B.3})$$

$$\mathcal{C}(\xi_{\mathcal{N}_i}^*) = \mathcal{C}^i_{\min} = \frac{\mathcal{C}_{\mathcal{N}_i}}{\omega}. \quad (\text{B.4})$$

The maximum or minimum threshold behaviour and argument minimisation follow identically as before. ■

Corollary B.2 *Consider Corollary B.1 and a single-edge channel property $\xi_{\mathbf{x}\mathbf{y}}$ for which the point-to-point capacity $\mathcal{C}(\xi_{\mathbf{x}\mathbf{y}})$ is monotonically bounded by lower and upper bounding functions $F \in \{F_l, F_u\}$ respectively. Then the threshold parameter $\xi_{\mathcal{N}}^*$ satisfies*

$$\xi_{F_j}^*(\delta) \leq \xi_{\mathcal{N}}^* \leq \xi_{F_k}^*(\delta), \quad (\text{B.5})$$

and the user-connected edge threshold parameter satisfies

$$\xi_{F_j}^*(\omega) \leq \xi_{\mathcal{N}_i}^* \leq \xi_{F_k}^*(\omega), \quad (\text{B.6})$$

where $j \neq k \in \{l, u\}$ label the order of the bounds and are the same as in Eq. (4.20).

Proof. This result is identical to Corollary 4.2 with the additional bounding of the user-connected threshold quantity $\xi_{\mathcal{N}_i}^*$ in order to guarantee optimal performance, which can be achieved in the same way. ■

B.2 Compound Thermal-Loss Channels

Consider a compound channel of N thermal-loss channels $\mathcal{E}_{\tau_j, \bar{n}_j}$ each with transmissivity τ_j and output thermal photon number \bar{n}_j for $j \in \{1, \dots, N\}$. A compound channel of these N thermal-loss channels can be summarised into a single thermal-loss channel with total transmissivity and thermal noise parameters τ_{tot} and \bar{n}_{tot} . More precisely,

$$\mathcal{E}_{\tau_{\text{tot}}, \bar{n}_{\text{tot}}} = \mathcal{E}_{\tau_N, \bar{n}_N} \circ \mathcal{E}_{\tau_{N-1}, \bar{n}_{N-1}} \circ \dots \circ \mathcal{E}_{\tau_1, \bar{n}_1}. \quad (\text{B.7})$$

We wish to determine the relationship between τ_{tot} , \bar{n}_{tot} and the individual sub-channel properties.

We can do this by inspecting the action of the compound channel on a single-mode Gaussian state, ρ . The relationship proves to be independent of the initial state considered, therefore this is sufficient without loss of generality. Let this initial single-mode state have the covariance matrix (with zero first moments) V_0^{in} which is a real, symmetric, positive definite matrix. A thermal-loss channel $\mathcal{E}_{\tau_j, \bar{n}_j}$ equates to the mixing of the input state on a τ_j transmissive beam-splitter with an environmental thermal state of $\bar{n}_j/(1 - \tau_j)$ mean photons. Hence, we can monitor the transformations induced by the compound channel in Eq. (B.7) by recursively applying each transformation to the initial state V_0^{in} .

After the first channel $\mathcal{E}_{\tau_1, \bar{n}_1}$, the output state has the covariance matrix

$$V_1^{\text{out}} = \tau_1 V_0^{\text{in}} + \epsilon_1 I_2, \quad (\text{B.8})$$

where I_2 is the 2×2 identity matrix, and we have defined the noise quantity for the j^{th} sub-channel

$$\epsilon_j := \left(\frac{\bar{n}_j}{|1 - \tau_j|} + \frac{1}{2} \right) |1 - \tau_j| = \bar{n}_j + \frac{1}{2} |1 - \tau_j|. \quad (\text{B.9})$$

All subsequent sub-channels result in the following sequence of transformations,

$$V_2^{\text{out}} = \tau_2 V_1^{\text{out}} + \epsilon_2 I_2, \quad (\text{B.10})$$

\vdots

$$V_{N-1}^{\text{out}} = \tau_{N-1} V_{N-2}^{\text{out}} + \epsilon_{N-1} I_2, \quad (\text{B.11})$$

$$V_N^{\text{out}} = \tau_N V_{N-1}^{\text{out}} + \epsilon_N I_2. \quad (\text{B.12})$$

By inspection of these recursive formula, it is clear that the final output state V_N^{out} can be rewritten in the form

$$V_N^{\text{out}} = \left(\prod_{j=1}^N \tau_j \right) V_0^{\text{in}} + \xi_N \cdot I_2, \quad (\text{B.13})$$

where we define $\xi_N \in \mathbb{R}$ is a thermal additive noise coefficient at the N^{th} channel output state. This function is recursive, as it compounds the noise properties from all previous channels. For $j = 1$, it takes the initial value $\xi_1 = (\bar{n}_1 + \frac{1}{2}) |1 - \tau_1|$, while for all further values $1 < j \leq N$ it takes the form

$$\xi_j = \tau_j \xi_{j-1} + \bar{n}_j + \frac{1}{2} |1 - \tau_j|. \quad (\text{B.14})$$

This can be easily solved via recursive techniques, and the final additive noise term after N thermal-loss channels is given by

$$\xi_N = \epsilon_N + \sum_{j=1}^{N-1} \left(\epsilon_{N-j} \prod_{i=N-j+1}^N \tau_i \right). \quad (\text{B.15})$$

It is then possible to derive relationships between τ_{tot} , \bar{n}_{tot} and all the individual τ_j , \bar{n}_j values by equating Eq. (B.13) to the transformation induced by a single thermal-loss channel with total loss/noise properties,

$$V_N^{\text{out}} = \tau_{\text{tot}} V_0^{\text{in}} + \left(\bar{n}_{\text{tot}} + \frac{1}{2} |1 - \tau_{\text{tot}}| \right) I_2. \quad (\text{B.16})$$

Then these parameters can be quickly identified. As a result, we find that they admit the forms

$$\tau_{\text{tot}} := \prod_{j=1}^N \tau_j, \quad (\text{B.17})$$

$$\bar{n}_{\text{tot}} := \xi_N - \frac{1}{2}|1 - \tau_{\text{tot}}|. \quad (\text{B.18})$$

These results can then be specifically applied to the compound channels in Chapter 4. They are clearly independent of the initial state and are thus universal for compound thermal-loss channels.

B.3 Compound Amplitude Damping Channels

An N length compound channel of amplitude damping channels each with damping probability p_j can be simply rewritten as a single amplitude damping channel with a total damping probability

$$\mathcal{E}_{p_{\text{tot}}} = \mathcal{E}_{p_N} \circ \mathcal{E}_{p_{N-1}} \circ \cdots \circ \mathcal{E}_{p_1}. \quad (\text{B.19})$$

The relationship between p_{tot} and the damping probability of each sub-channel can be easily derived. Indeed, let us consider a generic single qubit state as an input into the first channel,

$$\rho_0^{\text{in}} = \begin{pmatrix} 1 - \gamma & c^* \\ c & \gamma \end{pmatrix}, \quad (\text{B.20})$$

where $\gamma, c \in \mathbb{C}$. A single amplitude damping channel invokes the transformation,

$$\rho_1^{\text{out}} = \sum_{i=0,1} K_i \rho_0^{\text{in}} K_i^\dagger, \quad (\text{B.21})$$

$$= \begin{pmatrix} 1 - (1 - p_1)\gamma & c^* \sqrt{1 - p_1} \\ c \sqrt{1 - p_1} & \gamma(1 - p_1) \end{pmatrix}, \quad (\text{B.22})$$

where K_i are the Kraus operators of the channel, such that $K_0 = |0\rangle\langle 1| + \sqrt{1 - p} |1\rangle\langle 1|$ and $K_1 = \sqrt{p} |0\rangle\langle 1|$. It is clear we can rewrite this output state in the form,

$$\rho_1^{\text{out}} = \begin{pmatrix} 1 - \gamma_1 & c_1^* \\ c_1 & \gamma_1 \end{pmatrix}, \quad (\text{B.23})$$

where $\gamma_1 = 1 - p_1\gamma$ and $c_1 = c\sqrt{1 - p_1}$. Hence, the subsequent action of $N - 1$ further amplitude damping channels will result in the output state

$$\rho_N^{\text{out}} = \begin{pmatrix} 1 - \gamma_N & c_N^* \\ c_N & \gamma_N \end{pmatrix}, \quad (\text{B.24})$$

where the parameters of the N^{th} output state are

$$\gamma_N = \gamma \prod_{j=1}^N (1 - p_j), \quad c_N = c \prod_{j=1}^N \sqrt{1 - p_j}. \quad (\text{B.25})$$

As a result, we can equate the action of the N compound channel in Eq. (B.19) to that of a single amplitude damping channel with total damping probability,

$$p_{\text{tot}} := 1 - \prod_{j=1}^N (1 - p_j). \quad (\text{B.26})$$

Appendix C

Appendices for Chapter 5

C.1 General Aspects of Quantum Networks with Community Structure

In the main-text, we considered a specific modular network structure, using the idea of disjoint communities connected to a backbone quantum network. Here, using basic notions from graph and network theory [48, 49, 50, 51], we aim to generalise the concept of modular quantum networks, outlining a framework from which the ideal architecture in Definitions 5.1 and 5.3 emerge. In doing so, we derive general constraints which guarantee specific end-to-end performance bounds for communication between local community users and remote community users.

C.1.1 General Structure

Let us first consider general networks which display community structure. Consider a completely general architecture $\mathcal{N} = (P, E)$ such that P is the collection of all nodes, and E the set of all undirected edges. As discussed in the main text, it is possible to divide P into sub-collections of communities,

$$P = \bigcup_i P_{c_i}, \quad P_{c_i} \subset P. \quad (\text{C.1})$$

In general, the community structure on a given network is not unique, and the sets of community nodes can overlap, i.e. the subsets of nodes P_{c_i} are not necessarily pairwise disjoint, i.e. $P_{c_i} \cap P_{c_j} \neq \emptyset$, for all i, j . However, as we are physically motivated by separate communities connected via a backbone, we restrict our attention to the case in which each node can be uniquely assigned to a single community,

$$P = \bigcup_i P_{c_i}, \quad \text{s.t. } P_{c_i} \cap P_{c_j} = \emptyset, \forall i, j. \quad (\text{C.2})$$

This assumption is appropriate for large-scale communication networks, and applies for spatially modular networks [102], e.g. each community represents a separate metropolitan area. We make no further assumptions on the topology of the underlying communities.

The community structure additionally partitions the edges into distinct sets. The i^{th} community c_i has its own set of intracommunity edges,

$$E_{c_i} = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x}, \mathbf{y} \in P_{c_i}\}, \quad (\text{C.3})$$

while any two communities c_i and c_j are connected by a set of intercommunity edges

$$E_{c_i:c_j} = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x} \in P_{c_i}, \mathbf{y} \in P_{c_j}\}. \quad (\text{C.4})$$

Hence, for a network comprised of n communities we may define two global classes of edges: intracommunity and intercommunity edges respectively,

$$E_c := \bigcup_{i=1}^n E_{c_i}, \quad E_{c:c'} := \bigcup_{i \neq j=1}^n E_{c_i:c_j}. \quad (\text{C.5})$$

Using these notions we may introduce two related networks which will simplify our analysis. A *community sub-network* $\mathcal{N}_{c_i} = (P_{c_i}, E_{c_i})$ is defined as the graph consisting of all the nodes in the community c_i connected by the intracommunity edges E_{c_i} .

C.1.2 Simplified Quotient Network

Let \mathcal{N} be a network with an n -community structure. Since we consider only non-overlapping communities, we may define an equivalence relation R on the nodes of the network in the following way.

$$\mathbf{x} \sim \mathbf{y} \text{ iff } \mathbf{x}, \mathbf{y} \in P_{c_i}, \forall i \in [1, n]. \quad (\text{C.6})$$

Two nodes are equivalent if they are contained within the same community. This equivalence relation is a means of partitioning the network into a simplified form, such that nodes contained within equivalent classes (communities) are pooled and redefined as a unified, collective node. Then, R permits us to define a *quotient network*,

$$\mathcal{N}_Q := \mathcal{N}/R = (P_Q, E_Q), \quad (\text{C.7})$$

where P_Q is a set of quotient nodes and E_Q is a set of quotient edges. The set of quotient nodes is given by

$$P_Q := P/R = \{\mathbf{c}_Q^i\}_{i=1}^n, \text{ where } \mathbf{c}_Q^i = P_{c_i}, \quad (\text{C.8})$$

where by the equivalence relation R we have reduced the set of community nodes P_{c_i} into a single quotient node \mathbf{c}_Q^i . Meanwhile, there exists a quotient edge between the two community nodes \mathbf{c}_Q^i and \mathbf{c}_Q^j if there exists at least one intercommunity edge between a node $\mathbf{x} \in P_{c_i}$ and a node $\mathbf{y} \in P_{c_j}$. Therefore the set of edges on the quotient network E_Q is given by,

$$E_Q := E/R = \left\{ \left(\mathbf{c}_Q^i, \mathbf{c}_Q^j \right) \mid \exists (\mathbf{x}, \mathbf{y}) \in E_{c_i:c_j} \right\}_{i \neq j=1}^n. \quad (\text{C.9})$$

It is important to note that there may be more than one intercommunity edge between two given communities. Yet, our definition of the quotient network is still a simple graph. To account for this, the single-edge capacity of an edge in the quotient graph is actually defined as a multi-edge capacity from the original network. More precisely, the single-edge capacity of each quotient edge is equal to the sum of the capacities of the intercommunity edges,

$$\{ \mathcal{C}_{\mathbf{x}\mathbf{y}} \}_{(\mathbf{x}, \mathbf{y}) \in E_Q} = \{ \mathcal{C}_{c_i:c_j} \}_{i \neq j=1}^n, \quad (\text{C.10})$$

where we define the multi-edge capacity between communities,

$$\mathcal{C}_{c_i:c_j} := \sum_{(\mathbf{x}, \mathbf{y}) \in E_{c_i:c_j}} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (\text{C.11})$$

This community structure is extremely useful for simplifying investigations of end-to-end capacities. With this established, we can differentiate between two key scenarios for the end-to-end capacity: end-to-end communication in the same community, or between distinct communities.

C.1.3 Intra-Community Capacities

Let us focus on a pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ which are located within the same community, $\alpha, \beta \in P_{c_i}$. While it may be intuitive to assume that the flooding capacity for communication between these nodes is determined by a min-cut performed exclusively on \mathcal{N}_{c_i} , this is not always the case. Indeed, it is possible that a minimum cut will collect edges not only within the community \mathcal{N}_{c_i} , but also intercommunity edges, and edges from other communities. In general, we can write the following lemma:

Lemma C.1 *Consider two end-user nodes $\mathbf{i} = \{\alpha, \beta\}$ which are located within the same community $\mathcal{N}_{c_i} = (P_{c_i}, E_{c_i})$, such that $\alpha, \beta \in P_{c_i}$. Let \mathcal{C}_{c_i} be the end-to-end flooding capacity computed exclusively on the sub-network \mathcal{N}_{c_i} . Then the intracommunity flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N})$ is bounded by*

$$\mathcal{C}_{c_i} \leq \mathcal{C}_{cut}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c_i} + \mathcal{C}_{E \setminus c_i}, \quad (\text{C.12})$$

where $\mathcal{C}_{E \setminus c_i}$ is an additional capacity contribution associated with non-local community edges.

Proof. Consider the two end-user nodes $\mathbf{i} = \{\boldsymbol{\alpha}, \boldsymbol{\beta}\}$. We may exclusively investigate the flooding capacity of this induced sub-network, \mathcal{N}_{c_i} by ignoring all intercommunity edges. In this way, we can identify a minimum cut restricted on the community by minimising over all the local-community cuts,

$$\mathcal{C}_{c_i} = \min_{C_{c_i}} \mathcal{C}(C_{c_i}) = \min_{C_{c_i}} \sum_{(\mathbf{x}, \mathbf{y}) \in \tilde{C}_{c_i}} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (\text{C.13})$$

Here a local-community cut C_{c_i} is a cut performed exclusively on the community network, and \tilde{C}_{c_i} is its cut-set.

Now consider the addition of intercommunity edges $\{E_{c_i:c_j}\}_{j \neq i}$ which provide access to other remote communities. It is possible that these intercommunity edges will compromise the validity of a community cut C_{c_i} , since there may exist an end-to-end route that traverses the global network. In this scenario, it is necessary to cut additional edges from the rest of the network in order to consolidate the cut. We collect these additional edges within the following set $\tilde{C}_{E \setminus c_i} \subset E \setminus E_{c_i}$. More precisely, given a valid network cut C , we can always separate its cut-set into community and non-community edges

$$\tilde{C} = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x} \in \mathbf{A}, \mathbf{y} \in \mathbf{B}\}, \quad (\text{C.14})$$

$$= \tilde{C}_{c_i} \cup \tilde{C}_{E \setminus c_i}. \quad (\text{C.15})$$

We can then say that the community cut-set \tilde{C}_{c_i} is generated via a community cut C_{c_i} , while $\tilde{C}_{E \setminus c_i}$ is generated via an additional non-community cut $C_{E \setminus c_i}$. The network flooding capacity is thus generally given by

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \min_C \mathcal{C}_{\text{cut}}(C), \quad (\text{C.16})$$

$$= \min_C \left[\mathcal{C}_{\text{cut}}(C_{c_i}) + \mathcal{C}_{\text{cut}}(C_{E \setminus c_i}) \right]. \quad (\text{C.17})$$

The cut C_{c_i} always forms a valid partition of the user pair when we are restricted to the sub-network \mathcal{N}_{c_i} . Meanwhile, on its own, $\tilde{C}_{E \setminus c_i}$ is never a valid network cut between local end-users. Crucially, the addition of the non-community edges into the cut-set can never decrease the total flooding capacity between users, only increase it. Therefore we can

separate the minimisation in Eq. (C.17) and write

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \min_C \left[\mathcal{C}_{\text{cut}}(C_{c_i}) + \mathcal{C}_{\text{cut}}(C_{E \setminus c_i}) \right], \quad (\text{C.18})$$

$$\leq \min_{C_{c_i}} \mathcal{C}_{\text{cut}}(C_{c_i}) + \min_{C_{E \setminus c_i}} \mathcal{C}_{\text{cut}}(C_{E \setminus c_i}), \quad (\text{C.19})$$

$$= \mathcal{C}_{c_i} + \mathcal{C}_{E \setminus c_i}, \quad (\text{C.20})$$

where $\mathcal{C}_{E \setminus c_i}$ denotes the multi-edge capacity of the minimised non-community cut that validates the end-user partition.

It is then clear that we can write the following bounds on the global network flooding capacity,

$$\mathcal{C}_{c_i} \leq \mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c_i} + \mathcal{C}_{E \setminus c_i}. \quad (\text{C.21})$$

Here, the lower bound refers to the situation when non-community cuts are not required ($\mathcal{C}_{E \setminus c_i} = 0$), and the upper bound refers to when they are ($\mathcal{C}_{E \setminus c_i} > 0$). ■

Hence, the intracommunity capacity is always lower-bounded by the local-community capacity of a local network. The saturation of either the upper or lower bounds is completely determined via the network structure.

C.1.4 Intercommunity Capacities

We now turn our attention to the case in which the two end-users lie in distinct communities. This is the setting focussed on in Chapter 5, and is of most interest for relatively long-distance communication within large-scale, hybrid networks. Indeed, the intercommunity capacity depends more strongly on the interplay between sub-network properties, rendering its characterisation more difficult than the intracommunity capacity. Nonetheless, through the community structure developed in this appendix and the simplifications offered by the quotient graph representation, it is possible to glean conditions for which the end-to-end intercommunity capacity is analytically obtainable.

To achieve this, we must develop a number of helpful lemmas. We shall first show that any cut which collects an intracommunity edge automatically invokes a valid cut between the two nodes connected by that edge on a community sub-network.

Lemma C.2 *Consider two end-user nodes contained in remote communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$, and a cut C between them with a corresponding cut-set \tilde{C} . If \tilde{C} contains at least one intracommunity edge $(\mathbf{x}, \mathbf{y}) \in E_{c_i}$ from an arbitrary community c_i , then \tilde{C} contains a subset \tilde{C}' which is a valid cut between \mathbf{x} and \mathbf{y} on the induced sub-network \mathcal{N}_{c_i} .*

Proof. Consider a cut C such that the corresponding cutset \tilde{C} contains an intracommunity edge $(\mathbf{x}, \mathbf{y}) \in E_{c_i}$. Without loss of generality, the cut partitions the network nodes into two sets $\mathbf{A} = \{\alpha, \mathbf{x}, \dots\}$ and $\mathbf{B} = \{\beta, \mathbf{y}, \dots\}$. We can therefore identify two subsets $\mathbf{A}' \subseteq \mathbf{A}$ and $\mathbf{B}' \subseteq \mathbf{B}$ that consist solely of nodes that lie in the same community,

$$\mathbf{A}' = \{\mathbf{x} \mid \mathbf{x} \in \mathbf{A} \cap P_{c_i}\}, \quad \mathbf{B}' = \{\mathbf{y} \mid \mathbf{y} \in \mathbf{B} \cap P_{c_i}\}. \quad (\text{C.22})$$

It can be seen that this forms a bipartition for the nodes in P_{c_i} and thus forms a valid cut between the arbitrary nodes \mathbf{x} and \mathbf{y} on the community network \mathcal{N}_{c_i} . The corresponding cut-set \tilde{C}' may be formed as usual from these sets,

$$\tilde{C}' = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x} \in \mathbf{A}', \mathbf{y} \in \mathbf{B}'\}, \quad (\text{C.23})$$

Comparing this to the original cut-set

$$\tilde{C} = \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{x} \in \mathbf{A}, \mathbf{y} \in \mathbf{B}\}, \quad (\text{C.24})$$

it can clearly be seen that $\tilde{C}' \subseteq \tilde{C}$ since $\mathbf{A}' \subseteq \mathbf{A}$ and $\mathbf{B}' \subseteq \mathbf{B}$. ■

This is actually a very useful result. It tells us that a hybrid cut between two remote user-nodes α and β which collects edges from a network community will necessarily invoke a local-community cut between some arbitrary pair of nodes. Let us now make the following definition which will simplify our notation.

Definition C.1 (Min-local Community Capacity): *Consider a community sub-network given by \mathcal{N}_{c_i} . We define the minimum local-community capacity as the smallest flooding capacity that can be generated between any two nodes on community network,*

$$C_{c_i}^* := \min_{\mathbf{x} \neq \mathbf{y} \in P_{c_i}} \mathcal{C}(\{\mathbf{x}, \mathbf{y}\}, \mathcal{N}_{c_i}). \quad (\text{C.25})$$

As a result of the previous lemmas, we can present the following result which can be used to relate the intracommunity capacity with the minimum cut on the quotient network.

Lemma C.3 *Consider a quantum network \mathcal{N} with a disjoint community structure, and a pair of remote end-users $\mathbf{i} = \{\alpha, \beta\}$ which are located in distinct communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$. On the quotient graph \mathcal{N}_Q , we can equivalently consider the end-user-community pair $\mathbf{i}_Q = \{\mathbf{c}_Q^\alpha, \mathbf{c}_Q^\beta\}$. It follows that if all of the minimum local-community capacities are greater than the flooding capacity on the quotient network,*

$$\min_{c_i} C_{c_i}^* \geq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q), \quad (\text{C.26})$$

then the end-to-end flooding capacity between α and β is equal to

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q). \quad (\text{C.27})$$

Otherwise, the flooding capacity on the quotient network is an upper-bound on the true flooding capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q)$.

Proof. Since α and β lie in two different communities it is always possible to form cuts with cut-sets only containing intercommunity edges. These are exactly the same cuts as are possible on the quotient graph $\mathcal{N}_Q = \mathcal{N}/R$ where R is the equivalence relation partitioning the nodes into their communities. Hence, we can call these cuts *quotient cuts*, C_Q . Therefore we can obtain an initial bound for the multi-path capacity.

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q) = \min_C \mathcal{C}_{\text{cut}}(\mathbf{i}_Q, C_Q), \quad (\text{C.28})$$

where $\mathcal{C}_{\text{cut}}(\mathbf{i}_Q, C_Q)$ is the multi-edge capacity associated with a quotient cut partitioning the two communities in \mathbf{i}_Q . This is an upper bound since the cut taken on the quotient graph may not be a minimum cut.

Now consider an arbitrary cut C_0 between α and β , containing at least one intracommunity edge. From Lemma C.2 we have that the intracommunity edges form at least one valid cut between arbitrary nodes on an induced sub-network \mathcal{N}_{c_i} . Note that the corresponding cut-set will generally not correspond to a valid cut between α and β on \mathcal{N} . It is clear we can lower bound the capacity across C_0 by the minimum flooding capacity between any two nodes on P_{c_i} , that is

$$\mathcal{C}_{\text{cut}}(\mathbf{i}, C_0) \geq \mathcal{C}_{c_i}^*. \quad (\text{C.29})$$

Comparing this to the initial bound obtained on the quotient graph, we see that whenever the minimum flooding capacity between any two nodes on the community \mathcal{N}_{c_i} satisfies

$$\mathcal{C}_{c_i}^* \geq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q), \quad (\text{C.30})$$

then C_0 cannot be a minimum cut. Now since the intracommunity edge that C_0 collects is arbitrary, the left hand side must be minimised over all communities to ensure that no hybrid cut can ever be a minimum cut. Therefore, whenever

$$\min_{c_i} \mathcal{C}_{c_i}^* \geq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q), \quad (\text{C.31})$$

the minimum-cut must contain only intercommunity edges and $\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q)$. ■

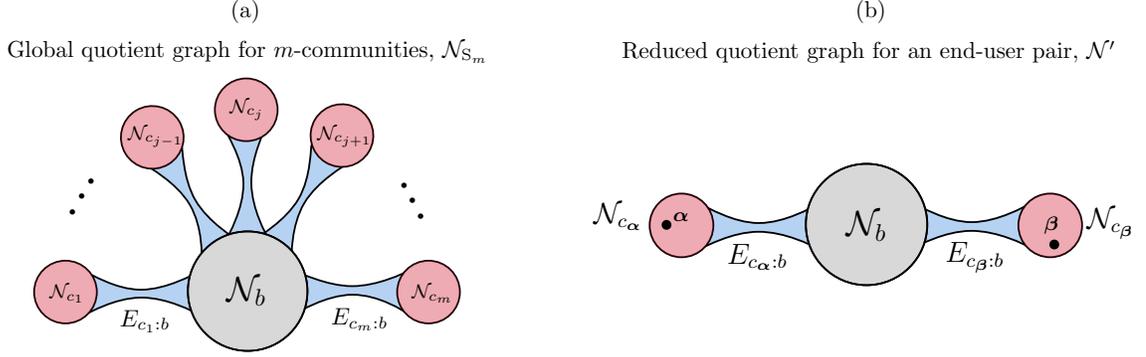


Figure C.1: (a) Quotient graph of a modular backbone network under the community equivalence relation for m -communities, resulting in a star network. (b) When considering the multi-path capacity between end-users α and β , the quotient graph can be simplified to a linear chain.

In general the condition given in Eq. (C.26) is fairly restrictive, as it places requirements on the minimum capacities between any two users in the same community. However, we shall see that in the case of communities connected to backbone networks (in which the quotient graph is simply a star network) the condition applies only to the two end-user community networks \mathcal{N}_{c_α} and \mathcal{N}_{c_β} .

C.1.5 Modular Networks with Backbone Structure

We can now turn our discussion to the modular networks as defined in Definitions 5.1 and 5.3, which are specific architectures with community structure. These are modular networks where all of the communities are disjoint and disconnected, but are all connected to a municipal backbone network. By imposing regularity (and thus high connectivity) on the backbone, we are able to study ideal modular networks. It is clear to see that the quotient network of this kind of modular architecture produces a star network. Let us denote a star network with m -children nodes and a central node by \mathcal{N}_{S_m} . Each community becomes a child node of the central backbone node, and we gather a very simple network structure. This is illustrated in Fig. C.1.

We find that when our modular network adopts this simple (yet very general) structure, then Lemma C.3 also simplifies significantly. It can be shown that the conditions in Lemma C.3 reduce to simple constraints only on community networks involved with the end-user pair; not on any other community sub-network. This result is captured in the following.

Lemma C.4 Consider a pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ and their associated pair of end-user

communities $\mathbf{i}_Q = \{\mathbf{c}_Q^\alpha, \mathbf{c}_Q^\beta\}$. The quotient graph \mathcal{N}_Q of the network under the community equivalence relation R is a star network. It then follows that if

$$\min_{x \in \{c_\alpha, c_\beta, b\}} \mathcal{C}_x^* \geq \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}, \quad (\text{C.32})$$

then the end-to-end flooding capacity between α and β is equal to

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}. \quad (\text{C.33})$$

Otherwise, the flooding capacity on the quotient network is an upper-bound on the true flooding capacity $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}$.

Proof. It is known that we can always perform a valid cut by community isolation, i.e. exclusively cutting the intercommunity edges between the backbone and either of the end-user communities. This type of cut equates to a cut on the quotient network, so that in general we can write the global-community capacity as an upper-bound on the flooding capacity

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}(\mathbf{i}_Q, \mathcal{N}_Q) = \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}. \quad (\text{C.34})$$

Now let us impose the condition in Eq. (C.32). This condition is similar to that which is proven in Lemma C.3 for more general networks. However in this setting, it is not necessary to consider communities which don't contain end-users. When Eq. (C.32) holds, this means that any cut which collects an edge from the sub-networks $x \in \{c_\alpha, c_\beta, b\}$ will not be a minimum cut. More precisely, by Lemma C.2, any cut which collects an edge from any of these sub-networks will automatically invoke a valid intracommunity cut between a pair of arbitrary local nodes. But per Eq. (C.32), the minimum local-community capacity is always larger than the global-community capacity, therefore this form of cut will never be the minimum cut.

We are now left to check that any cut which collects edges from other communities $c_i \notin \{c_\alpha, c_\beta, b\}$ will never be the minimum cut under these conditions. Consider a sub-graph of the original network $\mathcal{N}' = (P', E') \subset \mathcal{N}$ which consists solely of the communities c_α , c_β and c_b , each of the communities intracommunity edges and the corresponding intercommunity edges. Therefore the sets of sub-graph nodes and edges are

$$P' = P_{c_\alpha} \cup P_{c_\beta} \cup P_b, \quad (\text{C.35})$$

$$E' = (E_{c_\alpha} \cup E_{c_\beta} \cup E_b) \cup (E_{c_\alpha:b} \cup E_{c_\beta:b}). \quad (\text{C.36})$$

In general, the flooding capacity computed on the sub-network \mathcal{N}' will always be smaller than that computed on \mathcal{N} . The addition of extra communities can only ever increase the

number of end-to-end multi-path routes. As a result, we can write the lower-bound

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) \geq \mathcal{C}(\mathbf{i}, \mathcal{N}'). \quad (\text{C.37})$$

It is very important to note that $\mathcal{C}(\mathbf{i}, \mathcal{N}')$ is not necessarily a valid, end-to-end capacity. This is because the minimum cut which generates $\mathcal{C}(\mathbf{i}, \mathcal{N}')$ may not be a valid end-user partition on the global network \mathcal{N} . When the minimum cut which generates $\mathcal{C}(\mathbf{i}, \mathcal{N}')$ is also a valid cut on $\mathcal{C}(\mathbf{i}, \mathcal{N})$, then the above lower-bound saturates.

The quotient network of the sub-graph \mathcal{N}' can then be reduced to a simple linear chain, as shown in Fig. C.1(b). Now, thanks to the condition in Eq. (C.34), we can equate the flooding capacity on \mathcal{N}' to that computed on its quotient network,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}') = \mathcal{C}(\mathbf{i}_Q, \mathcal{N}'_Q) = \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}. \quad (\text{C.38})$$

Crucially, the minimum cut which generates this capacity is a valid cut on the global network \mathcal{N} , since it is always possible to partition the end-users via community isolation. As a result, when we combine this lower-bound with the upper-bound in Eq. (C.34), we gather that the end-to-end flooding capacity is given by

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\}, \quad (\text{C.39})$$

as required. If the condition Eq. (C.32) is violated, we re-gather the upper-bound in Eq. (C.34) since there may exist a cut that uses local community edges in $x \in \{c_\alpha, c_\beta, b\}$ to reduce the end-to-end capacity. This new cut *may* also be a valid cut on \mathcal{N}' , but it will not be achieved by community isolation, i.e. the lower-bound in Eq. (C.37) will still hold, but it will not be attributed to Eq. (C.38). ■

The technique used in the proof is actually rather more powerful than it may first appear. The key is to select a sub-graph whose quotient graph has exactly the same possible minimum cuts as the quotient graph of the overall network. When the condition in Eq. (C.26) on the sub-graph holds, this guarantees that the lower bound, found by asserting that the end-to-end capacity on the overall network must be greater than on the sub-graph, can be saturated on the overall network. This in turn allows the lower bound to match the upper and reduce the restrictiveness of the condition given in Eq. (C.26) to just minimising $\mathcal{C}_{c_i}^*$ over the communities that exist in the sub-graph. This highlights that the degree of simplification provided by Lemma C.3 depends on the underlying topology of the quotient network. Similar techniques can be applied to loosen the restrictions of Eq. (C.26) for other quotient network topologies, although we leave the exploration of these to future works.

C.1.6 Threshold Capacities of Modular Networks with Backbone Structure

Using the developments throughout this section, we can provide a concise proof of the main theorem in Chapter 5. This allows us to identify single-edge capacity thresholds for each of the end-user community networks and the backbone network, such that the end-to-end capacity is equal to the global-community capacity. As a result, we can identify unique physical constraints which can be used to motivate the construction of particular sub-networks, as was done in the main text. Here we restate the theorem for clarity:

Theorem 5.1 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3. Select any pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ contained in remote communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$. There exist single-edge threshold capacities on the communities $\mathcal{C}_{c_j}^{\min}$ and backbone \mathcal{C}_b^{\min} sub-networks for which the network flooding capacity is given by the global-community capacity,*

$$\left. \begin{aligned} \mathcal{C}_{\mathbf{x}\mathbf{y}} &\geq \mathcal{C}_{c_j}^{\min}, \forall (\mathbf{x}, \mathbf{y}) \in E_{c_j}, \\ \mathcal{C}_{\mathbf{x}\mathbf{y}} &\geq \mathcal{C}_b^{\min}, \forall (\mathbf{x}, \mathbf{y}) \in E_b, \end{aligned} \right\} \implies \mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{c:b}, \quad (\text{C.40})$$

for all $\mathbf{j} \in \{\alpha, \beta\}$. The threshold capacities are given by,

$$\mathcal{C}_{c_j}^{\min} := \frac{\mathcal{C}_{c:b}}{k_{c_j}}, \quad \mathcal{C}_b^{\min} := \frac{\mathcal{C}_{c:b}}{H_{\min}^*}, \quad (\text{C.41})$$

where H_{\min}^* is the minimum cut-set cardinality on the backbone network. If these threshold capacities are violated, then the global-community capacity becomes an upper-bound on the end-to-end capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c:b}$.

Proof. Any modular network following the form of Definition 5.3 admits a star network as its quotient graph. In order to assert that the global-community capacity between any two end-users located in remote communities $\mathbf{i} = \{\alpha, \beta\}$ is indeed the flooding capacity, we must reveal conditions for which all other possible cuts generate larger flooding capacities. Thanks to Lemma C.4 we know this condition is,

$$\min_{c \in \{c_\alpha, c_\beta, b\}} \mathcal{C}_c^* \geq \min\{\mathcal{C}_{c_\alpha:b}, \mathcal{C}_{c_\beta:b}\} = \mathcal{C}_{c:b}, \quad (\text{C.42})$$

where $\mathcal{C}_{c:b}$ is the global community capacity which implicitly performs the minimisation. To satisfy the condition in Eq. (C.42), it is sufficient to satisfy the set of equations,

$$\mathcal{C}_x^* \geq \mathcal{C}_{c:b}, \quad \forall x \in \{c_\alpha, c_\beta, b\}. \quad (\text{C.43})$$

Using this set of conditions, we are able to derive threshold capacities for each of the sub-networks of the modular structure to ensure the global community capacity is equal to the end-to-end capacity.

Let us first focus on satisfying this condition for the end-user communities, c_α and c_β . By definition, each of the communities in our idealised modular architecture adopt k_{c_j} -connectivity, $\mathbf{j} \in \{\alpha, \beta\}$. This means that the smallest possible cut between *any two nodes* on either of the community networks (which contain end-users) collects exactly k_{c_j} edges. Let us also assume that there exists a single-edge threshold capacity for each community $\mathcal{C}_{c_j}^{\min}$. Therefore, we can always say that the min-local community capacity $\mathcal{C}_{c_j}^*$ will never be smaller than that which is generated by cutting k_{c_j} edges each of which have a minimum threshold capacity $\mathcal{C}_{c_j}^{\min}$. That is, we can write

$$\mathcal{C}_{c_j}^* \geq k_{c_j} \mathcal{C}_{c_j}^{\min}, \quad \mathbf{j} \in \{\alpha, \beta\}, \quad (\text{C.44})$$

This lower-bound on the min-local capacity is achievable, since it is based on a valid cut on the communities. In order to satisfy Eq. (C.43) for each of communities, we must then demand

$$\mathcal{C}_{c_j}^* \geq k_{c_j} \mathcal{C}_{c_j}^{\min} \geq \mathcal{C}_{c:b}, \quad \mathbf{j} \in \{\alpha, \beta\}. \quad (\text{C.45})$$

As a result, we derive a single-edge threshold capacity for edges within the local communities,

$$\mathcal{C}_{c_j}^{\min} = \frac{\mathcal{C}_{c:b}}{k_{c_j}}, \quad \mathbf{j} \in \{\alpha, \beta\}. \quad (\text{C.46})$$

With this condition, we ensure that any valid cut performed exclusively on the local communities will always generate a larger multi-edge capacity than $\mathcal{C}_{c:b}$, and will not be the minimum cut.

We are now left to identify the single-edge constraint for the regular backbone network. While Eq. (C.42) will supply a sufficient condition for the backbone network to ensure it does not compromise the global community capacity, the property of regularity lets us determine a more specific constraint. The backbone may possess many connections from the communities, meaning that the minimum number of edges in a cut-set performed exclusively on the backbone as potentially very large. The minimum cut-set size of a backbone cut depends totally on the network regularity, and the distribution of intercommunity connections, i.e. the set of nodes $P_{b|c_j}$ which tell us where the community c_j is directly connected to the backbone (for $\mathbf{j} \in \{\alpha, \beta\}$). For regular networks, it is always possible to determine

this minimum cut-set size via collective node isolation. Hence, the minimum cardinality can be summarised by the function

$$|\tilde{\mathcal{C}}_b| \geq H_{\min}^* := \min_{j \in \{\alpha, \beta\}} H_{\min}(k_b, P_{b|c_j}) \quad (\text{C.47})$$

which chooses the minimum cut-set cardinality associated with either set of intercommunity connections. It then follows that, given some minimum single-edge capacity on the backbone network \mathcal{C}_b^{\min} , the minimum possible multi-edge backbone capacity is given by

$$\mathcal{C}_b \geq |\tilde{\mathcal{C}}_b| \mathcal{C}_b^{\min} \geq H_{\min}^* \mathcal{C}_b^{\min}. \quad (\text{C.48})$$

In order the global-community capacity to remain a minimum cut, it must always be smaller than this lower-bound. Hence, we assert that,

$$\mathcal{C}_{c:b} \leq H_{\min}^* \mathcal{C}_b^{\min} \quad (\text{C.49})$$

which leads to the required condition. ■

It is important to note that these conditions hold for any end-user pair in remote communities, even when the user nodes possesses direct connections to the backbone. When this is the case, there will never exist a valid end-user cut that is exclusively made up of local-community edges, since it is now necessary to also cut the direct connections to the backbone. Let $E_{\mathbf{x}} := \{(\mathbf{x}, \mathbf{y}) \in E \mid \mathbf{y} \in P\}$ be the set of all edges in the neighbourhood of a node \mathbf{x} . We can identify the intercommunity edges which provide direct connections from a node \mathbf{x} to the backbone via the edge set $E_{\mathbf{x}} \setminus E_{c_{\mathbf{x}}}$, i.e. all the directly connected edges to \mathbf{x} minus those which are community edges. Hence, we can never eliminate community-wide communication by means of a local community cut. If we impose the condition in Theorem 5.1 *anyway*, then this is sufficient to guarantee the global community capacity. Collecting k_{c_j} local community edges will automatically generate a multi-edge capacity which is at least as large as $\mathcal{C}_{c:b}$. Hence, the additional edges that one needs to collect to consolidate the cut can only increase this multi-edge capacity.

More precisely, the modification which minimises the number of extra edges collected is achieved by additionally collecting the edges which connect the user-node directly to the backbone. We can denote the multi-edge capacity associated with cutting the user-connected intercommunity edges as

$$\mathcal{C}_{j:b} := \sum_{(\mathbf{x}, \mathbf{y}) \in E_j \setminus E_{c_j}} \mathcal{C}_{\mathbf{x}\mathbf{y}}. \quad (\text{C.50})$$

The necessity of cutting additional intercommunity edges means that the min-local community capacity can *never* be the flooding capacity. It can never be a valid minimum cut on its own, since the direct backbone connection means there will remain a route to the backbone (and thus to the other end-user). Instead, we perform a cut of k_{c_j} edges on the local community and cut these direct backbone connections $E_{\mathbf{x}} \setminus E_{c_{\mathbf{x}}}$. This results in a multi-edge capacity of

$$\mathcal{C}_{c_j}^* + \mathcal{C}_{j:b} \geq k_{c_j} \mathcal{C}_{c_j}^{\min} + \mathcal{C}_{j:b}, \quad \forall j \in \{\alpha, \beta\}, \quad (\text{C.51})$$

To ensure that this cut is never the minimum cut, we ask that

$$k_{c_j} \mathcal{C}_{c_j}^{\min} + \mathcal{C}_{j:b} \geq \mathcal{C}_{c:b}, \quad \forall j \in \{\alpha, \beta\}, \quad (\text{C.52})$$

is always true. Hence the necessity of cutting additional intercommunity edges leads to the modified condition on the local community threshold capacities,

$$k_{c_j} \mathcal{C}_{c_j}^{\min} \geq \mathcal{C}_{c:b} - \mathcal{C}_{j:b}, \quad \forall j \in \{\alpha, \beta\}, \quad (\text{C.53})$$

which is clearly a looser condition than that in Theorem 5.1. Therefore, Theorem 5.1 holds regardless of if the end-users are directly connected to the backbone or not.

C.2 Application to Hybrid Quantum Networks

With the main theorem from the main text now proven, it is possible to elucidate the emergence of Corollaries 5.1 and 5.2. These are simply applications of Theorem 5.1 in the context of fibre/satellite modular quantum networks, and ground-based free-space/fibre architectures. To assist the reader, we restate each corollary before providing their proofs.

C.2.1 Fibre/Satellite Configuration

Corollary 5.1 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3 in the main text, and assume optical-fibre communities networks $\mathcal{N}_{c_{\alpha}}$, $\mathcal{N}_{c_{\beta}}$ and a satellite-based backbone \mathcal{N}_b . Select any pair of end-users $\{\alpha, \beta\}$ located in remote communities $\alpha \in P_{c_{\alpha}}$ and $\beta \in P_{c_{\beta}}$. There exists a maximum fibre-length in each community*

$$d_{c_j}^{\max} := -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c:b}/k_{c_j}} \right), \quad (\text{C.54})$$

and a maximum intersatellite separation in the backbone

$$z_b^{\max} := \arg \min_z \left| \log \left(\frac{H_{\min}^* \mathcal{B}_{F_{\sigma_p}}(\eta)}{\mathcal{C}_{c:b}} \right) \right|. \quad (\text{C.55})$$

for which the network flooding capacity is equal to the global-community capacity,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{c:b}. \quad (\text{C.56})$$

Otherwise, if any intersatellite links violate this condition $\exists z_{\mathbf{x}\mathbf{y}} > z_b^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_b$ or the local community links are in violation, $\exists d_{\mathbf{x}\mathbf{y}} > d_{c_j}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_{c_j}$, for either $\mathbf{j} \in \{\alpha, \beta\}$, then this becomes an upper-bound on the network flooding capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c:b}$.

Proof. The proof follows directly from the use of Theorem 5.1 and the direct substitution of single-edge capacity formulae into its results. As gathered from Theorem 5.1, for an ideal modular network of this form we can ensure that the end-to-end capacity between the end-users is equal to the global-community capacity if the following single-edge threshold capacities are satisfied: $\mathcal{C}_{c_j}^{\min} := \mathcal{C}_{c:b}/k_{c_j}$ and $\mathcal{C}_b^{\min} := \mathcal{C}_{c:b}/H_{\min}^*$.

Since the community sub-networks are consistent of fibre channels, we can equate the single-edge community threshold capacity to the precise expression of a bosonic pure-loss channel capacity (the PLOB bound). For fibre-channels, a minimum capacity threshold corresponds to a maximum fibre-length threshold, such that

$$\mathcal{C}_{c_j}^{\min} = -\log_2(1 - 10^{-\gamma d_{c_j}^{\max}}). \quad (\text{C.57})$$

where $\gamma \approx 0.2$ dB/km as used consistently throughout this thesis. This can be then be rearranged to determine the maximum permitted fibre-length within the community,

$$d_{c_j}^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c_j}^{\min}} \right) = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c:b}/k_{c_j}} \right). \quad (\text{C.58})$$

We may perform a similar procedure for the backbone network, which is constructed from intersatellite channels. Assuming negligible thermal contributions (see Section 5.3.8) but non-negligible pointing errors, then we can relate the threshold capacity \mathcal{C}_b^{\min} to the single-edge capacity expression for intersatellite channels from Eq. (5.37),

$$\mathcal{C}_b^{\min} = \frac{\mathcal{C}_{c:b}}{H_{\min}^*} \leq \mathcal{B}_{F_{\sigma_p}}[\eta_d(z)] = \frac{2a_R^2 \Delta(\eta_d(z), \sigma_p)}{w_d^2(z) \ln 2}. \quad (\text{C.59})$$

In general settings this is an upper-bound, as it is an extension of the PLOB bound to an ensemble of lossy channels where the convexity properties of the relative entropy of entanglement (REE) are exploited [87, 30, 31]. However, we reliably assume the intersatellite channels to be modeled as pure-loss channels, and thus can admit equality $\mathcal{C}_b^{\min} = \mathcal{B}_{F_{\sigma_p}}[\eta_d(z)]$. We are now in a position to compute the maximum tolerable intersatellite separation z_b^{\max} . This is the same as asking: for what channel length z does the following equality hold

$$\mathcal{C}_{c:b} = H_{\min}^* \mathcal{B}_{F_{\sigma_p}}[\eta_d(z)]. \quad (\text{C.60})$$

Due to the complicated nature of the capacity function $\mathcal{B}_{F_{\sigma_p}}[\eta_d(z)]$ this is not expedient analytically. However it is easy to compute numerically. Indeed, finding the maximum intersatellite separation equates to finding the minimum argument of

$$z_b^{\max} = \arg \min_z \left| \log \left(\frac{H_{\min}^* \mathcal{B}_{F_{\sigma_p}}(\eta)}{\mathcal{C}_{c:b}} \right) \right|. \quad (\text{C.61})$$

Here we use the absolute log-ratio to compare the right and left hand-side of Eq. (C.60) and determine for what channel length z_b^{\max} they are equivalent. ■

C.2.2 Ground-Based Free-Space Fibre Configuration

Corollary 5.2 *Consider an ideal modular network of the form \mathcal{N}^* introduced in Definition 5.3 in the main text, and assume free-space community networks \mathcal{N}_{c_α} , \mathcal{N}_{c_β} and an optical-fibre backbone \mathcal{N}_b . Select any pair of end-users $\mathbf{i} = \{\alpha, \beta\}$ located in remote communities $\alpha \in P_{c_\alpha}$ and $\beta \in P_{c_\beta}$. There exists a maximum free-space link length in each community*

$$z_{c_j}^{\max} \leq \arg \min_z \left| \log \left(\frac{k_{c_j} \mathcal{T}_{F_\sigma}(\eta, \bar{n}_j)}{\mathcal{C}_{c:b}} \right) \right|, \quad (\text{C.62})$$

and a maximum fibre length in the backbone

$$d_b^{\max} := -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c:b}/H_{\min}^*} \right), \quad (\text{C.63})$$

for which the network flooding capacity is equal to the global-community capacity,

$$\mathcal{C}(\mathbf{i}, \mathcal{N}) = \mathcal{C}_{c:b}. \quad (\text{C.64})$$

Otherwise, if any fibre links violate this condition $\exists d_{\mathbf{xy}} > d_b^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_b$ or the local community links are in violation, $\exists z_{\mathbf{xy}} > z_{c_j}^{\max}$, $(\mathbf{x}, \mathbf{y}) \in E_{c_j}$, for either $\mathbf{j} \in \{\alpha, \beta\}$, then this becomes an upper-bound on the network flooding capacity, $\mathcal{C}(\mathbf{i}, \mathcal{N}) \leq \mathcal{C}_{c:b}$.

Proof. Once again, a proof follows directly from Theorem 5.1 and the techniques used to prove the previous Corollary. For an ideal modular network of this form we can ensure that the end-to-end capacity between the end-users is equal to the global-community capacity if the following single-edge threshold capacities are satisfied: $\mathcal{C}_{c_j}^{\min} = \mathcal{C}_{c:b}/k_{c_j}$ and $\mathcal{C}_b^{\min} = \mathcal{C}_{c:b}/H_{\min}^*$.

In this setting, the community sub-networks are consistent of ground-based free-space quantum channels. We focus on the regime of weak turbulence, such that channel lengths are limited to $z \lesssim 1$ km. For a community containing an end-user $\mathbf{j} \in \{\alpha, \beta\}$ we can write

$$\mathcal{C}_{c_j}^{\min} = \frac{\mathcal{C}_{c:b}}{k_{c_j}} \leq \mathcal{T}_{F_\sigma}[\eta(z), \bar{n}_j], \quad (\text{C.65})$$

where $\mathcal{T}_{F_\sigma}[\eta(z), \bar{n}_j]$ is the single-edge capacity upper-bound associated with a ground-based free-space link, discussed in Section 5.3.4. This incorporates atmospheric fading dynamics, and free-space background noise \bar{n}_j which may be present in the community c_j . Hence, determining the maximum free-space link permitted in an end-user community is equivalent to finding the smallest channel length z for which the equality

$$\mathcal{C}_{c:b} = k_{c_j} \mathcal{T}_{F_\sigma}[\eta(z), \bar{n}_j], \quad (\text{C.66})$$

is satisfied. As before, this can be carried out numerically by finding the minimum argument

$$z_{c_j}^{\max} \leq \arg \min_z \left| \log \left(\frac{k_{c_j} \mathcal{T}_{F_\sigma}(\eta, \bar{n}_j)}{\mathcal{C}_{c:b}} \right) \right|. \quad (\text{C.67})$$

This is an upper-bound on $z_{c_j}^{\max}$, since it is not known whether $\mathcal{T}_{F_\sigma}(\eta, \bar{n}_j)$ is an achievable rate or not. Nonetheless, this single-edge upper bound has been shown to be tight, and therefore we can accurately utilise it in order to gain insight into the reliability of free-space links in a metropolitan network setting.

For the fibre-backbone, we possess exact expressions for single-edge capacities. Therefore, to find the maximum fibre-length we can simply compare the backbone threshold capacity to the PLOB bound and arrive at the result

$$d_b^{\max} = -\frac{1}{\gamma} \log_{10} \left(1 - 2^{-\mathcal{C}_{c:b}/H_{\min}^*} \right). \quad (\text{C.68})$$

This completes the proof. ■

C.3 Collective Node Isolation

C.3.1 Definition and Motivation

Consider a network based on an underlying undirected graph $\mathcal{N} = (P, E)$, and some collection of n -network nodes $\mathbf{I} = \{\mathbf{i}_1, \mathbf{i}_2, \dots, \mathbf{i}_n\} \subset P$ within it. Here, we will define \mathbf{I} as a set of *target-nodes* that we are interested in. We define the task of *Collective Node Isolation* as that of determining the smallest cut-set of edges $\tilde{C}_{\min} \subset E$ (that which collects the fewest number of edges) that need to be removed from the network in order to form a sub-graph $\mathcal{N}_{\mathbf{I}} = (P_{\mathbf{I}}, E_{\mathbf{I}})$ within which all the target nodes are contained, i.e. $\mathbf{I} \subseteq P_{\mathbf{I}}$. Importantly, this sub-graph need not be exclusively consistent of target nodes, but can also possess additional nodes. This question is relevant as it emerges within an unweighted minimum-cut problem for distant collections of nodes on a highly-connected network. That is, given some disjoint collections of sender nodes \mathbf{A} and receiver nodes \mathbf{B} , what is the minimum cut-set cardinality required to partition these collections of end-users?

Clearly, for a completely general network it is by no means obvious what this cut-set is. However, by asserting some form of connectivity constraints it is possible to gain some useful analytical insight. In particular we are interested in k -regular networks, relevant for the regular backbone networks studied within Chapter 5. The high level of connectivity guaranteed by regularity ensures that for given a pair of individual end-user nodes, the cut-set with the smallest cardinality (neglecting boundary effects) will always be found via nodal isolation. This is because regularity guarantees a high growth rate for cut-set sizes as one moves further away from either end-user; hence the closer one remains to either end-user, the smaller the cut-set will be. Regular networks with this property are defined as super-connected.

Hence, collective node isolation can be used to identify minimum-cut set sizes on a super-connected graph when it is necessary to isolate a number of particular nodes, \mathbf{I} . This is a generalisation of the work in Chapter 3 in which the focus is nodal isolation on weakly-regular graphs. In a modular network setting, collective node isolation is important for identifying minimum cut-set cardinalities when restricted to a particular sub-network of the global model. This is made clear via its application in the main body of this thesis. In the following, we devise the general result for the cut-set size of collective node isolation on regular networks.

C.3.2 Minimum Cut-Set Cardinality

Consider a k -regular network $\mathcal{N} = (P, E)$ and two specific disjoint collections of target nodes labelled $\{\mathbf{A}, \mathbf{B}\} \subset P$. These collections are used to represent end-user connected nodes on an intermediate sub-network within a modular structure. We wish to derive an expression for the minimum-cut set size required to completely partition the collections.

Consider either collection of target-connected nodes, $\mathbf{I} \in \{\mathbf{A}, \mathbf{B}\}$. A regular network is super-connected, hence the minimum cut-set cardinality is always achieved by neighbourhood isolation. In this way, the largest set generated by neighbourhood isolation occurs when all the users are sufficiently separated so that they do not share any edges or any neighbours. Then the cut-set has cardinality $|\tilde{C}| = k|\mathbf{I}|$. This is always an upper-bound on the minimum cut-set size. However, the potential for target-nodes sharing edges and sharing neighbours can diminish this cut-set size, since redundant edges may emerge. It's therefore possible to provide a corrective procedure for determining the minimum cut-set size.

Consider a set of N target nodes $\mathbf{I} = \{\mathbf{i}_1, \dots, \mathbf{i}_N\}$, and potential partitions of this set, i.e. $\mathbf{I} = \{\mathbf{i}_1, \mathbf{i}_2 | \mathbf{i}_3 | \mathbf{i}_4, \dots, \mathbf{i}_N\}$. Let us denote the space of all possible partitions via $\Gamma_{\mathbf{I}} \in \mathbf{\Gamma}_{\mathbf{I}}$,

such that $\Gamma_{\mathbf{I}} = \{\gamma_j\}_{j=1}^{|\Gamma_{\mathbf{I}}|}$ represents a single instance of a partition of the target-node set and $\gamma \in \Gamma_{\mathbf{I}}$ denotes the subset of target nodes within this partition.

As mentioned above, it is important to recognise that one can utilise the feature of edge-sharing and neighbour-sharing between target nodes that can reduce the number of edges that need to be cut. We can identify this potential through the following procedure: Let us define ω_γ as a path (sequence of edges) on the network that connects all the nodes within the target node set partition, $\mathbf{i} \in \gamma \in \Gamma_{\mathbf{I}}$. The path is free to form a closed cycle on the network, meaning that it may enclose a subset of network nodes within it (such that the path nodes form a boundary). Let us denote the set of nodes on the path and enclosed by it (if it is a cycle) by P_{ω_γ} . Note that if a particular partition consists of only one target node, e.g. $\gamma = \{\mathbf{i}\}$, $\mathbf{i} \in \mathbf{I}$, then there is no need to construct a path and $P_{\omega_\gamma} = \{\mathbf{i}\}$.

We then can conclude an expression for the minimum cut-set cardinality via the following logic: sometimes is cheaper to form paths between target nodes, and collect edges *around the edges along that path* than it is to simply isolate the target nodes. This is because these inter-target node edges can often be redundantly cut when they do not contribute to information flow outside of the target node partition. Furthermore, when these paths form cycles on the network, they identify enclosed nodes and edges that also do not benefit information flow outside of the partition. The minimum cut-set is then found by finding the set of shortest inter-target paths $\{\omega_{\gamma_j}\}_{j=1}^{|\Gamma_{\mathbf{I}}|}$ under a best case partition of the target nodes. Defining the following Kronecker-delta like function,

$$\delta(\mathbf{x}, \mathbf{y}) := \begin{cases} 1, & \text{if } (\mathbf{x}, \mathbf{y}) \in E, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{C.69})$$

then we arrive at the following quantity,

$$H_{\min}(k, \mathbf{I}) := \min_{\Gamma \in \Gamma_{\mathbf{I}}} \left[\sum_{\gamma \in \Gamma} \min_{\omega_\gamma} \left(k|P_{\omega_\gamma}| - \sum_{\mathbf{x} \neq \mathbf{y} \in P_{\omega_\gamma}} \delta(\mathbf{x}, \mathbf{y}) \right) \right]. \quad (\text{C.70})$$

Importantly, the quantity above recaptures the upper-bound via the target-node set partitioning $\Gamma_{\mathbf{I}} = \{\mathbf{i}_1 | \mathbf{i}_2 | \dots | \mathbf{i}_N\}$.

The challenging aspect is the minimisation performed over all target node partitions. The identification of shortest paths between nodes should be easy and efficient, but as the number of target nodes grows the number of possible partitions will scale super-exponentially. However, there should exist heuristics that assist in minimising the partition search space, e.g. locality of target nodes, since very distant nodes will be incredibly unlikely to benefit from alternative cuts from nodal isolation.

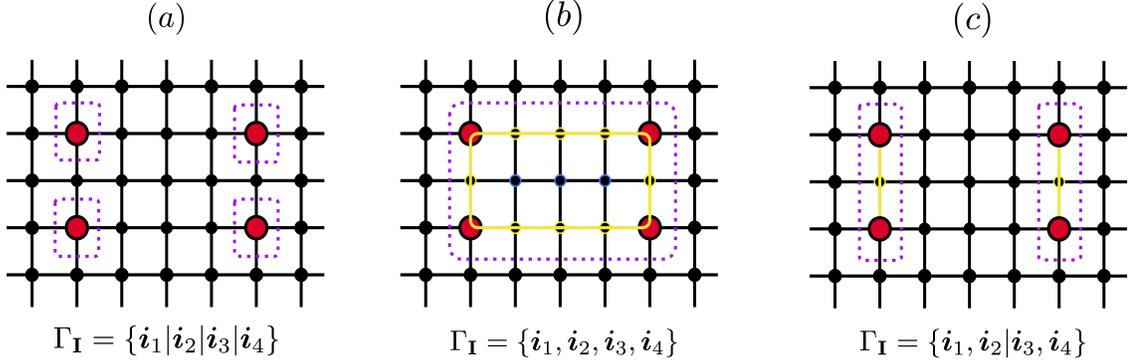


Figure C.2: Minimum-cut set cardinalities for collective node isolation. Panels (a), (b) and (c) depict different possible cuts that may emerge via the procedure of Section C.3.2 in which subsets of target nodes can be connected via paths to minimise their total cardinality. Panel (a) displays the natural case of target-node isolation, (b) shows how one can use closed paths, while (c) illustrates how multiple paths can be formed. In this instance, all cut-sets collect an equivalent number of edges.

C.3.3 Weakly-Regular Neighbourhood Isolation

We can use the example of WRNs from Chapter 3 to show the generality of the previous expression. In that work, it was important to determine the minimum cut-set cardinality that could be achieved when one is not permitted to cut neighbourhood edges of some potential end-users, α and β . In doing so, it was possible to derive conditions on the WRN for which the flooding capacity was always the minimum neighbourhood capacity between the users. Interestingly, this is equivalent to asking: what is the minimum cut-set size related to collectively isolating a user neighbourhood $\mathbf{A} = N_\alpha$ or $\mathbf{B} = N_\beta$ on a WRN? We can show that the result in Eq. (C.70) can reproduce the result found from this investigation.

Recall that a WRN is a network architecture based on an undirected graph $\mathcal{N}_{\text{WR}} = (P, E)$ which has the specific connectivity properties. In a $(k, \mathbf{\Lambda})$ -WRN, for any node \mathbf{x} , there is a multiset of values $\lambda_{\mathbf{x}}$ which collects the number of of common neighbours shared between \mathbf{x} and each $\mathbf{y} \in N_{\mathbf{x}}$. That is,

$$\lambda_{\mathbf{x}} := \{ |N_{\mathbf{x}} \cap N_{\mathbf{y}}| \mid \mathbf{y} \in N_{\mathbf{x}} \}. \quad (\text{C.71})$$

This is the *adjacent commonality multiset*. A network is $(k, \mathbf{\Lambda})$ -WR if each node is connected to exactly k other nodes, and each adjacent commonality multiset belongs to the superset $\mathbf{\Lambda}$ such that $\lambda_{\mathbf{x}} \in \mathbf{\Lambda}$ for all $\mathbf{x} \in P$.

Here, we focus on a scenario in which there is only one non-degenerate adjacent commonality multiset, i.e. $\mathbf{\Lambda} = \{\boldsymbol{\lambda}\}$, with the implicit understanding that this can be extended. Each node has k -neighbours, and the distribution of adjacent commonalities always follows $\boldsymbol{\lambda} = \{\lambda_1, \dots, \lambda_k\}$. As a result, every node \mathbf{i} has k -neighbours, and shares a unique number of common neighbours with each of them $\lambda_j \in \boldsymbol{\lambda}$. Consider performing collective isolation of the k -element neighbourhood of some node \mathbf{i} in this network. That is, our target nodes in this case become the set $N_{\mathbf{i}} = \{\mathbf{n}_1, \dots, \mathbf{n}_k\}$. The maximum cut-set size is of course

$$k|\mathbf{I}| = k|N_{\mathbf{i}}| = k^2, \quad (\text{C.72})$$

but this can be reduced due to edge-sharing and neighbour-sharing corrective factors.

We can immediately identify a cheaper upper-bound using the target-node partition set that isolates each target node. In that case, we need not build any paths, but can still identify the central node \mathbf{i} as enclosed within the boundary of the neighbourhood nodes, resulting in a corrective factor of k . That is, we can construct a smaller cut-set

$$H_{\min}(k, N_{\mathbf{i}}) \leq k^2 - k. \quad (\text{C.73})$$

Via the procedure outlined in the previous section, we could also identify a connective path $\omega^* = \omega_{\{\mathbf{n}_1, \dots, \mathbf{n}_k\}}$ that traverses the entire neighbourhood, connecting all of the target nodes. This might help us to identify a smaller cut-set, but only if the adjacent commonality between any two neighbourhood nodes $\lambda_j > 0$. If $\lambda_j = 0$ then know that it will take more than k edges to connect the path (since each neighbour \mathbf{n}_m shares no other neighbours with the target node \mathbf{i}). If $\lambda_j > 0$, one can be sure that the neighbourhood nodes can be connected via a path of length k edges; each neighbour shares a common neighbour with the target, so that the target node can be enclosed by a k -edge cycle.

Given that $\lambda(\mathbf{i}, \mathbf{n}) > 0$ for all $\mathbf{n} \in \{\mathbf{n}_1, \dots, \mathbf{n}_k\}$ then we can observe a further corrective factor to the minimum cut-set cardinality. Using Eq. (C.70) we will find that,

$$H_{\min}(k, N_{\mathbf{i}}) = k|P_{\omega^*}| - \sum_{\mathbf{x} \neq \mathbf{y} \in P_{\omega^*}} \delta(\mathbf{x}, \mathbf{y}), \quad (\text{C.74})$$

$$= k^2 - \sum_{\mathbf{x} \neq \mathbf{y} \in (P_{\omega^*} \setminus \{\mathbf{i}\})} \delta(\mathbf{x}, \mathbf{y}) - \sum_{\mathbf{x} \in N_{\mathbf{i}}} \delta(\mathbf{i}, \mathbf{y}). \quad (\text{C.75})$$

On the second line, we split the corrective contribution into two terms; a sum over the connections shared between nodes in the set $P_{\omega^*} \setminus \{\mathbf{i}\}$ that does not contain the target node, and a sum over the connections between the target node and those in P_{ω^*} . We find that the latter sum is simply equal to k , as it simply counts the number of nodes

in the neighbourhood. Meanwhile the former term counts the number of connections between neighbourhood nodes which are necessarily common connections to the target node. Consequently, this first sum is simply a sum over the adjacent commonalities,

$$\sum_{\mathbf{x} \neq \mathbf{y} \in (P_{\omega^*} \setminus \{i\})} \delta(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{x} \in N_i} \lambda(\mathbf{x}, i) = \sum_{j=1}^k \lambda_j. \quad (\text{C.76})$$

As a result, we arrive at the simplified min-cut set cardinality,

$$H_{\min}(k, N_i) = k|N_i| - \sum_{j=1}^k \lambda_j - k = \sum_{j=1}^k (k - \lambda_j - 1). \quad (\text{C.77})$$

which is the result reported in Chapter 3 where it was derived in a more direct fashion. As remarked, if $\lambda_j = 0$ for all j there is no additional minimisation possible due to path enclosures.

C.3.4 Bounds on Collective Node Isolation for Backbone Cuts

The function $H_{\min}(k, \mathbf{I})$ can be used for any general distribution of target-nodes. In the context of modular networks, it is very easy to write bounds on the minimum cut-set size for network cuts performed exclusively on the backbone. Given a k -regular backbone network, we can write the following bounds

$$k \leq H_{\min}(k, \mathbf{I}) \leq k|\mathbf{I}|. \quad (\text{C.78})$$

The lower bound corresponds to a situation where all the intercommunity edges are connected to a single node on the backbone. In this case it is sufficient to simply isolate the connected node on the backbone. The upper-bound refers to a situation where the intercommunity edges are connected to the backbone in such a way that the target-nodes do not share any edges or neighbours. The cut-set size for all other distributions \mathbf{I} fall within these bounds. Examples are illustrated for a Manhattan backbone network in Fig. 5.3.

Appendix D

Appendices for Chapter 7

In this appendix to Chapter 7, we briefly discuss how the developments made in this chapter can be extended to multicast protocols, and discuss how such concepts can be used to benchmark single and multiple quantum multicasting.

D.1 Extension to Quantum Multicasting

D.1.1 Multicast Communications

Let us consider the single-multicast communication setting. There exists a single sender node (Alice) \mathbf{a} which wishes to communicate with N_r receiver nodes (Bobs) $\{\mathbf{b}_j\}_{j=1}^{N_r}$. We can collect these quantities into an end-user *pool*,

$$\mathbf{i}^{N_r} := \left\{ \mathbf{a}, \{\mathbf{b}_j\}_{j=1}^{N_r} \right\}, \quad (\text{D.1})$$

where we maintain a similar notation for end-user *pairs* as end-user pools, utilising the superscript to denote the number of receivers within the pool. When exclusively considering end-user pairs $\mathbf{i} \equiv \mathbf{i}^1$ we may simply drop the superscript.

Alice may wish to share the same message with each of the receivers, or several different messages, which we can describe as her *message multiplicity*. We can quantitatively describe her message multiplicity via a multiplicity multiset,

$$\mathbf{m} := \{m_1, m_2, \dots, m_{N_r}\}, \quad m_j \in \{1, \dots, N_r\}, \quad (\text{D.2})$$

such that each $m_j \in \mathbf{m}$ denotes a label of the message she is sending to the j^{th} Bob. For example, if $\mathbf{m} = \{1, 1, \dots, 1\}$ then she is transmitting the same message to all Bobs, whereas $\mathbf{m} = \{1, 2, 3, \dots, N_r\}$ means she is transmitting unique messages to each receiver.

During single-multicast, the sender has access to all network resources in order to establish competent routes between \mathbf{a} and each of the receivers. In this multi-user configuration

it is clear that routing competition already exists. Each path from $\mathbf{a} \rightarrow \mathbf{b}_j$ must compete with every other sender to receiver path in order to reliably fulfil multicast communications. Dependent upon the objective of each receiver, one could also consider this as *routing collaboration* as it may be in the best interest of each receiver to help its counterparts construct high-rate routes. Regardless, it is clear that sophisticated routing strategies and TM techniques are necessary to ensure reliable multicast rates.

Multiple-multicast communications goes a step further and considers *many* senders communication with many receiver pools. That is, we may consider a collection of N_u end-user pools,

$$\mathcal{I}^{N_r} = \{i_1^{N_r^1}, \dots, i_{N_u}^{N_r^{N_u}}\} = \left\{ \left\{ \mathbf{a}_i, \{\mathbf{b}_j^i\}_{j=1}^{N_r^i} \right\}_{i=1}^{N_u} \right\}, \quad (\text{D.3})$$

where $N_r := \{N_r^i\}_{i=1}^{N_u}$ is a multiset that identifies N_r^i as the number of receiver nodes in the i^{th} receiver pool. Note that there is no need for these receiver pools to be unique, and different senders may wish to multicast to similar nodes and pools, i.e. in general,

$$\{\mathbf{b}_j^i\}_{j=1}^{N_r^i} \cap \{\mathbf{b}_j^k\}_{j=1}^{N_r^k} \neq \emptyset \quad \text{for } i \neq k. \quad (\text{D.4})$$

Routing competition is now even more intense than in multiple-unicast or single-multicast. Many senders must compete to establish paths between many (and potentially the same) receiver nodes.

D.1.2 Managing Single-Multicast Traffic

Let us consider single-multicast TM. Consider a quantum network $\mathcal{N} = (P, E)$ and a single EUP with N_r receivers,

$$i^{N_r} := \left\{ \mathbf{a}, \{\mathbf{b}_j\}_{j=1}^{N_r} \right\}. \quad (\text{D.5})$$

It is useful to translate this communication scenario into a pseudo-multiple-unicast format by considering a collection of N_r end-user *pairs*,

$$i^{N_r} \mapsto \mathcal{I} = \{i_j\}_{j=1}^{N_r} = \left\{ \{\mathbf{a}, \mathbf{b}_1\}, \{\mathbf{a}, \mathbf{b}_2\}, \dots, \{\mathbf{a}, \mathbf{b}_{N_r}\} \right\}. \quad (\text{D.6})$$

That is, multicasting can be thought of as an N_r multiple-unicast session in which all of the senders describes the same Alice node.

By thinking of single-multicast communications in this manner, we can recover all of the theoretical concepts regarding traffic management and network filters from the previous section. Each pair $i_j = \{\mathbf{a}, \mathbf{b}_j\}$ may employ a preferred, independent end-to-end routing protocol \mathcal{P}_{i_j} which describes how they would perform end-to-end routing in the absence

of routing competition. Once again \mathcal{P}_{i_j} is simply responsible for the mechanism by which routing is achieved; the opportunities and resources available to \mathcal{P}_{i_j} are controlled and influenced by the TM protocol \mathcal{T} . Each sender-receiver pair will compete for edges and establish routes dependent on the network filter that they have been exposed to throughout routing (dictated by the TM protocol). In this way, we can then state that the end-to-end rate between the sender and the j^{th} receiver will be,

$$K(i_j, \mathcal{N} | \mathcal{P}_j, \mathcal{T}) := \min_{\mathcal{C}} \sum_{(x,y) \in \tilde{\mathcal{C}}} \tilde{f}_{xy}^{i_j} \tilde{g}_{xy}^{i_j} K_{xy}. \quad (\text{D.7})$$

Multicasting invites the opportunity to replace routing competition with collaboration. Indeed, receivers may share a common objective, i.e. to achieve equal end-to-end rate across all receivers or to prioritise resources for particular receiver nodes. Fortunately, this notion is captured equivalently under the TM formalism. The TM protocol can monitor the multicast route and utilise feedback from the sender and receivers to optimise the network filters in line with their goal. A TM protocol can account for objectives desirable to both the network as a whole, and to particular communicators.

D.1.3 Managing Multiple-Multicast Traffic

As one might expect, an extension to multiple-multicast traffic is also immediate, since we can map any such communication configuration into a multiple-unicast scenario with repeated sender nodes. Given a N_u multicast sessions each with $N_r \in \mathbf{N}_r$ receivers that may take place simultaneously on the network \mathcal{I}^{N_r} , one can map this into multiple-unicast sessions described via the collection of $N_u \times N_r$ end-user pairs,

$$\mathcal{I}^{N_r} \mapsto \mathcal{I} = \{i_j\}_{j=1}^{N_u N_r} = \left\{ \{a_j, b_j^1\}, \dots, \{a_j, b_j^{N_r^j}\} \right\}_{j=1}^{N_u}. \quad (\text{D.8})$$

In this way, one can apply the machinery of network filters and TM protocols. such that each EUP is governed by an independent routing strategy \mathcal{P}_i which are then tied together via the overarching TM strategy, \mathcal{T} . End-to-end rates between senders and individual receivers are then computed via Eq. (D.7) once again, beyond which we can explicitly analyse capacity/rate regions.

D.1.4 Achievable Benchmarks

Through the translation of multicast sender/receiver configurations into multiple-unicast formats, it is immediately possible to study achievable benchmarks of quantum multicasting via naïve TM protocols. One can consider multicasting scenarios in which all EUPs aim

to communicate simultaneously such that each multicasting protocol between sender and receiver pools are governed by a naïve TM protocol, analogous to Section 7.3.2. Given the appropriate numerical tools to simulate routing protocols in quantum networks one can begin to compute achievable benchmarks for quantum multicasting performance.

Abbreviations

CV	Continuous Variable
DV	Discrete Variable
LO	Local Operations
CC	Classical Communications
QKD	Quantum Key Distribution
PLOB	Pirandola Laurenza Ottaviani Banchi
EPR	Einstein Podolsky Rosen
TMSV	Two Mode Squeezed Vacuum
REE	Relative Entropy of Entanglement
SQC	Satellite Quantum Communication
CI	Coherent Information
RCI	Reverse Coherent Information
LLO	Local Local Oscillator
TLO	Transmitted Local Oscillator
WR	Weakly Regular
WRN	Weakly Regular Network
DA	Dijkstra's Algorithm
MDPAlg	Multiple Disjoint Paths Algorithm
EUP	End-user pair
TM	Traffic Management
ED	Edge Disjoint
ND	Node Disjoint
QRD	Quantum Register Disjoint

Bibliography

- [1] C. Harney and S. Pirandola, “Analytical methods for high-rate global quantum networks,” *PRX Quantum* **3**, 010349 (2022).
- [2] C. Harney and S. Pirandola, “End-to-end capacities of imperfect-repeater quantum networks,” *Quantum Sci. Technol.* **7**, 045009 (2022).
- [3] C. Harney, A. I. Fletcher, and S. Pirandola, “End-to-end capacities of hybrid quantum networks,” *Phys. Rev. Applied* **18**, 014012 (2022).
- [4] J. Abbate, *Inventing the internet* (MIT Press, 2000).
- [5] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry,” *Rev. Mod. Phys.* **92**, 015003 (2020).
- [6] P. S. Emani, J. Warrell, A. Anticevic, S. Bekiranov, M. Gandal, M. J. McConnell, G. Sapiro, A. Aspuru-Guzik, J. T. Baker, M. Bastiani, J. D. Murray, S. N. Sotiropoulos, J. Taylor, G. Senthil, T. Lehner, M. B. Gerstein, and A. W. Harrow, “Quantum computing at the frontiers of biological sciences,” *Nature Methods* **18**, 701–709 (2021).
- [7] G. Spedalieri, L. Piersimoni, O. Laurino, S. L. Braunstein, and S. Pirandola, “Detecting and tracking bacteria with quantum light,” *Phys. Rev. Research* **2**, 043260 (2020).
- [8] Z. Huang and C. Lupo, “Quantum hypothesis testing for exoplanet detection,” *Phys. Rev. Lett.* **127**, 130502 (2021).
- [9] M. Tse, H. Yu, N. Kijbunchoo, A. Fernandez-Galiana, P. Dupej, L. Barsotti, C. D. Blair, D. D. Brown, S. E. Dwyer, A. Effler, M. Evans, *et al.*, “Quantum-enhanced

- advanced ligo detectors in the era of gravitational-wave astronomy,” *Phys. Rev. Lett.* **123**, 231107 (2019).
- [10] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM Journal on Computing* **26**, 1484–1509 (1997), <https://doi.org/10.1137/S0097539795293172> .
- [11] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” *Commun. ACM* **21**, 120–126 (1978).
- [12] V. S. Miller, “Use of elliptic curves in cryptography,” in *Advances in Cryptology — CRYPTO ’85 Proceedings*, edited by H. C. Williams (Springer Berlin Heidelberg, Berlin, Heidelberg, 1986) pp. 417–426.
- [13] N. Koblitz, “Elliptic curve cryptosystems,” *Mathematics of computation* **48**, 203–209 (1987).
- [14] W. Diffie and M. E. Hellman, “New directions in cryptography,” in *Democratizing Cryptography: The Work of Whitfield Diffie and Martin Hellman* (2022) pp. 365–390.
- [15] C. H. Bennett and G. Brassard, “Quantum cryptography: Public key distribution and coin tossing,” in *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing* (India, 1984) pp. 175–179.
- [16] S. L. Braunstein and P. van Loock, “Quantum information with continuous variables,” *Rev. Mod. Phys.* **77**, 513–577 (2005).
- [17] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, and et al., “Advances in quantum cryptography,” *Advances in Optics and Photonics* **12**, 1012 (2020).
- [18] J. I. Cirac, A. K. Ekert, S. F. Huelga, and C. Macchiavello, “Distributed quantum computation over noisy channels,” *Phys. Rev. A* **59**, 4249–4254 (1999).
- [19] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, “A quantum network of clocks,” *Nature Physics* **10**, 582–587 (2014).
- [20] D. Gottesman, T. Jennewein, and S. Croke, “Longer-baseline telescopes using quantum repeaters,” *Phys. Rev. Lett.* **109**, 070503 (2012).

-
- [21] H. J. Kimble, “The quantum internet,” *Nature* **453**, 1023–1030 (2008).
- [22] S. Pirandola and S. L. Braunstein, “Physics: Unite to build a quantum internet,” *Nature* **532**, 169–171 (2016).
- [23] M. Razavi, *An Introduction to Quantum Communications Networks*, 2053-2571 (Morgan & Claypool Publishers, 2018).
- [24] Y.-A. Chen, Q. Zhang, T.-Y. Chen, W.-Q. Cai, S.-K. Liao, J. Zhang, K. Chen, J. Yin, J.-G. Ren, Z. Chen, S.-L. Han, *et al.*, “An integrated space-to-ground quantum communication network over 4,600 kilometres,” *Nature* **589**, 214–219 (2021).
- [25] J. S. Sidhu, S. K. Joshi, M. Gundogan, T. Brougham, D. Lowndes, L. Mazzarella, M. Krutzik, S. Mohapatra, D. Dequal, G. Vallone, *et al.*, “Advances in space quantum communications,” *IET Quant. Comm.*, 1–36 (2021).
- [26] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, “Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels,” *Phys. Rev. Lett.* **70**, 1895–1899 (1993).
- [27] S. Pirandola, J. Eisert, C. Weedbrook, A. Furusawa, and S. L. Braunstein, “Advances in quantum teleportation,” *Nature Photonics* **9**, 641–652 (2015).
- [28] S. Pirandola, R. García-Patrón, S. L. Braunstein, and S. Lloyd, “Direct and reverse secret-key capacities of a quantum channel,” *Phys. Rev. Lett.* **102**, 050503 (2009).
- [29] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, “Fundamental limits of repeaterless quantum communications,” *Nature Communications* **8**, 15043 (2017).
- [30] S. Pirandola, “Limits and security of free-space quantum communications,” *Phys. Rev. Research* **3**, 013279 (2021).
- [31] S. Pirandola, “Satellite quantum communications: Fundamental bounds and practical security,” *Phys. Rev. Research* **3**, 023130 (2021).
- [32] S. Pirandola, “End-to-end capacities of a quantum communication network,” *Communications Physics* **2**, 51 (2019).
- [33] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, “Gaussian quantum information,” *Rev. Mod. Phys.* **84**, 621–669 (2012).

-
- [34] G. Adesso, S. Ragy, and A. R. Lee, “Continuous variable quantum information: Gaussian states and beyond,” *Open Systems and Information Dynamics* **21**, 1440001 (2014).
- [35] A. Ferraro, S. Olivares, and M. G. Paris, “Gaussian states in continuous variable quantum information,” arXiv:0503237 (2005).
- [36] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Taylor & Francis Group, 2017).
- [37] K. Stannigel, P. Komar, S. J. M. Habraken, S. D. Bennett, M. D. Lukin, P. Zoller, and P. Rabl, “Optomechanical quantum information processing with photons and phonons,” *Phys. Rev. Lett.* **109**, 013603 (2012).
- [38] W. Chen, Y. Lu, S. Zhang, K. Zhang, G. Huang, M. Qiao, X. Su, J. Zhang, J. Zhang, L. Bianchi, M. S. Kim, and K. Kim, “Scalable and programmable phononic network with trapped ions,” arXiv:2207.06115 (2022).
- [39] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, USA, 2011).
- [40] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, “Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels,” *Phys. Rev. Lett.* **70**, 1895–1899 (1993).
- [41] S. L. Braunstein and H. J. Kimble, “Teleportation of continuous quantum variables,” *Phys. Rev. Lett.* **80**, 869–872 (1998).
- [42] V. Vedral, M. B. Plenio, M. A. Rippin, and P. L. Knight, “Quantifying entanglement,” *Phys. Rev. Lett.* **78**, 2275–2279 (1997).
- [43] V. Vedral and M. B. Plenio, “Entanglement measures and purification procedures,” *Phys. Rev. A* **57**, 1619–1633 (1998).
- [44] V. Vedral, “The role of relative entropy in quantum information theory,” *Rev. Mod. Phys.* **74**, 197–234 (2002).
- [45] K. Horodecki, M. Horodecki, P. Horodecki, and J. Oppenheim, “Secure key from bound entanglement,” *Phys. Rev. Lett.* **94**, 160502 (2005).
- [46] B. Synak-Radtke and M. Horodecki, “On asymptotic continuity of functions of quantum states,” *Journal of Physics A: Mathematical and General* **39**, L423 (2006).

-
- [47] I. Devetak and A. Winter, “Distillation of secret key and entanglement from quantum states,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **461**, 207–235 (2005).
- [48] P. Slepian, *Mathematical Foundations of Network Analysis* (Springer-Verlag, New York, 1968).
- [49] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, New Jersey, 2006).
- [50] A. S. Tanenbaum and D. J. Wetherall, *Computer Networks*, 5th ed. (Pearson, 2010).
- [51] A. El Gamal and Y.-H. Kim, *Network Information Theory* (Cambridge University Press, 2011).
- [52] J. Borregaard, H. Pichler, T. Schröder, M. D. Lukin, P. Lodahl, and A. S. Sørensen, “One-way quantum repeater based on near-deterministic photon-emitter interfaces,” *Phys. Rev. X* **10**, 021071 (2020).
- [53] L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, “Fault-tolerant quantum communication based on solid-state photon emitters,” *Phys. Rev. Lett.* **96**, 070504 (2006).
- [54] C. Simon, H. de Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, “Quantum repeaters with photon pair sources and multimode memories,” *Phys. Rev. Lett.* **98**, 190503 (2007).
- [55] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische Mathematik* **1**, 269–271 (1959).
- [56] M. Pollack, “The maximum capacity through a network,” *Operations Research* **8**, 733–736 (1960).
- [57] T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 1990).
- [58] D. Medhi and K. Ramasamy, *Network routing: algorithms, protocols, and architectures*, 2nd ed. (Morgan Kaufmann, Cambridge, MA, 2018).
- [59] T. Harris and F. Ross, *Fundamentals of a method for evaluating rail net capacities*, Tech. Rep. (Rand Corporation, Santa Monica, CA, 1955).

-
- [60] L. R. Ford and D. R. Fulkerson, “Maximal flow through a network,” *Canadian Journal of Mathematics* **8**, 399–404 (1956).
- [61] J. Edmonds and R. M. Karp, “Theoretical improvements in algorithmic efficiency for network flow problems,” *J. ACM* **19**, 248–264 (1972).
- [62] J. B. Orlin, “Max flows in $o(nm)$ time, or better.” in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing, STOC’13* (2013) pp. 765–774.
- [63] J. Biamonte, M. Faccin, and M. De Domenico, “Complex networks from classical to quantum,” *Communications Physics* **2**, 53 (2019).
- [64] S. Brito, A. Canabarro, R. Chaves, and D. Cavalcanti, “Statistical properties of the quantum internet,” *Phys. Rev. Lett.* **124**, 210501 (2020).
- [65] Q. Zhuang and B. Zhang, “Quantum communication capacity transition of complex quantum networks,” *Phys. Rev. A* **104**, 022608 (2021).
- [66] B. Zhang and Q. Zhuang, “Quantum internet under random breakdowns and intentional attacks,” *Quantum Science and Technology* (2021).
- [67] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai, G.-B. Li, *et al.*, “Satellite-based entanglement distribution over 1200 kilometers,” *Science* **356**, 1140–1144 (2017).
- [68] J.-G. Ren, P. Xu, H.-L. Yong, L. Zhang, S.-K. Liao, J. Yin, W.-Y. Liu, W.-Q. Cai, M. Yang, L. Li, *et al.*, “Ground-to-satellite quantum teleportation,” *Nature* **549**, 70–73 (2017).
- [69] S.-K. Liao, J. Lin, J. Ren, W. Liu, J. Qiang, J. Yin, Y. Li, Q. Shen, L. Zhang, Y. Cao, *et al.*, “Space-to-Ground Quantum Key Distribution Using a Small-Sized Payload on Tiangong-2 Space Lab,” *Chinese Physics Letters* **34**, 090302 (2017).
- [70] A. Villar, A. Lohrmann, X. Bai, T. Vergoossen, R. Bedington, C. Perumangatt, H. Lim, T. Islam, A. Reezwana, Z. Tang, *et al.*, “Entanglement demonstration on board a nano-satellite,” *Optica* **7**, 734–737 (2020).
- [71] J. Yin, Y.-H. Li, S.-K. Liao, M. Yang, Y. Cao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, S.-L. Li, R. Shu, *et al.*, “Entanglement-based secure quantum cryptography over 1,120 kilometres,” *Nature* **582**, 501–505 (2020).

-
- [72] K. Menger, “Zur allgemeinen kurventheorie,” *Fundam. Math.* **10**, 96–115 (1927).
- [73] R. Aharoni and E. Berger, “Menger’s theorem for infinite graphs,” *Invent. Math.* **176**, 1–62 (2009).
- [74] R. Laurenza, N. Walk, J. Eisert, and S. Pirandola, “Rate limits in quantum networks with lossy repeaters,” *Phys. Rev. Research* **4**, 023158 (2022).
- [75] T. C. Ralph, “Continuous variable quantum cryptography,” *Phys. Rev. A* **61**, 010303 (1999).
- [76] B. Qi, P. Lougovski, R. Pooser, W. Grice, and M. Bobrek, “Generating the local oscillator “locally” in continuous-variable quantum key distribution based on coherent detection,” *Phys. Rev. X* **5**, 041009 (2015).
- [77] D. Huang, P. Huang, D. Lin, C. Wang, and G. Zeng, “High-speed continuous-variable quantum key distribution without sending a local oscillator,” *Opt. Lett.* **40**, 3695–3698 (2015).
- [78] J. Goodman, *Statistical Optics* (John Wiley & Sons, New York, 1985).
- [79] O. Svelto, *Principles of Lasers*, 5th ed. (Springer, New York, 2010).
- [80] C. Bohren and D. Huffman, *Absorption and scattering of light by small particles* (John Wiley & Sons, 2008).
- [81] V. Tatarskii, *The effects of the turbulent atmosphere on wave propagation* (Israel Program for Scientific Translations, Jerusalem, 1971).
- [82] A. Majumdar and J. Ricklin, *Free-Space Laser Communications* (Springer, New York, 2008).
- [83] H. Kaushal, V. Jain, and S. Kar, *Free Space Optical Communication* (Springer, New York, 2017).
- [84] L. Andrews and R. Phillips, *Laser Beam Propagation Through Random Medium* (SPIE, Bellingham, 2005).
- [85] V. C. Usenko, B. Heim, C. Peuntinger, C. Wittmann, C. Marquardt, G. Leuchs, and R. Filip, “Entanglement of gaussian states and the applicability to quantum key distribution over fading channels,” *New J. Phys.* **14**, 093048 (2012).

-
- [86] P. Papanastasiou, C. Weedbrook, and S. Pirandola, “Continuous-variable quantum key distribution in uniform fast-fading channels,” *Phys. Rev. A* **97**, 032311 (2018).
- [87] S. Pirandola, R. Laurenza, and L. Banchi, “Conditional channel simulation,” *Annals of Physics* **400**, 289–302 (2019).
- [88] R. García-Patrón, S. Pirandola, S. Lloyd, and J. H. Shapiro, “Reverse coherent information,” *Phys. Rev. Lett.* **102**, 210501 (2009).
- [89] S. Q. Duntley, “The reduction of apparent contrast by the atmosphere,” *J. Opt. Soc. Am.* **38**, 179–191 (1948).
- [90] D. Vasylyev, W. Vogel, and F. Moll, “Satellite-mediated quantum atmospheric links,” *Phys. Rev. A* **99**, 053830 (2019).
- [91] C. Liorni, H. Kampermann, and D. Bruß, “Satellite-based links for quantum key distribution: beam effects and weather dependence,” *New J. Phys.* **21**, 093055 (2019).
- [92] N. Jovanovic, C. Schwab, O. Guyon, J. Lozi, N. Cvetojevic, F. Martinache, S. Leon-Saval, B. Norris, S. Gross, D. Doughty, T. Currie, and N. Takato, “Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy,” *Astronomy & Astrophysics* **604**, A122 (2017).
- [93] R. Esposito, “Power scintillations due to the wandering of the laser beam,” *Proc. IEEE* **55**, 1533–1534 (1967).
- [94] D. L. Fried, “Statistics of laser beam fade induced by pointing jitter,” *Appl. Opt.* **12**, 422–423 (1973).
- [95] P. J. Titterton, “Power reduction and fluctuations caused by narrow laser beam motion in the far field,” *Appl. Opt.* **12**, 423–425 (1973).
- [96] R. Fante, “Electromagnetic beam propagation in turbulent media,” *Proc. IEEE* **63**, 1669–1692 (1975).
- [97] M. Ghalaii and S. Pirandola, “Quantum communications in a moderate-to-strong turbulent space,” *Communications Physics* **5**, 38 (2022).
- [98] R. Fante, “Electromagnetic beam propagation in turbulent media: An update,” *Proc. IEEE* **68**, 1424–1443 (1980).

-
- [99] H. T. Yura, “Short-term average optical-beam spread in a turbulent medium,” *J. Opt. Soc. Am.* **63**, 567–572 (1973).
- [100] J. A. Dowling and P. M. Livingston, “Behavior of focused beams in atmospheric turbulence: Measurements and comments on the theory,” *J. Opt. Soc. Am.* **63**, 846–858 (1973).
- [101] M. Er-long, H. Zheng-fu, G. Shun-sheng, Z. Tao, D. Da-sheng, and G. Guang-can, “Background noise of satellite-to-ground quantum key distribution,” *New J. Phys.* **7**, 215 (2005).
- [102] B. Gross, D. Vaknin, S. V. Buldyrev, and S. Havlin, “Two transitions in spatial modular networks,” *New J. Phys.* **22**, 053002 (2020).
- [103] D. Lopez-Pajares, E. Rojas, J. A. Carral, I. Martinez-Yelmo, and J. Alvarez-Horcajo, “The disjoint multipath challenge: Multiple disjoint paths guaranteeing scalability,” *IEEE Access* **9**, 74422–74436 (2021).
- [104] A. D. Broido and A. Clauset, “Scale-free networks are rare,” *Nature Communications* **10**, 1017 (2019).
- [105] S.-H. Yook, H. Jeong, and A.-L. Barabási, “Modeling the internet’s large-scale topology,” *Proceedings of the National Academy of Sciences* **99**, 13382–13386 (2002).
- [106] A.-L. Barabási and R. Albert, “Emergence of scaling in random networks,” *Science* **286**, 509–512 (1999).
- [107] J. O. Abe, H. A. Mantar, and A. Yayimli, “k -maximally disjoint path routing algorithms for sdn,” *2015 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery*, 499–508 (2015).
- [108] J. Tapolcai, G. Rétvári, P. Babarczi, and E. R. Bérczi-Kovács, “Scalable and efficient multipath routing via redundant trees,” *IEEE Journal on Selected Areas in Communications* **37**, 982–996 (2019).
- [109] O. Lemeshko, O. Yeremenko, M. Yevdokymenko, and B. Sleiman, “Enhanced solution of the disjoint paths set calculation for secure qos routing,” in *2019 IEEE International Conference on Advanced Trends in Information Theory (ATIT)* (2019) pp. 210–213.
- [110] B. Martín, A. Sánchez, C. Beltran-Royo, and A. Duarte, “Solving the edge-disjoint paths problem using a two-stage method,” *International Transactions in Operational Research* **27**, 435–457 (2020).

-
- [111] D. Lopez-Pajares, J. Alvarez-Horcajo, E. Rojas, J. A. Carral, and I. Martinez-Yelmo, “One-shot multiple disjoint path discovery protocol (1s-mdp),” *IEEE Communications Letters* **24**, 1660–1663 (2020).
 - [112] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, “Routing entanglement in the quantum internet,” *npj Quantum Information* **5**, 25 (2019).
 - [113] B. C. Coutinho, W. J. Munro, K. Nemoto, and Y. Omar, “Robustness of noisy quantum networks,” *Communications Physics* **5**, 105 (2022).
 - [114] S. Pirandola, “Bounds for multi-end communication over quantum networks,” *Quantum Science and Technology* **4**, 045006 (2019).
 - [115] S. Pirandola, “General upper bound for conferencing keys in arbitrary quantum networks,” *IET Quantum Communication* **1**, 22–25 (2020).