Energy Efficiency in Content Delivery Networks

Niemah Izzeldin Osman

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

School of Electronic and Electrical Engineering

October 2014

The candidate confirms that the work submitted is her own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 4 is based on the work from:

N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "Reduction of energy consumption of Video-on-Demand services using cache size optimization," in *2011 Eighth International Conference on Wireless and Optical Communications Networks*, 2011, pp. 1–5.

And:

N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "Caching in green IP over WDM networks," *J. High Speed Networks (Special Issue Green Netw. Comput.*, vol. 19, no. 1, pp. 33–53, 2013.

These papers have been published jointly with my PhD supervisor Prof. Jaafar Elmirghani and my PhD advisor Dr Taisir El-Gorashi.

Chapter 5 is based on the work from:

N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "The impact of content popularity distribution on energy efficient caching," in *2013 15th International Conference on Transparent Optical Networks (ICTON)*, 2013, pp. 1–6.

This paper has been published jointly with my PhD supervisor Prof. Jaafar Elmirghani and my PhD advisor Dr Taisir El-Gorashi.

Chapter 6 and 7 are based on the work from:

N. I. Osman, T. El-Gorashi, L. Krug, and J. M. H. Elmirghani, “Energy-Efficient Future High-Definition TV,” *J. Light. Technol.*, vol. 32, no. 13, pp. 2364–2381, 2014.

This paper has been published jointly with my PhD supervisor Prof. Jaafar Elmirghani, Ms. Louise Krug, industrial collaborator and my PhD advisor Dr Taisir El-Gorashi.

To my children, this PhD thesis is dedicated to you.

Acknowledgements

First and foremost, all praise to Almighty Allah, for life, for health and the countless blessings He gave me. I would like to acknowledge my supervisor, Professor Jaafar Elmirghani for his guidance and patience through my PhD journey. His insightful advice and constant encouragement is most appreciated.

I am very grateful and thankful to the Islamic Development Bank (IDB) for fully funding my PhD. Without this scholarship, my dream of a PhD from UK would have never been possible. I also thank Sudan University of Science and Technology for their financial support.

Many thanks and gratitude goes to my colleagues in the Communication Systems and Networks group in the School of Electronic and Electronic Engineering. It was a privilege to have known them. I thank them for their company, reassurance and fruitful discussions.

I am also indebted to my special friends Salima and Nahla. Their friendship comforted me and helped me escape many hurdles. I will always cherish the wonderful time I spent with them. I would also like to thank Nigel and Caroline, who were there for me when I most needed them.

My gratitude goes to the members of the IDB Scholars Association UK. I was very fortunate to have the opportunity to meet and collaborate with such exceptional role models from around the world. I truly thank them for inspiration and for uplifting my spirit.

A big thank you goes to the residents of the Unipol student accommodation in Shay Street/Holborn Terrace, all 56 flats. It was a delight to have known them. I will miss the picnics, tea parties and celebrations that we constantly enjoyed. I would like to especially thank Shilan and her family, for the friendship, the laughter and the joy they brought into my family's life during our stay at Leeds.

I would like to express my appreciation to my beloved family back home, my father, Professor Izzeldin Osman and my mother Dr Khadiga Karrar. I don't have enough words to thank them for supporting me in all possible ways. They are the reason behind my life achievements and have always shown me by example that everything is doable. I hope I made them proud. I would also like to thank my sisters, Dr Rasha Osman and Dr Ala Osman for always looking out for me and for their care and friendship.

Finally, all my thank-yous go to my young family. My children, the apple of my eye, little Malak and baby Izzeldin. Their smiles and innocent gazes have changed the definition of life for me. I thank them for making my life so fulfilled and rewarding and for making every day so special. I dedicate my final acknowledgements to the love of my life, my dear husband Hassan Daleel, without whose kindness, sacrifices, and patience I would not have been able to pursue my dreams. I truly appreciate him standing by my side and thank him for his love, support and attention.

Niemah

Abstract

The increasing popularity of bandwidth-intensive video Internet services has positioned Content Distribution Networks (CDNs) in the limelight as the emerging provider platforms for video delivery. The goal of CDNs is to maximise the availability of content in the network while maintaining the quality of experience expected by users. This is a challenging task due to the scattered nature of video content sources and destinations. Furthermore, the high energy consumption associated with content distribution calls for developing energy-efficient solutions able to cater for the future Internet. This thesis addresses the problem of content placement and update while considering energy consumption in CDNs.

First, this work contributed a new energy-efficient caching scheme that stores the most popular content at the edge of the core network and optimises the size of cached content to minimise energy usage. It takes into account the trend of daily traffic and recommends putting inactive segments of caches in sleep-mode during off-peak hours. Our results showed that power minimisation is achieved by deploying switch-off capable caches, and the trend of active cache segments over the time of day follows the trend of traffic.

Second, the study explores different content popularity distributions and determines their influence on power consumption. The distribution of content popularity dictates the resultant cache hit ratio achieved by storing a certain number of videos. Therefore, it directly influences the power consumption of the cache. The evaluation results

indicated that under video services where the popularity of content is very diverse, the optimum solution is to store the few most popular videos in caches. In contrast, when video popularities are similar, the most power efficient scheme is either to cache the whole library or to avoid caching completely depending on the size of the video library.

Third, this thesis contributed an evaluation of the power consumption of the network under real world TV data and considering standard and high definition TV programmes. We proposed a cache replacement algorithm based on the predictable nature of TV viewings. The time-driven proactive cache replacement algorithm replaces cache contents several times a day to minimise power consumption. The algorithm achieves major power savings on top of the power reductions introduced by caching.

CDNs are expected to continue to be the backbone for Internet video applications. This work has shown that storing the right amount of popular videos in core caches reduces from 42% to 72% of network power consumption considering a range of content popularity distributions. Maintaining up-to-date cache contents reduces up to 48% and 86% of power consumption considering fixed and sleep-mode capable caches, respectively. Reducing the energy consumption of CDNs provides a valuable contribution for future green video delivery.

Table of Contents

List of Figures

List of Tables

Glossary of Terms

4oD	4 on Demand
AMPL	A Mathematical Programming Language
BBC	British Broadcasting Corporation
BT	British Telecom
BT 21CN	British Telecom 21st Century Network
BT 21CN-MVMP	BT 21CN with Multiple Video head-ends with Minimum Power Locations
BT 21CN-SV4	BT 21CN with a Single Video head-end at Node 4
BT 21CN-SVMP	BT 21CN with a Single Video head-end with Minimum Power Location
CBGA	Constraint-Based Genetic Algorithm
CCN	Content-Centric Network
CDN	Content Distribution/Delivery Network
CO_2	Carbon Dioxide
CPU	Central Processing Unit
CuTV	Catch up Television
CRS-1	Carrier Routing System – 1st generation
DNS	Domain Name Server
DSL	Digital Subscriber Line

DSLAM	Digital Subscriber Line Access Multiplexer
EDFA	Erbium Doped Fibre Amplifier
GA	Genetic Algorithm
GB	Gigabyte
Gb/s	Gigabit per Second
GDSF	Greedy Dual Size Frequency
GHz	Giga Hertz
HD	High Definition
HDTV	High Definition Television
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
ICN	Information-Centric Network
ICT	Information and Communication Technology
ILP	Integer Linear Programming
IP	Internet Protocol
IPTV	Internet Protocol Television
ISP	Internet Service Provider
ITV	Independent Television
J/b	Joules per bit
J/GB	Joules per Gigabyte
LFU	Least Frequently Used

LP	Linear Programming
LRFU	Least Recently Frequently Used
LRU	Least Recently Used
MB	Megabyte
MB/s	Megabyte per Second
Mb/s	Megabit per Second
MC-OXC	Multicast-Capable Optical Cross Connect
MILP	Mixed Integer Linear Programming
MIT	Massachusetts Institute of Technology
MPEG-2	Motion Picture Experts Group 2
MPEG-4 AVC	Motion Picture Experts Group 4 - Advanced Video Coding
NGN	Next Generation Network
NSFNET	National Science Foundation Network
NSF-MVMP	NSFNET with Multiple Video head-ends with Minimum Power Locations
NSF-SVMD	NSFNET with a Single Video head-end with Minimum Delay Location
NSF-SVMP	NSFNET with a Single Video head-end with Minimum Power Location
OLT	Optical Line Terminal
ONU	Optical Network Unit
OXC	Optical Cross Connect

P2P	Peer-to-Peer
PC	Personal Computer
PON	Passive Optical Network
QoE	Quality of Experience
QoS	Quality of Service
RAM	Random Access Memory
RRS	Request-Routing System
SD	Standard Definition
SDTV	Standard Definition Television
SMS	Short Message Service
SRAM	Static Random Access Memory
SSD	Solid State Drive
STB	Set-Top-Box
UHDTV	Ultra High Definition Television
VBB	Viewers' Bias Based technique
VCR	Videocassette Recorder
VoD	Video-on-Demand
W	Watts
WDM	Wavelength Division Multiplexing
W/GB	Watts per Gigabyte

List of Symbols

a	Piecewise linear approximation equations coefficient
a_{ktr}	Replacements piecewise linear approximation equations coefficient
Amp_{ij}	Number of amplifiers used on each fibre on the physical link from node i to j, $Amp_{ij} = \lfloor D_{ij}/S - 1 \rfloor + 2$
AP_{it}	Aggregation ports at node i at time t
B	Capacity of a wavelength
b	Piecewise linear approximation equations coefficient
b_{ktr}	Replacements piecewise linear approximation equations coefficient
C_{xyt}	Wavelengths on the virtual link from node x to y at time t
C_Req	The number of requests served from the cache
D_{ij}	Distance from node i to j
$d(x, y)$	Demand from node x to y
f_{ij}	Fibres on the physical link from node i to j
G	Virtual topology

H	Cache hit ratio
H_{it}	Hit ratio of the cache deployed at node i at time t
H_{itr}	Cache hit ratio of node i at time t when r daily replacements are performed
H_{rt}	Cache hit ratio at time t when r daily replacements are performed
K	Set of equations that approximate the convex function describing the relationship between the cache and its hit ratio
M	Cache size in GB
M_{it}	Cache size in GB deployed at node i at time t
M_{itr}	Cache size of node i at time t when r daily replacements are performed
M_{rt}	Cache size in GB at time t when r daily replacements are performed
$mMin$	Pareto distribution minimum possible value of m
N	Set of nodes
Nm_i	Set of neighbouring nodes of node i
P_i	Popularity of object i
$P(i,a,b)$	Popularity of object i during the time window from a to b

Pa	Power consumption of an amplifier
Pc	Crossover rate
Pm	Mutation rate
Pmd	Power consumption of a multiplexer/demultiplexer
Po_{it}	Power consumption of optical switch i at time t
Pp	Power consumption of a router port
Pt	Power consumption of a transponder
R	Set of possible daily content replacement frequencies
r	Number of replacements performed
Rd	Server ratio of downlink demand to regular traffic
Req_i	The number of requests for object i
Req_{it}	The number of requests for object i at time t
$RPmax_x$	Maximum router ports available to node x
Ru	Server ratio of uplink demand to regular traffic
S	Span distance between two amplifiers
T	Set of points in time
Tot	Total number of videos in the library

Tot_Req	Total number of requests
Tot_Req_t	Total number of requests at time t
$Trep$	Time duration over which cache contents are updated
V	Cache size in videos
V_{it}	Cache size at node i at time t in number of videos $(= M_{it}/Vsize)$
$Vsize$	Video size in GB
W	Number of wavelengths in a fibre
w_{ijt}^{xy}	Wavelengths on the link from node i to j, part of the virtual link from node x to y at time t
w_{ijt}	Wavelengths on the physical link from node i to j at time t
α_{rt}	Is 1 if a replacement is performed at time t with r daily replacements, 0 otherwise
β	Popularity of the most popular video in the Pareto distribution
δ_i	Is 1 if node i has a video server, 0 otherwise, $\sum_{i \in N} \delta_i = u$, where u is the total number of servers in the network
λ^{xyt}	Demand from node x to y at time t

λ_{ijt}^{xy}	Regular traffic from node i to j, part of the virtual link from node x to y at time t
λd_{ijt}^{xy}	Downlink traffic from node i to j, part of the virtual link from node x to y at time t
λu_{ijt}^{xy}	Uplink traffic from node i to j, part of the virtual link from node x to y at time t
μ_{τ}	Bimodal distribution mean of content type τ
π_{itr}	Additional download traffic to be streamed to node i at time t when r daily replacements are performed
π_{rt}	Additional download traffic to be streamed to a node due to a replacement at time t when performing r replacements
σ^2	Bimodal distribution variance
σ_{τ}^2	Bimodal distribution variance of content type τ
τ	Programme type in the Bimodal distribution (example $\tau = 1$ represents news-type programmes and $\tau = 2$ represents drama-type programmes)
Φ	Cache power consumption factor in W/GB
ω	Time distribution factor for requests

Chapter 1 Introduction

The number of Internet users has grown to over 2.8 billion users [1] and Internet traffic is expected to exceed 1 Zettabytes (1 billion Terabytes) in 2016 with an annual growth rate of about 21% [2]. In a few years, video content is estimated to account for 80% to 90% of the total IP traffic, and on average one million minutes of video content is projected to cross the Internet every second [2]. Energy consumption is predicted to become the new Internet bottleneck of communication networks. Data centres which mange and provide content are a critical part of the Internet and consume significant energy—up to 70% of the total transmission energy [3]. These alarming figures depict the increasing energy consumption of the Information and Communication Technology (ICT) industry, thus implying increasing associated carbon dioxide (CO_2) emissions. ICT's CO_2 emissions are expected to increase from 0.5 billion tonnes in 2002 to 1.4 billion tonnes in 2020 [4], exceeding 3% of global emissions [5]. The possible environmental impacts of the Internet expansion have boosted a global movement towards reducing the CO_2 footprint of ICT.

One method gaining popularity for efficiently delivering Internet video traffic is Content Delivery/Distribution Networks (CDNs). A CDN is a collection of geographically distributed caches intended to enhance the performance of Internet video services and are forecasted to deliver around 67% of all Internet video traffic by 2018 [2]. CDNs are in charge of strategically distributing video content over the network to ensure content availability and reduce video delivery delay. Caching is an effective technique for replicating content in multiple network locations. The benefit of caching is maximised when the right content is stored at the right place in the network. At the same time, cache placement should be utilised with respect to energy consumption. Today, CDNs face great challenges due to the large number of content providers, increasing user demands and the bulkiness of high-quality video content. Consequently, designing and provisioning a CDN is a complex task, particularly when considering other equipment and network constraints.

The Internet is moving towards the content-oriented era [6], [7], thus triggering research aiming to reduce the energy consumption of CDNs. Even though some recent attempts have been proposed in the literature to reduce the energy consumption of CDNs, many issues and challenges in the area remain open. This thesis has addressed the issue of reducing the high energy consumption of delivering Internet video content. The intention is to decrease the energy consumption of video transport in a core network that deploys IP over WDM (Wavelength Division Multiplexing) and video caches. A Mixed Integer Linear Programming (MILP) model is developed to minimise the power consumption of the network by optimising the size of caches at the nodes. The model considers various video services and traffic volumes. Furthermore, this thesis investigates the optimisation of daily cache updates to minimise power usage. The existence of the most up-to-date content in caches minimises the traffic passing through the network, reducing its power consumption. Optimising the number of daily cache updates saves power lost in unnecessary downloads.

1.1 Research Objectives

The dynamic nature of content popularity and user viewing behaviour makes the reduction of the power consumption of CDNs a challenging task. The aim of this thesis is to investigate the energy efficiency of content caching in optical core networks and offload content servers by storing content locally. In order to fulfil the overall goal of this work, the following objectives were set:

1. Evaluate the direct impact of caching in the core network on power consumption and explore optimum cache sizes with respect to traffic volume and power consumption parameters. In addition, explore the effect of switching off unutilised sections of caches during off-peak hours and its influence on power usage.

2. Investigate the influence of content popularity distribution on power consumption. The goal here is to estimate the required cache sizes under different video services to minimise power usage and determine the optimum location for videos based on popularity.

3. Use real TV viewing data traces to minimise the power consumption of IPTV services considering standard definition and high definition TV. This provides an accurate insight on the potential applications of the proposed models. Moreover, produce generalised results for the models by comparing two core network topologies.

4. Highlight the significance of cache updates in relation to power consumption and explore frequent cache updates to find an optimum that maximises power efficiency. This provides an ideal strategy of keeping caches up-to-date in view of the additional power consumed in streaming content when a cache update occurs.

1.2 Original Contributions

To facilitate provisioning of CDN storage, identifying the ideal sizes and contents of deployed video caches is required. This work focuses on reducing the power consumption of video delivery in core topologies. First, the power efficiency of storing content in local caches to reduce the path to retrieve content is examined using a Mixed Integer Linear Programming (MILP) model. An extension to the model considers different content popularity distributions and optimises cache sizes under each distribution. Moreover, real TV viewing data is collected and the power consumption of the network is minimised considering standard and high definition TV. Finally, the work examines cache content replacements and proposes a MILP model to optimise daily cache content updates. Following are the specific contributions of this thesis.

1.2.1 Fixed and Variable Caching

Content can be stored locally to reduce power and delay and to maximise content availability. However, the additional power consumed on storage should not exceed the power to retrieve content remotely. The goal is to identify the content to store in caches so that the video service power consumption is minimised. A MILP model is developed to minimise the power consumption of a video service employing IP over WDM by optimising fixed and variable (sleep-capable) caches at the nodes. The MILP model considers a Zipf distribution, different traffic demands and has been validated using a simulation that optimises traffic routing and a genetic algorithm that optimises cache sizes. Caching popular content reduces up to 38% of power consumption using caches of fixed sizes, and up to 42% when sleep-mode capable caches are utilised. The impact of cache updates on power efficiency has also been evaluated by recalculating the power consumption of the network in the situation where the 10 most popular videos are not present in caches. The power consumption in this case was found to be higher by up to 20%.

1.2.2 Content Popularity Distributions

The popularity of a video determines the traffic associated with requests to that video. Different video services hold videos with different popularities. Consequently, the optimum content to store in caches depends on the popularity distribution of content. An improved MILP model has been implemented to consider different video services including YouTube-like services and Video-on-Demand (VoD). The optimum cache sizes that minimise power under each service have been determined. In addition, an energy efficiency sensitivity analysis has been provided. When the popularity of videos is highly diverse, storing the few most popular videos in caches minimises power consumption, with power savings of up to 72%. In contrast, when video popularities are similar the best power efficiency is achieved by maintaining variable caches in the network.

1.2.3 Caching in Future IPTV

Real TV viewing data traces have been analysed and used to acquire a content popularity distribution for a real IPTV service. The BT 21CN topology (the 21st Century Network is a next generation network implemented by British Telecom) has been considered and compared to the NSFNET topology in terms of power consumption and optimum cache sizes that minimise power. The minimum-power MILP model has been utilised to minimise the power required to deliver standard and high definition TV programmes. Since analyses for the BT 21CN topology are not publicly available, this work has explored a next-generation national core topology suitable for comparison with other topologies available in literature.

1.2.4 Cache Content Replacements

Performing frequent cache updates maximises the useful part of stored content, but consumes additional power. A MILP model has been developed to optimise cache update frequencies for power minimisation. The power consumption of the network considering a range of update frequencies has been compared and the benefit of maintaining up-to-date caches is highlighted. The model finds the optimum number of daily cache updates that lead to maximum content availability while taking into account the additional routing components that consume power to stream up-to-date content (for cache updates). In operator-controlled video services such as streaming TV, updating cache content up to 12 times a day minimises power consumption as TV programmes are very popular only during a certain time window. Maximum savings in power consumption under 12 cache updates are up to 48% and 86% considering fixed caches and sleep-mode capable caches, respectively. The impact of regular traffic on power efficiency has also been explored by considering a range of traffic mixtures and illustrating that savings due to caching are proportional to the video component in the traffic mixture.

1.3 List of Publications

During the course of my PhD, I have published the following journal and conference papers:

1. N. I. Osman, T. El-Gorashi, L. Krug, and J. M. H. Elmirghani, "Energy-Efficient Future High-Definition TV," *J. Light. Technol.*, vol. 32, no. 13, pp. 2364–2381, 2014.

2. N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "Caching in green IP over WDM networks," *J. High Speed Networks (Special Issue Green Netw. Comput.*, vol. 19, no. 1, pp. 33–53, 2013.
3. N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "The impact of content popularity distribution on energy efficient caching," in *2013 15th International Conference on Transparent Optical Networks (ICTON)*, 2013, pp. 1–6.
4. N. I. Osman, T. El-Gorashi, and J. M. H. Elmirghani, "Reduction of energy consumption of Video-on-Demand services using cache size optimisation," in *2011 Eighth International Conference on Wireless and Optical Communications Networks*, 2011, pp. 1–5.

The work in Chapter 4 throughout Chapter 7 is based on these publications.

1.4 Thesis Outline

Apart from the introduction in Chapter 1, this thesis is organised as follows: Chapter 2 reviews Content Distribution Networks and explains the video delivery process. It also highlights content popularity distributions and video services. In addition, the chapter explains the architecture of IP over WDM and IPTV.

Chapter 3 surveys cache placement and cache replacement algorithms. It also provides a description of relevant research related to energy reduction techniques in IPTV, CDNs and VoD. The chapter is concluded by a description of the optimisation methods used in this work.

In Chapter 4, the problem of high power consumption of video services is tackled by proposing caching video content in the core nodes. A full evaluation of the proposed MILP model is provided and discussed introducing fixed and variable caching. The results illustrate that there exists an optimum cache size which minimises power consumption. This optimum size depends on the content popularity distribution, traffic volume and power consumption parameters.

Chapter 5 examines the influence of content popularity distribution on power efficiency and optimum cache sizes. It discusses the amendments to the original MILP model required to consider different popularity distributions. It also highlights the influence of the most significant network parameters including cache power consumption, video sizes, router port power consumption and IP over WDM implementations on power efficiency. It demonstrates by results that only very popular videos should be present in caches when video popularities are highly diverse. It also demonstrates that variable caches better suit distributions with similar video popularities.

Chapter 6 sheds light on delivering high definition TV content in IPTV services. It explores the dynamics of TV viewing behaviour and investigates the benefit of caching on different network arrangements. The BT 21CN topology is introduced and compared to the NSFNET topology. Results reveal that larger cache sizes are required to store high definition content to achieve comparable power savings to those under standard definition. Results also show that similar power savings are attained considering both test networks due to the similarity in the average nodal degree and hop count.

In Chapter 7, using real TV viewing data, a proactive time-based cache replacement algorithm is proposed to deliver TV content with the minimal power consumption. The efficiency of the proposed scheme is tested considering various traffic mixtures and under current and future networks. We show that replacing cache contents several times a day does not compromise on power in a TV service environment. In addition, caching in the core network is most suitable with and without component power saving capabilities under the condition that optical multicast is not fully deployed in the network.

Chapter 8 summarises the contributions of this work and suggests possible directions for future work.

Chapter 2 Background

2.1 Introduction

The Internet began as a medium for text-based applications, e.g. email, file sharing, etc., thus the main goal of Internet Service Providers (ISPs) at the time was the efficiency of service. As network technologies improved, network bandwidth increased and service cost decreased leading to an increase in the number of Internet users and the widespread development of bandwidth-intensive multimedia applications. The rapid increase in Internet traffic has resulted in the degradation of Quality of Service (QoS) due to congestion.

An initial approach in addressing this problem was to enhance Web server hardware (processor, memory, disk space, etc.). However, this method was not flexible as it eventually would require replacing the whole server system [8]. The further increase in multimedia content applications, such as YouTube and Netflix, transformed the Internet from a textual information system to a multimedia information system [9] and thus the sender and/or location of the content became irrelevant. The current challenge has become to efficiently deliver a massive amount of multimedia content to a large number of geographically distributed users.

Content Delivery Networks (CDNs) are employed to aid in the efficient delivery of multimedia to a geographically diverse user base. A CDN is a group of geographically distributed content servers that deliver web content to users. CDNs stream multimedia content to requesting users, and therefore require high bandwidth for transmission and large storage for audio and video files. The bulky size of media files requires large storage compared to other forms of data. This exhausts the network bandwidth and raises the challenge of providing multimedia applications with an acceptable QoS at an affordable cost.

Content caching is an effective technique to reduce traffic on the long path between content servers and end users. In addition, caching can result in reducing the delay caused by congestion and content unavailability. Improving these factors has great influence on enhancing users' Quality of Experience (QoE). Moreover, reducing network traffic implies reducing cost and energy consumption.

This chapter provides a detailed overview of video services, content delivery networks, content caching and content popularity distributions. In addition, it describes the architectures and implementations of IP over WDM, CDNs and IPTV. The chapter is concluded by highlighting the current challenges facing CDNs.

2.2 Video Services

Video services are the Internet applications that involve upload and/or download of video. Today, a range of video services are delivered to residential and business users over Internet Protocol television (IPTV). These services include: broadcast TV, Video-on-Demand (VoD) and time-shifted TV. Other services such as Voice-over-IP (VoIP), cable television and Web/email are also delivered over IPTV [10]. Following is a brief overview of major video services.

2.2.1 Broadcast TV

Also referred to as over-the-top [11], broadcast TV streams content to the user's Set-Top-Box (STB) providing a cable TV-like experience with the intention of real-time consumption. The TV content provider has full control over the content and therefore a continuous stream from the provider to the user is guaranteed [12].

2.2.2 Video-on-Demand

Video-on-Demand (VoD) is a one-to-one service where the user selects a video from a list of available content to watch on a Personal Computer (PC), TV or other devices. Video-on-Demand offers videocassette recorder (VCR)-like functions such as pause, rewind and fast forward, with full control of the session [13]. For example, Netflix, a major VoD provider, is a video rental service featuring one of the largest streaming content libraries and has millions of subscribers. It provides users with old and new

TV programmes and movies to download and play. iTunes, is a leading media library in music that enables users to download and play audio and video including new releases.

2.2.3 YouTube-like Services

YouTube, with over 1 billion unique visitors a month [14], possesses a huge popularity in the content delivery market. YouTube and other YouTube-like services are considered a counterpart of VoD services as a platform for video download over IPTV. However, one major distinction is that they support two-way content streaming by enabling the upload of user generated content. Therefore, when considering YouTube-like services, video download traffic as well as video upload traffic should be taken into account.

2.2.4 Time-shifted TV

Time-shifted TV includes services such as Catch-up TV (Cu-TV) and start-over TV. Cu-TV is an Internet television service that allows users to watch previously broadcasted TV programmes. Start-over TV allows replaying the current TV programme from its beginning. The total number of videos in such services is relatively fewer compared to other video services..

2.2.5 Voice-over-IP

Voice-over-Internet Protocol (VoIP) is an Internet telephone service that allows audio and video communication over IP. Other supported services include Short Message Service (SMS), fax and voice-messaging. Popular VoIP providers today are Skype and Google+ Hangouts.

2.3 Caching and Video Popularity

A cache is a storage device placed in the network to store any type of data (web objects, images, audio, video, etc.) to offload servers. Another term for a cache, *replica*, is widely used in the literature to distinguish proxy caches from a PC cache (memory blocks for temporarily storing recently accessed data). Yet, the term *cache* has widely been used for content storage devices, as they have become a major component of Internetworking. In this work, the term *cache* is used to refer to network storage devices, and the term *caching* is used to mean storing in a cache.

The main objective of caching is to reduce traffic on the communication path between the server and the users by storing videos closer to the users [15]. As caches have limited storage capacity compared to servers, they should contain the most popular content. Having the right content stored in the cache increases content availability which reduces access latency and congestion at servers [16]. This also reduces cache updates and minimises required cache sizes which results in lowering the cost of caching, justifying the financial investments.

2.3.1 Video Delivery Process

Figure 2-1 shows the steps to deliver video content when the video service deploys caches. In a typical video service, a user requests a video from the server. If the video is available at the user's serving cache, the user receives the video from the cache. This outcome is known as a *cache hit*. If the video is not available at the cache, the request is forwarded to the remote server, where the video is sent to the user, representing a *cache miss*. The cache may update its contents and replace a video with the recently requested video, known as a *cache update* [15], which is governed by the cache replacement policy.

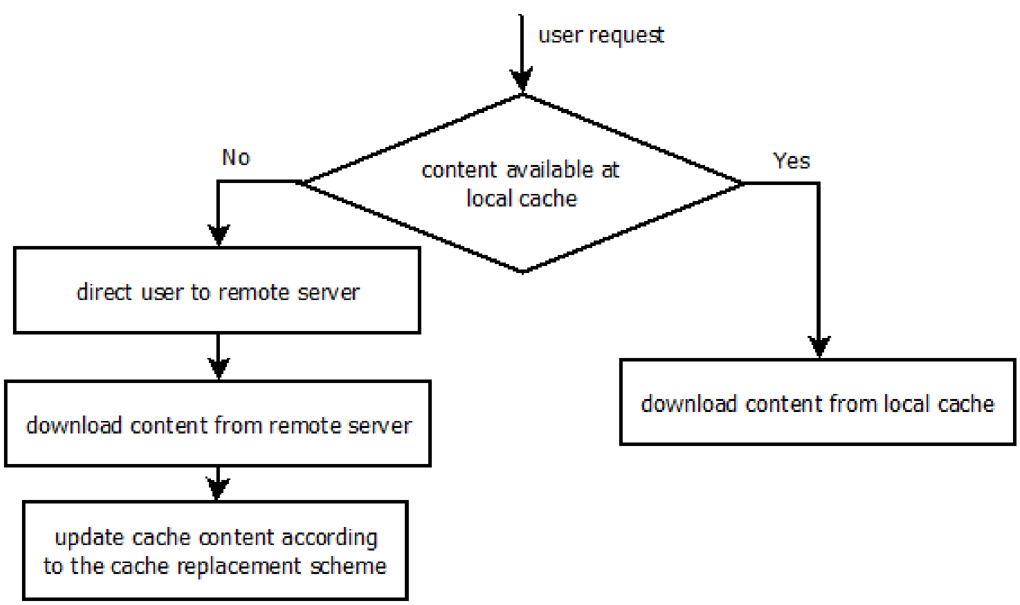

Figure 2-1: Content delivery process

The process described above is for a simple cache structure where a *cache miss* leads to direct communication with the video server. More complicated caching structures allow cache collaboration to maximise the benefit of caching. Figure 2-2 (a) shows an example of hierarchical caching where caches are arranged in a tree of multiple levels. When a cache miss occurs at a certain caching level, the request is forwarded to the next caching level (the cache in the parent node) until the requested item is found. The item is then transmitted down the request path to the user, and a copy of the item is made and stored in the caches on the path for future requests. This scheme suffers from an additional delay at each level as well as having multiple copies of a single item in multiple caches [17].

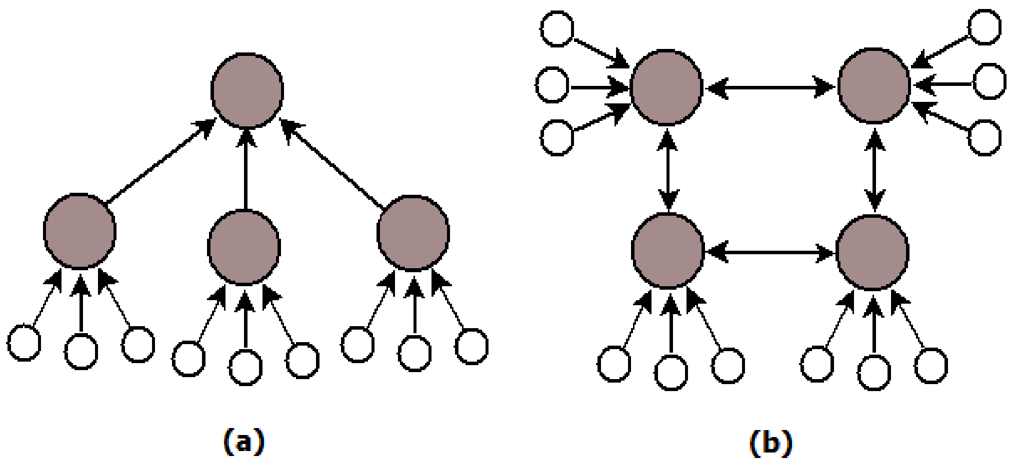

Figure 2-2 Caching structures: (a) hierarchical caching and (b) distributed caching

In a distributed caching architecture in Figure 2-2 (b), caches are allocated at the edge of the network and collaborate by distributing content among themselves and serving each other's misses. This scheme introduces additional communication and bandwidth overhead due to the necessary up-to-date cache update information that each cache needs in order to know about other cooperating caches [18]. An alternative caching structure is called en-routing, where each cache on the routing path of a request intercepts the request and delivers the requested item if found in the cache.

Otherwise, the request is forwarded along the routing path towards the video server [19]. However, due to storage capacity limitations this mechanism is used for web caching where HyperText Markup Language (HTML) pages and images are stored and not bulky multimedia files.

2.3.2 Caching Technologies

There is a large number of existing caching technologies and software that are employed by different web sites and content providers. Famous examples of these technologies are: Squid, Varnish and CoDeeN.

Squid Internet Object Cache is a group of proxy servers that cache web objects to reduce latency and congestion. Squid caches collaborate by searching neighbour caches first when a request is not found in the local cache before forwarding the request to the origin server. They also keep digests of each other's contents to reduce inter-cache signalling [20].

Varnish is an open-source HyperText Transfer Protocol (HTTP) cache that accelerates content delivery from web servers by load balancing and handling requests by separate threads. An overflow queue is provided to accommodate additional requests when the maximum active requests are reached. It is employed by high profile online newspaper sites including The Guardian and The New York Times as well as major social media sites such as Twitter and Facebook [21].

CoDeeN is an academic test bed Content Distribution Network developed at Princeton University. It consists of a network of high-performance proxy servers which act as both request redirectors and server surrogates. They collaborate with each other and provide fast and robust content delivery service to CoDeeN users [22].

2.3.3 Cache Hit Ratio

The cache hit ratio H is defined as the ratio of the number of requests served from the cache (*C_Req*) to the total number of requests (*Tot_Req*), or [23]:

$$H = C\_Req/Tot\_Req \qquad (2\text{-}1)$$

Each video in the service holds a certain value reflecting its popularity amongst other objects in the service. This popularity is calculated as [24]:

$$P_i = Req_i/Tot\_Req \qquad (2\text{-}2)$$

where Req_i is the number of requests for object i. The popularity of a video is also the probability that a request will be made to this video. The hit ratio can also be calculated from the summation of the popularities of cached videos, or [24]:

$$H = \sum_{i=1}^{V} P_i \qquad (2\text{-}3)$$

where P_i is the popularity of the i^{th} video in the cache, and V is the cache size in videos. The hit ratio of a cache depends on the capacity of the cache, the average size of videos and the popularity distribution of the content. The download demand between a node and a video server represents *1-H* of the total demand, making the cache hit ratio an influential parameter.

2.4 Content Popularity Distributions

Depending on the nature of the service, the popularity of videos follows different distributions. In this section, four content popularity distributions are discussed: the Zipf, Pareto, Bimodal and Equal Popularity distributions.

2.4.1 Zipf Distribution

The Zipf distribution is considered the best approximation that represents web access. Its long tail defines web objects where a very large number of objects exist, few of which are popular, and the remaining majority has little popularity. This pattern well describes the user access behaviour of YouTube with a limited number of hot videos experiencing high download rates and a long list of available videos with small number of hits [25]. The authors of [25] characterise the traffic of YouTube by analysing usage patterns, file properties and popularity and referencing characteristics. They found that YouTube video popularity follows a Zipf-like distribution. They also state that the use of caches can reduce network demands, as over 50% of video requests are for previously requested videos. Other studies on YouTube-like traffic follow their approach [26], [27]. Given that objects are arranged by popularity (the most popular object is number 1), the approximation in [28] defines the popularity P_i of the i^{th} object in the rank following a Zipf distribution by:

$$P_i = 1/(i \,.\ln Tot) \tag{2-4}$$

where Tot is the total number of objects.

2.4.2 Bimodal Distribution

In a CuTV service, TV programmes may be divided into two or more categories with respect to programme type. Each category has its own popularity rank and distribution with very popular programmes at the peak of each category. Such a service is best described by a Bimodal distribution. In a simple Bimodal distribution, representing the popularity of the j^{th} category by a Gaussian distribution, the popularity P_i of each TV programme i is given by [29]:

$$P_i = e^{-(i-\mu_j)^2/2\sigma^2} / \sqrt{2\pi\sigma_j^2} \qquad (2\text{-}5)$$

where μ_j is the mean of the j^{th} category, which defines the location of the peak, and σ_j^2 is the variance of the j^{th} category which gives the width of the distribution.

2.4.3 Pareto Distribution

The Pareto distribution is used to describe the popularity of content that is not as diverse in video popularities as the Bimodal distribution, but does not contain as long a tail as the Zipf distribution. The Pareto distribution can represent a VoD service where a long list of movies with different popularities is available for users on demand. The popularity P_i of a video i under the Pareto distribution is given as [30]:

$$P_i = \beta . iMin^{\beta} / i^{\beta+1}, \qquad \forall i > i_m \qquad (2\text{-}6)$$

where $iMin$ is the minimum possible value of i and β is the most popular video.

Both Zipf and Pareto distributions are widely used to describe access to web content and multimedia. Although the Zipf distribution is considered a discrete counterpart of the Pareto distribution [30] and both represent heavy-tailed distributions, interesting distinctions between the two distributions motivated this work to consider both. The popularity of videos for content where the popularity follows a Zipf distribution depends on the total number of videos in the library. Here the influence of the library size can be considered. The Pareto distribution explicitly allows for more control on the popularity of the most popular video. This comes handy in services where the popularity of the most popular video(s) is known, and consequently the effect of this value can be evaluated.

2.4.4 Equal Popularity Distribution

The Equal Popularity distribution implies having a set of videos where the probability of a request for any of the videos is equally likely. The popularity of each video under this distribution solely relies on the total number of videos in the library. Therefore, the popularity P_i of a video i is given as [29]:

$$P_i = 1/Tot \tag{2-7}$$

where Tot is the total number of videos in the library.

2.5 IP over WDM

Wave Division Multiplexing (WDM) allows transmitting data over multiple wavelengths on the same optical fibre. This feature increases the capacity of the optical network without laying more fibres. As a result, WDM is an attractive solution for optical network expansion and is suitable for accommodating increasing Internet traffic demands. On the other hand, IP is expected to continue to be the revenue generating layer as the choice for all end user communication. Therefore, IP over WDM is expected to play an important role in next generation optical Internet, offering low cost efficient service delivery for a large audience [31].

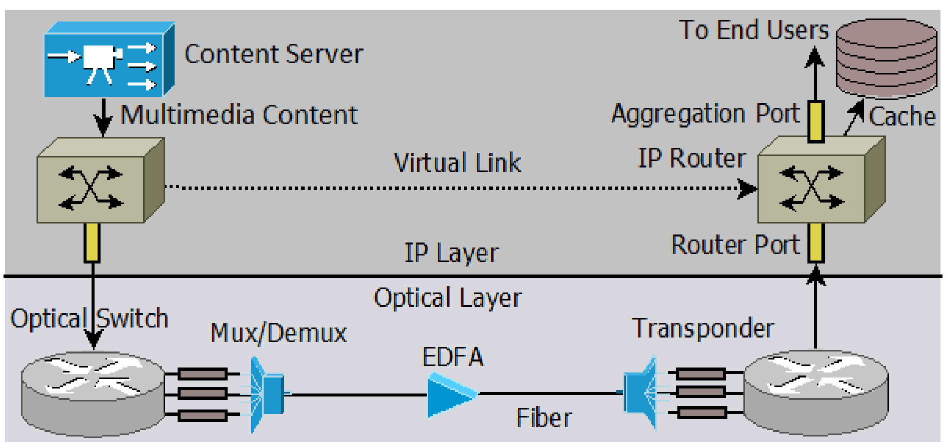

Figure 2-3: A Simple 2-node IP over WDM network architecture with content delivery

2.5.1 Network Architecture

A simple architecture of an IP over WDM network with content delivery consisting of two nodes is shown in Figure 2-3. A core router in the IP layer aggregates traffic demands generated at the access network which enter the core network through aggregation ports. Aggregation ports are utilised by the router for receiving and delivering user requests. A physical link between a pair of optical switches includes one or more fibres. Each fibre is equipped with a pair of multiplexer/demultiplexers to join/split wavelength signals in addition to a number of Erbium Doped Fibre Amplifiers (EDFAs) depending on the length of the link. The number of EDFAs required on each fibre is: $\lfloor D_{ij}/S - 1 \rfloor + 2$, where D_{ij} is the distance between node i and j, S is the EDFA spacing and $\lfloor \cdot \rfloor$ is the integer part [32]. Each occupied wavelength on a fibre requires a pair of transponders for wavelength modulation and regeneration.

2.5.2 Implementation

A lightpath is a point-to-point all optical wavelength channel connecting a source and a destination [33]. Considering the manner that a lightpath traverses IP routers, two strategies can be employed when implementing IP over WDM.

2.5.2.1 Lightpath non-Bypass

Under lightpath non-bypass, all the IP routers of intermediate nodes on the path are traversed, engaging IP router ports on each intermediate node. Most networks currently employ lightpath non-bypass routing.

2.5.2.2 Lightpath Bypass

When lightpath bypass is considered, traffic passing through an intermediate optical node is forwarded to the next optical node bypassing the IP router of the intermediate node. This approach is considered more energy efficient, as IP router ports consume relatively high power. Although lightpath bypass can result in higher power savings, it does not enable the operator to access data at intermediate nodes for security deep packet inspection and correction.

2.6 Overview of Content Delivery Networks

A Content Delivery/Distribution Network (CDN) is an overlay network that consists of a collection of strategically located servers that replicate and distribute content providing more reliability and better performance [8], [34]. The first CDN, Akamai [35], was founded in 1999 from a research project at Massachusetts Institute of Technology (MIT) aimed at solving the problem of flash crowds (server failure due to sudden increase in traffic). CDNs, such as Akamai and Amazon CloudFront [36] charge content providers for content delivery and allow them full control over how content is cached. The content is moved from the content provider and distributed over the CDN servers, and users access content from the CDN. CDNs represent a convincing solution for content providers as they take over the responsibility of hosting and distributing content. The large number of geographically distributed content servers owned by CDNs (thousands) offer high availability, easy access and less delay for users [37].

In a conventional network, a user request is forwarded to the Domain Name Server (DNS) which translates the website's name into its IP address and the request is redirected to the origin server which responds to the user request. In a CDN architecture, shown in Figure 2-4, the content distributer is responsible for distributing the content of the origin server between the CDN content servers. When a user request is received at the DNS, it is forwarded to the Request-Routing System (RRS) which is responsible for directing users to their corresponding CDN content server [8]. The selection of the content server depends on the availability of content, the distance between the user and the content server, delivery cost and load balancing. In order to optimise content delivery, the CDN performs network measurements to maintain up-to-date information about content location and network condition.

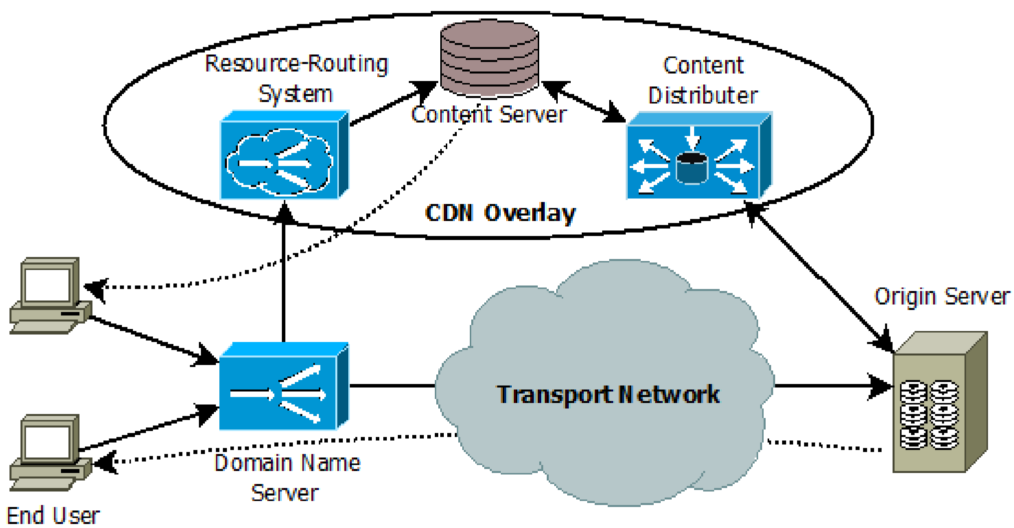

Figure 2-4: A CDN architecture

2.7 IPTV Streaming and Structure

IPTV is a system that offers streaming video content over IP networks. Typical services offered by IPTV include Broadcast TV, time-shifted TV and Video-on-Demand

(VoD). A typical IPTV network consists of three main parts: the content delivery network, the transport network and the access network. Following is an explanation of a typical IPTV network architecture, shown in Figure 2-5.

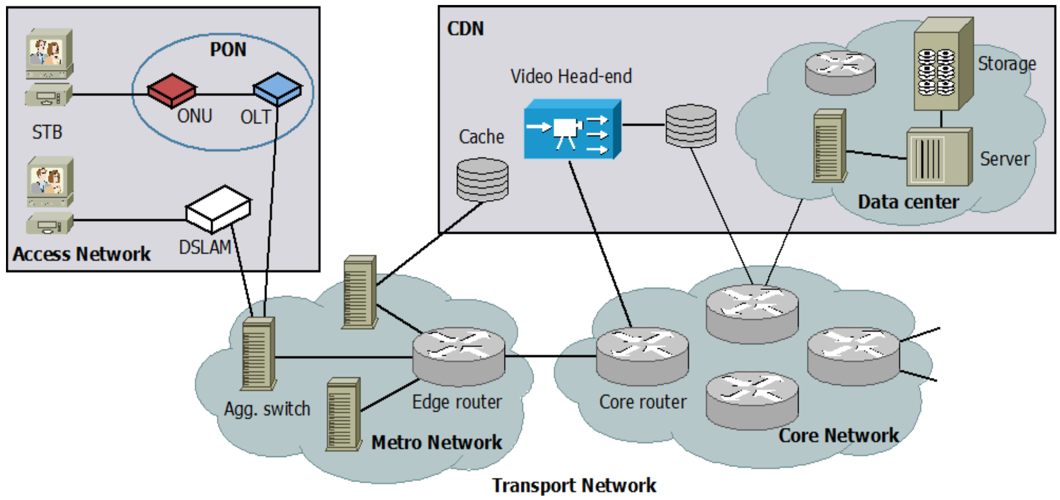

Figure 2-5: IPTV network architecture

2.7.1 Content Delivery Network

The content delivery network (CDN) includes the data centre which is made up of the content storage and server components. The CDN also includes the video head-end, responsible for transforming video streams into digital compressed streams that are encapsulated in IP packets and injected in the core network. These video streams are generally live TV broadcast and VoD libraries. One encoder is required at the video head-end for each TV channel, and possible output streams are standard definition, high definition and picture-in-picture [12], [38]. This part of the network also includes caches, which are deployed at one or more levels in the core, aggregation and/or access network.

2.7.2 Transport Network

All core switching, distribution and aggregation is considered a part of the transport network. The transport network consists of core and aggregation switches and edge routers. The backbone utilises WDM optical transport technology [12], [39].

2.7.3 Access Network

Passive Optical Networks (PONs) are becoming the leading choice for access as the popularity of Internet multimedia services grows. PONs provide attractive features including simplicity, cost and energy effectiveness and offer a bandwidth between 256 kb/s and over 1 Gb/s per user [40]. A PON is made up of the Optical Line Terminal (OLT), the Passive Optical Splitter and the Optical Network Unit (ONU). The OLT is the interface between the core network and the PON. It is connected to a number of ONUs through a splitter which evenly splits the signal from the input to the output. The ONU converts optical signals to electrical signals to be transmitted to individual users.

The Set-Top-Box (STB), connected to a TV, is one example of the end user equipment in an IPTV system. The STB has an embedded operating system, decodes the video to be viewed on TV and may include an Internet browser [12]. The STB can be connected directly to the PON or to a Digital Subscriber Line Access Multiplexer (DSLAM) through DSL (Digital Subscriber Line). The DSLAM connects multiple user DSL interfaces to a high-speed digital communication channel using multiplexing techniques.

2.8 CDN Challenges

CDNs face a number of challenges mainly due to the high demand for content and the dynamic nature of video delivery. Following is a summary of these challenges.

1. **Content delivery cost:** The massive commercial competition for content delivery has forced CDN operators to deploy a large number of content servers around the globe, increasing the cost of service. Consequently, the revenue of content delivery is monopolised by large CDN companies and only major content providers can enjoy CDN services [41]. The challenge is to bring down the cost of operating CDNs to expand the market.
2. **Content delivery energy consumption:** Network provisioning and operation is a difficult task for ISPs as the traffic fluctuates depending on the change in provided content and content popularity. As a result, ISPs provision for maximum traffic which consumes excessive energy. The challenge is to develop a model that accurately predicts network traffic and required network equipment so that minimum devices are utilised to minimise energy consumption.
3. **User mis-location:** In the CDN content delivery process, the Request-Routing System (RRS) receives requests from the underlying network Domain Name Server (DNS) and not from end users. The user is assumed to be located close to the requesting DNS (which is not always true) and therefore directed to a content server near the DNS. This results in a mis-location problem which results in content being delivered to the user from a

remote content server [41]. The challenge is to incorporate the user's IP address in the request forwarded to the RRS to improve the decision of the content server selection.

4. **Network bottlenecks:** Network bottlenecks represent a challenge facing CDNs as traffic flow is not always optimised. The reason is that despite the excessive network measurements that CDNs perform, it is a challenging task to maintain accurate up-to-date information regarding the network condition due to traffic fluctuations and distributed sources and destinations [42].

Today, the most popular CDN carries a significant part of the total daily Internet traffic, estimated between 15% and 30% [43]. Therefore addressing these challenges can have a great impact on the overall performance of video delivery over the Internet. An essential concern is the high energy consumption of CDNs which needs to be reduced in order to reduce CO_2 emissions and operating costs. This thesis addresses the high energy consumption of CDNs by caching content in the core network to minimise the energy consumption of video delivery. Chapter 4 to Chapter 7 explain in details the proposed models.

2.9 Summary

Content Delivery Networks have emerged to support the fast delivery of the growing demand for digital content to a large number of geographically distributed users. This chapter has introduced CDNs and the different video services it delivers and overviewed caching and content popularity distributions. It has also explained the importance of the cache hit ratio and its significance in estimating network performance. IP over WDM, CDN and IPTV architectures were discussed and the video delivery process was overviewed.

The next chapter surveys cache placement and cache replacement algorithms as well as the related work on energy reduction attempts in CDNs, IPTV and VoD in the literature.

Chapter 3 Related Work

3.1 Introduction

The work proposed in this thesis reduces the energy consumption of content delivery through optimum cache placement and content replacement. This chapter summarises previous research work related to caching and the energy efficiency of content delivery. The chapter opens with related work on cache placement and replacement algorithms. Afterwards, a section is devoted to green Internet and energy efficient optical networks. The following section surveys energy reduction strategies in content delivery. The chapter is concluded by a section describing the optimisation methods used in this thesis.

3.2 Cache Placement Algorithms

The objective of cache placement algorithms is location, size and content optimisation. Location optimisation specifies the location in which to implement caches in the network to meet certain criteria. Cache size optimisation decides how much capacity should be granted to each cache. Cache content optimisation identifies the best cache(s) to store each video. In the following, an overview of cache placement strategies proposed in the literature is presented.

3.2.1 Hierarchical Cache Optimisation

Hierarchical cache optimisation is concerned with optimising the size of caches implemented at each level in the caching hierarchy (core, access, etc.). There exists a trade-off between transport and caching costs and this trade-off is considered when optimising cache sizes. The cost of transport is related to the length of the path from the source to the destination. Therefore the influence of this trade-off is different at each network level.

The work in [44] and [45] evaluates this trade-off between transport and caching costs and its influence on optimum cache sizes at each network level. Having caches closer to clients reduces traffic on all links upstream of caches. Nevertheless, caching in the access network requires deploying a large number of caching equipment, increasing operating costs. The results in [44] and [45] show that the trade-off leads to having small caches closer to end users to regulate cost and larger caches are located further from the access network. Moreover, the smaller caches located closer to end users are populated with the most popular content to maximise traffic reduction due to caching, and less popular content is to be stored in caches that are higher in the caching hierarchy.

3.2.2 Combined CDN and Proxy Server Solution

The authors of [46] propose a caching method that increases the available storage area by making use of the caching space provided by proxy servers (intermediary servers between two networks) as well as the CDN. They propose a hybrid approach to reduce user delay. Their results show that the average response time is reduced by up to 40% compared to the stand-alone version due to the sufficient number of caches in the network containing the most popular content. They claim that their proposed hybrid solution does not introduce high administrative overhead for the CDN system. However, the hybrid scheme outperforms traditional caching by only 5%-10%, which is marginal considering the additional administrative overhead introduced by proxy-CDN collaboration.

3.2.3 Minimum Traversed Caches

The authors in [47] formulate the content distribution problem as an optimisation problem in which content is replicated in the caches such that the number of hops to retrieve requested items is minimised. This problem assumes collaboration and communication between CDN servers, leading to developing a model for a Peer-to-Peer (P2P) network, (A P2P network is a distributed architecture where peers provide and request content). The work considers four heuristics that utilise available data differently. These heuristics are:

1- Random: assign items to caches randomly
2- Popularity: populate each cache with the most popular content

3- Greedy-Single: use a cost function to cache the furthest items from each node with no collaboration between nodes

4- Greedy-Global: use a cost function to cache the furthest item from a node while using the new placement for the next cost function calculation

Results show that the Greedy-Global algorithm has improved performance by up to 24% compared to other heuristics. Content is made available at its closest possible cache for all the nodes through the cooperation between CDN servers provided by this method results.

Although this work provides a useful strategy to improve performance, it is only suitable when the number of objects is small since the problem becomes NP-complete with a large number of objects. In addition, the optimisation minimises the number of hops without considering the cost of storage against the cost of transport to optimise cache contents.

3.3 Cache Replacement Algorithms

The popularity of videos in a video service decays over time due to the release of new videos or the time-related viewing patterns of videos. As a result, the contents of caches become less popular and must be updated periodically to maintain the most popular videos. A cache replacement algorithm is the process in charge of selecting an item from the cache to be removed and substituted with a more popular item. The main goal of cache replacements is to maximise the cache hit ratio in order to improve other performance measurements.

The prior work in content replacements considers two main types of caching algorithms:

1. **Reactive**: In reactive cache replacements prior knowledge of future requests is not available. The work in [48] and [52] describes reactive caching algorithms that are invoked when the cache is full. A request for data not in the cache results in an item being evicted from the cache. The aim of this type of caching algorithm is to decide which item to evict from the cache. Examples of reactive cache replacement algorithms are the Viewers' Biased Based, LRFU, p-Based LRFU-k, Quality Based and Probabilistic In-Network caching algorithms, described below.
2. **Proactive**: Full or partial prior knowledge of requests can allow a caching algorithm to proactively update cache contents ahead of time [53], [54]. It is difficult to collect real request traces from video services and hard to predict future requests. Therefore, most prior works propose reactive algorithms. An example of a proactive cache replacement algorithm is the Reuse Time-Based algorithm, described below.

Cache replacement algorithms differ in the parameters used to select the item to be evicted from the cache and the way these parameters are applied. Following is an overview of the cache replacement algorithms applied in the literature.

3.3.1 Least Frequency Used

The Least Frequently Used (LFU) algorithm is a simple algorithm that has been used to manage computer memory. Its concept is that items which are accessed the most are the most likely to be requested again in the near future. The algorithm keeps track of the number of accesses for each item in the cache using a counter. When a new item is requested and the cache is full, the item with the smallest counter is removed from the cache [55].

The LFU is considered one of the simplest algorithms and is easy to implement, however it suffers from drawbacks. One problem is that considering a situation where an item was highly accessed in the past, its counter will hold a large value. Consequently, the item will be granted a place in the cache, although it is not being accessed at present. Another problem is that newly inserted items in the cache start with a small counter value, and therefore are most likely to be evicted first, even though they might experience many hits over a period of time.

3.3.2 Least Recently Used

The Least Recently Used (LRU) algorithm keeps track of the age of stored items. The item nominated to be evicted is the *oldest* item. LRU performance increases as the cache size increases, but suffers from its bias against items that are occasionally but consistently accessed [55]. Due to their simplicity, LRU and LFU algorithms are considered the base for many improved cache replacement algorithms and a reference for performance evaluation [48], [49], [50], [53].

3.3.3 Viewers' Bias Based Technique

Cache management algorithms can be improved by maximising the number of incorporated parameters used in the cache replacement decision. The Viewers' Bias Based technique (VBB) proposed in [48] considers the semantic content of the video such as fiction, horror, action, comedy, etc. as the most important feature influencing the viewers' decision to view a video. The semantic content of videos is measured at different times and places to make the algorithm adaptive. The VBB cache replacement algorithm steps are shown in Figure 3-1.

```
If (enough cache space for requested video)
    store in cache
Else
    Repeat
        calculate a priority for all videos
        mark video with lowest priority: can_be_deleted
    Until (enough space can be freed in the cache)
    If (enough space is freed in the cache)
        all videos marked can_be_deleted are removed
        the newly requested video is stored
    Else
        the newly requested video cannot be stored
        all marked videos are unmarked
```

Figure 3-1 Viewers' Bias Based algorithm steps

The performance of VBB is compared to LRU and LFU in [48] through simulation. For the evaluation, 400 videos are considered where 50 of the videos belong to the bias set. Results show that as the likelihood of a viewer choosing a video from the bias set increases, the cache hit ratio increases under all considered algorithms. When the likelihood is between 10% and 70%, the algorithm results in higher cache hit ratios, as it favours videos from the bias set to remain in the cache. However, the algorithm outperforms LFU and LRU only under a small mean interarrival time (<30 requests/second). The performance of VBB is similar to LFU and LRU for higher interarrival times.

3.3.4 LRFU and p-Based LRFU-k Algorithms

The authors of [49] propose the Least Recently/Frequently Used (LRFU) algorithm as an improvement to the LRU and LFU algorithms. The LRFU algorithm uses a weighted function to determine how much more favour is given to recently accessed items over older accessed items. Results show that the LRFU algorithm slightly increases the cache hit ratio compared to a number of LRU and LFU-based algorithms. Nevertheless, this algorithm requires infinite calculation memory as it calculates the weight for all past access history each time a cache replacement occurs.

The algorithm proposed in [50], p-based LRFU-k scheme, solves the LRFU problem by using the period of time span (p) and the number of time spans (k). The period of time span (p) determines how far back the algorithm keeps track of the number of references for each time span, while the number of time spans (k) allows adding more weight to recent time spans over older time spans.

The results show that a 30 minute-based LRFU-12 scheme results in the most increase in the cache hit ratio compared to LFU and LRU. However, the average hit ratio for a cache size of 100 items increased by only 6.5% and 4.5% compared to LFU and LRU algorithms, respectively.

3.3.5 Quality-Based Video Caching

In [51], the authors evaluate replacement strategies for quality-based video caches considering the LRU algorithm and a simplified Greedy Dual Size Frequency (GDSF) algorithm. They explore vertical and horizontal replacements where in vertical replacements all versions of the least popular video in the cache are replaced first. Under horizontal replacements the versions with the highest quality are removed first for all cached videos.

Simulations showed that horizontal replacements improve the cache hit-ratio by up to 22%. However, the performance of vertical replacements is similar to GDSF and slightly worse than LRU.

3.3.6 Reuse Time-Based Caching Algorithm

The authors in [53] treat movies as an ordered sequence of segments to predict future requests in a VoD service. They propose a caching algorithm that predicts the time that cached segments will be reused in the future. The segment with the furthest reuse time is evicted and replaced with a successive segment of a movie that is being viewed by active users.

The results in [53] show that the proposed algorithm increases the cache hit ratio compared to LFU and LRU. Nevertheless, they assume that once a user requests the first segment of a movie, he will continue watching the movie to the end and request all other movie segments in order. The probability of users skipping segments and interrupted viewing is not considered when calculating reuse times.

3.3.7 Probabilistic In-Network Caching for Information-Centric Networks

The work in [52] proposes a caching scheme, ProbCache, to reduce content redundancy in in-network caches (caching equipment provided at each content router in the network [56]). ProbCache is a probabilistic algorithm for distributed content delivered over a path of multiple caches in an ICN (Information-Centric Networking is a direction of moving the Internet to a content distribution architecture).

The algorithm makes the caching decision by considering other flows on the caching path and the amount of traffic associated with each caching request on the path. Based on this, content from different flows is fairly multiplexed in caches in the shared path. The proposed algorithm reduces up to 20% of requests to the origin server and cache eviction by 10% compared to universal caching.

The previously described cache placement and replacement algorithms do not consider the energy consumption of the network. They focus on maximising cache hit ratios to improve network performance. The energy efficiency of the network can be incorporated into these techniques by considering the power consumption of storage and transport. The objective is to minimise the network energy usage by optimising cache hit ratios. Other network factors including semantic video contents and video quality can be considered.

3.4 Green Internet and Optical Networks

This section reviews the academic and industrial efforts to reduce the energy consumption of the Internet as well as optical networks.

3.4.1 The Green Internet

The *Green Internet* concept was initiated in [57] in 2003 by proposing switching off unutilised network equipment and traffic aggregation. Other studies followed the lead and proposed more detailed energy reduction strategies. The authors of [58] presented a model to estimate the power consumption of the Internet. A thorough investigation of the energy efficiency of optical networks is provided in [59] reviewing energy-efficient protocols and network architectures. It recommends that existing optical network techniques including traffic grooming and protection should be re-evaluated with respect to energy consumption. With more focus on Wavelength Division Multiplexing (WDM), the works in [32], [60], [61] propose models to evaluate the energy consumption of optical core networks considering switch on/off techniques and optical bypass. Renewable energy sources such as wind and solar are deployed in networks as a green source of energy. The work in [62] and [63] focuses on minimising the non-renewable energy consumption of the network to reduce total CO_2 emissions.

The industry has significantly contributed to the *Green Internet* movement as well. British Telecom (BT) has recently become one of the largest companies in the world to source 100% of its electricity from renewable energy [64]. The GreenTouch initiative brings together industrial, academic and governmental bodies to reduce the energy consumption of networks by a thousand times by 2015 [65]. In addition, the Green Grid is an industry consortium of over 170 member companies including policy-makers, technology providers and utility companies. These companies collaborate to improve the resource efficiency of data centres, having power as a major goal [66].

3.4.2 Energy Efficient Optical Networks

A comprehensive survey in [59] on energy minimisation efforts in optical networks divides energy consumption approaches into four categories: (i) turning on/off network components, (ii) energy efficient network design, (iii) energy efficient IP packet forwarding, and (iv) green routing. When turning on/off network components, an entire node in the core network can selectively be turned off if it is unused or experiences low traffic. However, this approach results in additional control, management and operation overheads which can become significant. Similarly, inactive links, line cards and chassis can be shut down when traffic is low. Energy efficient network design implies developing energy efficient architectures during network design. This involves the selection of equipment with respect to capacity and energy consumption, placement of equipment and configuration.

Optimising the size of IP packets results in energy efficient IP packet forwarding, as larger IP packets consume less energy when transferred between routers compared to smaller packets. However, larger packets experience longer delays, and therefore packet size optimisation considers the trade-off between energy and delay. Pipeline forwarding, which is a time-driven packet switching scheme, is an approach for energy efficient IP packet forwarding. The optical implementation of pipeline forwarding allows packets to flow faster through the network, saving significant energy on transport. Green routing uses the energy consumption of network components as the optimisation objective. Traffic is routed through the most energy efficient route rather than the shortest route, and energy is saved by energy aware utilisation of line cards and chassis.

In what follows, the greening efforts for IPTV, CDNs and VoD are explored. The reader is referred to [59] for detailed information on energy-efficient attempts in optical networks in general.

3.5 Energy Reduction Strategies in Content Delivery

Continuous efforts have been made to reduce the energy consumption of the Internet and many issues in the field are yet to be solved. The rise in Internet video traffic increases the energy consumption of the network equipment that delivers that traffic [57], [58]. Many studies have proposed various approaches to reduce the energy consumption of the Internet at different network domains. Following is a survey of previously proposed energy reduction approaches in IPTV, CDNs and VoD.

3.5.1 Energy Reduction in IPTV

The design of IPTV networks has become a matter of increasing concern, as the high traffic growth rate of IPTV is coupled with the bandwidth consuming nature of the service. While capacity, cost and QoS remain important, energy consumption is expected to become an important driver and constraint in the design of the future Internet. Different attempts have been proposed in the literature to reduce the power consumption of IPTV services with particular focus on equipment energy efficiency and network traffic. Following is an overview of different methods for energy reduction in IPTV systems.

3.5.1.1 IPTV Caching Considering Viewer Behaviour

Caching can be provided at a single level in the network in the core, aggregation or the access network. Multi-level caching or hierarchical caching involves local storage of content at more than one part in the network. The authors of [67] optimise the number of network levels to employ caches depending on cache sizes and user viewing behaviour. The size of the caches determines their power consumption and the number of videos they can store, whereas the similarity and variation in viewer preference decides how much of user requests the cache can serve.

When viewing behaviour is similar, viewers tend to watch the same videos for comparable lengths of time. In this case, a single level of caching is sufficient to reduce the amount of energy used in the network substantially. As viewer preference becomes more diverse, the most energy efficient network design requires multi-level caching.

Although the results shown in [67] calculate the energy savings due to caching at different network levels, the higher power savings achieved assume a high similarity in viewer behaviour. The assumption states that the residents of a community are most likely to request the same video content over a similar session length. In reality, the similarity in viewer behaviour is expected to be low, and therefore the actual energy savings achieved using this technique would be moderate.

3.5.1.2 Statistics Based Caching

In [68], the caching decision is based on statistics from past requests where an algorithm calculates the optimum number of minutes from the start of the requested video that should be placed in the cache in order to minimise power consumption. The results show that the optimum chunk size to cache is larger for highly popular videos. They also show that when the number of requests increases, caching chunks from less popular videos becomes energy-efficient.

The results shown in [68] consider only a limited number of movies (1 to 100) and a small average request rate (a maximum of 18 requests/second). In order to produce generalised results, a larger movie library and higher request rates should be considered.

3.5.1.3 Selective Pre-Joining of TV Channels

In current IPTV systems, the operator distributes all TV channels to all DSLAMs to avoid channel switching delay. However, only a limited number of TV channels are watched by the vast majority of audience (80%-90% of channels [69], [70]). Therefore, pre-joining all channels is considered a waste of bandwidth (pre-joining a channel is including the channel in a specific multicast tree). The authors of [69] propose the pre-joining of a subset of channels (the most popular) instead of the full collection. This approach considerably reduces the network bandwidth, while increasing channel switching delay for only a small number of requests. Reductions in network bandwidth can be translated into energy savings by reducing the number of powered-on network components.

The results in [69] show that even though it is possible to reduce network bandwidth, the reduction in energy consumption due to selective pre-joining of channels is insignificant. They predict that savings in energy consumption will become much more significant when more channels become available in High Definition format (HDTV).

3.5.1.4 IP Flow Aggregation with Multiple Line Rates

Multicast IP over light-trees (a light-tree is a point-to-multipoint generalisation of a lightpath [33]) is an IP core network model that offers bandwidth reduction and improved performance. The work in [71] considers multicast IP flows for programme delivery in an IPTV network. The objective of the study is to optimise IP flows into light-trees and to find the appropriate line rate for each tree to minimise energy consumption. The results show that the high line rates are preferred by light-trees as long as the additional energy they consume is proportional to the extra rate they provide.

The work presented in [71] assumes optical multicast and therefore the complexity of the proposed model is high and increases with the increase in the number of established light-trees.

3.5.2 Energy Efficient CDNs

In CDNs and other related networks, a number of factors influence the energy consumption of video streaming. The distance between the cache and the user, the network domain at which caches are implemented and the cache configuration determine the energy consumption associated with streaming a video from the cache to the user. Continuous attempts are made to reduce the energy consumption of CDNs and related networks. Following is a summary of these attempts.

3.5.2.1 Energy Efficiency of Content Delivery Architectures

In [72], the energy consumption of several content delivery architectures is evaluated. These architectures include: a decentralised CDN, a Content-Centric Network (CCN) and a centralised CDN using optical bypass (CCN is a communications architecture that emphasises content rather than location [73]). The study takes into account the energy consumption of the core, edge and access network equipment and considers various library sizes.

The results recommend deploying CCNs to deliver small size libraries since they are more energy efficient compared to CDNs which consume less energy under larger libraries. Results also show that CCNs are more energy efficient to deliver the most popular content whereas optical bypass is more energy efficient when delivering less popular content.

3.5.2.2 In-Network Caching for Information Centric Networks

The study of energy efficient content delivery presented in [74] considers Information Centric Networks (ICNs) that support multicast and in-network caching. It offers two caching solutions: an offline solution based on prior knowledge of user requests and an online solution where the caching decision is based on energy minimisation.

The offline solution provides a lower bound for energy consumption, while the online solution consumes 67% less energy compared to LRU and LFU and 28% more energy than the lower bound.

Although the study considers minimizing the energy consumption of storage as well as transport, the reported results do not show the resulting cache hit ratio for the assumed fixed cache size. In addition, the study does not justify how caching 40% of content in each node provides the most energy efficient solution. Results should demonstrate the benefit of caching by showing the traffic reduction due to caching and the additional incurred costs of storage.

3.5.2.3 Energy Efficient Trade-offs among P2P and CDNs

The work in [39] compares the energy consumption of Internet TV in the transport network considering Peer-to-Peer (P2P) and CDN architectures. The results show that while P2P minimises the energy usage of the ISP, this reduced energy is migrated to users' Set-Top-Boxes. On the other hand, the CDN architecture results in the minimum overall energy consumption.

The comparison in [39] considers the energy consumption of downloading movies without taking into account the energy consumption of storing and updating movies. Considering the energy consumption of storage is necessary in the comparison as the CDN architecture consumes energy on a large number of caches whereas a P2P relies on user's storage capacities.

3.5.2.4 ISP and CP Cooperation

The authors of [75] propose a model to reduce the energy consumption of the backbone network based on full cooperation between the Internet Service Provider (ISP) and the Content Provider (CP). Results show that the proposed model can reduce the energy consumption by up to 71% on real ISP topologies. Nevertheless, they recognise the inconvenience of the idea, as ISPs and CPs do not openly share sensible data (network topology, server load, traffic, etc.). They therefore introduce another model in [76] to limit this cooperation. The results show that the loss in energy efficiency due to this limitation does not exceed 17% and hence full cooperation between the ISP and the CP is not necessary.

While the ISP and CP cooperation is limited, it still adds to the complexity and cost of the proposed model by employing third-party authorities to control this cooperation. In addition, the model assumes that user requests are made for very popular content and the influence of less popular content is negligible. Even though the most popular content accounts for a significant share of network traffic, less popular content contributes to a serious portion of network traffic and occupies a large part of storage capacities. As a result, less popular content accounts for a considerable amount of the energy consumption of the network.

3.5.2.5 Energy-Aware Load Balancing in CDNs

The authors of [77] propose reducing the energy consumption of CDNs by local load balancing within data centres and global load balancing between data centres. The results in [77] show that the energy consumption of the network is reduced by up to 51% through local balancing. However, only limited additional energy savings (4%-6%) can be achieved through global load balancing, as the load over the network is similar.

The authors claim that more savings can be achieved in scenarios where the load significantly increases in one cluster in the network; however, flash crowds are typically experienced in the whole network. In addition, the authors indicate that switching on as little as 10% of spare servers in data centres ensures acceptable content availability during load spikes. Nevertheless, continuously provisioning for an additional 10% of network traffic leads to increasing the energy consumption of the network. A trade-off between content availability and energy consumption needs to be considered.

3.5.3 Energy Efficient VoD

In VoD services, the total number of videos is larger compared to broadcast TV and CuTV. In addition, the popularity distribution of VoD content is heavy-tailed. As a result, it is a challenge to optimise content location, replacement and access patterns for VoD. The optimum solution depends on the type of service and other parameters such as the network topology and storage constraints. There is ongoing research that aims to minimise the energy consumption of VoD, and following is an overview of some of these attempts.

3.5.3.1 Proactive Server Provisioning

The authors of [78] propose a server provisioning strategy for a VoD service that turns on/off servers with respect to traffic load. They use a proactive online algorithm to calculate the predicted number of requests at each time. The data provided by the algorithm is utilised to turn servers on before they are needed, allowing them booting time to minimise user rejections. They compare their results to a non-prediction scheme that turns servers on/off with respect to instantaneous number of requests. Simulation results show that their proposed scheme reduces the number of rejected requests. In addition, it reduces the energy consumption by minimising the number of idle servers (which consume up to 66% of peak power [79]).

The drawback of this algorithm is that it results in increasing user reconnections. A reconnection occurs when the user's server is turned off and the user is transferred to an alternative active server, resulting in a short-term experience degradation.

3.5.3.2 Multi-level Content Placement Optimisation

Multi-level caching can potentially be provided at five strategic locations in the VoD network. Examples of possible content placement locations include at the head-end server (level 5) in the CDN network, the edge router (level 4) and aggregation switch (level 3) in the transport network, and at the DSLAM (level 2) and STB (level 1) in the access network. The work in [80] evaluates the energy consumption of retrieving content from each level and numerical analysis decides whether replicating a movie at a certain level is energy efficient. Movie request arrival rates play an important role in the accuracy of the decision, as replicating a movie stored at one level into a lower level is considered more energy efficient only if the overall user request arrival rate exceeds a certain threshold.

Results show that by evaluating energy consumption parameters it is feasible to determine which portions of content storages to switch off with respect to movie request arrival rate. Deploying the proposed multi-level content placement approach introduces significant savings in energy consumption, ranging from 25% at (level 4) to 97% at (level 1).

The authors propose varying the number of movies stored at each network level and switching off portions of caches during off-peak hours. This implies that a certain movie is served from five possible network locations depending on instantaneous traffic. As a result, additional communication overhead is introduced to the network as the location of movies has to frequently be updated in Domain Name Servers (DNSs).

The research reviewed in this section considered methods that can be used to reduce the energy consumption of video delivery by reducing network traffic or optimising the use of storage. To achieve that, the researchers used different strategies including load balancing, cache size optimisation and traffic aggregation techniques.

In this thesis, the power consumption of video delivered over an IP over WDM network is minimised by optimising cache sizes and cache updates. The evaluation utilises Mixed Integer Linear Programming (MILP) models, a Constrained-Based Genetic Algorithm (CBGA) and simulations. The following section describes these methods.

3.6 Optimisation Methods

Optimisation can be conducted using different methods including linear and nonlinear programming, simulation, game theory, genetic algorithms and many more. Following is a brief description of the optimisation methods used in this thesis.

3.6.1 Mixed Integer Linear Programming

Linear Programming (LP) is a form of constrained optimisation where the best solution is found while satisfying some constraints. Constrained optimisation problems consist of four main elements [81]:

1. **Objective function**: a mathematical expression that symbolises the goal. During the optimisation process this objective function is maximised or minimised.
2. **Variables**: the set of variables in the problem are adjusted during the optimisation process until the best value for the objective function is found.
3. **Constraints**: a set of mathematical expressions that set limits and determine the accuracy of possible solutions to the problem.
4. **Variable bounds**: a range of possible values that a variable can take is predefined to reduce the size of the problem.

In LP, the objective function and constraints are linear. Integer Linear Programming (ILP) is an LP problem where all of the solution variables are integer. This criteria makes them more difficult to solve [82].

Mixed Integer Linear Programming (MILP) is used to solve problems having some integer variables and some real variable, or a model with mixed variables. However, it can accept problems with any combination of integer, real and binary variables [81].

There exist a number of algorithms that are used for solving optimisation problems. Following is a brief description of the most popular in literature [83].

1. **Simplex Method**: this was the first method developed for solving LP problems. The simplex method states that the optimal solution is found by following the boundary of the feasible region (the region defined by the set of constraints). The method iterates from one vertex on the boundary of the feasible region to another

until an optimal solution is found. This method is considered the best for most LP problems; however, it is not suitable for MILP problems.

2. **Dual Simplex Method**: similar to the simplex method, this method relies on iterations until an optimal solution is found. It differs in the vertices it considers and the iteration method. While the simplex method starts at a feasible vertex and then iterates searching for an optimal solution, the dual simplex method starts at an optimal solution (that is not feasible) and iterates until a feasible solution is found. The dual simplex method solves MILP problems and therefore was ideal for the MILP models proposed in this thesis.
3. **Newton Barrier Method**: this method differs from the simplex methods in that it iterates through solutions that are not on the boundary of the feasible region. Therefore, only an approximate solution is found. The number of iterations is determined by the level of proximity required by the optimal solution. Consequently, the number of iterations is similar irrespective of the size of the problem.
4. **Branch and Bound Method**: this method is most useful when no feasible solution is found by other methods. It works by dividing (branching) the problem into sub-problems and using LP relaxation (the LP problem resulting from dropping the integer constraint from variables) to find feasible solutions for the original problem. The bounding part of the method is to estimate how good a solution can be found for all sub-problems.

The IBM ILOG CPLEX Optimisation Studio is a popular optimisation software package that solves complex MILP problems. It is accessed through AMPL (A

Mathematical Programming Language) [84] and provides interfaces to a number of programming languages including C++, C# and Java. The AMPL/CPLEX solver was used to solve the proposed MILP models on a Pentium(R) Dual – Core CPU at 2.8GHz with 3GB RAM using the dual simplex method. Each model defines an objective function, 7 to 11 constraints and between 1 million and 3.6 million variables.

3.6.2 Constraint-Based Genetic Algorithms

A Genetic algorithm (GA) is an evolutionary algorithm that is used to solve optimisation and search problems. The components of a GA are [85]:

1. **Population:** is made up of a number of individuals called chromosomes that hold information (genes) representing a possible configuration of the problem. Genes can hold binary, integer or real values and are initialised either randomly or by seeding likely optimum values.
2. **Fitness function:** is an equation representing the target of the optimisation problem. The fitness function of a chromosome is calculated using the values of the genes.
3. **Parent selection:** a group of the existing population is selected to reproduce the next generation based on their fitness. Fit individuals are more likely to produce a fitter offspring, bringing the problem closer to the optimum solution.
4. **Crossover:** is the recombination method that produces an offspring using the chromosome information of two parents.

5. **Mutation:** is an operation where the value of one or more gene is altered. This alteration results in avoiding local minima and improving results. Mutation is performed by a given rate which should not be too high to avoid transforming the problem into a random search.
6. **Termination criteria:** the GA process iterates by selecting the group of parents for the iteration, producing the next generation and calculating the fitness function for each offspring. This process is repeated until the fitness function target is reached or additional repetitions do not improve fitness.

A Constraint-Based Genetic algorithms (CBGA) is a GA that describes an optimisation problems having one or more constraints to satisfy [86]. In CBGAs, the conventional GA selection, crossover and mutation processes are performed to obtain fitter population individuals. A CBGA is more complex compared to a GA, as the values of genes of new chromosomes generated by the GA process must satisfy one or more constraints. Since the GA process does not consider these constraints during execution, chromosomes are likely to violate them. The CBGA can overcome these violations either by repairing each chromosome once generated or by applying a penalty value to the chromosome’s fitness. Repairing chromosomes becomes challenging when the problem contains several constraints and/or the chromosome consists of a large number of genes.

No prior work in network energy optimisation is carried out using CBGAs. In this work a CBGA is used to optimise cache sizes to minimise power usage. The number of genes and capacity and flow constrains to satisfy in the proposed model is large (over

150,000 genes and 11 constraints). Therefore, the CBGA penalises chromosomes that violate constraints by attaching a high penalty to their fitness function (power consumption). The optimum cache sizes found by the CBGA were used to validate those found by the MILP model.

3.6.3 Simulation

A computer simulation is a software reproduction of a system implemented to evaluate the performance of the system. Simulations can be classified into two types with respect to the manner that they track events. Discrete-event simulation is used when the system contains a list of discrete events and simulation starting and ending points are predefined. An example for a discrete-event simulation is a queuing system as the system is evaluated whenever a customer enters or leaves the queue. Continuous simulation evaluates the state of variables and events at equal time intervals and variables change continuously over time. Network traffic routing is an example of this type since the flow of traffic is continuous.

In this work, continuous simulations are used to route network traffic between nodes. The state of the network is evaluated at two-hour time intervals where the traffic demand is input and the power consumption of the network is calculated. Simulations were used to validate power consumption results produced by the MILP models.

3.7 Summary

This chapter has provided an overview of cache placement and replacement algorithms. It has also surveyed different approaches in reducing the energy consumption of IPTV, CDNs and VoD. The chapter was concluded by a description of the optimisation methods used in this work.

The energy consumption of CDNs is important in deciding the future directions of video applications delivered over the Internet. The ongoing research covered in this chapter tries to minimise the energy consumption associated with CDNs in an effort to improve the feasibility of future CDNs. In general, these attempts offer a valuable insight into network energy efficiency by highlighting the network areas where energy reduction is possible. Some of these studies compare the energy efficiency of different content delivery techniques including P2P and CDNs while others are considered with reducing the energy consumption of storage. However, none of the studies evaluates the energy efficiency of different video services delivered by CDNs. In addition, they do not incorporate sleep-capable equipment to optimise utilised network equipment. In this thesis, minimum-power optimisation Mixed Integer Programming (MILP) models are developed to optimise the size of stored content in network nodes to minimise power usage. In these models network resources are provisioned for current traffic requirements and network equipment is provided with sleep-mode capabilities for maximum power efficiency.

Chapter 4 Energy Efficient Video-on-Demand with Fixed and Variable Caches

4.1 Introduction

Due to the high energy consumption associated with the storage and delivery of bulky multimedia files, Video-on-Demand (VoD) is an energy consuming service [80]. Storing the most popular content towards the edge of the network is an effective strategy to reduce the energy consumption of video services. This Chapter evaluates the energy consumption of delivering video over an IP over WDM network. It explores the power savings introduced by utilising caches at core nodes. A Mixed Integer Linear Programming (MILP) model is developed to optimise the cache size for each node in the network to minimise power consumption. The MILP model is extended to consider variable caches (caches which are equipped with sleep-mode capabilities allowing them to power down inactive sections when traffic is low). The model finds the optimum cache size for each node at different times of the day that achieve the most power savings. The MILP model is validated by simulation as well as a Constraint-Based Genetic Algorithm (CBGA).

4.2 Energy-Efficient Cache Size Optimisation

Caching content locally results in shorter routes to content and hence lower power consumption. This strategy however results in increased equipment power consumption through the deployment of local caches. A trade-off has to be struck therefore where the optimum cache size is a function of the two above drivers. This evaluation aims to minimise the power consumption of a video service by optimising the sizes of caches deployed at the nodes. It takes into account a VoD service deploying an IP over WDM network with the network architecture described in Figure 2-3 in Chapter 2.

4.2.1 Fixed Cache MILP Model

This MILP model finds the optimum cache size to be deployed at each node in the network to minimise power consumption. Caches are considered having a fixed size that is fully operated for the whole day. The model declares a number of sets, parameters and variables as follows:

***Sets*:**

N	Set of nodes

Nm_i	Set of neighbouring nodes of node i
T	Set of points in time
K	Set of equations that approximate the convex function describing the relationship between the cache and its hit ratio.

Parameters:

Pp	Power consumption of a router port
Po_{it}	Power consumption of optical switch i at time t
Pt	Power consumption of a transponder
Pa	Power consumption of an amplifier
Pmd	Power consumption of a multiplexer/demultiplexer
B	Capacity of a wavelength
W	Number of wavelengths in a fibre
D_{ij}	Distance from node i to j
S	Span distance between two amplifiers
Amp_{ij}	Number of amplifiers used on each fibre on the physical link from node i to j, $Amp_{ij} = \lfloor D_{ij}/S - 1 \rfloor + 2$
$RPmax_x$	Maximum router ports available to node x
λ^{xyt}	Demand from node x to y at time t

δ_i	Is 1 if node i has a video server, 0 otherwise, $\sum_{i \in N} \delta_i = u$, where u is the total number of servers in the network
Ru	Server ratio of uplink demand to regular traffic
Rd	Server ratio of downlink demand to regular traffic
Φ	Cache power consumption factor in W/GB
a, b	Approximation vectors

Variables:

f_{ij}	Fibres on the physical link from node i to j
λ_{ijt}^{xy}	Regular traffic from node i to j, part of the virtual link from node x to y at time t
λu_{ijt}^{xy}	Uplink traffic from node i to j, part of the virtual link from node x to y at time t
λd_{ijt}^{xy}	Downlink traffic from node i to j, part of the virtual link from node x to y at time t
w_{ijt}^{xy}	Wavelengths on the link from node i to j, part of the virtual link from node x to y at time t
w_{ijt}	Wavelengths on the physical link from node i to j at time t
C_{xyt}	Wavelengths on the virtual link from node x to y at time t

AP_{it} Aggregation ports at node i at time t

H Cache hit ratio

M Cache size in GB

Under lightpath bypass, the power consumption of the network consists of the power consumption of the following components:

1. Router ports at time t, where a port is required for each occupied wavelength:

$$\sum_{i \in N} Pp \left(AP_{it} + \sum_{j \in Nm_i : i \neq j} C_{ijt} \right)$$

2. Optical switches at time t:

$$\sum_{i \in N} Po_{it}$$

3. Transponders at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pt \cdot w_{ijt}$$

4. Amplifiers at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pa \cdot Amp_{ij} \cdot f_{ij}$$

5. Multiplexers/demultiplexers at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pmd \cdot f_{ij}$$

6. Deployed caches at time t:

$$\sum_{i \in N} \emptyset M$$

It is worth mentioning that the model does not assume a simple symmetric case, and therefore, the number of lightpaths from node i to j can be different to the number of lightpaths in the reverse direction. Mainly, f_{ij} , w_{ijt} and C_{ijt} are not necessarily equal to f_{ji}, w_{jit} and C_{jit}, respectively. Note that uplink traffic is the video traffic uploaded from nodes to video servers, downlink traffic is the video traffic downloaded from video servers to nodes and regular traffic is other non-cacheable traffic (email, live video, dynamic content, etc.).

The goal of the proposed MILP model is to minimise the network total daily power consumption while satisfying a number of flow and capacity constraints. The complete MILP model is defined as:

Objective: minimise

$$\sum_{t \in T} \left(\begin{array}{c} \sum_{i \in N} Pp \left(AP_{it} + \sum_{j \in Nm_i : i \neq j} C_{ijt} \right) + \sum_{i \in N} Po_{it} + \\ \sum_{i \in N} \sum_{j \in Nm_i} Pt \cdot w_{ijt} + \sum_{i \in N} \sum_{j \in Nm_i} Pa \cdot Amp_{ij} \cdot f_{ij} + \\ \sum_{i \in N} \sum_{j \in Nm_i} Pmd \cdot f_{ij} + \sum_{i \in N} \emptyset M \end{array} \right) \qquad (4\text{-}1)$$

Subject to:

$$\sum_{y \in N} C_{xyt} + AP_{xt} \leq RPmax_x \quad (4\text{-}2)$$
$$\forall\, x \in N, \forall\, t \in T$$

$$\sum_{y \in N} C_{yxt} + AP_{xt} \leq RPmax_x \quad (4\text{-}3)$$
$$\forall\, x \in N, \forall\, t \in T$$

$$\sum_{x \epsilon N} \sum_{y \epsilon N : x \neq y} w_{ijt}^{xy} \leq W \cdot f_{ij} \quad (4\text{-}4)$$
$$\forall\, i \in N, j \in Nm_i, \forall\, t \in T$$

$$\sum_{x \epsilon N} \sum_{y \epsilon N : x \neq y} w_{ijt}^{xy} = w_{ijt} \quad (4\text{-}5)$$
$$\forall\, i \in N, j \in Nm_i, \forall\, t \in T$$

$$\sum_{j \epsilon Nm_i} w_{ijt}^{xy} - \sum_{j \epsilon Nm_i} w_{jit}^{xy} = \begin{cases} C_{xyt} & i = x \\ -C_{xyt} & i = y \\ 0 & otherwise \end{cases} \quad (4\text{-}6)$$
$$\forall\, i, x, y \in N, \forall\, t \in T$$

$$\sum_{j \epsilon N : i \neq j} \lambda_{ijt}^{xy} - \sum_{j \epsilon N : i \neq j} \lambda_{jit}^{xy} = \begin{cases} \lambda^{xyt} & i = x \\ -\lambda^{xyt} & i = y \\ 0 & otherwise \end{cases} \quad (4\text{-}7)$$
$$\forall\, i, x, y \in N, \forall\, t \in T$$

$$\sum_{j \epsilon N : i \neq j} \lambda u_{ijt}^{xy} - \sum_{j \epsilon N : i \neq j} \lambda u_{jit}^{xy} = \begin{cases} \lambda^{xyt} \cdot\ Ru \cdot \delta_y & i = x \\ -\lambda^{xyt} \cdot\ Ru \cdot \delta_y & i = y \\ 0 & otherwise \end{cases} \quad (4\text{-}8)$$
$$\forall\, i, x, y \in N, \forall\, t \in T$$

$$\sum_{j \epsilon N : i \neq j} \lambda d_{ijt}^{xy} - \sum_{j \epsilon N : i \neq j} \lambda d_{jit}^{xy} = \begin{cases} \lambda^{xyt} \cdot Rd \cdot \delta_x \cdot (1 - H) & i = x \\ -\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot (1 - H) & i = y \\ 0 & otherwise \end{cases} \quad (4\text{-}9)$$
$$\forall\, i, x, y \in N, \forall\, t \in T, \forall\, r \in R$$

$$\sum_{x \in N} \sum_{y \in N: x \neq y} \left(\lambda_{ijt}^{xy} + \lambda u_{ijt}^{xy} + \lambda d_{ijt}^{xy} \right) \leq C_{ijt} \cdot B \tag{4-10}$$

$$\forall\, i, j \in N, \forall\, t \in T$$

$$AP_{it} = \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot (1 - H) \right) / B \tag{4-11}$$

$$\forall\, i \in N, \forall\, t \in T$$

$$M \geq a_k \cdot H + b_k \qquad \forall\, k \in K \tag{4-12}$$

Objective (4-1) calculates the power consumption of the network by summing up the power consumption of different network components at each time point. Constraints (4-2) and (4-3) limit the number of occupied router ports at each node to its maximum. Constraint (4-4) and (4-5) are the capacity constraints for the optical layer. Constraint (4-6) is the flow conservation constraint in the optical layer. Constraint (4-7) is the flow conservation constraint for regular traffic in the IP layer. Constraints (4-8) and (4-9) differ from Constraint (4-7) by considering the uplink and downlink traffic terminating and originating at nodes equipped with a video server, respectively. Constraint (4-10) ensures that the total regular, uplink and downlink traffic carried by a lightpath does not exceed its capacity. Constraint (4-11) calculates the number of required aggregation ports. Finally, Constraint (4-12) is the piecewise linear approximation utilised to find the cache size M from its hit ratio H.

4.2.2 Variable Cache MILP Model

The variable cache MILP model finds the optimum cache size of each node varied over the time of the day. The model assumes that caches are equipped with sleep-mode capabilities such that inactive parts of a cache can go to sleep. The goal is to explore the potential additional power savings on top of the use of fixed caches and to analyse the variation in optimum cache sizes that minimise power consumption with the variation in network traffic.

The variable cache MILP model defines the same sets, parameters and variables defined for the fixed cache MILP model. Since cache sizes are variable for each node at each time of the day, the cache size variable M and its hit ratio H are modified as follows:

H_{it} Hit ratio of the cache deployed at node i at time t

M_{it} Cache size in GB deployed at node i at time t

Therefore the objective function becomes:

$$\sum_{t\in T}\left(\begin{array}{c}\sum_{i\in N} Pp\left(AP_{it}+\sum_{j\in Nm_i: i\neq j} C_{ijt}\right)+\sum_{i\in N} Po_{it}+\\ \sum_{i\in N}\sum_{j\in Nm_i} Pt\cdot w_{ijt}+\sum_{i\in N}\sum_{j\in Nm_i} Pa\cdot Amp_{ij}\cdot f_{ij}+\\ \sum_{i\in N}\sum_{j\in Nm_i} Pmd\cdot f_{ij}+\sum_{i\in N}\emptyset M_{it}\end{array}\right) \tag{4-13}$$

In addition, Constraints (4-9), (4-11) and (4-12) are modified as follows:

$$\sum_{j\epsilon N: i\neq j}\lambda d_{ijt}^{xy}-\sum_{j\epsilon N: i\neq j}\lambda d_{jit}^{xy}=\begin{cases}\lambda^{xyt}\cdot Rd\cdot\delta_x\cdot(1-H_{yt}) & i=x\\ -\lambda^{xyt}\cdot Rd\cdot\delta_x\cdot(1-H_{yt}) & i=y\\ 0 & otherwise\end{cases} \tag{4-14}$$
$$\forall\, i,x,y\in N,\forall\, t\in T,\forall\, r\in R$$

$$AP_{it} = \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot (1 - H_{it}) \right) / B \qquad (4\text{-}15)$$

$$\forall\, i \in N, \forall\, t \in T$$

$$M_{it} \geq a_k \cdot H_{it} + b_k \qquad (4\text{-}16)$$

$$\forall\, k \in K$$

4.3 Constraint-Based Genetic Algorithm and Simulation Schemes

A Constraint-Based Genetic algorithm (CBGA) and a simulation are developed to validate the proposed MILP models, explained below.

4.3.1 CBGA for Cache Size Optimisation

To validate the results of the variable MILP model, a Constraint-Based Genetic algorithm (CBGA) is developed having objective (4-13) as the fitness function. The genetic algorithm defines a population of individuals or chromosomes, and their structure is shown in Figure 4-1. Each chromosome represents one possible network configuration made up of the number of wavelengths in the optical and virtual layer, the amount of (regular, uplink and downlink) traffic carried by each link and the cache size at each node. The fitness of a chromosome is the power consumed using the configuration given by the values of the chromosome genes.

The CBGA process, illustrated in Figure 4-2, starts by initialising chromosomes using random numbers. The fittest individuals are selected to reproduce the offspring by undergoing the crossover process with a crossover rate *Pc* and the mutation process with a probability *Pm*. All members of the offspring go through a screening process where constraints are enforced. Violating members are assigned a penalty that increases their power consumption and therefore decreases their fitness which gives them a higher expectancy to die out in the next generation.

C_{ij}	w_{ij}^{xy}	λ_{ij}^{xy}	$\lambda\_u_{ij}^{xy}$	$\lambda\_d_{ij}^{xy}$	H_i

Figure 4-1: The CBGA chromosome structure

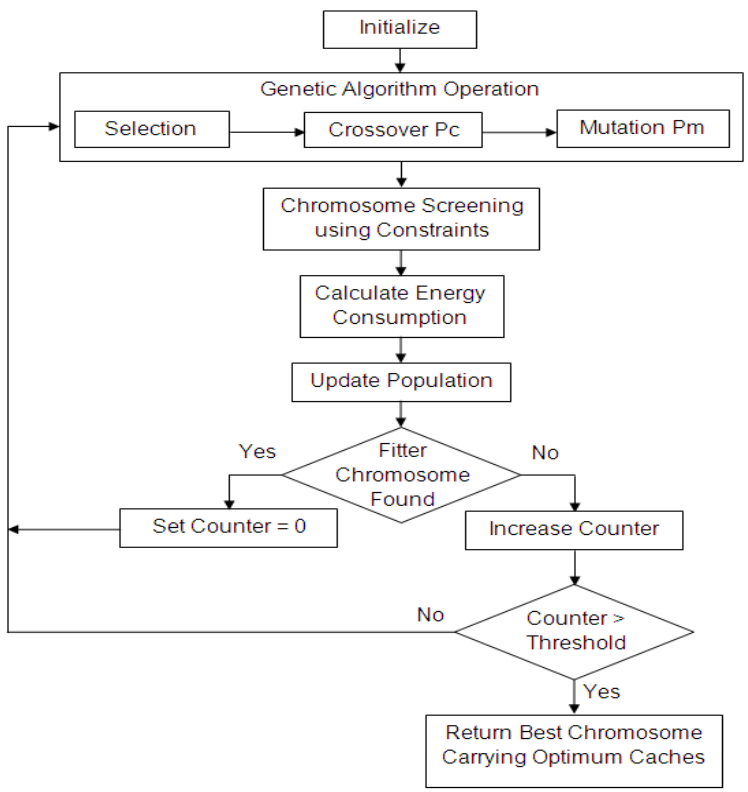

Figure 4-2: The constraint-based genetic algorithm process

Since the objective is to minimise power consumption, chromosomes with low values of the fitness function are considered to be fit individuals. This procedure is repeated until no reduction in network power consumption is observed for numerous iterations. In other words, the process terminates when the number of iterations *Count* exceeds a large number *Threshold*, with no reduction in power consumption. Finally, the optimum individual representing the most energy-efficient configuration is obtained, carrying the optimum cache size for each node in the network.

4.3.2 Minimum Power Simulation

As a second validation technique for the MILP models, a simulation is developed based on the lightpath bypass heuristic proposed in [32]. The algorithm, illustrated in Figure 4-3, makes use of the optimum cache sizes produced by the MILP model by assuming the deployment of the optimum cache size at each node in the network at each considered time of the day. The downlink traffic demand from a server to a node is directly influenced by the size of the cache deployed at the node (becomes $1-H$ of the original downlink traffic, where H is the cache hit ratio). Therefore, the remaining downlink traffic and consequently the total traffic (regular, uplink and downlink traffic) is calculated for each node after assuming a certain cache size for a node.

Node pairs are then arranged in a descending order starting with the node pair having the highest demand to ensure that the algorithm accommodates high demands on virtual links first and hopefully accommodate lower demands on existing links. An empty topology G is created to track established links and their capacities. The node pair (x, y) with the highest demand $d(x, y)$ is selected and the algorithm attempts to route $d(x, y)$ over existing virtual links. If this process is successful, the remaining capacities of G are updated. Otherwise, (capacities of existing virtual links in G are insufficient to accommodate some or all $d(x, y)$), a new virtual link connecting x and y is created and added to G. The demand $d(x, y)$ is routed over the new link and the remaining capacities of G are updated. The algorithm continues by repeating the selection of node pairs until all demands are routed.

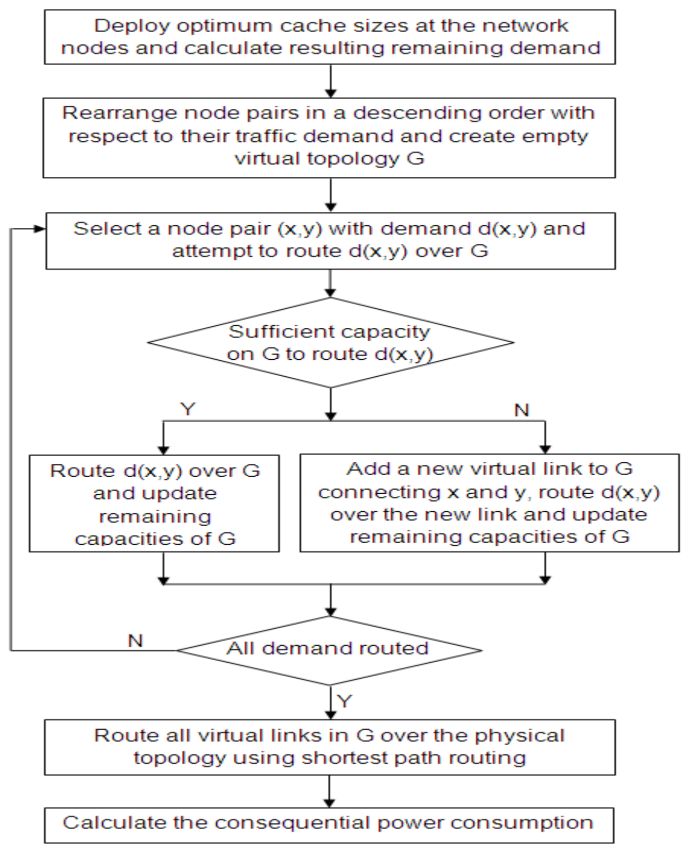

Figure 4-3: The simulation flowchart

At this point, G represents the set of lightpaths to be routed over the physical topology in the optical layer. The simulation uses the shortest path routing algorithm to route the lightpaths over the physical topology, as the shortest path provides the minimum usage of EDFAs, transponders and multiplexer/demultiplexers, and therefore minimises the power consumption of the optical layer. Note that minimum-hop routing can result in lower power consumption if the IP layer (at nodes) power consumption dominates. (A shorter route may involve many intermediate nodes). However, shortest path routing is implemented here as it minimises delay, and the simulation results obtained are in any case close to the results found using the MILP models.

The power consumption of the IP layer is determined using the capacities of the virtual links to calculate the number of required ports and their power consumption, and the total network power consumption is found. The algorithm improves virtual link utilisation by allowing more than one demand to be routed on the same virtual link. This feature results in decreasing the number of established virtual links and consequently utilising less IP router ports, the network major power consuming component, leading in turn to overall network power reduction.

4.4 Power Consumption Evaluation

This section demonstrates by results the outcomes of the proposed power-minimised MILP models. The power consumption of an IP over WDM network is evaluated and the optimum fixed and variable cache sizes to achieve the most power efficiency are found.

4.4.1 Test Network and Input Parameters

A network architecture having multiple servers where content is replicated and/or distributed among servers influences the traffic behaviour, as demand to a certain node is retrieved from multiple servers. This behaviour better describes traffic generated by famous services such as Google, which has 19 data centres in the US alone [87] and YouTube which has 6 data centres excluding the CDN [88]. To evaluate the proposed power-minimisation MILP model, the NSFNET topology of 14 nodes and 21 bidirectional links is considered. The network is considered having 7 video servers where their optimum locations (locations that minimise the network power consumption) are obtained from the MILP model assuming caches of fixed sizes. Results show that the optimum locations of video servers are independent of cache sizes, and are given as nodes 1, 3, 5, 8, 10, 12 and 14. Figure 4-4 shows the NSFNET topology with distances between nodes in kilometres and locations of optimum servers.

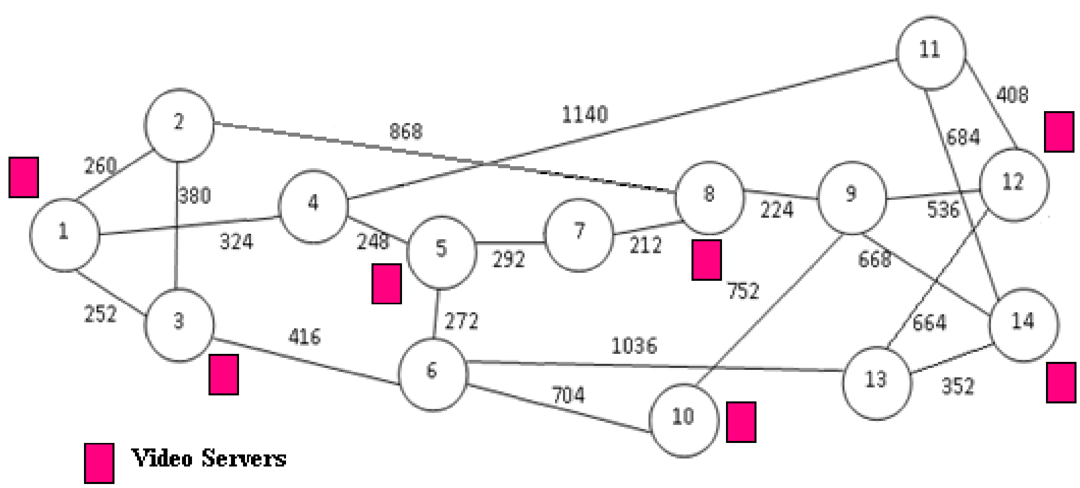

Figure 4-4: The NSFNET topology with video servers and fibre lengths (km) [32]

The average demand between node pairs ranges from 20Gb/s to 120Gb/s and the peak occurs at 23:00. The regular traffic demand between each node pair is generated using a random function having a uniform distribution on the interval (10Gb/s, 230Gb/s][63]. The presence of servers creates a hot node scenario where more traffic originates from and terminates at servers. Therefore uplink and downlink traffic demand between nodes and video servers is considered and is generated based on regular traffic demand between nodes. Three different values for the uplink to regular traffic demand ratio Ru and the downlink to regular traffic demand ratio Rd are considered: (1) Rd = 1.5 and Ru = 0.2, (2) Rd = 4.5 and Ru = 0.6, and (3) Rd = 7.5 and Ru = 1.0. These values match the input and output rates of a typical video server [89] and reflect the expected growth rate in Internet video traffic [2]. The considered regular and total network traffic at each time of the day is shown in Figure 4-5.

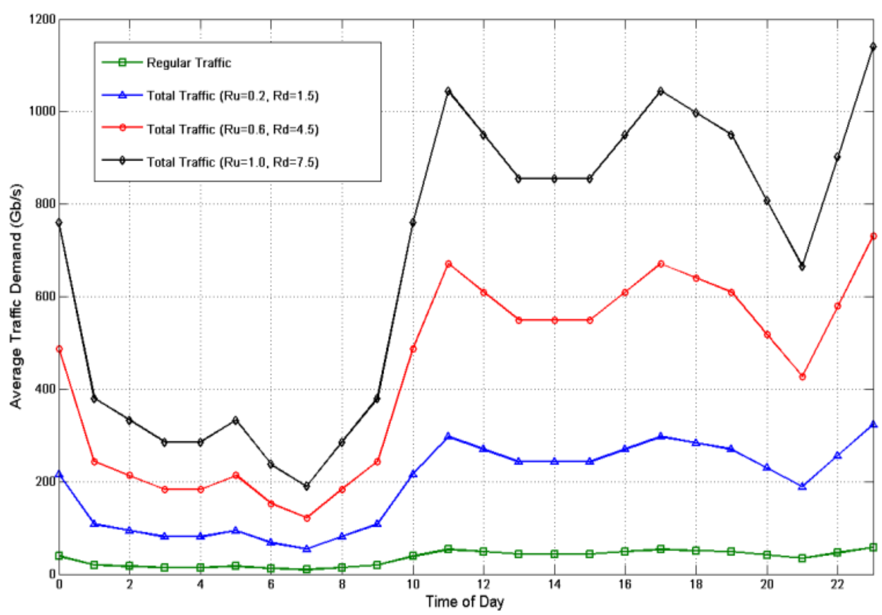

Figure 4-5: Network traffic demand considering uplink and downlink traffic

The input values for power consumption parameters used in the model are shown in Table 4-1. The power consumption of an 8-slot CRS-1 is 8000W including backplane power [90], and therefore the power consumption of a 40Gb/s port is estimated at 1000W [32]. The power consumption of a Cisco ONS 15454 Optical Filter Card is 16W [91], and a Cisco ONS 15501 Optical Amplifier is 8W [92].

Table 4-1: Input data for the MILP model

Distance between two neighbouring EDFAs (*S*)	80 (km)
Number of wavelengths in a fibre (*W*)	16 [32]
Capacity of a wavelength (*B*)	40 (Gb/s)
Power consumption of a router port (*Pp*)	1000 (W) [90]
Power consumption of a transponder (*Pt*)	73 (W) [32]
Power consumption of an EDFA (*Pa*)	8 (W) [92]
Power consumption of an optical switch (*Po*)	85 (W) [93]
Power consumption of a MUX/DEMUX (*Pmd*)	16 (W) [91]
Cache power consumption factor (∅)	7.4 (W/GB)

The energy consumption of streaming 1 bit (READ/WRITE operation) is given as 211×10^{-9}J [94]. This value is converted into the power in Watts consumed to stream 1GB of data over a given time duration, to utilise in the MILP model. The energy consumption of streaming 1GB of data is: $211\times10^{-9}\times10^{9}\times8=1688$J. The typical maximum access data rate for a fibre channel hard disk is up to 100MB/s. However, the typical average data rate is around 50% of that value and is up to 50MB/s [95]. Higher hard disk access speeds reduce the cache update time, but lead to higher power consumption. A lower hard disk data rate of 35Mb/s is considered to achieve reasonable power consumption. The time required to stream 1GB becomes: 10^{3}

×8/35=228 seconds. Consequently the power consumption of caching 1GB is 1688/228=7.4W/GB. Each node in the network is assumed to have a cache containing the most popular objects and therefore a total of 14 caches are deployed in the network. The Zipf distribution explained in Section 2.4.1 in Chapter 2 is considered for the popularity of video content. The MILP model utilises the piecewise linear approximation equation given in Equation (4-12) to define the relationship between the cache size and its hit ratio.

To run the proposed MILP model, the described NSFNET topology in Figure 4-4 is assumed (14 nodes and 21 links) having an IP over WDM architecture, where the details of the architecture are shown in Figure 2-3 in Chapter 2 considering 2 nodes only. Considering the traffic, content distribution and power consumption parameters explained above, a typical run of the model requires between 1 and 5 hours for its different states using the solver and computer specified in Section 3.6.1 in Chapter 3. The model defines over 1.8 million variables and uses dual simplex iterations to find the optimum solution.

4.4.2 Optimum Fixed Caches

The MILP model, CBGA and simulation are utilised to evaluate the power consumption of the fixed cache size approach and compare it to the power consumption of the network when no caches are deployed at the nodes. A library of 2 million objects of the same size of 0.2GB (a typical YouTube video size) is considered. To evaluate the influence of deploying caches of fixed sizes on the network energy efficiency, a range of cache sizes are deployed at the network nodes and the corresponding energy consumption of the network is calculated. Figure 4-6 shows the energy consumption of the network when deploying caches of sizes ranging from 20GB to 4000GB when Rd = 1.5, 4.5 and 7.5. The addition of caches increases the power consumption of network components, however content can be accessed locally in the presence of caches which reduces the power used to access an object. The net effect of these two trends is the existence of an optima in Figure 4-6 where a particular cache size minimises the overall power consumption of the network (200, 1000 and 2000GB when Rd = 1.5, 4.5 and 7.5, respectively).

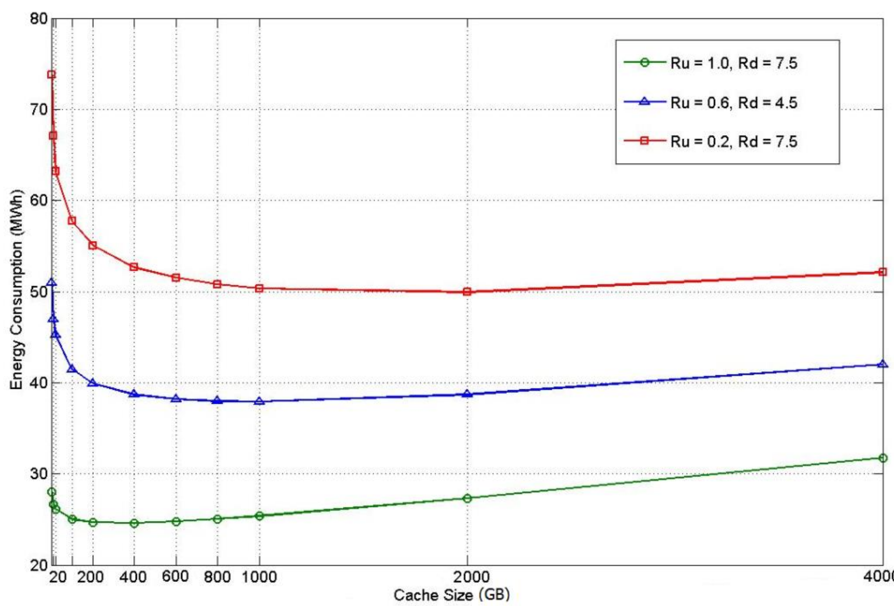

Figure 4-6: Energy consumption considering fixed cache sizes

4.4.3 Optimum Variable Caches

Caching can be provided through a number of hard disks, and unused capacity can go to sleep mode. Future implementations can use flash memory and therefore provide finer granularity in terms of the number of memory units powered at a given point in time. The MILP model is used to find an optimum variable cache size that minimises power consumption varied over network nodes and time of the day. Figure 4-7 depicts the optimum cache size for each node in the network at different times of the day. For each node, the optimum cache size varies with the traffic passing through the node at a certain time. Therefore when considering one row along the *time of day* axis, the optimum cache size for a certain node over 24 hours follows the trend of the traffic shown in Figure 4-5. When considering a column showing the optimum cache sizes of all nodes at a certain time of day, cache sizes are comparable, as the traffic demand of nodes congregates around the average network traffic demand at that time of day.

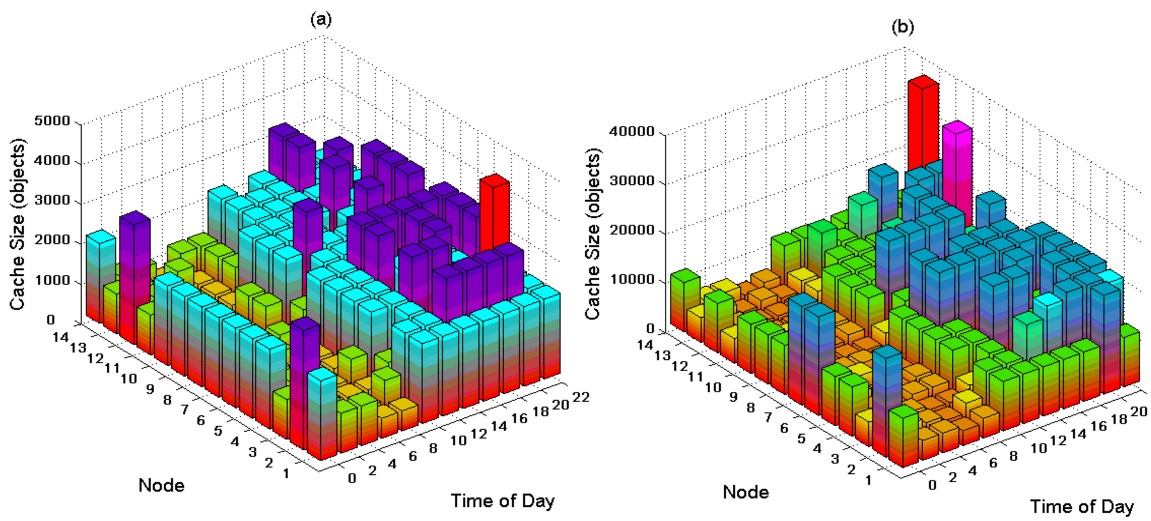

Figure 4-7: Optimum cache sizes for each node at different times of the day (a) *Ru* = 0.2 and *Rd* = 1.5, (b) *Ru* = 1.0 and *Rd* = 7.5

Observing cache sizes in Figure 4-7, there exists a step between one cache size and another. This is caused by the piecewise linear approximation equation given in

Equation (4-16), as the optimum cache size is approximated to the closest value provided by the approximation. Improved granularity can be achieved by including additional linear equations in the approximation. Optimum cache sizes are relatively high throughout the network when the traffic is high and vice versa. Note that the cache power consumption is a function of the desired hit ratio. A higher hit ratio requires the storage of more objects and higher cache power consumption. For a 100% hit ratio, all the objects are stored at the cache and the power consumption of the cache is at its maximum. Additional effects can be considered beyond the current work. For example, the transmission mechanism used may call for small local caches to be attached to each link so that the main cache is shared between all links and is used to update the small transmission caches in each outgoing link. Such refinements warrant further investigation.

4.4.3.1 Power Evaluation using the CBGA

To validate the optimum cache sizes determined by the MILP model, the CBGA is employed to find the optimum cache size for each node at different times of the day. The considered test network (NSFNET topology), traffic and power consumption parameters are similar to those applied to the MILP model. The algorithm defines a population of 240 individuals to form a generation having 0.5 as the crossover point. The values of the mutation rate influences the rate at which the algorithm converges. A slow convergence allows the algorithm to perform more iterations searching for the best possible fitness. Initially, three values for the mutation rate are assumed (0.02, 0.1 and 0.2) to evaluate the effect of the mutation rate on the optimum cache sizes found by the CBGA. The insert in Figure 4-8 shows that applying a large mutation rate results in a slower convergence of the algorithm, and consequently a longer execution time for the CBGA. While the slow convergence might be considered as a downside, this larger value for the mutation rate resulted in the least error rate compared to the MILP model results. The value of 0.2 results in an acceptable error rate, and was therefore used to attain the optimum cache sizes for the considered duration. The CBGA was developed using Java programming language and the algorithm takes over 6 hours to converge per time of day. Figure 4-8 shows the average of the optimum cache sizes obtained by the MILP model and the CBGA taken over 24 hours. The optimum cache sizes produced by the two techniques are comparable and follow the trend of the traffic demand shown in Figure 4-5.

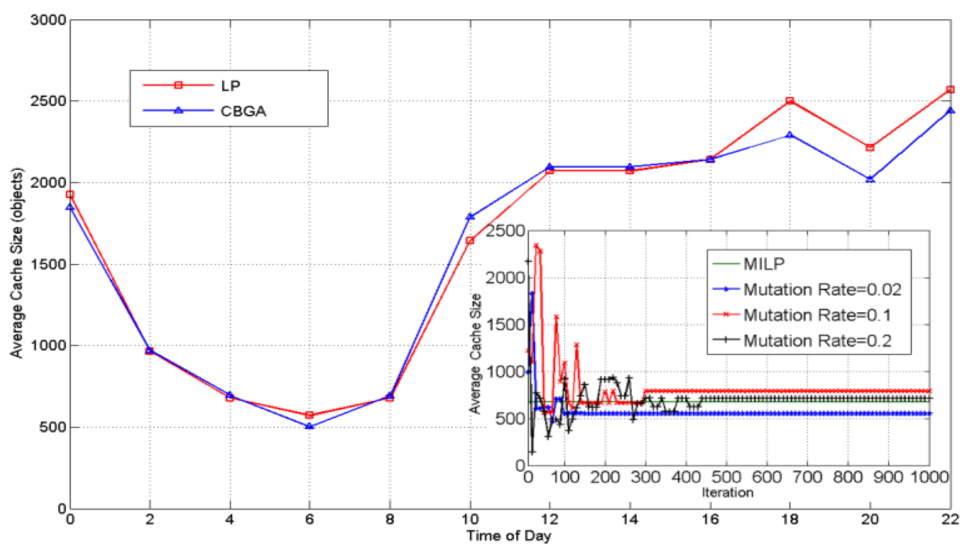

Figure 4-8: The average of the optimum cache sizes of 14 nodes taken over 24 hours obtained by the MILP model and the CBGA. The insert figure shows the average of the optimum cache sizes at 4:00 using different mutation rates

4.4.3.2 Power Evaluation using Simulation

The power consumption of routing the traffic demand over the physical topology configuration obtained by the MILP model is validated by implementing the minimum-power simulation. The optimum cache sizes found by the MILP model and validated by the CBGA are used as input cache sizes to the simulation. All other power consumption parameters and assumptions are similar to those used in the MILP model.

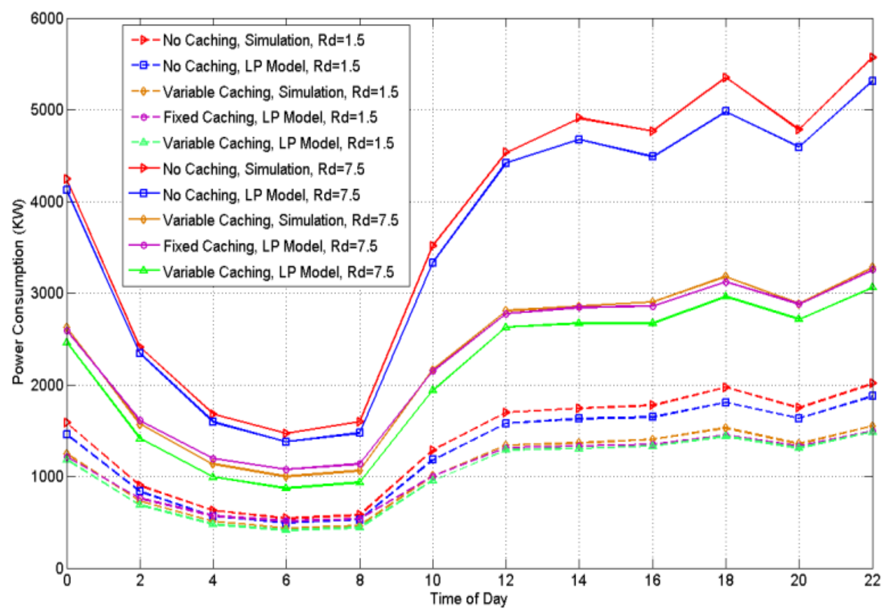

Figure 4-9: Network power consumption at different times of the day using fixed and variable cache sizes obtained by the MILP model and simulation

As can be seen in the simulation flowchart in Figure 4-3, the algorithm performs one iteration per traffic demand, and therefore the simulation execution time is negligible. The simulation considers the situation of network dimensioning where extra resources are added where required. As a result, traffic is always routed over the shortest path in the physical layer. In the case where a maximum physical link capacity is identified, some traffic will have to be routed over longer routes, consuming more power. A more extreme situation to consider is where demands may be rejected due to lack of resources. Such scenarios require further investigation.

Figure 4-9 shows the power consumption of the network taken over 24 hours using the MILP model and the simulation when Rd = 1.5, Ru = 0.2 and Rd = 7.5, Ru = 1.0. The figure shows the power consumption of the network when no caches are deployed, using fixed caches, and when variable caches are deployed at the nodes.

The results produced by the MILP model, CBGA and simulation show consistency in the power consumption figures and optimum cache sizes, giving confidence in the developed methods. In addition to the two considered validation approaches, the power consumption figures obtained using the MILP model considering no caching are consistent with the results in [63] when similar number of servers and input values for traffic and power consumption parameters are considered.

4.4.3.3 Power Saving Figures

Using the MILP model and caches of fixed sizes, reductions in network power consumption of up to 19% (average of 8%) and up to 38% (average of 30%) when Rd = 1.5, Ru = 0.2 and Rd = 7.5, Ru = 1.0 are achieved, respectively. It is also observed that more significant power is saved when the traffic is high compared to when the traffic is low. However, using optimum fixed size caches under low traffic results in higher power consumption compared to when no caches are deployed in the network. This is because optimum caches of fixed sizes are larger than required when the traffic is low, as optimum cache sizes are found for all nodes at all times considering overall network traffic. As a result, unnecessary power is consumed for storage.

To achieve the maximum power efficiency, the MILP model identifies the optimum cache size for each node for each considered time of day, shown in Figure 4-7. Variable cache sizes are optimised with respect to traffic conditions of each time of the day. Using the power minimising MILP model under caches of variable sizes, the power consumption of the network is reduced by a maximum of 19% (average of 16%) and a maximum of 42% (average of 37%) when Rd = 1.5, Ru = 0.2 and Rd = 7.5, Ru = 1.0, respectively. Varying the cache size with respect to instantaneous traffic achieves more power savings compared to caches of fixed sizes, although limited. A maximum reduction in power consumption of 20% (average of 6%) and 14% (average of 6%) is obtained when Rd = 1.5, Ru = 0.2 and Rd = 7.5, Ru = 1.0, respectively. The improvement in power efficiency introduced by utilising caches of variable sizes is higher when traffic experiences steep fluctuations during the course of the day. In other words, the use of caches of variable sizes becomes more desirable when the difference between the minimum and maximum traffic passing through a node is large. Network additional savings of 6% suggest that the use of fixed caches saves a significant amount of power without the need for traffic measurements and storage devices that have sleep-mode capabilities.

All the previously reported power saving figures are for the power consumption of transport as well as the additional power consumed on storage. To analyse the achieved power savings, the net power savings on transport and the additional cache power consumption are also reported, shown in Table 4-2.

Table 4-2: The Breakdown of network power savings

	Low Traffic				**High Traffic**			
	Fixed caching		**Variable caching**		**Fixed caching**		**Variable caching**	
	max	**Avg.**	**max**	**Avg.**	**max**	**Avg.**	**max**	**Avg.**
Power savings on transport %	22	12	22	18	43	37	49	42
Cache power consumption %	5	4	3	2	8	7	7	5
Total (overall) power savings %	19	8	19	16	38	30	42	37

The power consumption of caching is added to the power consumption of transport to attain the overall network power consumption. The power consumption of transport is the power consumption of the network excluding the power consumption of caching, and therefore the power saving figures are higher. The optimum cache sizes found under variable caching are smaller than when considering fixed cache sizes, contributing however to larger overall power savings. The additional power consumed on storage (2% – 8% of network power) is hence a worthwhile power investment, leading to considerable overall power savings (42% maximum).

4.5 Summary

Energy efficiency has become a key factor in the design and implementation of communication networks. This Chapter has evaluated the energy efficiency associated with the introduction of caching in IP over WDM networks. A MILP model has been developed with the objective of minimising the network power consumption by optimising the cache size deployed at network nodes. An extended MILP has been developed considering optimising the size of the cache for each node varied over the time of the day. A simulation and a Constraint-Based Genetic Algorithm (CBGA) have also been developed to validate the results obtained by MILP models.

The results reveal that optimising the cache size for each node at different times of the day introduces power savings of up to 42% considering a Zipf distribution for content popularity. Varying the size of the cache at each node during the day can save up to 20% power compared to utilising a cache of a fixed size all the time at each node in the network. These savings depend on the level of fluctuation that the traffic experiences during the day. However, the additional average network power savings are 6% and therefore, finding the optimum fixed cache size to deploy in the core network is adequate to introduce considerable power savings.

Chapter 5 Impact of Content Popularity Distributions and Network Parameters

5.1 Introduction

The previous chapter has investigated the impact of caching the most popular videos towards the edge of the core network to minimise power consumption. It has incorporated a Zipf distribution for content popularity and illustrated by results the optimum cache sizes that minimise power usage, and the power saving figures obtained. The popularity distribution of content plays an important role in the caching decision since it provides an accurate indication of the hits associated with each video and therefore the cache hit ratio. Consequently, it is an important measure for optimising cache sizes that minimise power usage. This chapter investigates the impact of content popularity distribution on power efficiency by considering and comparing four different distributions for content popularity: Zipf, Bimodal, Pareto and Equal Popularity distributions. The MILP model proposed in Chapter 4 is modified to adapt to the variations introduced by these distributions.

In addition, this chapter provides a sensitivity analysis of power consumption using the most influential network parameters. The size of the video in GB influences optimum cache sizes as the video size is linked with the power consumption of caching a video and the availability of storage space. Moreover, the power consumption of caches and IP router ports indicates the benefit of caching. The IP over WDM implementation also effects the power consumption by dictating the number of occupied router ports. Results show the impact each of these parameters has on the power consumption of the network.

5.2 Content Popularity Distributions and the MILP Model

This section evaluates the influence of content popularity distribution on power consumption. In addition, this section generalises the original minimum-power MILP model required to evaluate the power consumption of the network under the considered popularity distributions.

5.2.1 The Influence of Content Popularity Distributions

The content popularity distribution of videos provides an accurate measurement to calculate the number of hits each video receives in the video service. To carry out the power consumption evaluation, four distributions are considered: Zipf, Bimodal, Pareto and Equal Popularity distributions. These distributions have been explored in Chapter 2.

Video popularities can be very diverse as in a Bimodal distribution, or can be alike as in an Equal Popularity distribution. The objective of this evaluation is to find out the best caching strategy to achieve the most power efficiency and how this strategy varies with the variance in video popularities. The optimum number of videos to store in caches and resultant cache hit ratios are utilised to calculate the cache power consumption and the remaining network traffic which is used in turn to compute the video delivery power consumption. Therefore, the content popularity distribution highly influences the power consumption of the network and optimum cache sizes that minimise power.

5.2.2 Minimum Power MILP Model

In order to carry out this evaluation, the MILP model introduced in Chapter 4 is taken into account. To apply the four considered content popularity distributions to the MILP model, an independent piecewise linear approximation equation describing the relationship between a cache size and its hit ratio is required under each distribution.

Here the parameters appended to the original MILP model are introduced and the modifications applied to the equations are explained.

In addition to the parameters that the original MILP model defines, the modified MILP model utilises the following parameters:

τ	Programme type in the Bimodal distribution (example $\tau = 1$ represents news-type programmes and $\tau = 2$ represents drama-type programmes)
μ_τ	Bimodal distribution mean of content type τ
σ_τ^2	Bimodal distribution variance of content type τ
$mMin$	Pareto distribution minimum possible value of m
β	Popularity of the most popular video in the Pareto distribution
$Vsize$	Video size in GB
V_{it}	Cache size at node i at time t in number of videos ($= M_{it}/Vsize$)
Tot	Total number of videos in the library

Constraints (4-14) and (4-15) in Chapter 4 are generalised to include four different content popularity distributions as follows:

$$
\sum_{j\epsilon N:i\neq j} \lambda_{d\,ijt}^{\;xy} - \sum_{j\epsilon N:i\neq j} \lambda_{d\,jit}^{\;xy}
= \begin{cases}
\begin{cases}
\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} \frac{1}{(m \cdot lnTot)} \right) & i = x \\
-\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} \frac{1}{(m \cdot lnTot)} \right) & i = y \\
0 & otherwise
\end{cases} & Zipf \\
\begin{cases}
\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{k=1}^{2} \left(2 \cdot \left(\sum_{m=\mu_k+1}^{\mu_k+(V_{it}-2)/4} \frac{e^{\frac{-(m-\mu_k)^2}{2\sigma_k^2}}}{2\sqrt{2\pi\sigma_k^2}} \right) + \frac{1}{2\sqrt{2\pi\sigma_k^2}} \right) \right) & i = x \\
-\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{k=1}^{2} \left(2 \cdot \left(\sum_{m=\mu_k+1}^{\mu_k+(V_{it}-2)/4} \frac{e^{\frac{-(m-\mu_k)^2}{2\sigma_k^2}}}{2\sqrt{2\pi\sigma_k^2}} \right) + \frac{1}{2\sqrt{2\pi\sigma_k^2}} \right) \right) & i = y \\
0 & otherwise
\end{cases} & Bimodal \\
\begin{cases}
\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} \beta \cdot mMin^{\beta} / m^{\beta+1} \right) & i = x \\
-\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} \beta \cdot mMin^{\beta} / m^{\beta+1} \right) & i = y \\
0 & otherwise
\end{cases} & Pareto \\
\begin{cases}
\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} 1/Tot \right) & i = x \\
-\lambda^{xyt} \cdot Rd \cdot \delta_x \cdot \left(1 - \sum_{m=1}^{V_{yt}} 1/Tot \right) & i = y \\
0 & otherwise
\end{cases} & Equal
\end{cases}
\quad \forall\, i, x, y \in N, \forall\, t \in T \tag{5-1}
$$

$$AP_{it} = \begin{cases} \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot \left(1 - \sum_{m=1}^{V_{it}} 1/(m \cdot lnTot) \right) \right) / B & Zipf \\ \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot \left(1 - \sum_{k=1}^{2} \left(2 \cdot \left(\sum_{m=\mu_k+1}^{\mu_k + (V_{it}-2)/4} \frac{e^{\frac{-(m-\mu_k)^2}{2\sigma_k^2}}}{2\sqrt{2\pi\sigma_k^2}} \right) + \frac{1}{2\sqrt{2\pi\sigma_k^2}} \right) \right) \right) / B & Bimodal \\ \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot \left(1 - \sum_{m=1}^{V_{it}} \beta \cdot mMin^{\beta} / m^{\beta+1} \right) \right) / B & Pareto \\ \sum_{y \in N: y \neq i} \left(\lambda^{iyt} + \lambda^{iyt} \cdot Ru \cdot \delta_y + \lambda^{yit} \cdot Rd \cdot \delta_y \cdot \left(1 - \sum_{m=1}^{V_{it}} 1/Tot \right) \right) / B & Equal \end{cases}$$

$$\forall\, i \in N, \forall\, t \in T \tag{5-2}$$

Constraint (5-1) is the flow conservation constraint for download traffic considering different content popularity distributions. Constraint (5-2) calculates the number of required aggregation ports at each node under various content popularity distributions.

5.3 Power Consumption Evaluation of Popularity Distributions

In order to evaluate the influence of different content popularity distributions, the modified MILP model is run using different input values for the associated parameters. All other input values are similar to those used in Chapter 4.

Four distributions are considered in the evaluation: Zipf, Bimodal, Pareto and Equal Popularity distributions. The popularities of videos under each distribution are shown in Figure 5-1. The next sections explain the input values applied to the MILP model and power consumption results under each distribution.

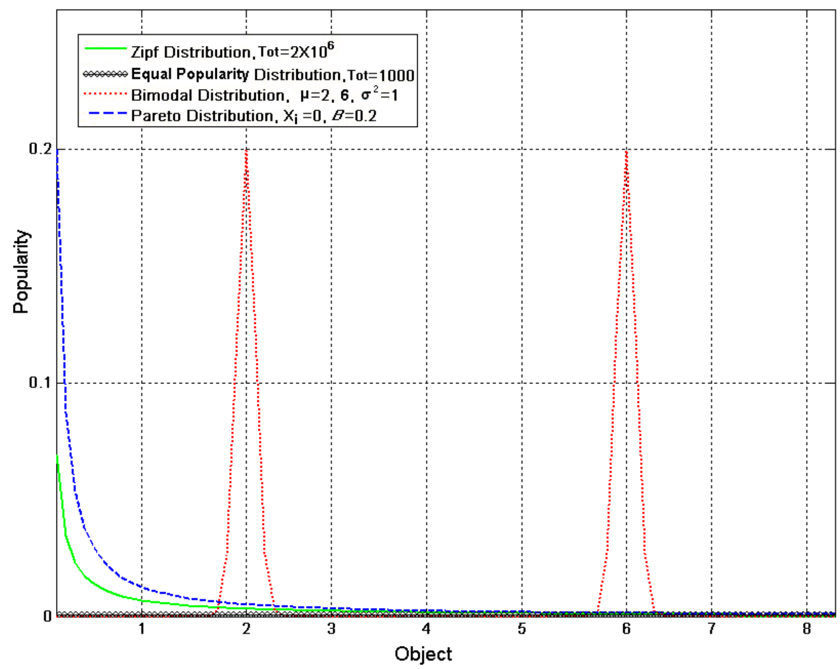

Figure 5-1: The popularity of videos under the Zipf, Bimodal, Pareto and Equal Popularity distributions

5.3.1 Input Parameters

A Zipf distribution represents a YouTube-like service having a total of 2 million videos in the library. For simplicity, the library is used without any interim updates during the considered 24 hours. The Bimodal distribution signifies a Catch-up TV service that stores highly watched TV programmes available to users all the time. The number of videos in such a service is relatively smaller compared to other services. A library of 1000 videos is assumed divided into two categories, each category has a Gaussian distribution with the same variance, but a different mean. Three values of σ^2

are considered to examine the impact of video popularity diversity on the optimum cache size: 1) σ^2 =0.2, 2) σ^2 =1 and 3) σ^2 =5.

The Pareto distribution describes a VoD service assuming a library of 1000 videos taking into account three different values of the parameter β representing three different popularity values for the most popular video. These values are: 1) β=0.2, 2) β=0.4, and 3) β=0.6. The Equal Popularity distribution exemplifies a service where all videos are of the same popularity. Three services are assumed having a library of: 1) Tot =1000, 2) Tot =10000 and 3) Tot =2 million videos to evaluate the influence of the total number of videos on the popularity of videos and on power efficiency.

5.3.2 Power Consumption and Optimum Cache Sizes

Table 5-1 shows the popularity of the most popular video under different distributions considering different values of the relevant distribution parameters as well as the optimum cache sizes found by the MILP model. The videos are assumed to be of the same size of 200MB and the power consumption of deployed caches is 7.4W/GB. Figure 5-2 shows the power consumption of the network at different times of the day under different content popularity distributions.

Table 5-1 Popularity values and optimum cache sizes under different content popularity distributions

Distribution	Parameter Value	Popularity of the Most Popular Video	Optimum Cache Size
Bimodal	$\sigma^2 = 0.2$	0.44	5
	$\sigma^2 = 1$	0.19	25
	$\sigma^2 = 5$	0.08	50
Pareto	$\beta = 0.2$	0.2	variable
	$\beta = 0.4$	0.4	40
	$\beta = 0.6$	0.6	10
Zipf	Tot = 2 million	0.068	variable
Equal Popularity	Tot = 1000	0.001	1000
	Tot = 10000	0.0001	variable
	Tot = 2 million	0.0000005	0

5.3.3 Analyses of Results under Each Popularity Distribution

In the following, the obtained outcomes for power consumption and optimum cache sizes are analysed under each considered content popularity distribution.

5.3.3.1 Zipf Distribution

The power consumption under the Zipf distribution is higher compared to the other distributions as popularities of the most popular videos are low. This requires either storing a large number of videos in caches to achieve high cache hit ratios which consumes high power for storage, or caching less videos and consequently consuming more power on streaming content from servers. The resulting optimum caches that minimise power are variable during the day and their size depends on network traffic.

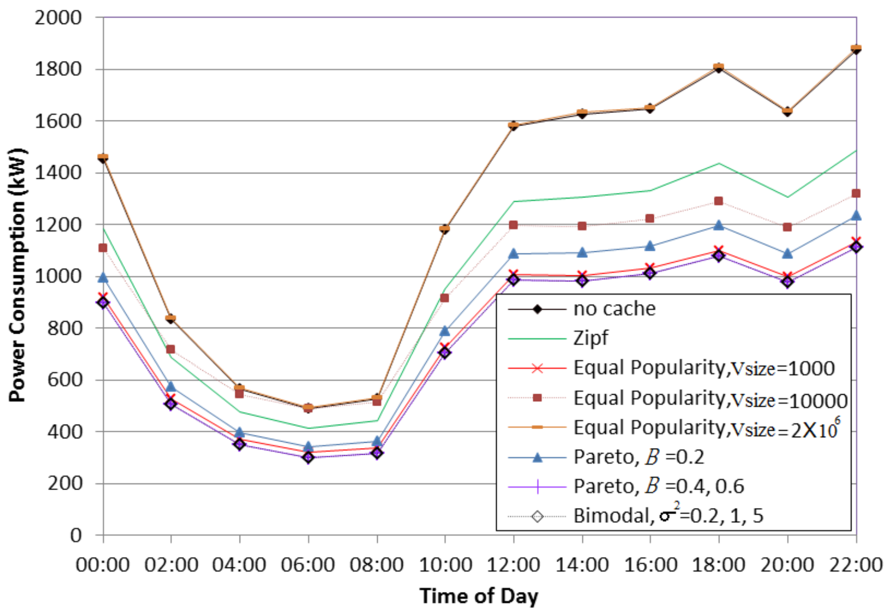

Figure 5-2: The power consumption of the network considering no caching and under Zipf, Bimodal, Pareto and Equal Popularity distributions with different parameter values

5.3.3.2 Bimodal Distribution

In contrast to the Zipf distribution, the Bimodal distribution results in the largest power savings due to the fact that the most popular videos have very high popularity values. Therefore storing a small number of videos achieves a high cache hit ratio. Under all values of σ^2, the power consumed at different times of the day is the same when the other network parameters are fixed. Optimising the cache sizes under the Bimodal distribution resulted in cache hit ratios of approximately 1. A fixed cache size for all the nodes in the network of 5, 25 and 50 videos achieves the maximum power efficiency for $\sigma^2 = 0.2$, 1 and 5, respectively.

5.3.3.3 Pareto Distribution

The Pareto distribution results in similar power consumption values to those of the Bimodal distribution under β=0.6 and β=0.4. For lower values of β, the power consumption of the network increases, since lower popularity values result in relatively larger caches and lower cache hit ratios making it more power-efficient to retrieve more content directly from servers. The maximum savings in power consumption are reached by deploying a fixed cache size of 10 and 40 when β=0.6 and β=0.4, respectively. Similar to a Zipf distribution, under a Pareto distribution having β=0.2 the maximum power savings are achieved by deploying variable cache sizes at the nodes.

5.3.3.4 Equal Popularity Distribution

Under Equal Popularity distribution, with a library of 1000 videos, maintaining a copy of the whole library at each node (cache hit ratios of 1) results in the minimum power consumption. However for a library of 2 million videos, the power consumption of the network is minimised when all content is retrieved directly from servers, i.e. when no caches are deployed. This is because storing a limited number of videos of low popularity (0.5×10^{-6}) will not achieve a significant cache hit ratio. The resultant power consumption is similar to that under no caching. The power consumption ranges between the two considered extremes as the total number of videos in the library rises from 1000 to 2 million. Minimising the power consumption when the library is made up of 10000 videos is realised by the deployment of variable cache sizes at the nodes. When the traffic is high, the optimum configuration is to store part of the content in caches whereas under low traffic (videos experience small number of hits) the optimum solution is no caching.

5.3.4 Power Saving Figures

Results reveal that optimising the cache size for each node at different times of the day under a downlink traffic demand ratio of *Rd*=1.5 and an uplink traffic ratio of *Ru*=0.2, introduces power savings of 20%, 34%, 40% and 39% for the Zipf, Pareto, Bimodal and Equal Popularity distributions, respectively. These savings increase to 42%, 68%, 72% and 71% under *Rd*=7.5 and *Ru*=1 considering the same distributions, respectively. The results also show that under the Bimodal distribution the network can achieve the maximum power savings by caching the 5 most popular objects.

5.4 Network Parameters and MILP Modifications

In this section, the video and network parameters considered in the sensitivity analyses are highlighted and the influence of each of these parameters on network power consumption is evaluated. The required changes to the original MILP model to perform the power consumption evaluation are also given.

5.4.1 Network Parameters and Power Consumption

In this section the significance of each considered network parameter is evaluated in terms of its impact on the power consumption of the network.

5.4.1.1 Video Sizes

The importance of studying the influence of the size of video objects on power efficiency is due to a number of reasons. Different services use videos of different sizes, i.e. a Catch up TV service has large 1-hour programmes, while downloadable videos of a social network are mostly short 3-minute clips. In addition, for some services the amount of storage required to store videos depends on the data rate, resulting in many possible video sizes with respect to data rate. Moreover, the advances in video delivery technologies have resulted in one video having several versions in Standard Definition SD and High Definition HD for example, of different sizes in GBs. Another possible reason is that using different data compression techniques in a service produces different video sizes, and switching from one technique to another affects the average video size.

5.4.1.2 Cache Power Consumption

The power consumed to store 1GB of data in a cache depends on the type of storage implementation whether it is hard disk arrays or a flash memory for example. In addition, there is ongoing research aiming at reducing the power consumption of network equipment including storage devises. Caching technology enhancements can improve the power efficiency of caching equipment resulting in possible changes in the feasibility of caching, and therefore caching decisions. The aim in this evaluation is to examine the extent of influence the cache power consumption has on the total network power consumption and the selection of the optimum cache sizes.

5.4.1.3 Power Consumption of a Router Port

The power consumption of a router port is the decisive factor of whether to stream a video remotely or to store it in a local cache since the router port is the most power consuming element in the IP over WDM network. The benefit of caching increases when the power consumption of streaming content from servers is high. The power consumption of other components in the network such as amplifiers, transponders, switches and multiplexing equipment also influence the power consumption of the network and the caching decision. However, since the IP router port consumes considerably more power compared to other network components (1000W per router port operating at 40Gb/s versus 8-85W for other components), this evaluation only considers router ports.

5.4.1.4 IP over WDM Implementation

Under lightpath non-bypass, IP router ports are engaged at each intermediate node that traffic traverses. In contrast, under lightpath bypass traffic is forwarded to the next optical node on the path without occupying router ports at that node. This difference implies that no power is consumed by router ports at intermediate nodes under lightpath bypass. Since router ports are major power consumers, the difference in the IP over WDM implementation is bound to influence the power consumption of the network.

5.4.2 Modifications to the MILP Model

To carry out the power consumption evaluation, the MILP model introduced in Chapter 4 is considered. However, the equation to calculate the power consumption of router ports is different under each IP over WDM implementation. Under lightpath bypass, router ports are considered only at the source and destination nodes and their power consumption is given in Objective equation (4-1) in Chapter 4. Under lightpath non-bypass, the power consumption of router ports is considered at each IP node connected to an optical node on the path that the traffic traverses, and therefore becomes:

$$\sum_{i \in N} Pp. \left(AP_{it} + \sum_{j \in Nm_i : i \neq j} w_{ijt} \right) \tag{5-3}$$

Constraints (4-2) and (4-3) in Chapter 4 that restrain the number of occupied ports at each node to a maximum considering lightpath bypass are modified to consider lightpath non-bypath and are defined as:

$$\sum_{j \in N} w_{ijt} + AP_{it} \leq RPmax_i \qquad \forall\, i \in N, \forall\, t \in T \tag{5-4}$$

$$\sum_{j \in N} w_{jit} + AP_{it} \leq RPmax_i \qquad \forall\, i \in N, \forall\, t \in T \tag{5-5}$$

5.5 Power Consumption Evaluation of Network Parameters

This section illustrates by results the influence of considered network parameters on power consumption.

5.5.1 Video Sizes

To investigate the impact of varying the average video size on power consumption, the cache sizes are optimised assuming three average video sizes: 1) *Vsize*=0.2GB, 2) *Vsize*=1GB and 3) *Vsize*=2GB.

Figure 5-3 illustrates the power consumption of the network using different video sizes. Here, the results under the four considered content popularity distributions are shown to generalise the results. The Bimodal distribution and the Pareto distribution are considered having a library of 1000 videos with σ^2=5 and β=0.2, respectively. The Equal Popularity distribution is assumed to have a library of 10000 videos. Under all distributions the power consumption increases as the size of videos increases. The increase is negligible under the Bimodal distribution as the number of videos stored in caches is limited. The largest variation in power consumption from one video size to another is observed under the Equal Popularity distribution. The reason behind this is the large number of videos that have to be stored in caches in order to make an influential increase in cache hit ratios.

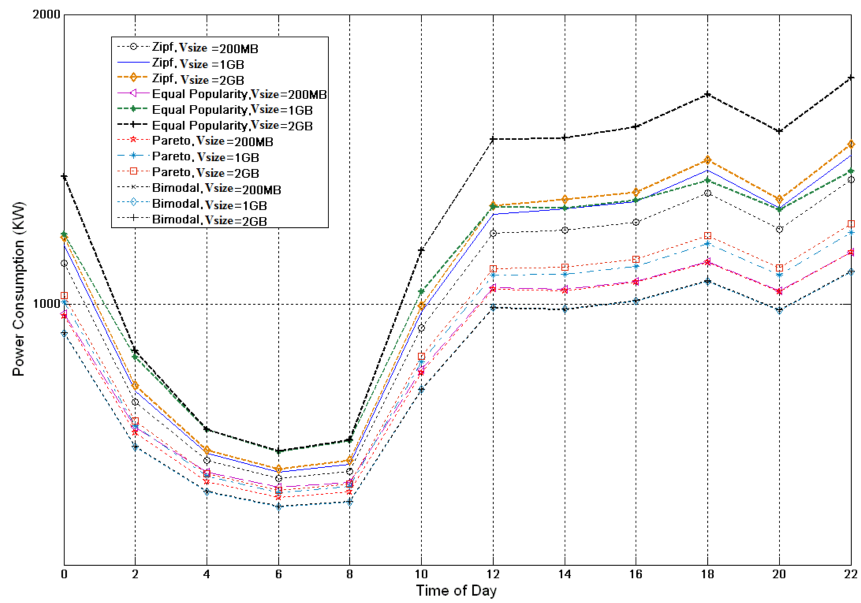

Figure 5-3: Power Consumption of the network at different times of the day considering different video sizes under different content popularity distributions

5.5.2 Cache Power Consumption

The power that a cache consumes to store a unit of data depends on the cache technology. Considering massive numbers of large videos, the power efficiency of a cache plays a significant role in the network overall power consumption. This study investigates the impact of the cache power efficiency considering three different values for the cache power consumption: 1) Φ =2.5W/GB, 2) Φ =5W/GB and 3) Φ =7.5W/GB.

The power consumption of the network is evaluated under four popularity distributions deploying caches of different power efficiencies. Video sizes are assumed to be 200MB and the parameter values of the distributions are similar to those in the previous section. As depicted in Figure 5-4 the power consumption of the network increases under all distributions as the power consumption of caching 1GB of data increases.

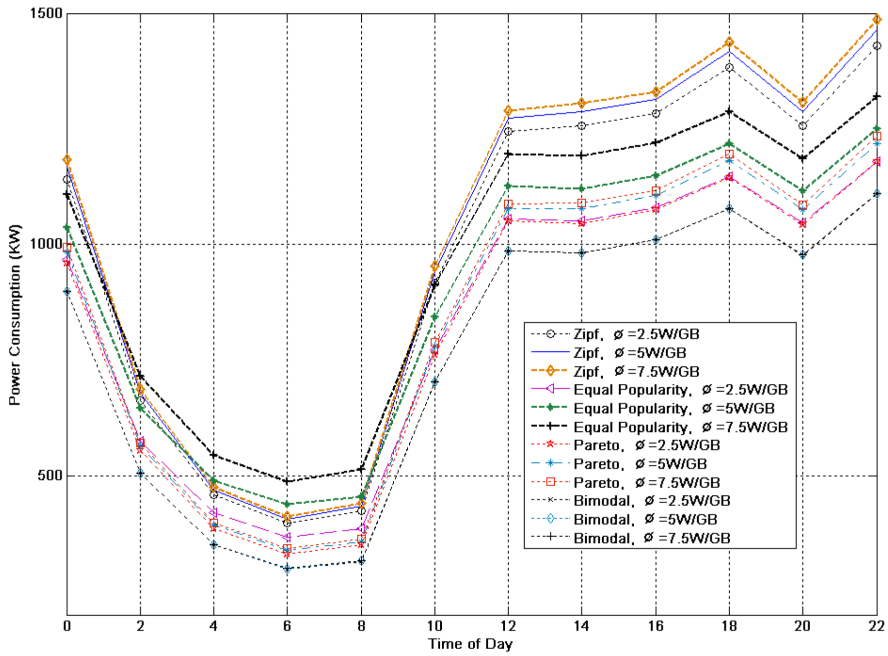

Figure 5-4: Power Consumption of the network at different times of the day considering caches of different power consumption under different content popularity distributions

Under the Equal Popularity distribution the optimum cache size averaged over the time of the day decreases from 10000 to 8450 videos when the cache power consumption rises from 2.5 to 7.5W/GB since the storage of videos becomes less power-efficient. Under the same rise in cache power consumption considering the Zipf distribution, the averaged optimum cache size for the network decrease from 5120 to 1670 videos. Under the Bimodal and Pareto distributions the average cache size in the network remains the same. When a library of 2 million videos is assumed under the Equal Popularity distribution, the network power consumption is not affected as no videos are stored in caches. The power consumed in this case is similar to the situation where no caches are deployed in the network. The Zipf distribution results in the highest power consumption compared to other popularity distributions.

5.5.3 Router Port Power Consumption

The power consumption of an 8-slot CRS-1 was reported at 8000W at the start of this work in 2011. Due to technology enhancements this value was reduced to 4834W in 2012. As a result, the estimated power consumption of a single 40Gb/s port *Pp* was reduced from 1000W to 604W. Figure 5-5 shows the influence of the router port power consumption *Pp* on the network power consumption with no caching and under variable caches. The input values for the cache power consumption and video size are 7.4W/GB and 200MB, respectively. The popularity distribution of content is considered to follow a Zipf distribution.

Since router ports are the major power consumers, the power consumption of the network reduces when router ports are more power efficient under no caching and variable caching. The power consumption of the network falls by 36% and 35% (maximum and average) under no caching and variable caching, respectively. Under variable caching, the optimum cache sizes that minimise power consumption averaged over the time of the day drop from 1660 to 1080 videos when *Pp* falls from 1000W to 604W. When *Pp*=1000W the power consumption to download videos from servers is high, making it more power efficient to store more videos at caches deployed in local nodes. As this value reduces the feasibility of consuming more power on caching reduces. Nevertheless, the total network power savings due to caching are similar considering the two values of *Pp* (19% maximum and 16% average). Consequently, for the remainder of the thesis the most recent value for the power consumption of an 8-slot CRS-1 (*Pp*=604W) is taken into account to produce more up-to-date results.

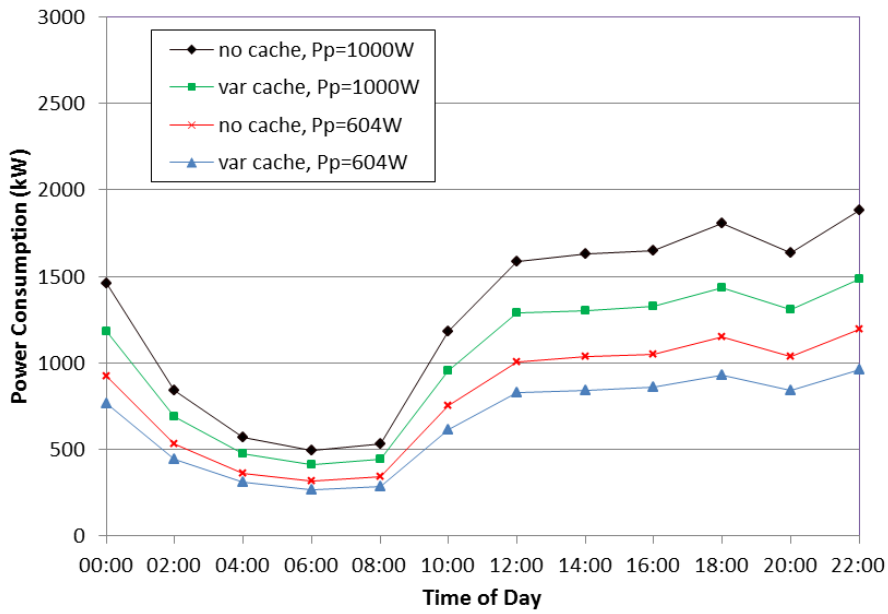

Figure 5-5: Network power consumption with no caching and under variable caching considering previous and current values for the power consumption of a router port

5.5.4 IP over WDM Implementation

Under lightpath non-bypass, the traffic passing through an intermediate node in the optical layer is forwarded to the corresponding intermediate node in the IP layer. This results in additional power being consumed in the IP layer (including router ports which are the main power consumers in the network). Figure 5-6 shows the power consumption of the network under the two IP over WDM implementations considering no caching and variable caching. The influence of the IP over WDM implementation on power consumption increases when the average number of hops required to deliver content is high. In other terms, implementing lightpath bypass reduces more power when the average number of bypassed intermediate nodes is high.

Savings in power consumption are up to 21% (19% average) and 19% (16% average) under lightpath non-bypass and lightpath bypass, respectively. When variable caches are considered, the average optimum variable cache sizes that minimise power consumption are 2600 and 1660 videos under lightpath non-bypass and lightpath bypass, respectively. The reason behind this is that lightpath bypass consumes less power on transport, favouring more videos to be delivered remotely.

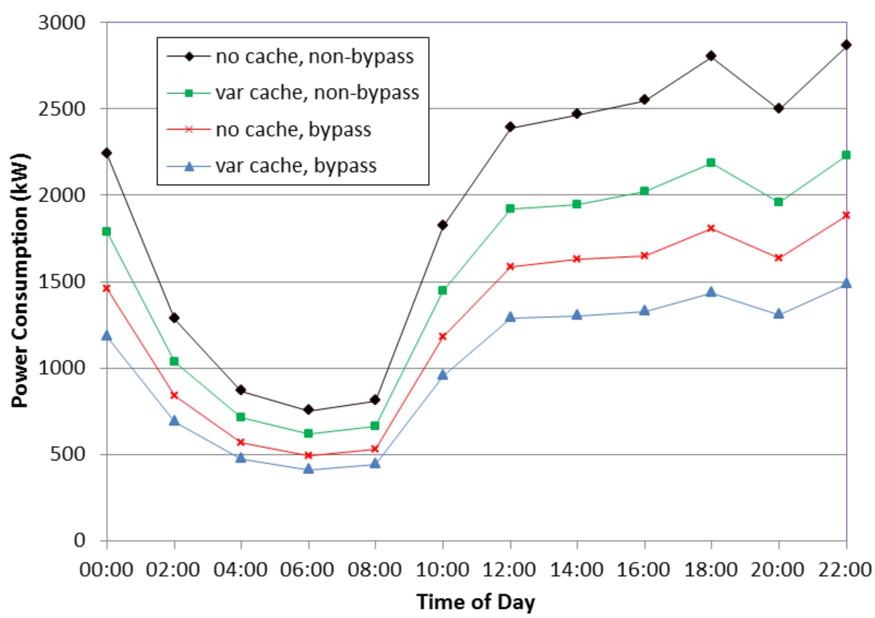

Figure 5-6: Network power consumption with no caching and under variable caching considering lightpath non-bypass and lightpath bypass

5.6 Summary

This chapter has explored the influence of content popularity distribution on the network power consumption. The popularity distribution which content follows decides the popularity of each video in the service. Therefore, it directly determines the cache hit ratio from the number of stored videos. When video popularities are highly diverse, the most power efficient solution is to store the few most popular videos in caches reducing up to 72% of network power consumption. Caches of variable sizes are deployed to minimise power consumption when video popularities are less diverse. Specifying the amount of content to cache under different video services increases the diversity of potential application areas of this work.

The chapter has also evaluated the influence of significant parameters on the network power consumption. The power consumption of the network rises following increase in video sizes, power consumption of caches and power consumption of IP router ports. If technology advances in the transport network are not met by matching power reduction strategies in storage devices, caching is expected to become less favourable. Bypassing intermediate IP nodes on the path that traffic traverses is an effective technique for power reduction.

Chapter 6 Energy Efficient Future IPTV

6.1 Introduction

The rapidly growing IPTV market has resulted in increased traffic volumes raising concerns over Internet energy consumption. In this chapter, the dynamics of TV viewing behaviour and programme popularity are explored in order to devise a strategy to minimise energy usage. The power consumption of IPTV delivered over an IP over WDM network is calculated. The evaluation considers standard definition TV (SDTV) as the most common video delivery technology for today's TV channels and high definition TV (HDTV) as the emerging technology rapidly replacing SDTV. The state of art BT 21CN topology is introduced under different scenarios and compared to the NSFNET topology. The aim of the work in this chapter is to evaluate the power consumption of watching SDTV and HDTV programmes considering real TV data and two network topologies. Fixed and variable caches are employed in the network to investigate the savings in power consumption due to caching TV programmes. The results illustrate these power savings as well as identify optimum fixed and variable cache sizes.

6.2 TV Programme Popularity

In Chapter 4 and Chapter 5, the power consumption of different video services was minimised using Zipf, Pareto, Bimodal and Equal Popularity distributions to describe the popularity of video content. These distributions allow a close-to-real environment to evaluate the power consumption of content delivery networks and enable comparisons between different video services. However, using real data to calculate traffic demand and video popularities in a video service gives a better insight into actual input values for these services and provides a more accurate evaluation of the network.

One day TV viewing data from Friday the 9th November 2012 is obtained from [96] [97]. The number of viewers of the most popular TV programmes is mapped against the total number of viewers at each time of the day. The number of viewers for each TV programme and the times these programmes aired were collected from [96] and [97], respectively. A viewer of a TV programme is considered among the programme audience if the viewer watches 3 minutes or more of the programme. The number of viewers of the most popular TV programmes is based on a sample of viewers. The total number of viewers is the average daily viewing of all TV channels measured over

November 2012 [96], [98]. The number of viewers per TV programme is used to calculate the popularity of the programme and hence its potential contribution to the network traffic due to user requests.

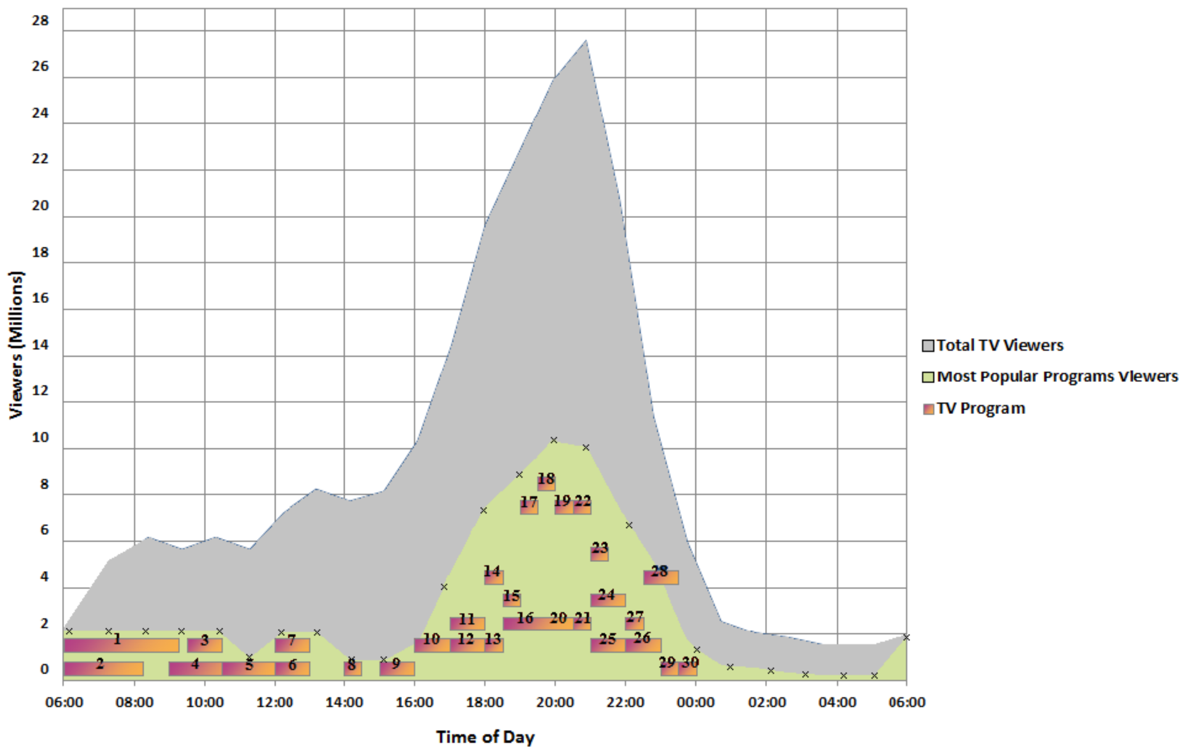

1 Breakfast (BBC1)
2 Daybreak (ITV1)
3 The Jeremy Kyle Show (ITV1)
4 The Wright Stuff (Channel 5)
5 This Morning (ITV1)
6 Daily Politics (BBC2)
7 Loose Women (ITV1)
8 Neighbours (Channel 5)
9 Countdown (Channel 4)
10 Deal or no Deal (Channel 4)
11 Antiques Road Trip (BBC2)
12 Neighbours (Channel 5), Come Dine with Me (Channel 4)
13 Eggheads (BBC2), The Simpsons (Channel 4)
14 Six O'Clock News (BBC1)
15 ITV News (ITV1)
16 Strictly Come Dancing - It Takes Two (BBC2)
17 Emmerdale (ITV1)
18 Coronation Street (ITV1)
19 Eastenders (BBC1)
20 Mastermind (BBC2)
21 Gardeners' World (BBC2)
22 Coronation Street (ITV1)
23 Have I Got News for You (BBC1)
24 Who Wants to Be a Millionaire Specials (ITV1)
25 The Mentalist (Channel 5)
26 Alan Carr: Chatty Man (Channel 4), Castle (Channel 5)
27 QI (BBC2)
28 The Graham Norton Show (BBC1)
29 Family Guy (BBC3)
30 Family Guy (BBC3)

Figure 6-1 The average daily TV viewing figures, the most popular TV programmes and their number of viewers at each time of the day viewed on Friday November the 9th 2012

Figure 6-1 lists the most popular TV programmes and their corresponding broadcasting TV channels as well as the number of viewers for each programme in the UK. It also plots the total number of TV viewers during the day, and therefore can be used to estimate the amount of traffic devoted to the most popular programmes. The most popular TV programmes would account for a significant portion of network traffic, particularly at primetime. This observation justifies the use of caching the most popular TV programmes to reduce the duplicate traffic and therefore reduce the power consumption of transport. Moreover, requests for other TV Catch-up services (CuTV) such as the BBC iPlayer and Channel 4 on Demand (4oD) can be served from content already stored in caches.

It is worth mentioning that mainstream TV today is usually broadcast over the air and where scheduled TV is delivered over IP, IP multicast is the technology of choice. However, two recent trends motivated us to consider cache solutions:

1. Firstly, whilst viewers prefer to watch scheduled TV, they increasingly expect more control over the timing of their viewing – to be able to start a programme slightly late (a need that is driving the +1 channels) or to pause a viewing temporarily [99].
2. Secondly, there is a rise in emerging rich streaming space served by YouTube live streams, Twitch.tv, etc. which have already attracted millions of viewers.

Little is known yet about user behaviour and content popularity of these services, leaving mainstream TV data the closest available alternative. Mainstream TV viewing patterns can be a good reflection of future IPTV viewing patterns as, despite the availability of CuTV, most users prefer to watch TV close to schedule. The use of social media to communicate between friends whilst watching TV may in part be driving this trend. However, certain types of programmes may be less suited to a cache solution (e.g. live sports). Therefore, it is important to estimate the proportion of live content to decide if caching is beneficial. The number of TV programmes broadcasted on 91 major UK channels are obtained from [100] and the number of live and non-live TV programmes shown between the 14th and the 20th of September 2013 are calculated. The methodology includes basic channels in addition to entertainment, documentary, lifestyle, films, sport and children channels. In total, 14323 TV programmes were aired during the time considered. Of these, 322 were live shows, accounting for only 2.25% of broadcast TV programmes.

6.3 Test Networks

The study considers two physical topologies as test networks. The NSFNET topology is selected as a communication infrastructure widely used in network studies, and the UK BT 21CN topology is introduced as a Next Generation Network (NGN) network. Including the NSFNET enables other researchers to compare their work to the outcomes of this work. The UK BT 21CN topology is also included since the TV viewing data used in the evaluation is based on UK audience. Original analysis on the BT 21CN topology is provided as not much research has been reported on this network.

Considering these two topologies allows comparison between a European network where population densities are high and an American network with large distances between content locations.

6.3.1 The NSFNET Topology

The NSFNET topology of 14 nodes and 21 links is shown in Figure 6-2, with distances between nodes given in kilometres. Each node in the network is assumed to have a video cache. With respect to the number and location of video head-ends in the network, three options are considered:

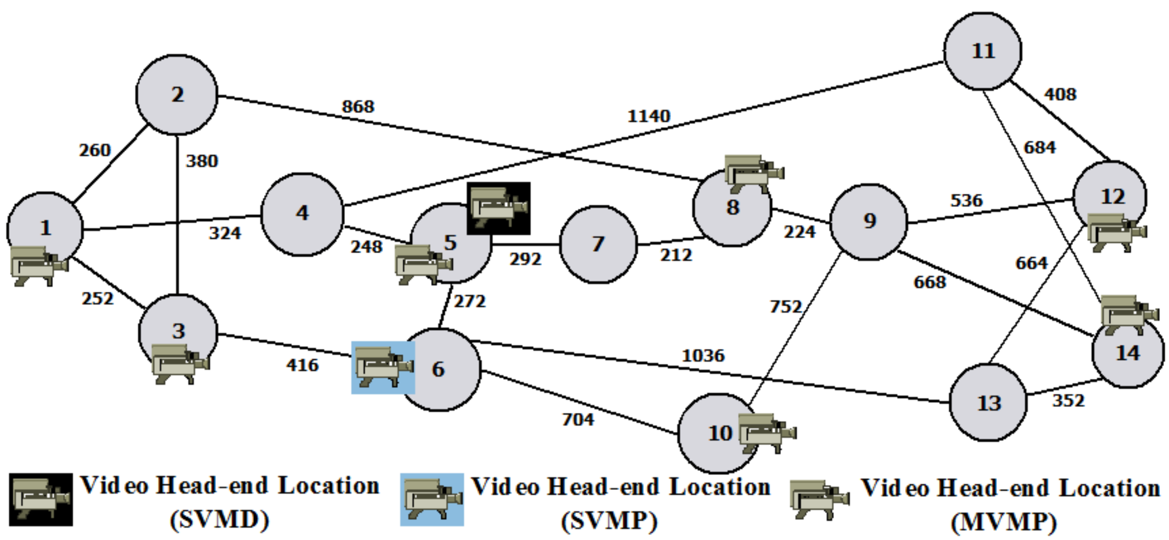

Figure 6-2: The NSFNET topology with video head-end locations under SVMD, SVMP and MVMP and fibre length in kilometres

6.3.1.1 A Single Video Head-end with Minimum Delay Location (NSF-SVMD)

The location of the video head-end is optimised such that the propagation delay to end users is minimised. Other end-to-end delays are not considered here as they are proportional to the amount of streamed video traffic. The location of the video head-end under NSF-SVMD is node 5. The optimum locations of the video head-ends assuming no caching were obtained using the MILP model proposed in Chapter 4. This required that the server location indicator is set as a variable rather than a parameter. The MILP model was re-run to find the optimum server locations with respect to the total number of servers in the network and the target parameter (delay here and power consumption for the next two options).

6.3.1.2 A Single Video Head-end with Minimum Power Location (NSF-SVMP)

The location of the video head-end is optimised to minimise the power consumption of the network. Unlike in SVMD where the distance between nodes is the dominant factor in finding the optimum location of the video head-end, here minimising the number of hops that content traverses from source to destination is the key element under SVMP. This is explained by the fact that the most power consuming network components are routers which are utilised at each hop in the communication path. The optimum location of the video head-end under NSF-SVMP is node 6.

6.3.1.3 Multiple Video Head-ends with Minimum Power Locations (NSF-MVMP)

A further evolution is considered where 7 video head-ends are assumed, each injecting unique content into the core network. This is the most complex yet realistic scenario[87], [88]. The locations of video head-ends are optimised to minimise the overall network power consumption. Since no content replication is considered among video head-ends, the comparison of NSF-SVMD, NSF-SVMP and NSF-MVMP is not directly possible.

6.3.2 The BT 21CN Topology

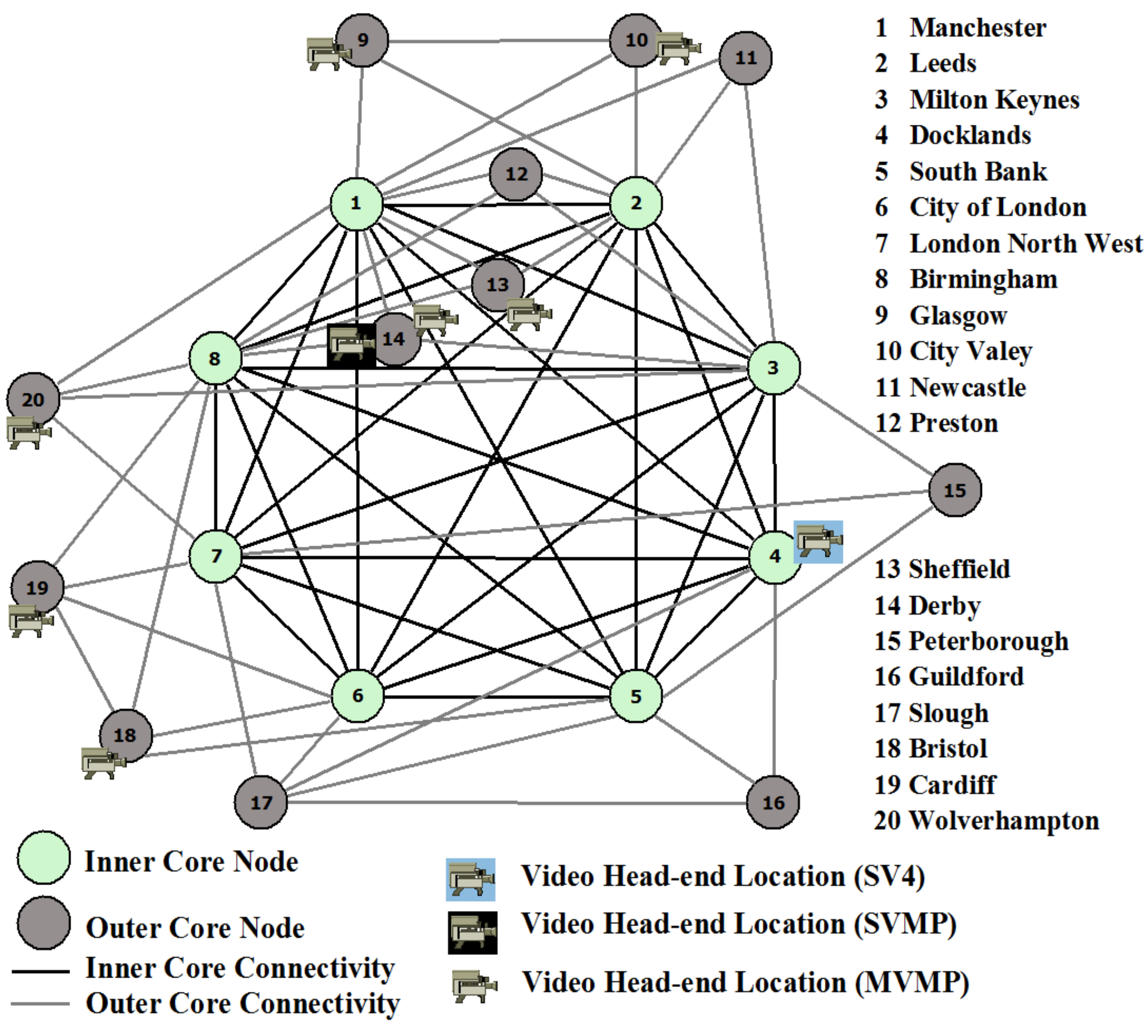

Figure 6-3: The BT 21CN topology with inner and outer core node locations

The BT 21CN (British Telecom 21st Century Network) is a NGN implemented by BT. Its core topology shown in Figure 6-3 consists of 20 nodes and 68 links. Core nodes are divided into inner core nodes (8 nodes) which are fully meshed and outer core nodes (12 nodes) which are connected to at least three other core nodes [101]. In order to determine the topology information required to carry out the investigation, the network core node connectivity is obtained from [101] and the core node locations from [102]. Three topologies under BT 21CN are considered with respect to server location:

6.3.2.1 A Single Video Head-end at Node 4 (BT 21CN-SV4)

This approach represents the current network situation as Telehouse located in Docklands (node 4) is the major peering location for the UK.

6.3.2.2 A Single Video Head-end with Minimum Power Location (BT 21CN-SVMP)

Similar to the method employed under the NSFNET topology, the location of the single video head-end that minimises the power consumption is found, which in this case is node 14.

6.3.2.3 Multiple Video Head-ends with Minimum Power Locations (BT 21CN-MVMP)

The locations of 7 video head-ends in the network that minimise overall power are found using the same approach used under the NSFNET topology. These locations are node 9, 10, 16, 14, 18, 19 and 20. The locations of video head-ends under all considered topologies are shown in Figure 6-3.

The two network topologies considered are small. Nevertheless, note that this work is concerned with core networks and they are typically of this size. Heuristics based on the MILP model insights can be developed and applied to larger networks.

Table 6-1 compares the NSFNET and BT 21CN topologies in terms of coverage, average hop count and average nodal degree.

Table 6-1: Comparison of the NSFNET and BT 21CN topologies

Property	**NSFNET**	**BT 21CN**
Coverage	Continental	National
Average hop count	1.7	1.8
Average nodal degree	3	6.6

6.4 Input Parameters

Table 6-2: Input data for the MILP model

Distance between two neighbouring EDFAs (*S*)	80 (km)
Number of wavelengths in a fibre (*W*)	16 [32]
Capacity of a wavelength (*B*)	40 (Gb/s)
Power consumption of a router port (*Pp*)	604 (W) [90]
Power consumption of a transponder (*Pt*)	73 (W) [32]
Power consumption of an EDFA (*Pa*)	8 (W) [92]
Power consumption of an optical switch (*Po*)	85 (W) [93]
Power consumption of a multiplexer/demultiplexer (*Pmd*)	16 (W) [91]
Cache power consumption factor (∅)	7.4 (W/GB)

The power consumption parameters used in the evaluation are shown in Table 6-2. The power consumption of an 8-slot CRS-1 is 4834W[90]. Each slot can contain one 40Gb/s port and therefore the power consumption of a 40Gb/s router port is estimated at 604W.

Chapter 4 explained how the energy consumption of streaming data is converted from J/b to W/GB to utilise in the MILP model. The time required to stream 1GB was 228 seconds resulting in a cache power consumption of 7.4W/GB. Note that the 228 seconds streaming time results in one 1-hour TV programme being delivered in approximately 2.5 minutes and 17 minutes under SDTV and HDTV, respectively. To reduce this time for example to 1 minute under HDTV, 17 parallel disks are required, however the hard disk access speeds continue to improve and access speeds below 1 minute in this case may become possible in the near future.

To generate the network video traffic demand, all the TV demand of Figure 6-1 is assumed to be carried over an IP over WDM network, at standard or high definition rates. For simplicity, user requests are assumed to be equally distributed among the nodes. A more complex and more accurate approach is to distribute user requests with respect to population density which requires further investigation. The IP over WDM implementation considered in the evaluation is lightpath non-bypass where router ports are occupied in each intermediate node allowing the operator more security and packet correction.

Assuming the compression standard MPEG-4 AVC (Motion Picture Experts Group 4 - Advanced Video Coding), the Digital Subscriber Line Forum (DSL Forum) recommends bit rates between 1.5Mb/s - 3Mb/s for SDTV and 8Mb/s - 12Mb/s for HDTV [103]. Here 1.5Mb/s is assumed for SDTV and 10Mb/s for HDTV producing 1-hour TV programmes of size 675MB and 4.5GB under SDTV and HDTV, respectively. Although MPEG-4 AVC is not as widely deployed as MPEG-2, it is expected to become more commonplace as it delivers up to 50% bit rate savings, offering better bandwidth utilisation, deployment of advanced HDTV services and lower power consumption due to reduced download times.

The number of viewers of each TV programme and the total viewers are utilised to calculate programme popularities. Note that the most popular TV programmes shown in Figure 6-1 account for only a trivial number of the entire TV programmes broadcasted on the day (30 out of over 2000 programmes). The MILP model finds optimum cache sizes that minimize power consumption, and are expected to expand beyond 30 programs. As a result, the popularity of remaining TV programmes must be found. Since these values are not available, an extrapolation is used to estimate these values based on the popularity of the known 30 most popular programmes. The popularity of TV programs generated by the extrapolation are utilised in the model in the piecewise linear approximation equation. Since the popularity of less popular programmes are insignificant compared to the 30 most popular programmes, the consequent error due to the estimation is negligible.

The utilised TV viewing data is based on UK audience. US TV programme popularity and viewing data are not available in public domain to the best of our knowledge. Therefore, the same TV traffic demand derived from Figure 6-1 is applied to both network topologies (BT 21CN and NSFNET).

6.5 Energy Efficiency of Caching in Future IPTV

Requests for TV programmes broadcasted from a video head-end are generated at different nodes in the network. Dedicating a stream for each programme request, results in consuming high power in transport. Therefore as traffic increases it becomes more power-efficient to store popular programmes in local caches. Generally, deploying caches at network nodes reduces the traffic passing through the network and the amount of reduction in traffic at each time of the day depends on the volume of traffic. It would be valuable to apply the obtained real TV viewing data to the minimum-power MILP model explained in Chapter 4 (considering fixed and variable caches) to explore the potential power savings and optimum cache sizes for the scenarios considered.

6.5.1 Caches of Fixed Sizes

To explore the power savings introduced by caching popular content locally, caches of fixed sizes are assumed at each node in the network. Applying the obtained TV viewing profile, the optimum fixed cache sizes that minimise power consumption under SDTV and HDTV are 200GB and 650GB corresponding to approximately 37 and 18 1-hour TV programmes, respectively. Figure 6-4 shows the power consumption of watching SDTV when no caching is considered and when each node in the network is allocated its optimum cache size under NSFNET and BT 21CN.

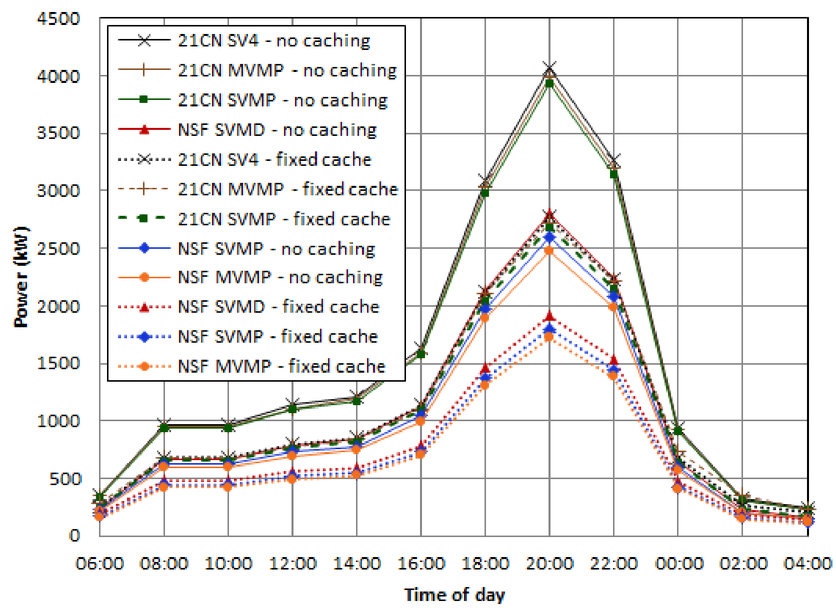

Figure 6-4: Power consumption of watching SDTV programmes with no caching and considering caches of a fixed size of 200GB under SVMD, SVMP, MVMP (NSFNET) and SV4, SVMP and MVMP (BT 21CN)

The power consumption of watching HDTV under the same scenarios is shown in Figure 6-5. One observation is the higher network power consumed under HDTV which is explained by the larger sizes of HDTV programmes which consume more power when streamed through the core network. Deploying caches of fixed sizes achieves similar power savings at each time of the day which results in the largest power reduction gained when the traffic is high. Figure 6-8 displays maximum and average power savings achieved by deploying caches of fixed sizes at the nodes. The power savings introduced by caches of fixed sizes are comparable under SDTV and HDTV considering all network topologies, as maximum savings range between 31%

and 38%. Slightly higher power savings are achieved assuming the BT 21CN topology under HDTV.

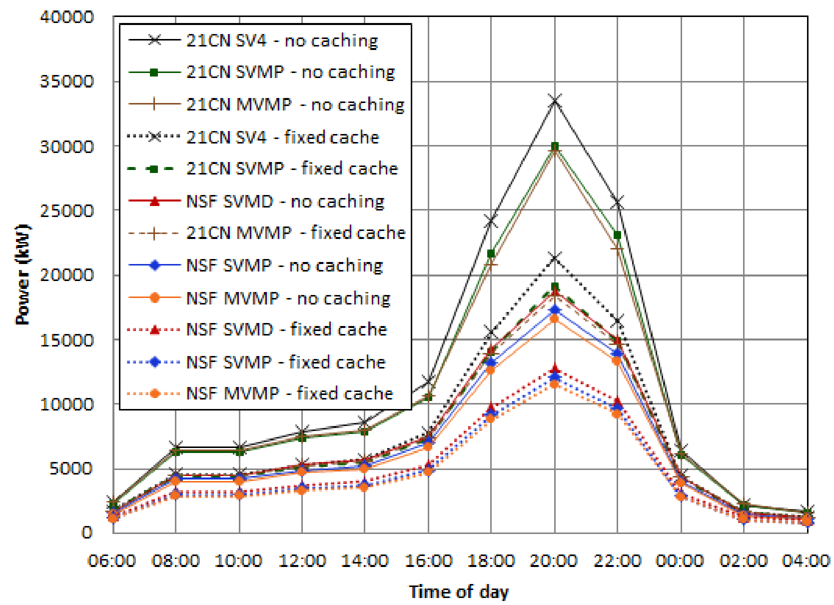

Figure 6-5: Power consumption of watching HDTV programmes with no caching and considering caches of a fixed size of 650GB under SVMD, SVMP, MVMP (NSFNET) and SV4, SVMP and MVMP (BT 21CN)

6.5.2 Caches of Variable Sizes

Operating large caches during off-peak hours consumes unnecessary power, as it is more power-efficient to stream programmes from video head-ends directly. TV viewing data is applied to the variable cache minimum-power MILP model proposed in Chapter 4 to find the optimum cache size for each node at each time of the day that minimises power consumption. The results confirm that using caches of larger sizes is power-efficient only during primetime whereas the network power consumption during off-peak hours is minimised when smaller caches are utilised.

Figure 6-6 and Figure 6-7 show the power consumption of watching SDTV and HDTV programmes under variable cache sizes considering NSFNET-SVMP and BT 21CN-SV4, respectively. They also show the power consumption of the network under fixed caches for comparison. The power consumption is similar under all network topology scenarios, therefore for clarity, one scenario for each topology is shown. The maximum and average savings in network power consumption with caches of variable sizes considering all network topologies are shown in Figure 6-8. Maximum savings in power consumption are up to 40% under HDTV BT 21CN-MVMP. The improvement in power efficiency of caches of variable sizes over the power savings achieved by caches of fixed sizes is limited. Nevertheless, it is worth investigating the situation where cache sizes are optimised over the time of the day to minimise power consumption in order to observe the variance in optimum cache sizes with respect to traffic trend, network topology and the number of video head-ends in the network.

The optimum sizes of caches during the hours of the day follow the trend of programme requests shown in Figure 6-1. Figure 6-9 to Figure 6-14 show the optimum cache sizes for each node in the network at each considered time of the day under SDTV BT 21CN-SV4, HDTV BT 21CN-SV4, SDTV NSF-SVMP, HDTV NSF-SVMP, SDTV NSF-MVMP and HDTV NSF-MVMP, respectively. Optimum cache sizes increase as the traffic volume increases reaching their peak at primetime. At a given time of the day, nodes where video head-ends are allocated require caches of smaller sizes, since programmes requested from those nodes are served locally. Therefore when a single video head-end is assumed as under BT 21CN-SV4 and NSF-SVMP (Figure 6-9 to Figure 6-12), the node equipped with the video head-end does not require a cache, as content is available at the node (node 4 and 5, respectively). Under HDTV, optimum cache sizes are considerably larger compared to SDTV due to larger sizes of HDTV programmes.

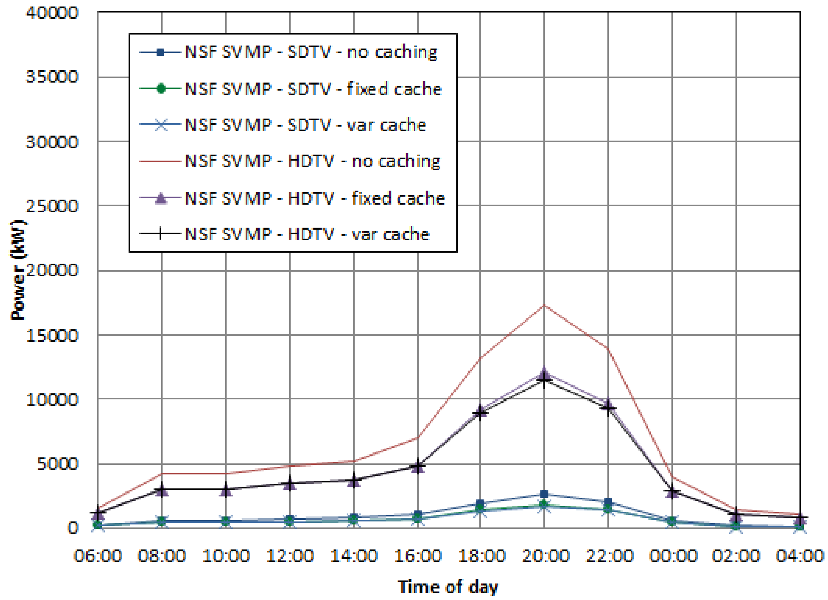

Figure 6-6: Power consumption of watching SDTV and HDTV programmes with no caching and considering caches of fixed and variable sizes under NSFNET-SVMP

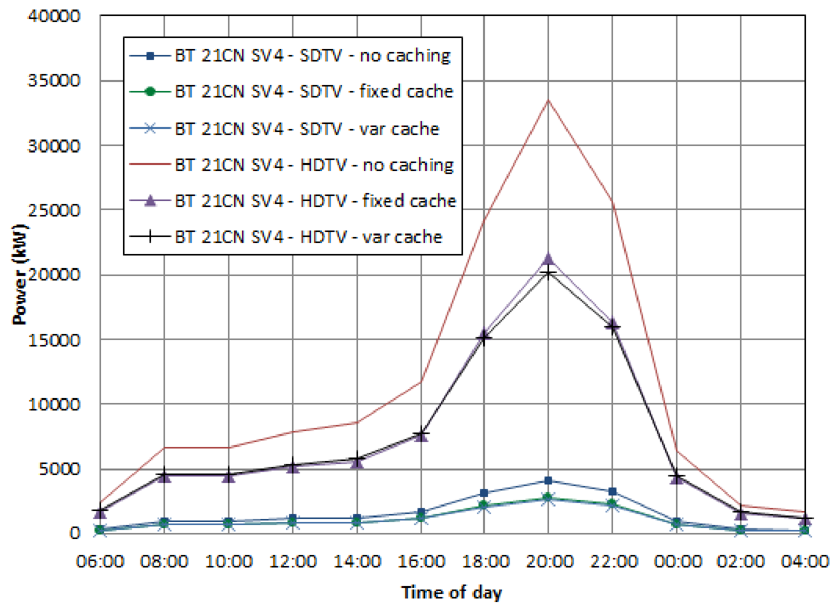

Figure 6-7: Power consumption of watching SDTV and HDTV programmes with no caching and considering caches of fixed and variable sizes under BT 21CN-SV4

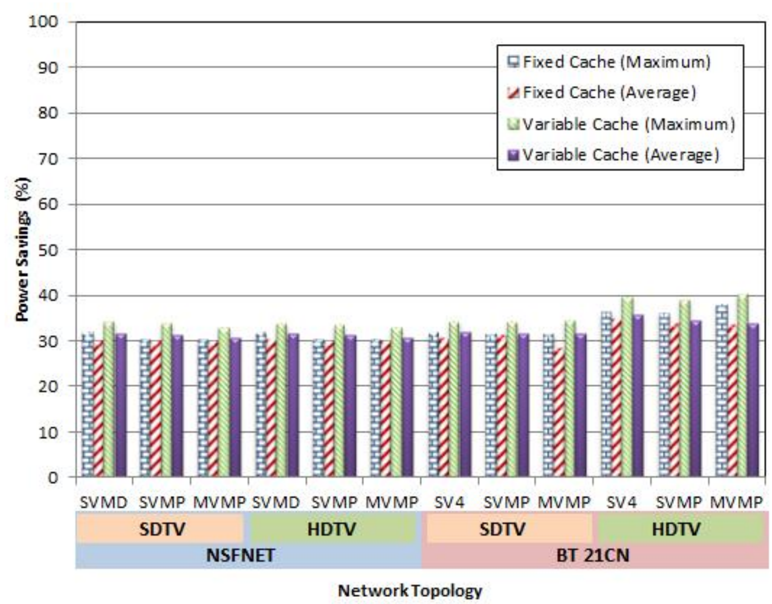

Figure 6-8: Maximum and average power savings of fixed and variable caching under SDTV and HDTV considering SVMD, SVMP and MVMP (NSFNET) and SV4, SVMP and MVMP (BT 21CN)

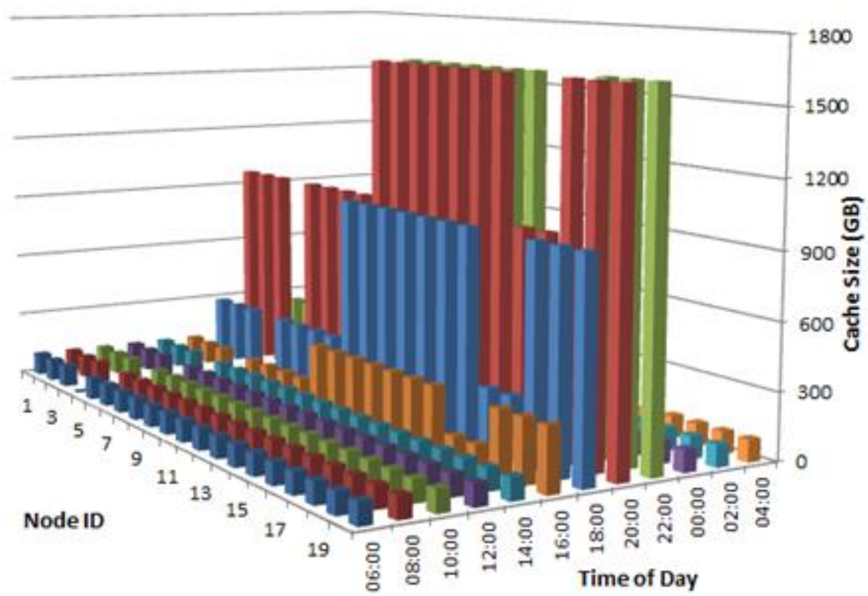

Figure 6-9: Optimum cache sizes at each node varied over the time of day under SDTV BT 21CN-SV4

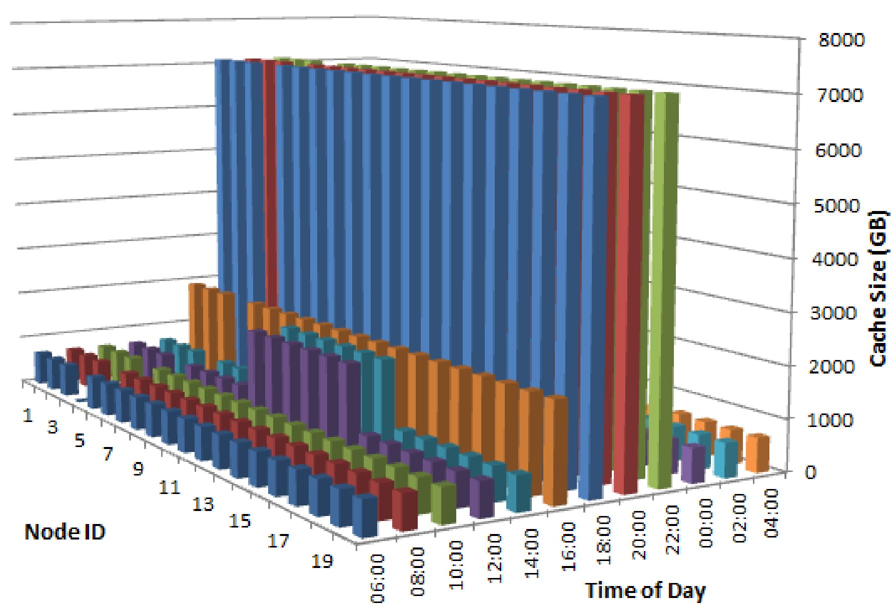

Figure 6-10: Optimum cache sizes at each node varied over the time of day under HDTV BT 21CN-SV4

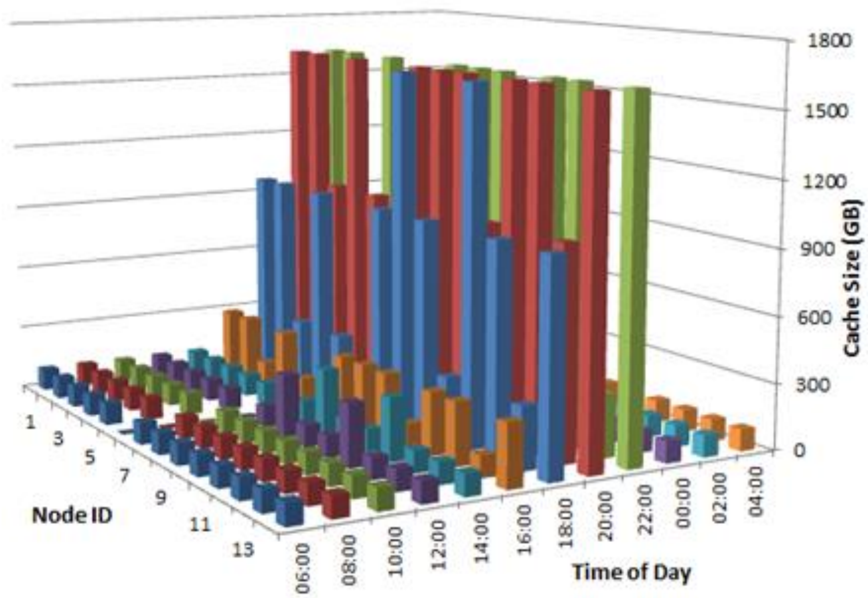

Figure 6-11: Optimum cache sizes at each node varied over the time of day under SDTV NSF-SVMP

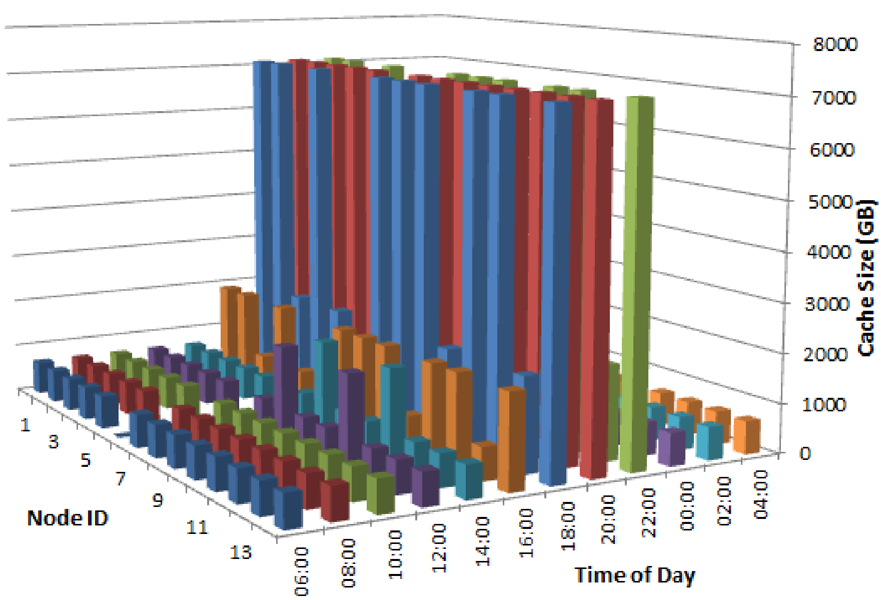

Figure 6-12: Optimum cache sizes at each node varied over the time of day under HDTV NSF-SVMP

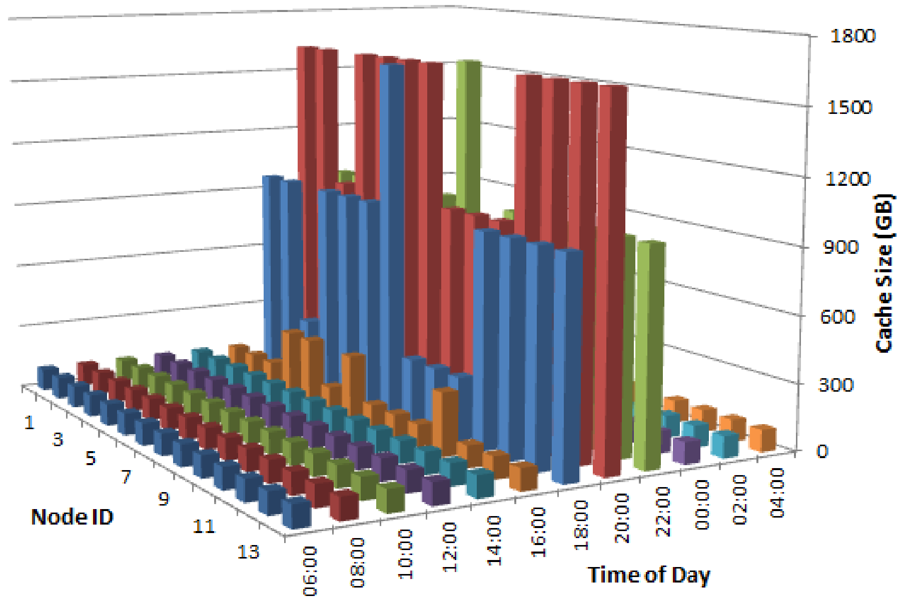

Figure 6-13: Optimum cache sizes at each node varied over the time of day under SDTV NSF-MVMP

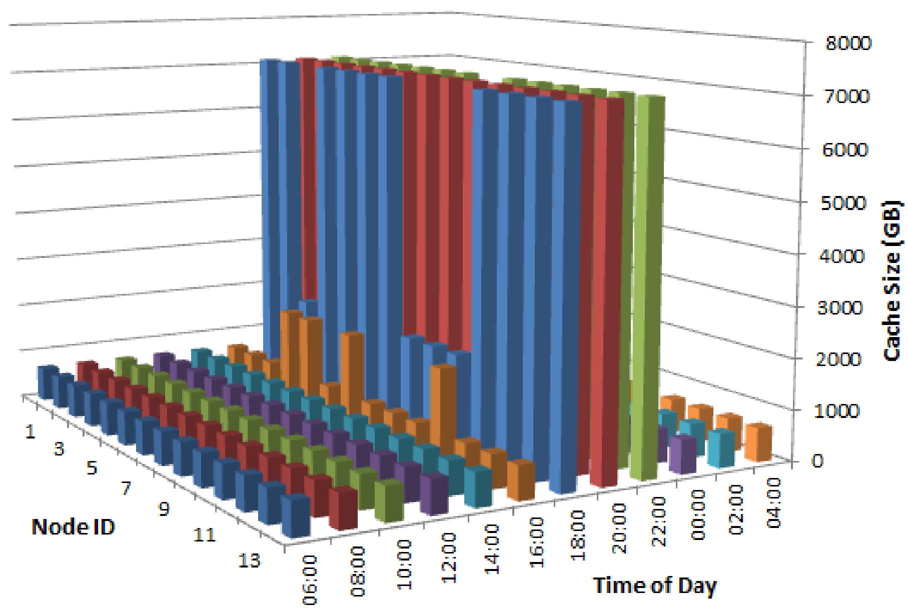

Figure 6-14: Optimum cache sizes at each node varied over the time of day under HDTV NSF-MVMP

6.6 Summary

This chapter has investigated the power consumption associated with the delivery of TV programmes over an IP over WDM network. It has explored the dynamics of watching TV including viewing figures and programme popularities. This work is significant due to the high power consumption associated with high quality TV during storage and transport. The state of the art BT 21CN network has been introduced considering three topologies with respect to video head-end location. The BT 21CN topology is similar to the NSFNET topology in terms of average hop count, but different in terms of coverage area and average nodal degree. The real TV data obtained in this chapter has been applied to the minimum-power optimum cache MILP model. The power consumption of delivering SDTV and HDTV programmes with caches of fixed and variable sizes has been evaluated.

The traffic demand generated by the most popular TV programmes accounts for a significant share of the total daily TV traffic. Only 2.25% of broadcast TV programmes were found to be live shows. These findings support the use of caching in the core network as a power reduction strategy. Storing popular TV programmes towards the edge of the network achieves an instantaneous power reduction of up to 36% under fixed caches. These savings are up to 40% when variable caches are considered. The optimum variable cache sizes that minimise power consumption follow the trend of TV traffic, and are larger under high definition TV. Similar power saving figures have been obtained under the two network topologies considered.

Chapter 7 Energy Efficient IPTV with Cache Content Replacements

7.1 Introduction

Chapter 4 to Chapter 6 have shown that caching data can enable reductions in overall energy use. This energy optimisation occurs when the cache hit ratio is high (reducing the energy of the transport network) and the cache size is small (constraining the energy of the cache). Chapter 6 has explored TV viewing figures and TV programme popularities. Measuring TV viewing and its distribution over a day provides information about likely IPTV traffic demand and its variation over time. The similarity of TV viewing is higher when more viewers prefer to view the same programme. This measurement is useful to estimate the potential reduction in network traffic due to caching, as the popularity of each programme stored in the cache is used to find the cache hit ratio. TV viewing behaviour can also be used to calculate the required cache capacities to achieve a desired cache hit ratio.

In this chapter, a novel time-based content replacement algorithm is proposed to maintain high cache hit ratios and small cache sizes as programme popularities change. The algorithm is based on the fact that programme popularities are both time dependent and predictable. Cache content replacements are therefore performed with respect to time-based programme popularities to maximise cache hit ratios whilst minimising the required cache size. The influence of regular traffic on power savings is explored where a number of regular and TV traffic mixtures are considered. The significance of the proposed cache management techniques is highlighted with respect to current and future network technologies.

7.2 Time Driven Cache Hit Ratio

For a given TV programme, we are interested in three quantities with respect to viewing popularity:

1. The programme average popularity over all time which determines the significance of caching the programme.
2. The instantaneous popularity that shows the dynamics of viewing the programme over time.
3. The programme popularity over a time window which is useful when cache content replacements are considered.

Assuming a popular programme which is, for example, primarily viewed in the morning and hardly viewed during other hours of the day, the average popularity P_i of this programme is given as:

$$P_i = Req_i / Tot\_Req \qquad (7\text{-}1)$$

where Req_i is the number of requests to programme i and Tot_Req are the total requests. Note that Req_i and Tot_Req are measured over the same time duration. If this time duration is long and the process is stationary then P_i converges to its actual probability of that event.

Let us next consider the popularity of a programme over a time window from time $t=a$ to time $t=b$ to signify a duration in time over which cache contents remain the same. The popularity of programme i during the time from a to b $P(i,a,b)$ is calculated as the ratio of the sum of requests for the programme throughout the time duration from a to b to the total requests taking place during the same time duration,

$$P(i,a,b) = \sum_{t=a}^{b} Req_{it} \Big/ \sum_{t=a}^{b} Tot\_Req_t \qquad (7\text{-}2)$$

Let $\omega(i,a,b)$ be a factor that specifies how requests for programme i are distributed over time and is given as the ratio of the number of requests for programme i over a time window to the overall number of requests, or:

$$\omega(i,a,b) = \sum_{t=a}^{b} Req_{it} / Req_i \quad 0 \leq \omega \leq 1 \qquad (7\text{-}3)$$

where $\omega(i,a,b)$=1 implies that all requests for programme i occur in the considered time window and the value of $\sum_{t=a}^{b} Req_{it}$ reaches its maximum $\sum_{t=a}^{b} Req_{it} = Req_i$. Consequently programme i must be stored in the cache for the time duration from time t=a to time t=b (assuming that the programme's time-driven popularity ranks at the top of all programmes' popularity list). In the other extreme when the value of $\omega(i,a,b)$ is 0, programme i is never requested in the considered time duration, since $\sum_{t=a}^{b} Req_{it}$=0. Although this programme might have a high global popularity, having the programme in the cache during the particular duration of time from a to b is not useful. The number of daily content replacements to be performed influences the resultant time-driven programme popularity $P(i,a,b)$ by determining the lengths of considered time windows. Note that time t is a continuous variable and therefore Equation (7-2) and (7-3) can be calculated using integrals. In this evaluation however, the number of programme requests are grouped at the start of each hour of the day, thus a summation provides an accurate approximation.

As previously explained in Chapter 2, the cache hit ratio H is the ratio of the number of requests served from the cache to the total number of requests,

$$H = C\_Req / Tot\_Req \qquad (7\text{-}4)$$

where *C_Req* is the number of requests served from the cache. The cache hit ratio is also calculated from the summation of the popularities of programmes stored in the cache,

$$H = \sum_{i=1}^{V} P_i \qquad (7\text{-}5)$$

where V is the cache size in programmes. Considering content replacements, the time-driven cache hit ratio during a time window from time $t=a$ to time $t=b$, $H(a,b)$, is derived from the sum of the time-driven popularities of programmes $P(i,a,b)$ stored in the cache during that time window,

$$H(a,b) = \sum_{i=1}^{V} P(i,a,b) \qquad (7\text{-}6)$$

The cache hit ratio can be considerably increased by performing content replacements, since a replacement populates the cache with programmes which are highly relevant during the considered time window. Such programmes replace programmes which are hardly viewed during that time window thus removing programmes with lower values of $P(i,a,b)$. This strategy acts on shorter time spans and is therefore more effective (in terms of the required cache size and therefore its power consumption) than conventional approaches which use the global popularity P_i.

7.3 Content Replacements MILP Models

Two MILP models are developed to find the optimum number of daily replacements in order to minimise power consumption, one for fixed and one for variable caches. These models are applied to two test network topologies (NFSNET and BT 21CN) considering both SDTV and HDTV.

7.3.1 Content Replacements MILP Model Assuming Caches of Fixed Sizes

The optimum number of daily replacements that achieves the best power efficiency is a trade-off. Infrequent replacements lead to low cache hit ratios and so high power consumption associated with viewing TV from the head-end. At the same time, frequent cache updates waste energy through populating the cache unnecessarily. The MILP model finds the optimum number of cache updates, using the IP over WDM network architecture described in Figure 2-3 in Chapter 2 and the programme popularities obtained in Chapter 6. The model assumes that the cache size is fixed and defines sets, parameters and variables as follows:

Sets:

N	Set of nodes
Nm_i	Set of neighbouring nodes of node i
T	Set of points in time
R	Set of possible daily content replacement frequencies

Parameters:

Pp	Power consumption of a router port
Po_{it}	Power consumption of optical switch i at time t
Pt	Power consumption of a transponder
Pa	Power consumption of an amplifier
Pmd	Power consumption of a multiplexer/demultiplexer
B	Capacity of a wavelength

W Number of wavelengths in a fibre

D_{ij} Distance between nodes i and j

S Span distance between two amplifiers

Amp_{ij} Number of amplifiers used on each fibre on the physical link from node i to j, $Amp_{ij} = \lfloor D_{ij}/S - 1 \rfloor + 2$

$RPmax_x$ Maximum router ports available to node x

λ^{xyt} Demand from node x to y at time t

δ_i Is 1 if node i has a video head-end, 0 otherwise, $\sum_{i \in N} \delta_i = u$, where u is the total number of servers in the network

H_{rt} Cache hit ratio at time t when r daily replacements are performed

M_{rt} Cache size in GB at time t when r daily replacements are performed

Φ Cache power consumption factor in W/GB

α_{rt} Is 1 if a replacement is performed at time t with r daily replacements, 0 otherwise

π_{rt} Additional download traffic to be streamed to a node due to a replacement at time t when performing r replacements

Variables:

f_{ij} Fibres on the physical link from node i to j

λ_{ijt}^{xy} Traffic from node i to j, part of the virtual link from node x to y at time t

w_{ijt}^{xy} Wavelengths on the link from node i to j, part of the virtual link from node x to y at time t

w_{ijt} Wavelengths on the physical link from node i to j at time t

C_{xyt} Wavelengths on the virtual link from node x to y at time t

AP_{it} Aggregation ports at node i at time t

r Number of replacements performed

The power consumption of the network consists of the power consumption of the following components:

1. Router ports at time t, where a port is required for each occupied wavelength:

$$\sum_{i \in N} Pp \left(AP_{it} + \sum_{j \in Nm_i : i \neq j} w_{ijt} \right)$$

2. Optical switches at time t:

$$\sum_{i \in N} Po_{it}$$

3. Transponders at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pt \cdot w_{ijt}$$

4. Amplifiers at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pa \cdot Amp_{ij} \cdot f_{ij}$$

5. Multiplexers/demultiplexers at time t:

$$\sum_{i \in N} \sum_{j \in Nm_i} Pmd \cdot f_{ij}$$

6. Deployed caches at time *t*:

$$\sum_{i \in N} \emptyset M_{rt}$$

Note that the number of lightpaths from node i to j is allowed to be different to the number of lightpaths in the reverse direction. Thus, f_{ij} and w_{ijt} are not necessarily equal to f_{ji} and w_{jit} respectively, since the model does not assume a simple symmetric case.

The goal of the proposed power-minimised content replacements MILP model is to minimise the network overall daily power consumption, and therefore the objective function is defined as:

Objective: minimise

$$\sum_{t \in T} \left(\begin{array}{c} \sum_{i \in N} Pp \left(AP_{it} + \sum_{j \in Nm_i : i \neq j} w_{ijt} \right) + \sum_{i \in N} Po_{it} + \\ \sum_{i \in N} \sum_{j \in Nm_i} Pt \cdot w_{ijt} + \sum_{i \in N} \sum_{j \in Nm_i} Pa \cdot Amp_{ij} \cdot f_{ij} + \\ \sum_{i \in N} \sum_{j \in Nm_i} Pmd \cdot f_{ij} + \sum_{i \in N} \emptyset M_{rt} \end{array} \right) \tag{7-7}$$

The model specifies a number of capacity and flow conservation constraints that must be satisfied, as follows:

Subject to:

$$\sum_{x\epsilon N}\sum_{y\epsilon N:x\neq y} w_{ijt}^{xy} \leq W \cdot f_{ij} \qquad (7\text{-}8)$$
$$\forall\, i \in N, j \in Nm_i, \forall\, t \in T$$

$$\sum_{x\epsilon N}\sum_{y\epsilon N:x\neq y} w_{ijt}^{xy} = w_{ijt} \qquad (7\text{-}9)$$
$$\forall\, i \in N, j \in Nm_i, \forall\, t \in T$$

$$\sum_{x\in N}\sum_{y\in N:x\neq y} \lambda_{ijt}^{xy} \leq\ C_{ijt} \cdot B \qquad (7\text{-}10)$$
$$\forall\, i, j \in N, \forall\, t \in T$$

$$\sum_{j\in N} w_{ijt} + AP_{it} \leq RPmax_i \qquad (7\text{-}11)$$
$$\forall\, i \in N, \forall\, t \in T$$

$$AP_{xt} \geq \sum_{y\in N} \left(\lambda^{yxt} \cdot \delta_y\right) / \ B \qquad (7\text{-}12)$$
$$\forall\, x \in N, \forall\, t \in T$$

$$\sum_{j\epsilon Nm_i} w_{ijt}^{xy} - \sum_{j\epsilon Nm_i} w_{jit}^{xy} = \begin{cases} C_{xyt} & i = x \\ -C_{xyt} & i = y \\ 0\ otherwise & \end{cases} \qquad (7\text{-}13)$$
$$\forall\, i, x, y \in N, \forall\, t \in T$$

$$\sum_{j\epsilon N:i\neq j} \lambda_{ijt}^{xy} - \sum_{j\epsilon N:i\neq j} \lambda_{jit}^{xy} = \begin{cases} \delta_x \cdot \left(\lambda^{xyt} \cdot (1 - H_{rt}) + (\propto_{rt} \cdot \pi_{rt})\right) & i = x \\ -\delta_x \cdot \left(\lambda^{xyt} \cdot (1 - H_{rt}) + (\propto_{rt} \cdot \pi_{rt})\right) & i = y \\ 0 & otherwise \end{cases} \qquad (7\text{-}14)$$
$$\forall\, i, x, y \in N, \forall\, t \in T, \forall\, r \in R$$

Objective (7-7) specifies the power consumption of the network by considering the power consumption of occupied network components at each time of the day. Constraints (7-8) and (7-9) are the capacity constraints of the physical layer.

Constraint (7-10) is the lightpath capacity constraint. Constraint (7-11) ensures that the total router ports used at a node do not exceed its maximum. Constraint (7-12) calculates required aggregation ports. Constraint (7-13) is the flow conservation constraint in the optical layer. Constraint (7-14) is the flow conservation constraint for traffic originating at nodes equipped with a video head-end. Note that the traffic increases at the times of day when a replacement is performed, but with the advantage of having higher cache hit ratios throughout the day.

7.3.2 Content Replacements MILP Model with Variable Cache Sizes

This model extends the first model by introducing caches that can reduce their active capacity. A smaller cache will use less power but will be able to store less data. The relationship between the number of programmes stored in the cache and its hit ratio is obtained. This relationship is represented by a convex function. However, cache hit ratios vary with the difference in the number of daily content replacements, requiring a different convex function for each considered number of replacements. Therefore, the goal is to find the optimum cache size for each node at each time of the day under each number of daily replacements.

The MILP model declares the number of replacements as a parameter in order to maintain the linearity of the model. In addition, having the number of replacements as an input is required to construct the equations of the piecewise linear approximation. Note that the amount of content to be replaced at each replacement is different for each node at each time of the day under each considered number of replacements. The sets, parameters and variables declared in the original content replacements MILP model are utilised, in addition to the following amendments:

Sets:

K	Set of equations that approximate the convex function describing the relationship between the cache size and its hit ratio

Parameters:

r	Number of replacements performed
a_{ktr}, b_{ktr}	Piece-wise linear approximation equations coefficients, three dimensional vectors
$Trep$	Time duration over which cache contents are updated

Variables:

H_{itr}	Cache hit ratio of node i at time t when r daily replacements are performed
M_{itr}	Cache size of node i at time t when r daily replacements are performed
π_{itr}	Additional download traffic to be streamed to node i at time t when r daily replacements are performed

The objective of the model is to minimise the overall network power consumption over time t, and is given as:

Objective: minimise

$$\sum_{t\in T}\left(\begin{array}{c}\sum_{i\in N} Pp\left(AP_{it}+\sum_{j\in Nm_i:i\neq j} w_{ijt}\right)+\sum_{i\in N} Po_{it}+\\ \sum_{i\in N}\sum_{j\in Nm_i} Pt\cdot w_{ijt}+\sum_{i\in N}\sum_{j\in Nm_i} Pa\cdot Amp_{ij}\cdot f_{ij}+\\ \sum_{i\in N}\sum_{j\in Nm_i} Pmd\cdot f_{ij}+\sum_{i\in N} \emptyset M_{itr}\end{array}\right) \tag{7-15}$$

Subject to:

The model satisfies constraints (7-8), (7-9), (7-10), (7-11), (7-12) and (7-13) in the original model. Constraint (7-14) is replaced with

$$\sum_{j\epsilon N:i\neq j}\lambda_{ijt}^{xy}-\sum_{j\epsilon N:i\neq j}\lambda_{jit}^{xy}=\begin{cases}\delta_x\cdot\left(\lambda^{xyt}\cdot\left(1-H_{ytr}\right)+\left(\propto_{rt}\cdot\pi_{ytr}\right)\right) & i=x\\ -\delta_x\cdot\left(\lambda^{xyt}\cdot\left(1-H_{ytr}\right)+\left(\propto_{rt}\cdot\pi_{ytr}\right)\right) & i=y\\ 0 & otherwise\end{cases}$$
$$\forall\, i,x,y\in N,\forall\, t\in T,\forall\, r\in R \tag{7-16}$$

In addition, the model satisfies the following constraints

$$M_{itr}\geq\ a_{ktr}\cdot H_{itr}+b_{ktr}$$
$$\forall\, i\in N,\forall\, t\in T,\forall\, r\in R,\forall\, k\in K \tag{7-17}$$

$$\pi_{itr}=\ M_{itr}\cdot\ 8/Trep$$
$$\forall\, i\in N,\forall\, t\in T,\forall\, r\in R \tag{7-18}$$

Objective (7-15) is the power consumption of the network made up of the power consumption of network components at each considered time of the day. Constraint (7-16) is the flow conservation constraint for downlink traffic. Note that the different cache sizes for each node at each time of the day under each considered number of daily replacements governs the amount of additional download traffic due to a performed replacement. Constraint (7-17) is the set of convex equations of the piecewise linear approximation utilised to convert a cache hit ratio into its corresponding cache size with respect to the number of daily replacements. Constraint (7-18) calculates the additional traffic to update cache contents by converting the cache size (GB) into Gb/s.

7.4 Content Replacements Simulation Approach

In order to validate the proposed MILP model, a simulation is developed to calculate the power consumption of the network. The routing algorithm of the simulation is based on the heuristic proposed in [32] where traffic is routed over an IP over WDM network considering lightpath bypass. This algorithm is extended to include caches and consider lightpath non-bypass.

7.4.1 Demand Calculation

The content replacement process, shown in Figure 7-1 starts by obtaining the hit ratios of deployed caches using the optimum number of replacements found by the content replacements MILP model. Using these cache hit ratios, the network traffic demand is calculated. The algorithm then checks if content is to be replaced at the current time of day according to the number of replacements. If so, the additional traffic demand due to content replacements is added to the total traffic entering each node.

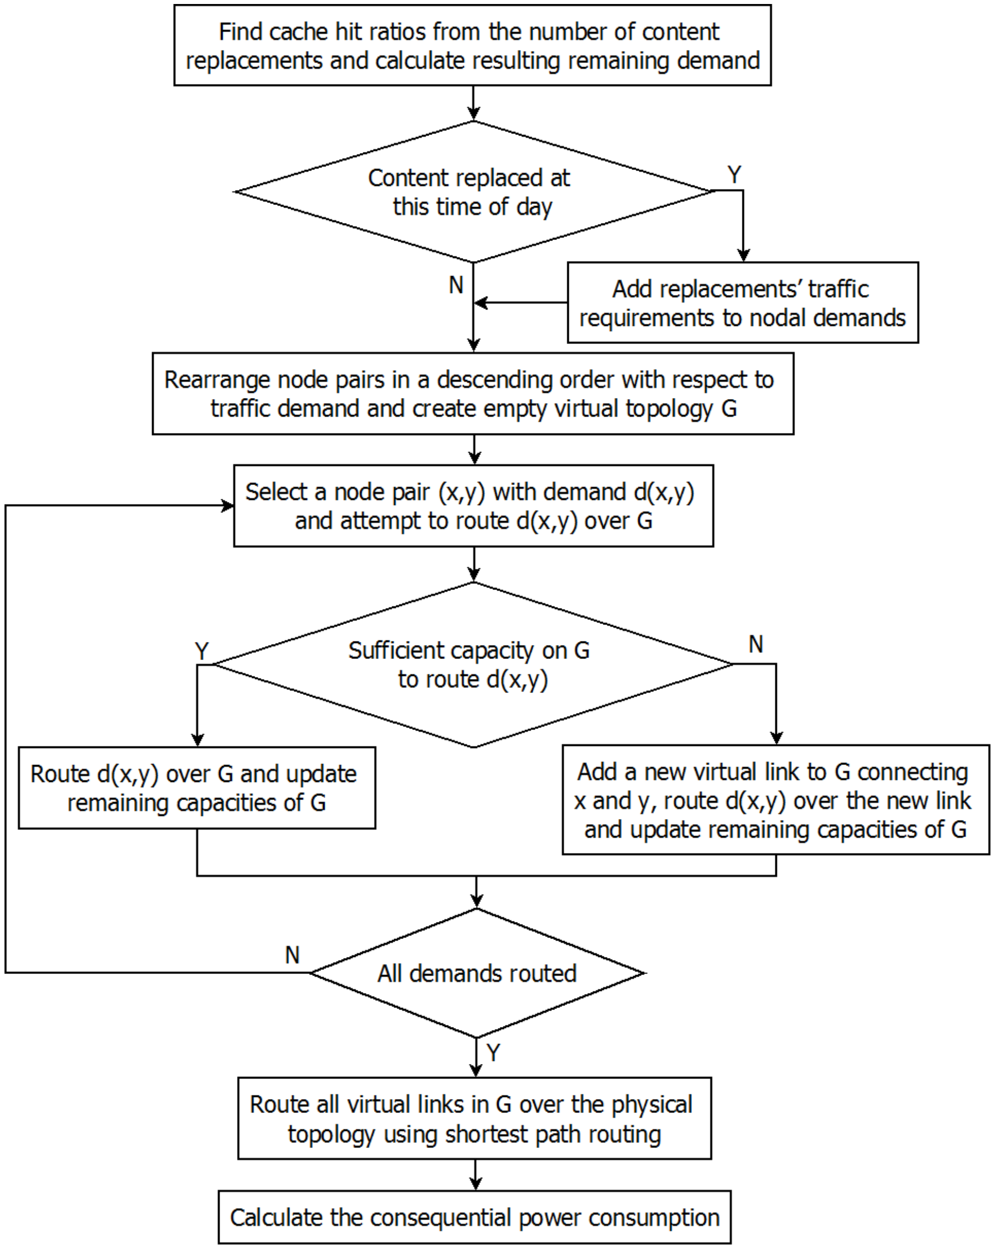

Figure 7-1: The content replacements algorithm flow chart

7.4.2 Routing Traffic over the Virtual Layer

The algorithm continues by arranging node pairs in a descending order starting with the node pair having the highest demand to accommodate high demands on virtual links first and try to accommodate lower demands on the same existing links. An empty topology G is created to record established links and their capacities. The node pair (x,y) with the highest demand $d(x,y)$ is selected and the algorithm attempts to route $d(x,y)$ over existing virtual links. If the capacities of existing virtual links in G are sufficient to accommodate $d(x,y)$, the process is successful and the remaining capacities of G are updated. Otherwise, a new virtual link connecting x and y is created and added to G. The demand $d(x,y)$ is routed over the new link and the remaining capacities of G are updated. The selection of node pairs is repeated until all demands are routed.

7.4.3 Routing Traffic over the Physical Layer

When iterations are complete, G holds the set of lightpaths to be routed over the physical topology in the optical layer. The simulation uses the shortest path routing algorithm to route the lightpaths over the path with the shortest physical distance in the physical topology. The shortest physical distance minimises delay but does not necessarily result in the minimum usage of network components, and consequently minimum overall power consumption. For example the path with shortest physical distance may traverse more intermediate nodes compared to a slightly longer path and therefore consumes more power in this case. It minimises delay however (and in general is a good choice for power minimisation) and is therefore chosen here [62], [63]. The alternative minimum hop routing approach may result in lower power consumption, but higher delay and is therefore not adopted in this study. In the IP layer, the number of ports required to accommodate the capacities of lightpaths is calculated, and the power consumption of the IP layer is found.

The algorithm allows traffic grooming where more than one demand may be routed on the same virtual link, and therefore improves virtual link utilisation. This feature results in decreasing the number of established virtual links and hence fewer IP router ports are required. Since router ports are the network major power consuming components, the overall power consumption is reduced.

7.5 Content Replacements Power Consumption Evaluation

To evaluate the power consumption of content replacements, the TV viewing data obtained in Chapter 6 is considered. The evaluation also considers the same NSFNET and BT 21CN network topologies as in Chapter 6 depicted in Figure 6-2 and Figure 6-3, respectively. The power consumption of different network components, EDFA spacing, wavelengths per fibre and wavelength capacity are the same as Table 6-2 in Chapter 6. These inputs are applied to the previously explained MILP models and simulation to calculate the power consumption of delivering TV programmes under content replacements.

Caches are in the READ state as long as no replacements are performed. During a replacement, the portion of the cache being replaced enters the WRITE state for a very short time (a few minutes versus 24 hours). The power consumption of a READ or WRITE state is therefore considered to be the same (7.4W/GB).

The MILP models use dual simplex iterations to find the optimum and define around 1 million variables and over 3.6 million variables in the NSFNET and the BT 21CN topologies respectively. A typical run of the content replacements MILP model required between 2.4 and 6 hours for the NSFNET and 48 hours for the BT 21CN network using the solver and computer specified in Section 3.6.1 in Chapter 3. In this section the results are presented.

7.5.1 Content Replacements with Fixed Size Caches

Chapter 6 has shown that the optimum cache sizes that minimise power consumption using TV programme popularities are 200GB and 650GB (fixed-size caches) for SDTV and HDTV respectively. Here the popularity of each programme is calculated with respect to time windows using Equation (7-2) and the traffic profile in Figure 6-1 in Chapter 6. The values for programme popularities are used to find the cache hit ratio from Equation (7-6). Each programme is assumed to have a fixed number of viewers for the programme duration and to have no viewers outside broadcasting time. To include consideration of CuTV or time-shifted viewing, a non-zero popularity could be assigned to each TV programme after the end of broadcasting. Since the contents of the cache may not always be entirely replaced, this approach is conservative as CuTV would lead to higher cache hit ratios than calculated here. However, this traffic is currently small compared to TV traffic [98]. The influence of CuTV viewing on cache hit ratios will become visible when CuTV traffic reaches considerable levels.

When a content replacement is performed, the new cache content is streamed from the video head-end to the cache resulting in additional traffic passing through the network. The amount of traffic is a function of the percentage of the cache size that is replaced. It is possible that only a fraction of the cache is updated with each replacement as some programmes may remain popular through two or more time windows. Here the entire cache contents are replaced each time. This assumption has the advantage of achieving a high cache hit ratio as the cache is occupied with the most popular programmes for the considered time duration. However, the shortcoming is that more traffic passes through the network to fill the cache when a replacement occurs. This assumption allows the evaluation of the full effect of additional traffic introduced by content replacements and therefore the reported power savings are conservative.

It takes time to update the cache contents, and this time delay is used to calculate the additional network traffic in Gb/s. The cache update time should be much shorter than the minimum replacement interval (2 hours). In addition, the resulting data rate should not overload the network. Therefore, a cache update time of 1 minute is selected, leading to cache update traffic of 26.6Gb/s and 86.6Gb/s under SDTV and HDTV, respectively. The cache update traffic is much smaller than the overall network traffic (40Tb/s). A more sophisticated scenario for cache updates may be considered where only sections of the cache are updated. For example, it might be more power-efficient to replace only 20% of the programmes during content replacements performed in the early hours of the morning, whereas updating 100% of the cache becomes necessary at primetime to achieve maximum power savings. A more practical approach can also be considered where the list of the most popular programmes during the next time window is compared to the list of the most popular programmes currently stored in the cache. Only the TV programmes which are not already in the cache are streamed from the video head-end. This strategy results in communication overhead where knowledge of the up-to-date cache contents is required. Such complex refinements warrant further investigation.

7.5.1.1 Power Consumption Evaluation

In order to evaluate the power consumption of the network under various daily replacements, replacement frequencies of 2, 3, 4, 6 and 12 per day are compared. The MILP model is solved considering the NSFNET topology under SVMD, SVMP and MVMP and for the BT 21CN topology under SV4, SVMP and MVMP. Figure 7-2 shows the resultant power consumption of content replacements under BT 21CN-SV4, NSF-MVMP and NSF-SVMP compared to the power consumed without caching and having fixed caches with no replacements. The largest savings in power consumption are achieved when 12 content replacements are performed. Nevertheless with a fixed cache size, content replacements give little benefit over a cache that is populated once a day as there is only a marginal increase in cache hit ratios. The largest instantaneous power savings of 48% were seen for the BT 21CN topology under HDTV and MVMP.

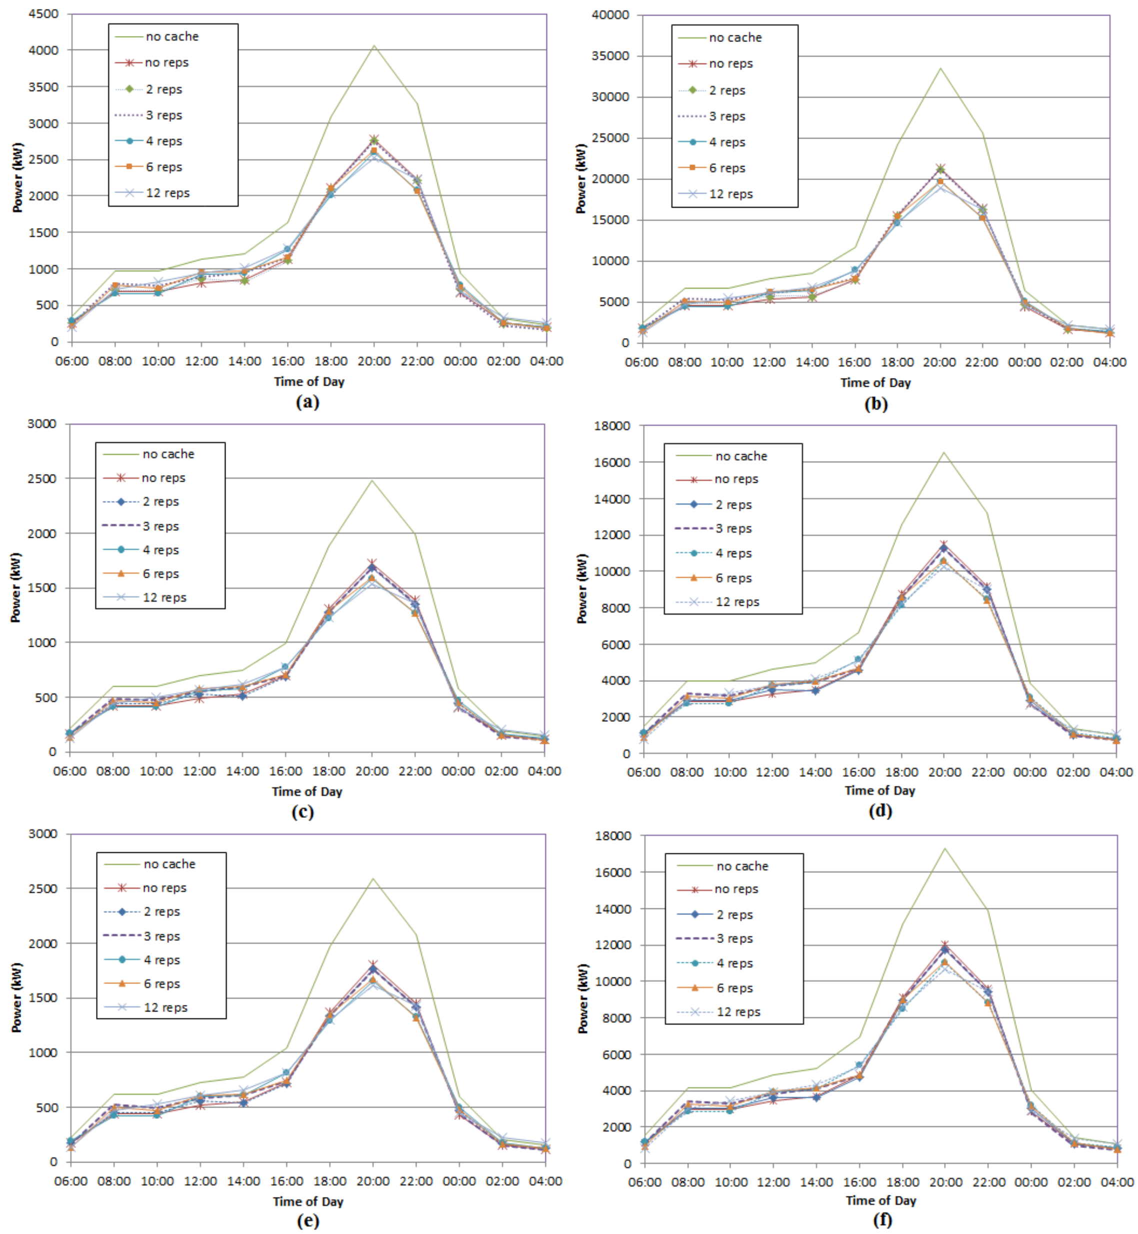

Figure 7-2: Power consumption of watching TV programmes with no caching, fixed optimum caches and under 2, 3, 4, 6 and 12 daily content replacements considering (a) SDTV BT 21CN-SV4 (b) HDTV BT 21CN-SV4 (c) SDTV NSF-MVMP, (d) HDTVNSF-MVMP, (e) SDTV NSF-SVMP, and (f) HDTV NSF-SVMP

7.5.2 Content Replacements under Variable Caches

The goal here is to investigate the additional savings possible by using variable size caches.

7.5.2.1 Cache Hit Ratio Dynamics

To carry out the investigation, the popularity of TV programmes is found to obtain the relationship between the number of programmes stored in the cache and its hit ratio. The popularity of each programme is different under each of the considered daily replacement schemes and at each time of the day, so cache hit ratios will vary with time. Each equation that describes the relationship between the number of programmes stored in the cache and its hit ratio under a certain replacement frequency is a concave function. The combination of all equations representing all considered number of daily replacements form a surface.

Figure 7-3 shows two examples of this relationship at time 6:00 and 20:00. From each resulting surface, a piecewise linear approximation is calculated under each considered daily replacement frequency. All piecewise linear approximation equations are input to the MILP model to calculate cache sizes from optimum cache hit ratios rather than finding the cache hit ratio from the cache size as in Figure 7-3. Therefore the relationship described in the MILP model is a convex function.

The resulting surfaces shown in Figure 7-3 (a) and (b) show that the hit ratios of variable caches increase with the increase in the number of content replacements performed. When no content replacements are performed, increasing the cache size leads to increasing the cache hit ratio. When further increasing the cache size, the cache hit ratio saturates as programmes become less popular.

If the cache is of a fixed size, then at any point in time it is storing programmes which are not needed at that time. Under content replacements, the cache stores programmes which are viewed during the current time window and thus the whole cache is usable. This explains the few dips in cache hit ratio that can be observed in Figure 7-3 where at a number of points the cache hit ratio under replacements becomes lower than the cache hit ratio under fixed cache sizes, yet entirely effective.

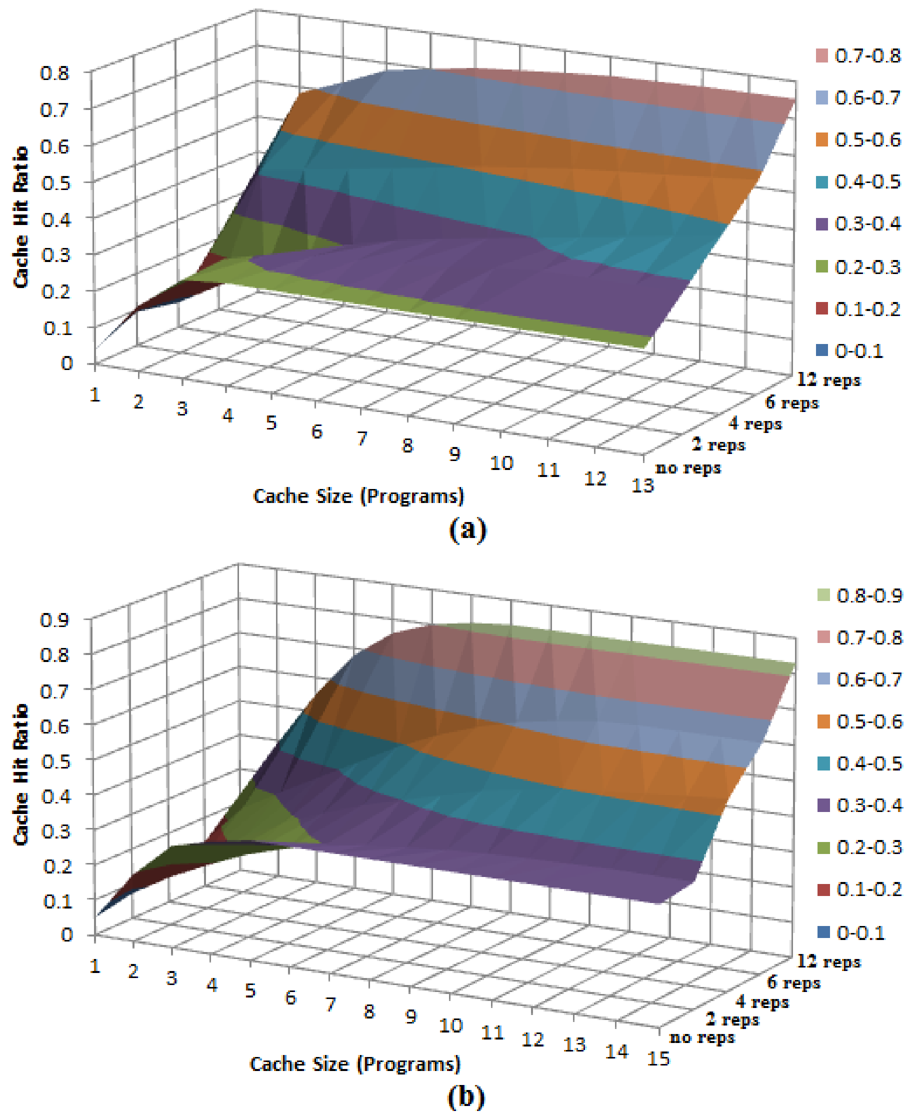

Figure 7-3: The relationship between the number of TV programmes stored in the cache, the number of daily content replacements and the cache hit ratio at time: (a) 6:00 and (b) 20:00

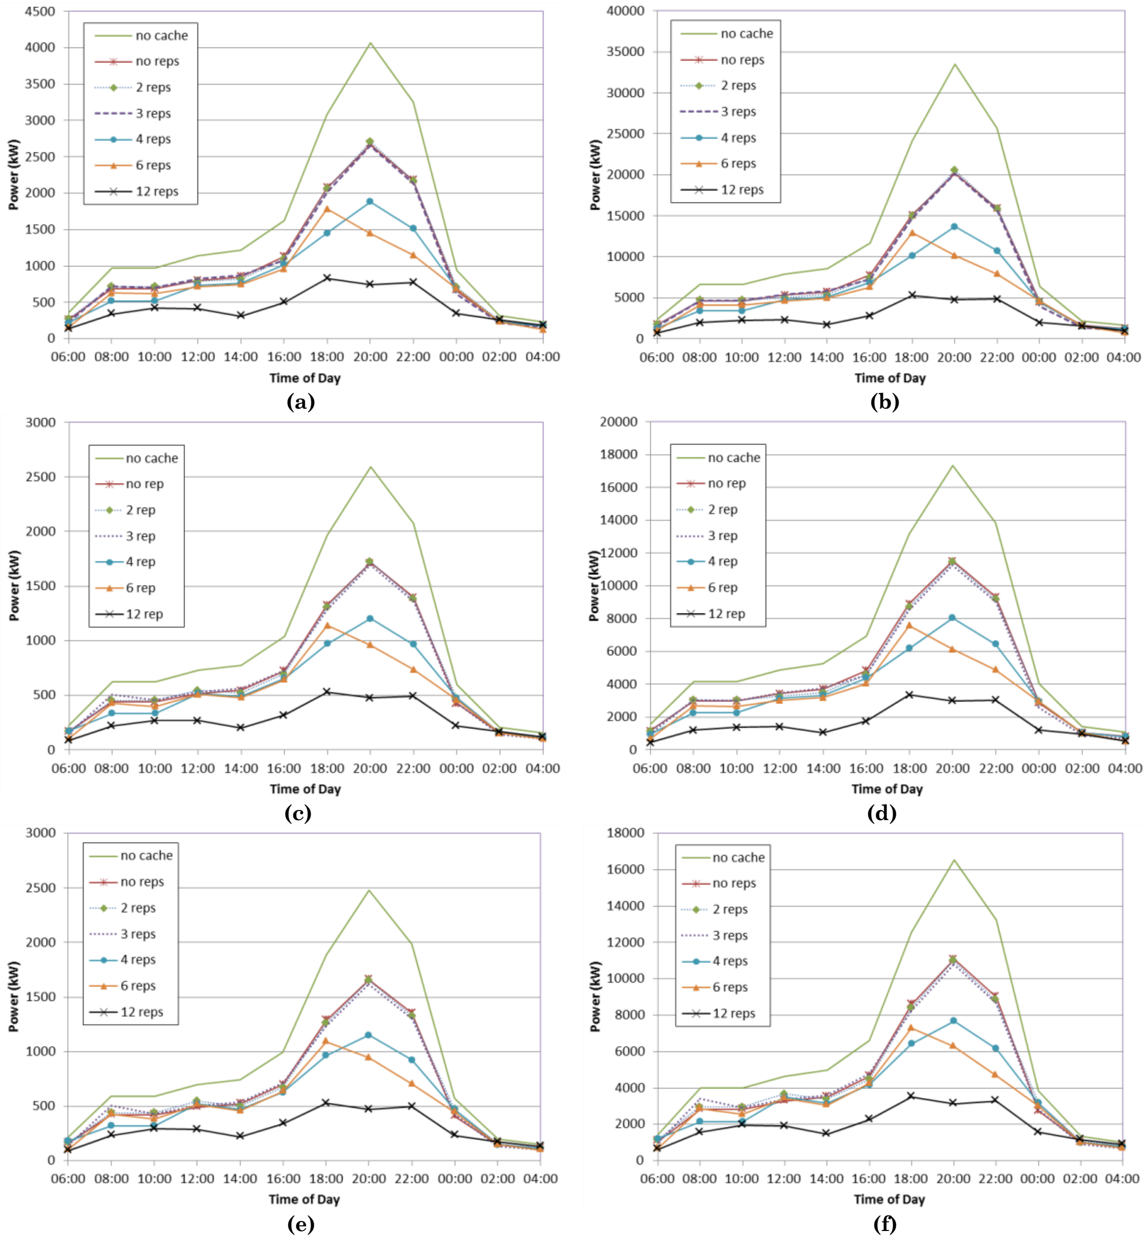

Figure 7-4: Power consumption of watching TV programmes with no caching, variable cache sizes with no content replacements and under 2, 3, 4, 6 and 12 daily content replacements considering (a) SDTV BT 21CN-SV4, (b) HDTV BT 21CN-SV4, (c) SDTV NSF-SVMP, (d) HDTV NSF-SVMP, (e) SDTV NSF-MVMP and (f) HDTV NSF-MVMP

7.5.2.2 Power Consumption Evaluation

Figure 7-4 shows the power consumption of the network under different topologies generated by the MILP model. This shows that, with variable cache sizes, there is now a significant benefit in having frequent content replacements. The savings are greatest during primetime when the traffic volume is high and there are only a few very popular programmes.

The peak and average savings in power consumption due to content replacements performed on caches with fixed and variable sizes are shown in Figure 7-5. The maximum savings under variable caches range between 80% and 82% for the NSFNET topology and up to 86% for the BT 21CN topology. These savings are significantly greater than the savings due to content replacements under fixed cache sizes (47% and 48% for the NSFNET and BT 21CN topology, respectively).

The following observations are made:

1. The greatest savings in energy consumption are achieved by performing 12 daily replacements on cache contents (see Figure 7-4).
2. Under time-based content replacements the cache hit ratio varies over the day, resulting in varying savings in power consumption (see Figure 7-7).
3. Variable cache sizes are beneficial as the optimal cache size is different depending on the traffic volume. As a result, when performing replacements on variable cache sizes, the maximum power savings achieved are much higher than the average power savings. The greatest power savings are

achieved at primetime where traffic is high and a small number of programmes are very popular (see Figure 7-2 and Figure 7-4).

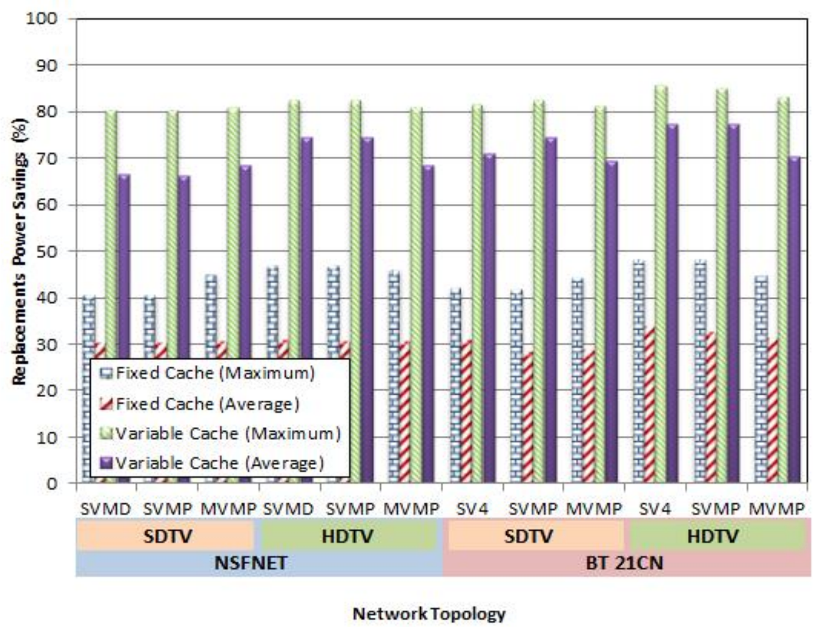

Figure 7-5: Maximum and average power savings of content replacements using fixed and variable caching under SDTV and HDTV considering SVMD, SVMP and MVMP (NSFNET) and SV4, SVMP and MVMP (BT 21CN)

It is worth pointing out that the TV data utilized in this evaluation is one day viewing figures obtained on a weekday (Friday) in November 2012. TV data is expected to differ when other dates are considered. The total number of viewers changes from one year to another. The most popular TV shows are mainly seasonal and do not run throughout the year [96]. Consequently, the most popular TV shows and their corresponding number of viewers are different from one month to another. In addition, daily viewing figures are different in weekends compared to weekdays [104],[99]. Figure 7-6 shows the total number of TV viewers during weekdays and weekends in 2011 and 2012. As depicted in Figure 7-6 the difference in the total viewers is marginal and the trend of viewing during the day is identical. This can also be observed when comparing viewing figures of different months of the same year [105]-[107],[98]. In addition, the difference between the number of viewers in weekdays and weekends is insignificant considering that the trend and peak are similar.

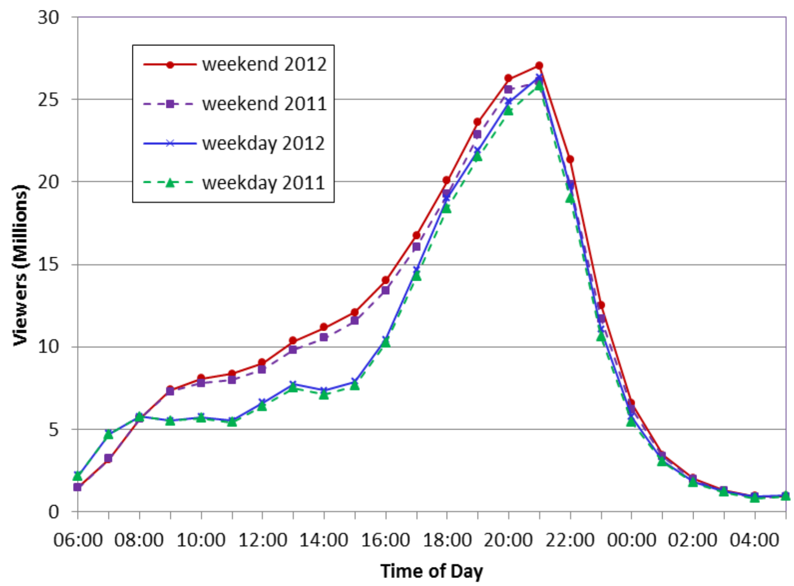

Figure 7-6: Average 2011/2012 TV viewers during weekdays and weekends [104],[99]

The combined popularity of the 5 most popular TV programs viewed on Friday 9th and Saturday 10th November 2012 is 0.34 and 0.36, respectively [96]. The variance in the number of viewers and the popularity of the most popular programs results in obtaining different power saving figures from the replacements MILP model. However, given the facts above, these differences are expected to be insignificant. An exception to this is when a major event is broadcasted causing considerable changes in viewing figures. One example is the London 2012 Olympic Games where live TV viewing and BBC iPlayer requests hit record breaking figures [105].

7.5.2.3 Optimum Variable Cache Sizes

Figure 7-7 shows how the cache hit ratios change due to content replacements when cache sizes can be varied over the day. The increase in the number of daily content replacements leads to an increased cache hit ratio. The average cache hit ratio increases from 0.32 to 0.69 and 0.7 under SDTV and HDTV, respectively. These increases in cache hit ratios are the primary driver for the reduction in power consumption.

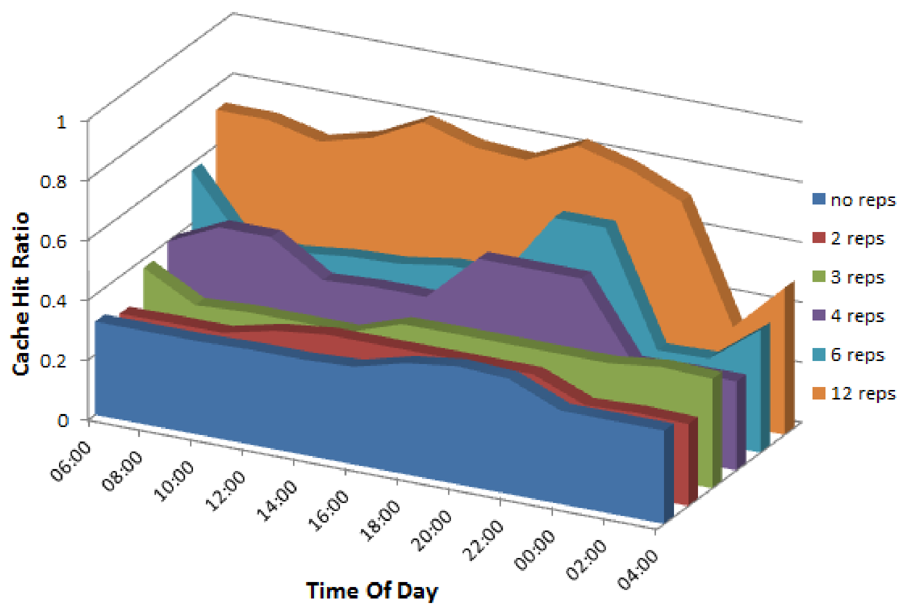

Figure 7-7: Cache hit ratios of optimum cache sizes averaged over network nodes with no content replacements and considering 2, 3, 4, 6, and 12 daily content replacements under SDTV NSFNET with a single video head-end

It is worth mentioning that it is not guaranteed that the cache hit ratio will increase at each time slot with the increase in replacement frequencies. For instance, the average cache hit ratio at 18:00 when performing 4 and 6 replacements is 0.5 and 0.39, respectively. This results (as shown in Figure 7-4) in the network power consumption with 6 replacements being higher than that with 4 replacements for that particular point in time. However, the overall energy consumption of the network reduces with increase in the replacements frequency, including that additional power consumed to replace content.

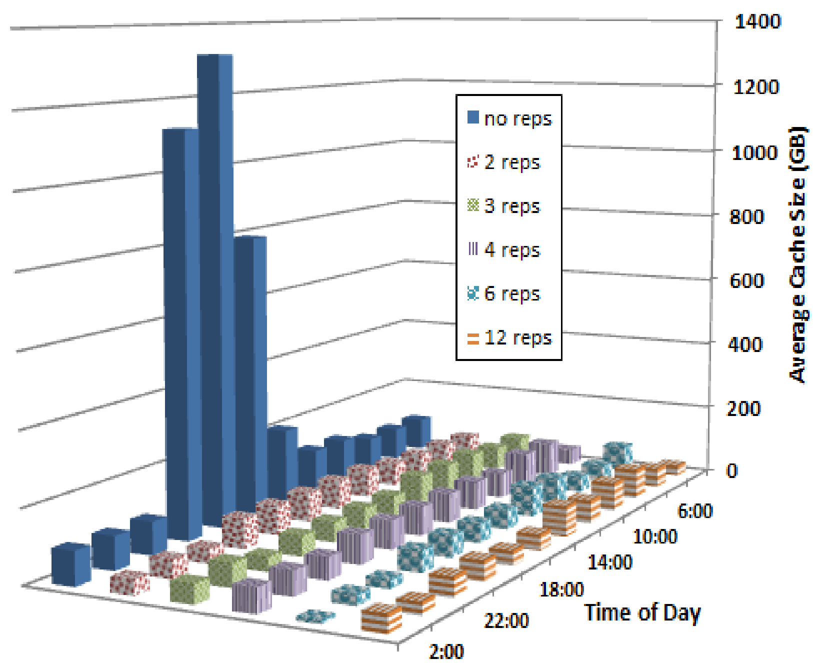

Figure 7-8: Optimum variable cache sizes averaged over network nodes with no content replacements and considering 2, 3, 4, 6, and 12 daily content replacements under SDTV NSFNET with a single video head-end

The optimum variable cache sizes over the time of the day found by the MILP model follow the daily trend of input traffic. These cache sizes are averaged over network nodes and are shown in Figure 7-8. Performing content replacements on smaller caches achieves higher cache hit ratios than deploying variable cache sizes with no content replacements. When a single video head-end is considered in the network the average cache size falls from 320GB to 51GB and from 2058GB to 443GB (variable cache sizes) under SDTV and HDTV, respectively. In other words, the power savings in caching are an average of 84% and 78.4% under SDTV and HDTV, respectively.

The results on cache savings are only accurate when the number of replacements is high (e.g. 12 replacements). Considering variable caches sizes, it is assumed that by caching the most popular programmes and by obtaining a cache hit ratio of *H,* this cache hit ratio remains effective all day. In reality, the most popular programmes may be evening programmes and in real time TV there is no interest in them in the morning. Therefore in Figure 7-8 for example, the variable cache size at 2:00 should be zero since the most popular programmes are evening programmes that are cached, but are not relevant in the morning. With the increase in the adoption of CuTV and iPlayer type services, this approximation becomes valid as popular programmes remain popular and available to play for extended hours and may be for a few days (one week typically for iPlayer). For the approximation to be accurate, popular programmes must remain popular all the time, with no real time TV effect. In addition, programmes need to be available to stream at any time of the day.

7.6 Comparison of the Two Network Topologies and Simulation Results

This section compares the power savings achieved by content caching and content replacements under the two considered network topologies (NSFNET and BT 21CN).

It also shows the network energy savings obtained by the simulation compared to those obtained by the MILP model.

7.6.1 Comparison of the NSFNET and BT 21CN Topologies

The NSFNET and BT 21CN topologies are similar with respect to hop count, but are different in terms of the coverage area and nodal degree (see Table 6-1 in Chapter 6). Average power savings of caching (with no content replacements) are 30% – 31% and 33% – 35% under NSFNET and BT 21CN, respectively. The average power savings introduced by content replacements are 68% – 74% and 70% – 77% under NSFNET and BT 21CN, respectively. The coverage area of the network does not influence the power consumption much, as the additional distance between nodes requires extra EDFAs which do not consume significant power (8W). Consequently, the power savings are comparable under all cache management techniques in the two topologies. The slight further improvement in power savings under the BT 21CN topology is due to the higher nodal degree which provides more possible paths through which traffic can be routed without underutilising resources (in the MILP model the path that achieves the most power saving is selected). The consideration of two network topologies (the NSFNET and BT 21CN) in the evaluation helped generalise the results. Although there are slight differences, there is good consistency in the results giving confidence in the methods developed.

7.6.2 Simulation Results

Figure 7-9 shows the overall energy consumption of the network under SDTV and HDTV using both MILP models and simulation.

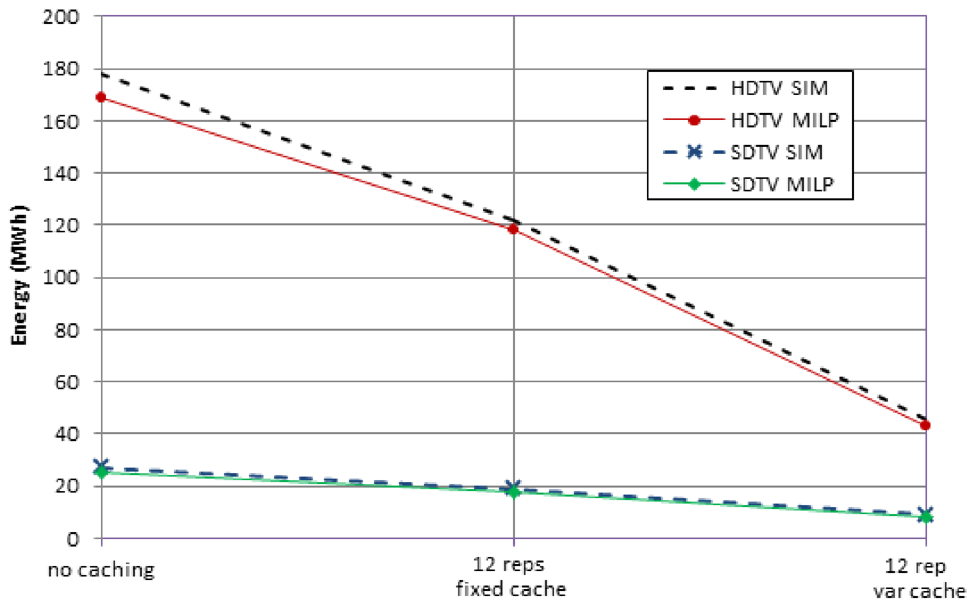

Figure 7-9: The daily energy consumption of watching SDTV and HDTV programmes when no caches are deployed and when 12 daily replacements are performed on fixed and variable caches under NSF-SVMD using the MILP model and simulation

To run the simulation considering content replacements on caches of fixed sizes, the cache hit ratios are calculated considering the number of daily replacements, and these cache hit ratios are used to calculate the remaining network traffic to be routed from the video head-end to each node. The simulation considers the additional traffic due to replacements and is included to find the total traffic to route at each time of the day. Considering replacements on variable cache sizes, the variable cache sizes and cache hit ratios are found from the MILP model for each replacement frequency and are used in the simulation. Traffic demands are routed over the shortest path as network dimensioning is assumed where additional network resources are allocated when demanded. However, physical link maximum capacities are not considered where some demands are forced to be routed over longer routes or rejected due to unavailable resources. All other parameters used in the simulation are similar to the input values used in the MILP models. The simulation is run to find the network energy consumption considering SDTV and HDTV when no caches are deployed in the network and when 12 daily content replacements are performed on caches of fixed and variable sizes.

The results show close agreement between the MILP models and the simulation. They also clearly show that content replacements are more beneficial under HDTV, supporting expected future developments in video delivery technologies.

7.7 The Influence of Regular Traffic on Power Efficiency

The results presented in Chapter 6 and Chapter 7 so far have considered TV video traffic downloaded from the video head-end to network nodes, but have not included regular traffic. However, some IPTV service providers cater for both video and non-video services including web, email, data, gaming and interactive TV. In this section, some of the previous scenarios are re-evaluated assuming that the network traffic comprises both regular and TV video traffic.

7.7.1 Traffic Mixtures

In addition to the formerly evaluated network traffic having only TV video traffic, four further traffic mixtures are considered:

1. 10 – 90: Internet traffic reports forecast that by 2016 Internet video traffic will account for about 80%-90% of total Internet traffic [2].

2. 30 – 70: This is based on the fact that the 80%-90% share of video in the total traffic is made up of various types of video including TV along with VoD and Peer-to-Peer. Consequently, this mixture considers the situation where regular, VoD and Peer-to-Peer traffic represent 30% of the traffic and the remaining 70% of network traffic is TV video.

3. 50 – 50: This traffic mixture represents a service having equal amounts of regular and TV traffic.

4. 70 – 30: The traffic mixture considered here is that of a service provider whose main service is not TV but still carries some TV video content having 70% regular traffic and 30% TV traffic.

7.7.2 Power Consumption Evaluation

The intention is to investigate how the presence of regular traffic along with TV traffic influences the power consumption. The evaluation considers three network schemes that were evaluated in the previous sections: deploying caches of fixed sizes at the network nodes, performing time-based content replacements on the contents of caches of fixed sizes and replacing the contents of variable caches. The fixed-cache

MILP model proposed in Chapter 4 and content replacements MILP models are run after including regular traffic matrices in the input data.

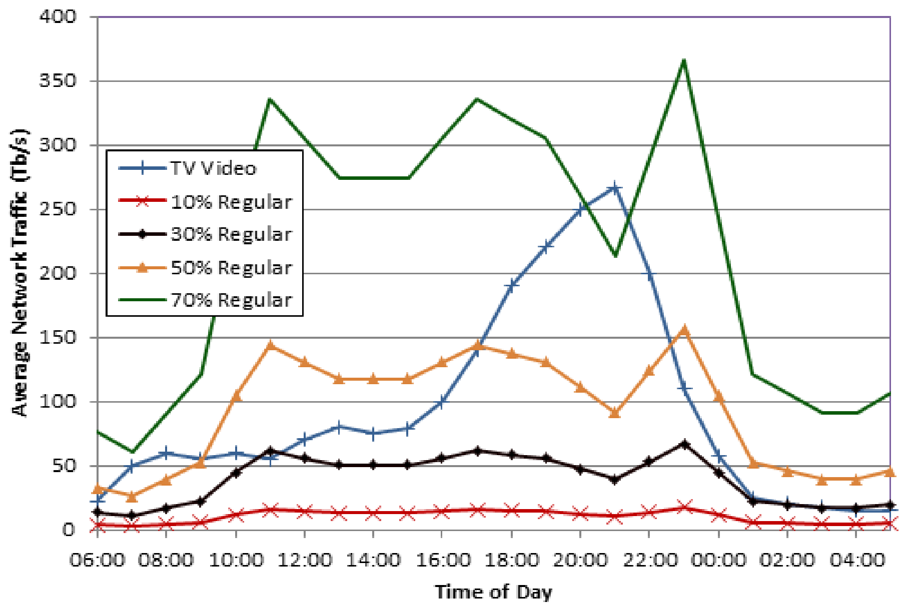

Figure 7-10: Average network traffic demand considering TV video traffic and regular traffic where regular traffic is: 10%, 30%, 50% and 70% of the total (regular + TV video) traffic

Under each considered traffic mixture, the total daily regular traffic is obtained from the total TV traffic using the traffic ratios that apply to each case. The regular traffic between each node pair is then generated using a random function with mean values compliant with the considered traffic mixtures. The trend and volume of regular traffic under each considered case are shown in Figure 7-10. The curves in Figure 7-10 show TV video and regular traffic components. Therefore the total traffic mixture of 10 – 90 for example can be calculated by adding the TV video traffic curve to the 10% regular traffic curve. The traffic volume is calculated from the average daily TV viewing figures in Figure 6-1 in Chapter 6 assuming that all TV programmes are delivered using HDTV. Therefore a peak traffic of over 260Tb/s can be observed which is moderate compared to future busy-hour Internet traffic that is expected to reach 720Tb/s in 2016 [108].

Table 7-1: Network maximum and average power savings (%) with different traffic mixtures under HDTV NSF-SVMP

Traffic Mixtures	Power Savings (%)					
	Fixed Caching		Replacements on Fixed Caches		Replacements on Variable Caches	
	Max.	Avg.	Max.	Avg.	Max.	Avg.
TV Video	31%	30%	47%	31%	83%	74%
10 – 90	30%	27%	36%	26%	78%	66%
30 – 70	26%	21%	32%	20%	68%	50%
50 – 50	21%	14%	26%	14%	54%	34%
70 – 30	15%	8%	18%	8%	38%	21%

Table 7-2: Network maximum and average power savings (%) with different traffic mixtures under HDTV BT 21CN-SV4

Traffic Mixtures	Power Savings (%)					
	Fixed Caching		Replacements on Fixed Caches		Replacements on Variable Caches	
	Max.	Avg.	Max.	Avg.	Max.	Avg.
TV Video	36%	34%	48%	34%	86%	77%
10 – 90	36%	33%	44%	32%	84%	73%
30 – 70	28%	25%	35%	24%	74%	59%
50 – 50	25%	19%	30%	18%	64%	45%
70 – 30	19%	12%	23%	12%	44%	21%

Table 7-1 and Table 7-2 show the maximum and average savings in power consumption considering TV traffic only and the four assumed traffic mixtures when deploying caches of fixed sizes and when performing content replacements on caches of fixed and variable sizes under NSF-SVMP and BT 21CN-SV4, respectively. Deploying caches in the network reduces the traffic by storing popular TV programmes locally. The presence of caches however does not reduce regular traffic passing through the network since the objects related to this traffic type are not stored in caches. Since the MILP models are linear, savings in power consumption are likely to be proportional to the portion of TV traffic in the traffic mixture. As can be inferred from Table 7-1 and Table 7-2, overall power savings are relative to the TV video component in the network traffic since maximum savings are attained when the traffic is made up of only TV and less power savings are achieved as the percentage of regular traffic increases in the traffic mixture. The linear property of the MILP models allows an estimated calculation of network power savings for any traffic mixture as long as the portion of traffic that will benefit from deployed caches is known.

Figure 7-11 and Figure 7-12 show the percentage of power savings over the time of the day when deploying caches of fixed sizes and when performing 12 content replacements on caches of fixed and variable sizes in the BT 21CN-SV4 topology. Figure 7-11 shows the power savings with traffic mixtures 10 – 90 and 30 – 70 while Figure 7-12 considers the traffic mixtures 50 – 50 and 70 – 30. The peaks of regular traffic and TV traffic are not aligned (see Figure 7-10), resulting in a different trend of power consumption over the time of the day under each considered traffic mixture. If caches of fixed sizes are deployed in the network, moderate and comparable power savings are achieved over the time of the day. These savings become more diverse as the amount of regular traffic increases in the traffic mixture.

When 12 content replacements are performed on the contents of fixed caches, the amount of power savings vary over the time of the day since popularities of TV programmes are different over the time of the day resulting in different cache hit ratios. Nevertheless, the resultant average daily power savings are similar to those assuming fixed caches with no content replacements. The maximum power savings are attained under 12 content replacements with variable caches under all traffic mixtures. The combined influence of varying the size of the cache with respect to traffic and maximising cache hit ratios due to content replacements results in the greatest power savings compared to other methods.

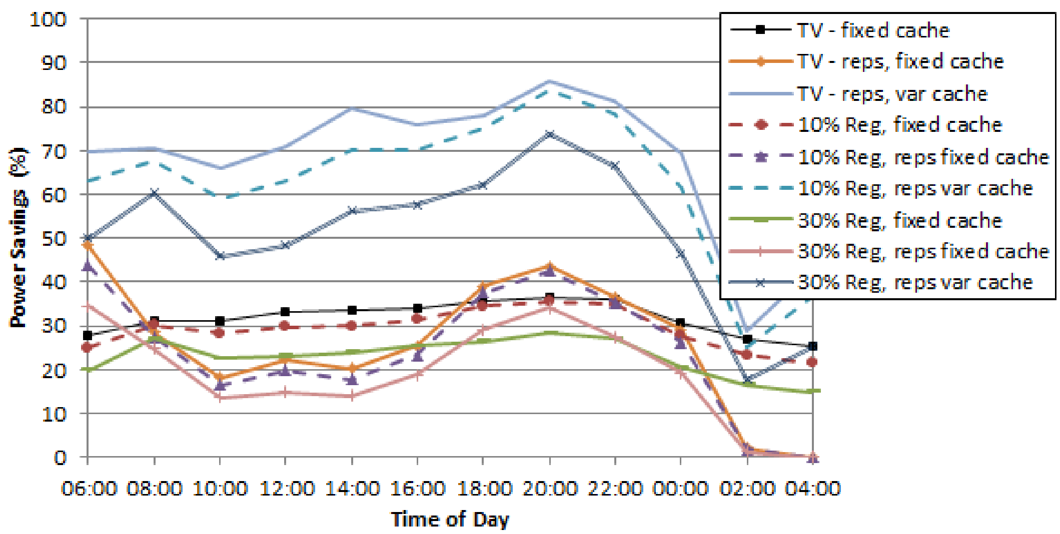

Figure 7-11: Power savings (%) over the time of the day with fixed caching and when 12 content replacements are performed on fixed and variable size caches under BT 21CN-SV4 with traffic mixtures of TV video only, 10 – 90 and 30 – 70

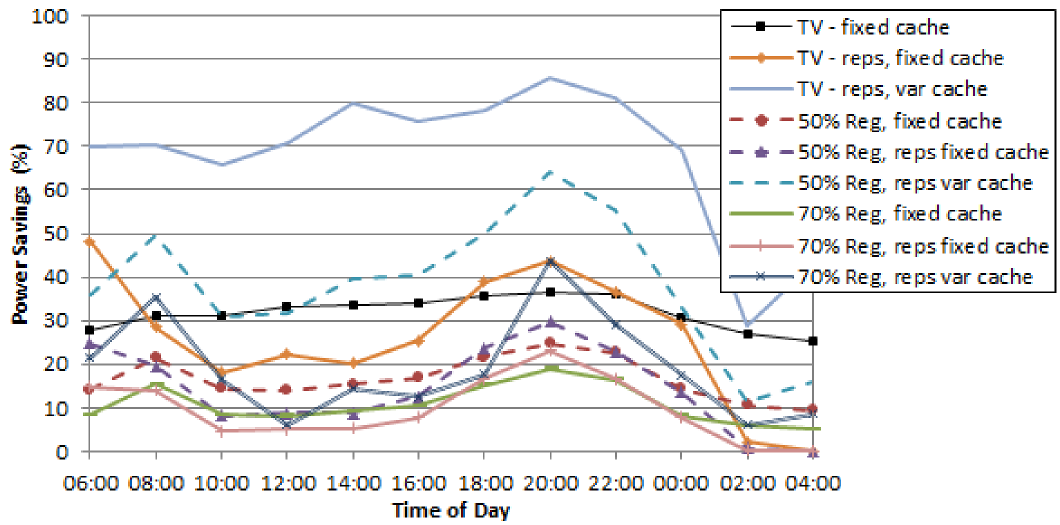

Figure 7-12: Power savings (%) over the time of the day with fixed caching and when 12 content replacements are performed on fixed and variable size caches under BT 21CN-SV4 with traffic mixtures of TV video only, 50 - 50 and 70 – 30

7.8 Current and Future Networks

In this section the significance of the proposed cache management techniques are highlighted with respect to current and future network technologies.

7.8.1 Adaptation

Current networks do not support resource adaptation and networks are provisioned for peak load. As a result, current networks consume constant power, proportional to the peak traffic. Caching significantly reduces peak traffic which could therefore lead to reduction in power consumption. However, the greatest benefits will not be obtained until variable sized caches can be deployed.

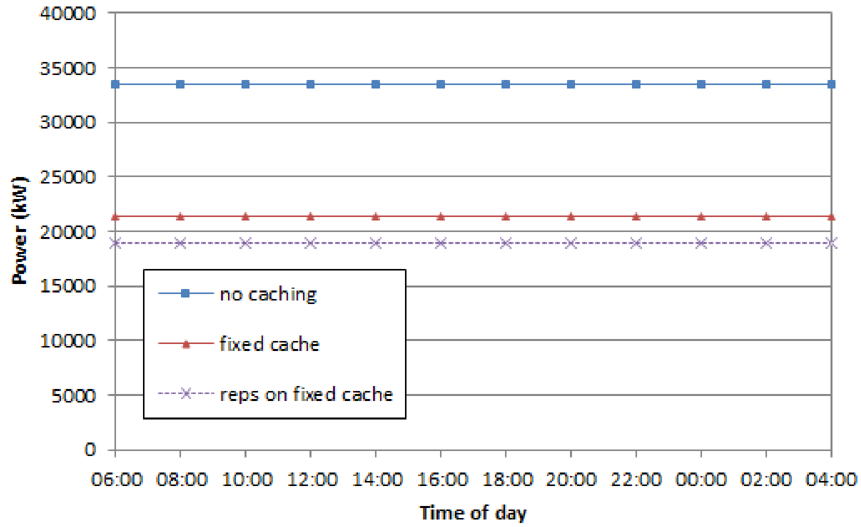

Figure 7-13: Power consumption of watching HDTV programmes assuming BT 21CN-SV4 considering no caching, caches of fixed sizes, variable cache sizes, content replacements on caches of fixed sizes and content replacements on variable cache sizes assuming the network with no resource adaptation facilities

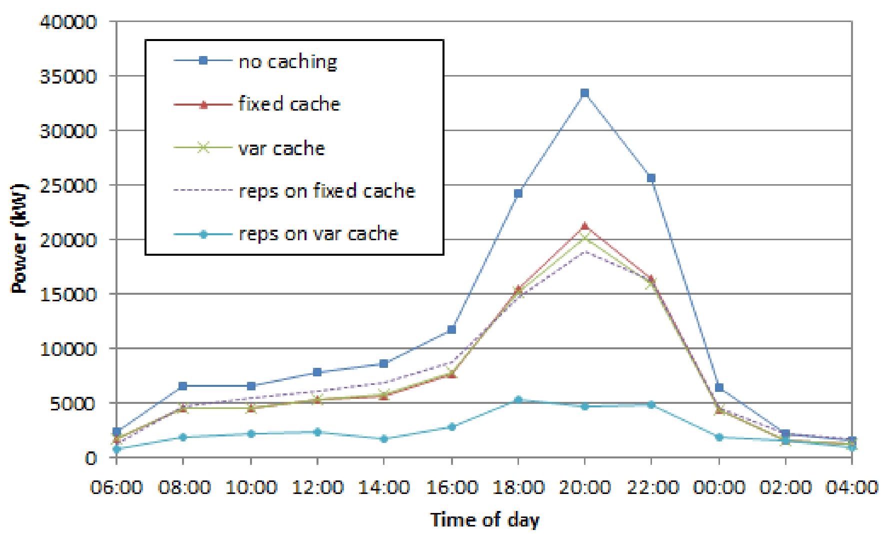

Figure 7-14: Power consumption of watching HDTV programmes assuming BT 21CN-SV4 considering no caching, caches of fixed sizes, variable cache sizes, content replacements on caches of fixed sizes and content replacements on variable cache sizes assuming the network with resource adaptation facilities

Figure 7-13 shows the power savings that would be possible with current technology caching HDTV content in the BT 21CN-SV4 topology. The maximum power savings in the network are 36% and 48% for fixed size caches without and with content replacements respectively.

Future networks are expected to be able to adapt their resources by switching off any components that are not in use. This would reduce the power consumption when the traffic is low and therefore the network consumes the maximum power only when the traffic is at its peak. The overall power consumption under this assumption follows the trend of traffic. The proposed cache management techniques reduce network traffic resulting in lower overall power consumption. Maximum power savings are achieved when the traffic peaks.

Figure 7-14 shows the power consumption of the cache management techniques considering a future network with resource adaptation capabilities. The maximum power savings are 36%, 40%, 48% and 86% with caches of fixed sizes, variable cache sizes, content replacements performed on caches of fixed sizes and replacements on variable cache sizes, respectively. The daily-averaged power savings under these schemes are 34%, 36%, 34% and 77%, respectively. As can be observed in Figure 7-13 and Figure 7-14, the proposed cache management techniques are suitable for network power consumption reduction considering both current and future network capabilities.

7.8.2 Alternative Technologies

Cache management can reduce video traffic generated by the growing number of Internet video services. An alternative technique that can effectively reduce the number of video replications in the network is multicasting. Multicast routing supports the simultaneous delivery of one copy of a video to multiple recipients.

Currently, multicast support is provided at the IP routing layer. In order to minimise energy consumption, it would be preferred to implement this functionality within the optical layer. In order for an optical core network to support multicasting, it requires the deployment of Multicast-Capable Optical Cross Connects (MC-OXCs) equipped with light splitters. However, the implementation of WDM multicast in reality is difficult for many reasons:

1. The high cost of the MC-OXC construction.
2. Multicast algorithms require a large number of wavelengths which cannot be supported by current optical device technology [109].
3. Designing optical multicast algorithms is complex as two multicast trees cannot be assigned the same wavelength if they traverse common links [110].

The benefit of multicasting over caching is that no caches are required at the nodes, and therefore no power is consumed for caching. Consequently, multicasting may be more power efficient than cache management if all OXCs in the core network could support multicast routing.

Multicast has a limitation in that viewers have no control (on demand, pause, rewind, etc.) on the video stream. Caching provides an added degree of freedom that allows delay and differentiated viewing times most suitable for emerging TV and time-shifted services.

7.9 Summary

This chapter has utilised the TV viewing data explored in Chapter 6 and proposed a cache content replacement scheme to reduce the power consumption of TV delivery. It has considered the time-driven popularity of TV programmes, taking into consideration the large number of requests associated with each programme during a time window with the programme losing its popularity during the rest of the day. Time-based popularities have been used to form equations to calculate the resultant cache hit ratios. A MILP model has been developed and validated by simulation to minimise the power consumption by optimising the number of daily content replacements. In addition, the model has been extended to perform content replacements while varying the sizes of caches at each node at each time of the day. The chapter has also studied the influence of regular traffic on percentage power savings by evaluating the power consumption assuming various traffic mixtures containing different shares of TV and regular traffic. The NSFNET and BT 21CN core topologies have been considered with both single and multiple video head-ends. Finally, the power savings achieved by caching techniques have been examined with respect to current and future networks.

When the power consumption of caching is low and fixed caches are deployed, the optimum cache sizes that minimise power consumption are sufficient to store the popular content for the whole day. Therefore, content replacements during the day are not beneficial. However, content replacements save significant power when user viewing has high variance over time with up to 48% instantaneous power savings in transmission and caching power consumption. These savings are maximised and are up to 86% when variable caches are considered. The significance of caching TV programmes rises as demand for IPTV Catch-up services grows and viewing is not real time, and is expected to surpass the power efficiency of IP multicast.

Chapter 8 Summary of Contributions and Future Directions

This chapter summarises the work that has been accomplished and presented in this thesis as well as its original contributions and findings. Furthermore, it suggests possible directions for future research in the area.

8.1 Summary of Contributions

This thesis addresses the problem of high energy usage in content delivery networks. The increasing demand for bandwidth-intensive content introduces routing and storage challenges for content providers and distributors. The energy consumption associated with such services is a major factor in designing the future Internet infrastructure.

Motivated by these challenges, the first contribution in this thesis was to evaluate the energy efficiency of caching video content delivered over an IP over WDM network. Storing popular content at the edge of the network is proposed by equipping each node with limited storage capacity in the form of disk arrays. This strategy allowed streaming one copy from the video server to each cache instead of a dedicated stream for each request. In addition, the optimum cache sizes that minimise energy usage when deployed at nodes are found based on traffic demand and power consumption parameters. Furthermore, the benefit of deploying caches with sleep-mode capabilities and their variation with traffic is examined. The main finding of this study was that caching the most popular content at the edge of the network reduces traffic on the path from servers to end users and therefore energy consumption. There exists an optimum cache size that minimises energy consumption, this size follows the increase in traffic demand. When sleep-mode capable caches are considered, the trend of the active section of caches during the day follows the trend of traffic. Using this strategy, further reductions in energy consumption are possible. These additional energy savings are more significant when the traffic experiences high fluctuations during the day.

The second contribution in this thesis is the analysis of the influence of content popularity distribution on energy consumption. The energy consumption of the network is evaluated considering a number of distributions for content popularity exemplifying a range of popular video services. While a YouTube-like service has a massive number of videos with a long tail of videos with low popularity, a CuTV service enjoys a limited number of videos, some of which are highly popular. These variances in popularity influence the resulting cache hit ratio achieved under each distribution. The main outcome of this study is that when a small number of very popular videos exist in a service, the most energy-efficient solution is to store these videos in caches, and the rest of the content is efficiently delivered remotely. Furthermore, there is a situation where not caching at all is energy-efficient, and that is when the probability of requesting a video from a huge catalogue is equally likely. A size-adaptable cache provides the most energy savings for popularity distributions that fall between the two extremes.

One day viewing figures for British TV channels were acquired and utilised to find the popularity of TV programmes for the day. This data was applied to the caching MILP model to minimise energy consumption of IPTV services considering standard and high definition streaming. Also, the topology of the BT 21CN was introduced and included in the evaluation for comparison. The main conclusion was that caching provided significant energy savings under high definition TV, supporting future video enhancement technologies. In addition, similar savings were achieved when comparing two major core topologies (NSFNET and BT 21CN), as the average hop count was comparable.

The predictable patterns of TV viewing lead to the fourth contribution of the thesis. The work contributed a proactive time-based method for cache management that updates the contents of caches a number of times a day to maintain the most popular content in caches. Updating cache contents by predicting which programmes will be requested in the next time slot increases the cache hit ratios, which reduces the traffic traversing the network, reducing in turn the power consumption. The evaluation considered fully replacing cache contents 2, 3, 4, 6 and 12 times a day and compared the energy consumption in each case to that when no replacements are performed. The main conclusion is that time-based content replacements reduce substantial energy for these services. The additional traffic due to the cache content replacements is justified by the conclusive energy savings achieved.

All proposed schemes in this thesis were evaluated using MILP models validated by simulation and a genetic algorithm. Developing generic models gave a better insight of predicting energy savings for different services. The use of real data added to the credibility of results and this data can be utilised in further research. The models introduced in this thesis support the recent green Internet movement, and are expected to contribute to further research in the area and enhance the implementation of caching in CDNs.

Note that the work in this thesis was completed prior to the increased dominance of alternative non-spinning disk technologies such as Solid State Drive (SSD) and Static Random Access Memory (SRAM). These alternative caching technologies are expected to provide higher speed and more power efficient caching.

8.2 Future Directions

This thesis has tackled the challenging task of reducing the power consumption of bandwidth-consuming video services. The findings achieved during the study are key to more energy-efficient content delivery and lead to the following future research directions.

8.2.1 Energy-Efficient Caching in the Access Network

The energy efficiency of caching in the access network can be considered, as energy savings due to caching increase with the size of the network. Also, deploying caches closer to end users minimises the journey between the source and destination much more than caches in the core network. In addition, the connectivity of the access network is denser compared to the core network. Therefore, a network topology with an increased nodal degree and average hops can be evaluated.

The work in [80] and [67] considers the access network and optimises the location where each video in the service is to be cached to minimise power usage. A caching algorithm for PONs is proposed in [111] that responds to the dynamics of video popularity distribution. It maintains the most popular videos at any time in the cache by assessing the most recent interarrival requests for videos. Additional models considering caching in the access network can be explored. The energy-efficient trade-off between having a large number of small caches located closer to users and deploying a small number of large caches further away from users can be optimised. Different emerging content delivery architectures including CCNs and P2P can be applied and compared. Analysing power savings in the access network considering the influence of video service, access network topology and content popularity distribution is worth investigating.

8.2.2 Energy-Efficient Cache Collaboration in CDNs

In the work presented in this thesis, caches are only accessed by the local node. This assumption is valid here, since the caches are considered at the core network where the average hop count is small (1.7 – 1.8). Therefore, when content is not provided by the local cache it can be retrieve from the origin server without compromising the power efficiency. Nevertheless, this setting might not be optimum in highly connected networks with large average hop counts (the access network for example) or when the location of a single origin server is not optimum.

Cache collaboration is already implemented in existing caching solutions as described in Section 2.3.2 in Chapter 2, and is an active field of research where a number of schemes have been proposed in the literature [112]-[115]. The authors of [112] propose a mechanism that locates multiple nearby copies of requested content for efficient delivery. They demonstrate that their proposed scheme delivers content faster than other schemes and reduces network traffic. In [113], a cooperative cache management algorithm is developed to minimise bandwidth cost and maximise traffic served from caches. The work classifies cache collaboration into intra-level, where content is only retrieved from leaf nodes and not from the parent node, and inter-level cache collaboration, where content is only fetched from the parent node. Results show that the proposed algorithm is guaranteed to perform within a constant factor from the globally optimal performance with better worst-case scenarios than previous work. The content caching scheme presented in [114], WAVE, adjusts the number of file chunks to be cached with respect to the popularity of the file. Each file request results in exponentially increasing the number of chunks to cache at each node on the path between the origin server and the end host. The proposed scheme reduces the average hop count of content delivery and achieves higher hit ratios compared to other schemes. The work in [115] investigates cooperative caching in content delivery networks. It evaluates the influence of request rates and the size of locally stored objects on the local cache hit rate and the cluster hit rate. The results reveal that devoting a small portion of storage area for local popular content ensures high local and cluster cache hit rates.

All these proposed mechanisms do not take into consideration the energy consumption of the network. While minimising cost and traffic has a significant influence on reducing the energy consumption of the network, incorporating energy in the optimisation process affects cache configuration and optimum content placement without compromising other performance measurements. A distributed caching environment can be investigated where cache collaboration is managed based on energy minimisation. It would be interesting to compare collaborative cache optimisation models based on geographical distance, minimum hop, cluster size and content popularity distributions. It addition, a cache content optimisation model can be used to determine a popularity threshold where all videos with popularities higher than the threshold are replicated in all caches, and the remaining are distributed over the cluster.

8.2.3 Energy Evaluation of Caching-Aided WDM Multicast

The rapidly growing market of video-rich applications calls for the utilisation of the high bandwidth provided by WDM. The importance of multicast arises from the increasing popularity of services that require the delivery of data from one source to multiple destinations. Multicasting is more beneficial in such bandwidth-intensive services including video conferencing and Internet TV. To date, full multicast in the optical layer is not viable due to cost, complexity and optical device technology reasons [109], [110]. Previous work has considered proposing models to construct multicast light trees and traffic routing over WDM networks [116]-[118]. Other studies focused on wavelength assignment strategies [119]-[122]. While some studies have considered caching-aided multicast schemes [115], [123], [124], only a few considered video traffic [125]-[127]. Nevertheless, all previous work considered multicasting in the IP network and therefore does not address the complexity of multicasting in the optical network. In addition, none of the previous work addresses the energy consumption perspective of caching-aided multicast. These facts make this topic an open area for research and extensions. Tackling the issue of the energy efficiency of caching-aided multicast in WDM has great value due to the expected increase in video traffic coupled with the success of IP over WDM for Internet traffic delivery. On top of that, proposing models that consider full optical multicast networks can provide valuable insight to expected energy reductions when all-optical multicast is realised. A very important statement would be to prove whether the deployment of all-optical multicast actually eliminates the practicability of caching in the core network.

8.2.4 Minimising Energy Usage for CDNs using Network Coding

Network coding optimises the use of bandwidth by coding multiple flows at intermediate nodes on the route of traffic. It has been widely utilised in wireless communication to increase the resilience of the network and save energy [128]-[130]. Merging the benefits of network coding with emerging network architectures such as ICNs and P2P can produce promising outcomes suitable for catering for the future Internet video demand [6], [123], [131], [132]. For instance, in [133], the location of content in an ICN is optimised and network coding is used to reduce network bandwidth and cost. Nevertheless, the energy efficiency of network coding considering CDNs and optical networks is a new research topic. Recent work in [134] and [135] evaluated the energy-efficiency of network coding in PONs. It is worth investigating the influence of introducing caching in this model on energy consumption. Another method is to evaluate coded storage at the caches. This technique can lead to savings in storage, reducing in turn energy consumption. The energy efficiency of network coding can be studied on different network levels. For instance, coding in the optical network can be coupled with coding in the IP network to maximise energy savings. Developing network coded content delivery architectures is a wide area of research where ICNs, CCNs, P2P, etc. architectures can be considered. Also, a possible direction is to evaluate the impact of physical topology and the dynamics of video content (traffic trend, location and number of sources and destinations, etc.) on energy efficiency.

References

[1] "Internet world stats." [Online]. Available: http://www.internetworldstats.com/stats.htm. [Accessed: 05-Sep-2014].

[2] Cisco white papers, "Cisco Visual Networking Index: Forecast and Methodology, 2013-2018," 2014. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html. [Accessed: 28-Aug-2014].

[3] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64–76, 2014.

[4] "SMART 2020: Enabling the low carbon economy in the information age," 2008. [Online]. Available: http://www.smart2020.org/_assets/files/02_smart2020Report.pdf.

[5] Parliamentary office of science and technology, "ICT and CO2 emissions," *Postnote no. 319*. 2008.

[6] M.-J. Montpetit, C. Westphal, and D. Trossen, "Network Coding Meets Information-Centric Networking," in *Proceedings of the 1st ACM workshop on Emerging Name-Oriented Mobile Networking Design - Architecture, Algorithms, and Applications*, 2012, pp. 31–36.

[7] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A Survey of Information-Centric Networking Research," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.

[8] R. Buyya, M. Pathan, and A. Vakali, *Content Delivery Networks*. Springer Science & Business Media, 2008.

[9] T. Plagemann, V. Goebel, A. Mauthe, L. Mathy, T. Turletti, and G. Urvoy-Keller, "From content distribution networks to content networks—issues and challenges," *Computer Communications*, vol. 29, no. 5, 2006.

[10] Y. Xiao, X. Du, J. Zhang, F. Hu, and S. Guizani, "Internet Protocol Television (IPTV): The Killer Application for the Next-Generation Internet," *IEEE Communications Magazine*, vol. 45, no. 11, pp. 126–134, 2007.

[11] F. M. V. Ramos, "Green IPTV: a resource and energy efficient network for IPTV," PhD Thesis, Cambridge University, 2012.

[12] "Network Management Solution: Optimize Infrastructure for IPTV Services." [Online]. Available: http://www.cisco.com/en/US/technologies/tk869/tk769/technologies_white_paper0900aecd80730d28.html. [Accessed: 02-Nov-2012].

[13] F. Thouin and M. Coates, "Video-on-Demand Networks: Design Approaches and Future Challenges," *IEEE Network*, vol. 21, no. 2, pp. 42–48, 2007.

[14] "YouTube Statistics." [Online]. Available: https://www.youtube.com/yt/press/statistics.html. [Accessed: 12-Sep-2014].

[15] D. De Vleeschauwer and K. Laevens, "Performance of Caching Algorithms for IPTV On-Demand Services," *IEEE Transactions on Broadcasting*, vol. 55, no. 2, pp. 491–501, 2009.

[16] J. Angel, "Caching in with content delivery," *NPN: New Public Network Magazine*. 2000.

[17] A. Jiang and J. Bruck, "Optimal content placement for en-route web caching," in *Second IEEE International Symposium on Network Computing and Applications*, 2003, pp. 9–16.

[18] P. Rodriguez, C. Spanner, and E. W. Biersack, "Analysis of Web Caching Architectures: Hierarchical and Distributed Caching," *IEEE/ACM Transactions on Networking*, vol. 9, no. 4, pp. 404–418, 2001.

[19] S. T. Chanson, "Coordinated en-route Web caching," *IEEE Transactions on Computers*, vol. 51, no. 6, pp. 595–607, 2002.

[20] "Squid Internet Object Cache -- Advanced caching proxy server for Unix." [Online]. Available: http://www.serverwatch.com/server-reviews/article.php/1376561/squid-internet-object-cache--advanced-caching-proxy-server-for-unix.htm. [Accessed: 15-Jan-2015].

[21] "Varnish Software wins Red Herring 100 Global Award | Redpill Linpro." [Online]. Available: http://www.redpill-linpro.com/varnish-software-wins-red-herring-100-global-award. [Accessed: 15-Jan-2015].

[22] "CoDeeN -- A CDN on PlanetLab." [Online]. Available: http://codeen.cs.princeton.edu/. [Accessed: 15-Jan-2015].

[23] D. K. Krishnappa, S. Khemmarat, L. Gao, and M. Zink, “On the feasibility of prefetching and caching for online TV services: a measurement study on Hulu,” in *12th International Conference on Passive and Active Measurement*, 2011, pp. 72–80.

[24] M. Hefeeda and O. Saleh, “Traffic modeling and proportional partial caching for peer-to-peer systems,” *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1447–1460, 2008.

[25] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, “YouTube Traffic Characterization: A View From the Edge,” in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007, pp. 15–28 .

[26] J. Araujo, F. Giroire, Y. Liu, R. Modrzejewski, and J. Moulierac, “Energy efficient content distribution,” in *2013 IEEE International Conference on Communications (ICC)*, 2013, pp. 4233–4238.

[27] X. Cheng, C. Dale, and J. Liu, “Statistics and Social Network of YouTube Videos,” in *2008 16th Interntional Workshop on Quality of Service*, 2008, pp. 229–238.

[28] D. N. Serpanos, G. Karakostas, and W. H. Wolf, “Effective caching of Web objects using Zipf’s law,” in *IEEE International Conference In Multimedia and Expo (ICME)*, 2000, vol. 2, pp. 727–730.

[29] M. Baron, *Probability and Statistics for Computer Scientists, Second Edition*. CRC Press, 2013.

[30] Barry C. Arnold, *Pareto distribution*. International Co-operative Publishing House, 1983.

[31] S. Dixit, *IP over WDM: building the next-generation optical Internet*. New York, USA: John Wiley & Sons, 2003.

[32] G. Shen and R. S. Tucker, “Energy-Minimized Design for IP Over WDM Networks,” *Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, 2009.

[33] L. H. Sahasrabuddhe and B. Mukherjee, “Light trees: optical multicasting for improved performance in wavelength routed networks,” *IEEE Communications Magazine*, vol. 37, no. 2, pp. 67–73, 1999.

[34] G. Fortino and C. E. Palau, *Next Generation Content Delivery Infrastructures: Emerging Paradigms and Technologies - Powell’s Books*. Information Science Reference, 2012.

[35] "Akamai." [Online]. Available: http://www.akamai.com/? [Accessed: 07-Sep-2014].

[36] "Amazon Cloudfront." [Online]. Available: aws.amazon.com/cloudfront/. [Accessed: 07-Sep-2014].

[37] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, "An Analysis of Internet Content Delivery Systems," in *The 5th Symposium on Operating Systems Design and Implementation*, 2002, vol. 36, pp. 315–327.

[38] "Global IPTV Forecast 2011-2015 Multimedia Research Group." [Online]. Available: http://www.mrgco.com/reports/amy-title-you-want-2/. [Accessed: 05-Jul-2011].

[39] A. Feldmann, A. Gladisch, M. Kind, C. Lange, G. Smaragdakis, and F.-J. Westphal, "Energy trade-offs among content delivery architectures," in *9th Conference of Telecommunication, Media and Internet*, 2010, pp. 1–6.

[40] W. Lim, P. Kourtessis, M. Milosavljevic, and J. M. Senior, "Dynamic Subcarrier Allocation for 100 Gbps, 40 km OFDMA-PONs With SLA and CoS," *Journal of Lightwave Technology*, vol. 31, no. 7, pp. 1055–1062, Apr. 2013.

[41] J. P. Mulerikkal and I. Khalil, "An Architecture for Distributed Content Delivery Network," in *15th IEEE International Conference on Networks*, 2007, pp. 359–364.

[42] B. Frank, I. Poese, G. Smaragdakis, A. Feldmann, B. M. Maggs, S. Uhlig, V. Aggarwal, and F. Schneider, "Collaboration Opportunities for Content Delivery and Network Infrastructures," *ACM SIGCOMM eBook: Recent Advances in Networking*, vol. 1, 2013.

[43] "The Akamai Internet." [Online]. Available: http://www.akamai.com/html/riverbed/akamai_internet.html. [Accessed: 07-Sep-2014].

[44] B. Krogfoss, L. Sofman, and A. Agrawal, "Caching architectures and optimization strategies for IPTV networks," *Bell Labs Technical Journal*, vol. 13, no. 3, pp. 13–28, 2008.

[45] L. B. Sofman and B. Krogfoss, "Analytical Model for Hierarchical Cache Optimization in IPTV Network," *IEEE Transactions on Broadcasting*, vol. 55, no. 1, pp. 62–70, 2009.

[46] S. Bakiras and T. Loukopoulos, “Combining replica placement and caching techniques in content distribution networks,” *Computer communications*, vol. 28, no. 9. pp. 1062–1073, 2005.

[47] J. Kangasharju, J. Roberts, and K. W. Ross, “Object replication strategies in content distribution networks,” *Computer Communications*, vol. 25, no. 4, pp. 376–383, 2002.

[48] B. Sonah and M. B. Ito, “New adaptive object replacement policy for video-on-demand systems,” in *Proceedings. Sixth International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (Cat. No.98TB100247)*, 1998, pp. 13–18.

[49] S. H. Noh, “LRFU: a spectrum of policies that subsumes the least recently used and least frequently used policies,” *IEEE Transactions on Computers*, vol. 50, no. 12, pp. 1352–1361, 2001.

[50] H. J. Roh, Y. Kim, K. S. Ko, and Y. I. Eom, “Design of a Content Replacement Scheme using the p-based LRFU-k algorithm in Contents Delivery Networks,” in *2008 10th International Conference on Advanced Communication Technology*, 2008, vol. 3, pp. 2067–2070.

[51] S. Podlipnig and L. Boszormenyi, “Replacement strategies for quality based video caching,” in *Proceedings. IEEE International Conference on Multimedia and Expo*, 2002, vol. 2, pp. 49–52.

[52] I. Psaras, W. K. Chai, and G. Pavlou, “Probabilistic in-network caching for information-centric networks,” in *ACM Workshop on Information- Centric Networking (ICN)*, 2012.

[53] T. Wu, K. De Schepper, W. Van Leekwijck, and D. De Vleeschauwer, “Reuse time based caching policy for video streaming,” in *IEEE Consumer Communications and Networking Conference (CCNC)*, 2012, pp. 89–93.

[54] Q. Zhu and Y. Zhou, “Power-Aware Storage Cache Management,” *IEEE Transactions on Computers*, vol. 54, no. 5, pp. 587–602, 2005.

[55] A. Silberschatz, P. B. Galvin, and G. Gagne, *Operating System Concepts*, 9th ed. Wiley, 2012.

[56] C. Fang, F. R. Yu, T. Huang, J. Liu, and Y. Liu, “Energy-efficient distributed in-network caching for Content-Centric Networks,” in *2014 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2014, pp. 91–96.

[57] M. Gupta and S. Singh, “Greening of the internet,” in *Proceedings of the ACM conference on Applications, technologies, architectures, and protocols for computer communications*, 2003, p. 19.

[58] J. Baliga, K. Hinton, and R. S. Tucker, “Energy Consumption of the Internet,” in *COIN-ACOFT 2007 - Joint International Conference on the Optical Internet and the 32nd Australian Conference on Optical Fibre Technology*, 2007, pp. 1–3.

[59] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, “Energy Efficiency in Telecom Optical Networks,” *IEEE Communications Surveys & Tutorials*, vol. 12, no. 4, pp. 441–458, 2010.

[60] F. Vismara, V. Grkovic, F. Musumeci, M. Tornatore, and S. Bregni, “On the energy efficiency of IP-over-WDM networks,” in *2010 IEEE Latin-American Conference on Communications*, 2010, pp. 1–6.

[61] M. Caria, M. Chamania, and A. Jukan, “To switch on or off: A simple case study on energy efficiency in IP-over-WDM networks,” in *2011 IEEE 12th International Conference on High Performance Switching and Routing*, 2011, pp. 70–76.

[62] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, “IP Over WDM Networks Employing Renewable Energy Sources,” *Journal of Lightwave Technology*, vol. 29, no. 1, pp. 3–14, 2011.

[63] D. Xiaowen, T. El-Gorashi, J. M. H. Elmirghani, and X. Dong, “Green IP Over WDM Networks With Data Centers,” *Journal of Lightwave Technology*, vol. 29, no. 12, pp. 1861–1880, 2011.

[64] “BT sources 100% renewable electricity from npower,” 2013. [Online]. Available: http://www.btplc.com/News/Articles/Showarticle.cfm?ArticleID=C5B8C544-A575-485C-8123-81F67CD9AC07. [Accessed: 25-Sep-2014].

[65] “GreenTouch.” [Online]. Available: www.GreenTouch.org. [Accessed: 25-Sep-2014].

[66] “The green grid.” [Online]. Available: http://www.thegreengrid.org. [Accessed: 25-Sep-2014].

[67] C. A. Chan, E. Wong, and A. Nirmalathas, “Energy savings dependency of IPTV caching systems on similarity in user behavior,” in *37th European Conference and Exhibition on Optical Communication (ECOC)*, 2011, pp. 1–3.

[68] A. Dewangan and D. Jalihal, “Statistics based energy efficient caching decisions for IPTV services,” in *2013 National Conference on Communications (NCC)*, 2013, pp. 1–5.

[69] F. Ramos, P. Rodriguez, R. Gibbens, J. Crowcroft, F. Song, and I. White, “Reducing Energy Consumption in IPTV Networks by Selective Pre-Joining of Channels,” in *Proceedings of the first ACM SIGCOMM workshop on Green networking*, 2010, pp. 47–52.

[70] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, “Watching television over an IP network,” in *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement conference - IMC ’08*, 2008, pp. 71–84.

[71] Y. Zhu and J. P. Jue, “Energy-efficient flow aggregation for IPTV program delivery in optical backbone networks with multiple line rates,” in *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference* , 2011, pp. 1–3.

[72] K. Guan, G. Atkinson, D. C. Kilper, and E. Gulsen, “On the Energy Efficiency of Content Delivery Architectures,” in *2011 IEEE International Conference on Communications Workshops (ICC)*, 2011, pp. 1–6.

[73] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, “Networking named content,” in *Proceedings of the 5th international conference on Emerging networking experiments and technologies - CoNEXT ’09*, 2009, pp. 1–12.

[74] J. Llorca, A. M. Tulino, K. Guan, J. Esteban, M. Varvello, N. Choi, and D. C. Kilper, “Dynamic in-network caching for energy efficient content delivery,” in *2013 Proceedings IEEE INFOCOM*, 2013, pp. 245–249.

[75] L. Chiaraviglio and I. Matta, “GreenCoop: Cooperative Green Routing with Energy-efficient Servers,” in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking - e-Energy ’10*, 2010, p. 191.

[76] L. Chiaraviglio and I. Matta, “An energy-aware distributed approach for content and network management,” in *2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2011, pp. 337–342.

[77] V. Mathew, R. K. Sitaraman, and P. Shenoy, “Energy-aware load balancing in content delivery networks,” in *2012 Proceedings IEEE INFOCOM*, 2012, pp. 954–962.

[78] K. Zeng and J. Yang, “Energy-Aware Server Provisioning in Large Scale Video-On-Demand Systems,” in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, 2010, pp. 1–5.

[79] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao, “Energy-Aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services,” in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, 2008, pp. 337–350.

[80] C. Jayasundara, A. Nirmalathas, E. Wong, and C. A. Chan, “Improving Energy Efficiency of Video on Demand Services,” *Journal of Optical Communications and Networking*, vol. 3, no. 11, p. 870, 2011.

[81] “Practical Optimization: A Gentle Introduction.” [Online]. Available: http://www.sce.carleton.ca/faculty/chinneck/po.html. [Accessed: 02-Sep-2014].

[82] “Introduction to Integer Linear Programming.” [Online]. Available: http://wpweb2.tepper.cmu.edu/fmargot/introILP.html. [Accessed: 02-Sep-2014].

[83] “Xpress-optimizer reference manual,” *FICO Xpress optimization suite*, 2009. [Online]. Available: http://brblog.typepad.com/files/optimizer-1.pdf. [Accessed: 02-Sep-2014].

[84] “AMPL.” [Online]. Available: http://ampl.com/. [Accessed: 29-Sep-2014].

[85] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing*. Springer Science & Business Media, 2003.

[86] C. Chiu and P.-L. Hsu, “A Constraint-Based Genetic Algorithm Approach for Mining Classification Rules,” *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, vol. 35, no. 2, pp. 205–220, 2005.

[87] “Map of all Google Data Center Locations.” [Online]. Available: http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/. [Accessed: 18-Jan-2011].

[88] “YouTube Architecture.” [Online]. Available: http://highscalability.com/youtube-architecture. [Accessed: 22-Apr-2012].

[89] “YouTube fact sheet, traffic and stats,” 2012. [Online]. Available: http://www.youtube.com/t/fact_sheet. [Accessed: 12-Apr-2012].

[90] “Cisco CRS-1 8-Slot Single-Shelf System.” [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/routers/crs-1-8-slot-single-shelf-system/product_data_sheet0900aecd801d53a1.html. [Accessed: 22-Mar-2013].

[91] "Cisco ONS 15454 Optical Filter Cards." [Online]. Available: http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/product_data_sheet09186a00801a5572_ps5791_Products_Data_Sheet.html. [Accessed: 22-Mar-2011].

[92] "Cisco ONS15501 Erbium Doped Fiber Amplifier." [Online]. Available: http://www.cisco.com/en/US/products/hw/optical/ps2011/products_data_sheet09186a008008870d.html. [Accessed: 22-Mar-2011].

[93] "Glimmerglass Intelligent Optical System 500." [Online]. Available: http://www.glimmerglass.com/products/intelligent-optical-system-500/. [Accessed: 22-Mar-2011].

[94] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez, "Greening the internet with nano data centers," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies - CoNEXT '09*, 2009, pp. 37–48 .

[95] J. Elerath and M. Pecht, "A Highly Accurate Method for Assessing Reliability of Redundant Arrays of Inexpensive Disks (RAID)," *IEEE Transactions on Computers*, vol. 58, no. 3, pp. 289–299, 2009.

[96] "Broadcasters' Audience Research Board (BARB)." [Online]. Available: http://www.barb.co.uk/whats-new/weekly-top-30? [Accessed: 20-Nov-2012].

[97] "UK TV Listings." [Online]. Available: http://onthebox.com/#/. [Accessed: 09-Nov-2012].

[98] "BBC iPlayer Performance Pack - November 2012." [Online]. Available: http://www.bbc.co.uk/mediacentre/latestnews/2012/iplayer-performance-nov12.html. [Accessed: 24-Dec-2012].

[99] "Ofcom Communications Market Report 2013." [Online]. Available: http://stakeholders.ofcom.org.uk/binaries/research/cmr/cmr13/2013_UK_CMR.pdf. [Accessed: 27-Mar-2014].

[100] "TV and Satellite week Magazine, 14-20 September 2013."

[101] BT Wholesale, "Current Thinking on Colossus & 21CN Connectivity," 2006.

[102] "BT 21CN - Network Topology & Technology." [Online]. Available: http://www.kitz.co.uk/adsl/21cn_network.htm. [Accessed: 18-Mar-2013].

[103] "The Digital Subscriber Line Forum", “Triple-play Services Quality of Experience (QoE) Requirements.” [Online]. Available: http://www.broadband-forum.org/technical/download/TR-126.pdf. [Accessed: 15-Feb-2012].

[104] “Ofcom, Communications Market Report 2012.” [Online]. Available: http://stakeholders.ofcom.org.uk/binaries/research/cmr/cmr12/CMR_UK_2012.pdf. [Accessed: 19-Jan-2015].

[105] “BBC iPlayer Performance Pack - June to August 2012.” [Online]. Available: http://downloads.bbc.co.uk/mediacentre/iplayer/iplayer-performance-jun-aug-12.pdf. [Accessed: 19-Jan-2015].

[106] “BBC iPlayer Performance Pack - September 2012.” [Online]. Available: http://downloads.bbc.co.uk/mediacentre/iplayer/iplayer-performance-sep12.pdf. [Accessed: 19-Jan-2015].

[107] “BBC iPlayer Performance Pack - October 2012.” [Online]. Available: http://downloads.bbc.co.uk/mediacentre/iplayer/iplayer-performance-oct12.pdf. [Accessed: 19-Jan-2015].

[108] “Cisco Visual Networking Index: The Zettabyte Era, Cisco White Papers.” [Online]. Available: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/VNI_Hyperconnectivity_WP.html. [Accessed: 08-May-2013].

[109] D. Danh Le, M. Molnar, and J. Palaysi, “An improved multicast routing algorithm in sparse splitting WDM networks,” in *2013 International Conference on Computing, Management and Telecommunications (ComManTel)*, 2013, pp. 99–104.

[110] N. K. Singhal, L. H. Sahasrabuddhe, and B. Mukherjee, “Optimal Multicasting of Multiple Light-Trees of Different Bandwidth Granularities in a WDM Mesh Network With Sparse Splitting Capabilities,” *IEEE/ACM Transactions on Networking*, vol. 14, no. 5, pp. 1104–1117, Oct. 2006.

[111] C. Jayasundara, A. Nirmalathas, E. Wong, and N. Nadarajah, “Popularity-Aware Caching Algorithm for Video-on-Demand Delivery over Broadband Access Networks,” in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, 2010, pp. 1–5.

[112] M. Lee, K. Cho, K. Park, T. Kwon, and Y. Choi, “SCAN: Scalable Content Routing for Content-Aware Networking,” in *2011 IEEE International Conference on Communications (ICC)*, 2011, pp. 1–5.

[113] S. Borst, V. Gupta, and A. Walid, "Distributed Caching Algorithms for Content Distribution Networks," in *2010 Proceedings IEEE INFOCOM*, 2010, pp. 1–9.

[114] K. Cho, M. Lee, K. Park, T. T. Kwon, and Y. Choi, "WAVE: Popularity-based and collaborative in-network caching for content-oriented networks," in *2012 Proceedings IEEE INFOCOM Workshops*, 2012, pp. 316–321.

[115] D. H. K. Tsang and J. Ni, "Large-scale cooperative caching and application-level multicast in multimedia content delivery networks," *IEEE Communications Magazine*, vol. 43, no. 5, pp. 98–105, 2005.

[116] J. Y. Wei, "Constrained multicast routing in WDM networks with sparse light splitting," *Journal of Lightwave Technology*, vol. 18, no. 12, pp. 1917–1927, 2000.

[117] H. T. Mouftah, "Dynamic Constrained Multicast Routing in WDM Networks: Blocking Probability, QoS and Traffic Engineering," in *11th IEEE Symposium on Computers and Communications (ISCC'06)*, 2006, pp. 974–980.

[118] H.-H. Yen, S. S. W. Lee, and B. Mukherjee, "Traffic Grooming and Delay Constrained Multicast Routing in IP over WDM Networks," in *2008 IEEE International Conference on Communications*, 2008, pp. 5246–5251.

[119] K.-M. Yong, T.-H. Cheng, and G.-S. Poo, "Dynamic Multicast Routing and Wavelength Assignment with Minimal Conversions in Delay-Constrained WDM Networks," in *2009 Proceedings of 18th International Conference on Computer Communications and Networks*, 2009, pp. 1–6.

[120] Y. Zhou and Gee-Swee Poo, "A new multiwavelength multicast wavelength assignment (MMWA) algorithm in wavelength-routed WDM networks," in *2004 IEEE International Conference on Communications*, 2004, vol. 3, pp. 1786–1790.

[121] J. Wang, X. Qi, and B. Chen, "Wavelength assignment for multicast in all-optical WDM networks with splitting constraints," *IEEE/ACM Transactions on Networking*, vol. 14, no. 1, pp. 169–182, 2006.

[122] B. Chen and J. Wang, "Efficient routing and wavelength assignment for multicast in WDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 1, pp. 97–109, 2002.

[123] J. Llorca, A. M. Tulino, K. Guan, and D. C. Kilper, "Network-coded caching-aided multicast for efficient content delivery," in *2013 IEEE International Conference on Communications (ICC)*, 2013, pp. 3557–3562.

[124] J. Zhang, Z. Li, and L. Chen, "Dynamic Cache Allocation Algorithm and Replacement Policy for Reliable Multicast Network," in *2009 5th International Conference on Wireless Communications, Networking and Mobile Computing*, 2009, pp. 1–5.

[125] K. Katsaros, G. Xylomenos, and G. C. Polyzos, "A Hybrid Overlay Multicast and Caching Scheme for Information-Centric Networking," in *2010 INFOCOM IEEE Conference on Computer Communications Workshops*, 2010, pp. 1–6.

[126] S. Ramesh and K. Guo, "Multicast with cache (Mcache): an adaptive zero-delay video-on-demand service," in *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213)*, 2001, vol. 1, pp. 85–94.

[127] J. Y. Kim, G. M. Lee, and J. K. Choi, "Efficient Multicast Schemes Using In-Network Caching for Optimal Content Delivery," *IEEE Communications Letters*, vol. 17, no. 5, pp. 1048–1051, 2013.

[128] C. Fragouli, J. Widmer, and J.-Y. Le Boudec, "A Network Coding Approach to Energy Efficient Broadcasting: From Theory to Practice," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, 2006, pp. 1–11.

[129] P. Poocharoen and M. Magana, "Energy efficient multi-hop communication using partial network coding with cooperation," *IEEE Latin America Transactions*, vol. 10, no. 4, pp. 1874–1880, 2012.

[130] J. Chen, V. Lee, and E. Chan, "Network coding-aware cache replacement policy in on-demand broadcast environments," in *2011 6th International ICST Conference on Communications and Networking in China (CHINACOM)*, 2011, pp. 691–697.

[131] H. Wang, Y. Zhang, P. Li, and Z. Jiang, "The Benefits of Network Coding in Distributed Caching in Large-Scale P2P-VoD Systems," in *2010 IEEE Global Telecommunications Conference GLOBECOM*, 2010, pp. 1–6.

[132] C. Gkantsidis and P. R. Rodriguez, "Network coding for large scale content distribution," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, 2005, vol. 4, pp. 2235–2245.

[133] J. J. Wang, J. Ren, K. Lu, S. Liu, and C. Westphal, "An optimal Cache management framework for information-centric networks with network coding," in *2014 IFIP Networking Conference*, 2014, pp. 1–9.

[134] T. Ichikawa, M. Tadokoro, T. Kubo, T. Yamada, K.-I. Suzuki, N. Yoshimoto, and R. Kubo, "Energy-efficient peer-to-peer communication with network coding over WDM/TDM-PON," in *2013 IEEE 2nd Global Conference on Consumer Electronics (GCCE)*, 2013, pp. 481–482.

[135] X. Liu, K. Fouli, R. Kang, and M. Maier, "Network-Coding-Based Energy Management for Next-Generation Passive Optical Networks," *Journal of Lightwave Technology*, vol. 30, no. 6, pp. 864–875, 2012.