A Green Analysis of the Content Centric Networking Architecture

Muhammad Rizwan Butt

Department of Electrical & Computer Engineering
McGill University
Montreal, Canada

November 2012

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Master of Engineering.

2012/11/01

Abstract

Content-Centric Networking (CCN) is a recently proposed networking architecture that can potentially lead to reduced bandwidth usage and better scalability and security as compared to the current IP-based architecture. In this thesis, we conduct a green analysis of content-centric networking and IP-based networking for a video streaming scenario. We consider two types of energy consumption: the energy required to manufacture the network devices and the energy required for operation. We perform simulations of content centric networking over three different network topologies (i.e., general tree, Content Distribution Network (CDN) tree and partial mesh) to assess the traffic rate reductions achieved by CCN's insertion of caches at routers. We generated two different types of traffic demands (Zipf and Uniformly distributed) to perform our analysis. Although CCN network devices have a higher intrinsic energy consumption compared to the IP-based devices because of the presence of additional memory, by exploiting their caching capabilities it is possible to reduce the overall energy consumption of the network. Content caching at the routers present on lower levels of the network (near clients) results in reducing traffic on the links which are close to the server (content source). We exploit this feature of CCN-based network by using rate adaptation to achieve energy benefits. We consider both the incorporation of an on-line rate adaptation mechanism as well as a static network provisioning approach and observe that these approaches can lead to a substantial reduction in energy consumption for CCN. On the other hand, an IP-based network cannot benefit from rate adaptation due to the absence of the cache capable routers.

Sommaire

CCN (Content Centric Networks) est une architecture réseau récemment proposée. Elle peut potentiellement réduire l'utilisation de bande passante et améliorer l'extensibilité et la sécurité du réseau par rapport á l'architecture IP existante. Dans cette thèse, nous conduisons une analyse énergétique comparative des CCN et des réseaux IP dans le cas du streaming vidéo. Nous considérons deux types de consommation d'énergie: celle requise pour construire les éléments du réseau et celle requise pour le fonctionnement du réseau. Nous réalisons des simulations de CCN sur trois topologies réseaux différentes (réseau en arbre, réseau de distribution et maillage partiel) afin de vérifier la réduction du traffic obtenue avec l'introduction d'un cache aux niveau des routeurs. Nous générons deux types de demandes de traffic (Zipf et distribution uniforme) pour réaliser cette analyse. Bien que les éléments d'un réseau CCN aient une plus grande consommation d'énergie par rapport á leur équivalent des réseaux IP et qui sont dues á la présence de mémoire supplémentaire, l'exploitation de leur capacité de cache permet de réduire la consommation d'énergie totale du réseau. Contenu de mise en cache au niveau des routeurs prsents sur les niveaux inférieurs du réseau (clients prés) se traduit par la réduction du trafic sur les liens qui sont á proximité du serveur (source de contenu). Nous exploitons cette caractéristique du CCN á base de réseau á l'aide d'adaptation de débit pour obtenir des avantages de l'énergie. Nous considérons á la fois l'incorporation d'un mécanisme en ligne taux d'adaptation ainsi que d'un réseau statique approche de provisionnement et d'observer que ces approches peuvent conduire á une réduction substantielle de la consommation d'énergie pour les CCN. D'autre part, un réseau basé sur IP ne peut pas bénéficier de l'adaptation du débit en raison de l'absence des routeurs capables de cache.

Acknowledgments

In the first place, I am truly indebted and thankful to my supervisor, Prof. Mark Coates, for his excellent guidance and advice in completion of this thesis. I was highly motivated by the challenging research environment provided by him in the lab and it encouraged me immensely for my research.

I am grateful to my dear parents for their love and constant prayers for me throughout my studies abroad. My self-belief is a gift from my parents which helped me all the way through my life. I am obliged to my fellow colleagues of the Computer Networks Research Lab for their companionship and assistance, particularly to Santosh, Zhe Yao, Deniz, Alok, Milad, Oscar, Haani, Nilofar, Rodrigo and Arslan. I would also like to thank Prof. Michael Rabbat for his precious comments during my lab presentations which helped me to improve my work. Besides I would like to thank Benjamin for translating the abstract into French and Carrie Serban for her help in the administrative matters.

A special thanks to my wife for her continuous love and support in the two years of our married life.

Most of all, I am thankful to Allah for His countless blessings on me in the most difficult circumstances. He bestowed me with determination and will power which helped me every day.

Contents

List of Figures

List of Tables

List of Acronyms

CCN	Content Centric Network
CDN	Content Dissemination Network
IP	Internet Protocol
VoD	Video on Demand
Mbps	Mega bits per second
MB	Mega Byte
GB	Mega Byte
TB	Tera Byte

Chapter 1

Introduction

1.1 Motivation

The demand for distribution of content through the Internet is increasing every day. A major portion of the Internet traffic is generated by video streaming and Video on Demand (VoD) applications. The host oriented IP-based architecture is not well-matched to these types of applications. To handle this issue different approaches have been adopted in the recent past including peer-to-peer [2] and content-distribution networks (CDNs) [3]. These solutions are very effective but cannot be applied throughout the network. This motivates the development of a new networking paradigm. Substantial research efforts have been made in recent years to explore data-centric networking architectures. Ghodsi et al. in [4] have provided a review of the developments in this field and also shed light on commonalities and differences of the proposals. One of the prominent approaches in this regard is content centric networking (CCN) [5]. Numerous projects have been initiated to address the research issues related to the architectural technicalities of CCN. One important aspect that has received little attention is the energy efficiency of the content centric networking approach. Research on energy-efficient communication networks, especially with respect to environmental concerns is commonly referred to as "green networking". Prem et al. in [6] provided a detailed survey on green networking research and argue that working towards a green network can result in reduced operational costs, reduced power usage and increased network and power efficiency.

1.2 Thesis problem statement

The existing Internet consists of thousands of network devices (i.e., routers, switches, hubs) deployed all over the world. Introducing a new networking architecture would involve replacement or upgradation of many of these devices at a global level. In order to commit to a change of such large scale, the research and industrial community would need to be convinced that the new networking paradigm is superior to the existing one in many respects and has no major deficiencies. The need to reduce the emissions of greenhouse gases (i.e., carbon footprint) for environmental protection and rising energy costs have turned energy efficiency into an important consideration.

The problem is to find a reasonable mechanism for comparing the energy and environmental aspects of CCN with the existing Internet architecture. CCN is still in its initial stages of development. According to our knowledge, devices having CCN capabilities are not currently available. The problem is then twofold: (i) model a CCN and implement its basic functionalities, (ii) determine the energy and carbon footprint of CCN and host-oriented network architectures by running a similar application on both networks. Our goal is to analyze whether the CCN is greener than the IP-based Internet.

1.3 Thesis contribution and organization

In this thesis, we provide a comparative study of CCN and IP-based networks in terms of energy consumption. As it constitutes one of the more compelling applications for the CCN architecture, we consider a dedicated video streaming application. Our analysis includes the energy consumption during both manufacturing and operation of the network devices.

We use the notion of emergy [7] to embody the latter expenditure. The network routers in the CCN architecture are provisioned with additional memory and have the capability to cache content. The manufacturing and operation of this memory is an additional energy expenditure compared to the IP network. This implies that if both network architectures are operated in the same fashion, the CCN approach will consume more energy to deploy and operate. The caching of content reduces the traffic on many links, suggesting that links and routers in the CCN architecture can operate at a lower rate compared to the IP-based implementation. This can lead to a decreased energy consumption [8–10], under the reasonable assumption that devices operating at lower rates require less energy. We investigate the impact of two different strategies on energy consumption. The first involves on-line rate adaptation, in which routers are capable of adjusting their operating rate (and hence energy consumption) according to the bandwidth demand. The other is a static approach, where links and routers are statically provisioned according to the anticipated load. We have implemented two simulators for content-dissemination in CCNs and we use them to estimate the traffic rate on each link during operation of the VoD application. This allows us to calculate energy consumption for the two proposed strategies. Due to the absence of cache capable routers IP-based networks cannot take advantage of the proposed rate adaptation strategies.

1.4 Published Work

Some parts of this thesis have appeared in following publications. The work has been performed in collaboration with Oscar Delgado and Prof. Mark Coates.

1.4.1 Energy assessment of CCN (Chapters 3 and 4)

- M. Butt, O. Delgado, and M. Coates,“An Energy-efficiency Assessment of Content Centric Networking (CCN)”, in *Proc.* IEEE 25th Annual Canadian Conference on

Electrical and Computer Engineering, Montreal, Canada, Mar. 2012.

In this paper, Oscar Delgado helped in the implementation of CCN simulator *CCNSim* in MATLAB and producing some simulation results.

- M. Butt, M. Coates,“An Energy-efficiency Assessment of Content Centric Networking (CCN)”, Colloquium Entretiens Jacques Cartier Information and Communications Technologies: Are they Green?, Montreal, Canada Oct. 2011. (Poster)

Chapter 2

Background and Literature Review

In this chapter, we provide a brief discussion of the literature available on content-based networking. Initially, we discuss the different projects which are based on developing a content-based networking architecture. In the second section, we focus on one particular project, CCNx [11], and explain the merits of content centric networking. Section 2.3 discusses some of the work on modeling of a content centric network. Section 2.4 and 2.5 provide brief overviews of green networking and dynamic rate adaptation respectively.

2.1 Content-based networking architectures

A significant fraction of the Internet traffic is composed of multi-media content [12, 13]. It requires high speed servers and mass storage to accommodate user demand. This makes the networking paradigm more complex and difficult to manage for the existing host-oriented network architecture. At present three projects are being developed in parallel to address this issue using different networking paradigms. Content centric networking (CCN) [5] is a content based architecture which is implementable over the existing IP-based infrastructure. It uses named content[1] rather than host identifiers as its central abstraction, which enables

[1] CCN names the chunks of content at the packet level of the network, similar to web's uniform resource identifier (URI) hierarchical naming structure.

it to retain the simplicity and scalability of IP and potentially provide much better security, delivery efficiency, and disruption tolerance.

Network of Information (NetInf) [14] uses a publish/subscribe data dissemination scheme. NetInf is an information-centric global networking architecture which delivers dissemination and non-dissemination objects only to the consenting nodes and as a result reduces unwanted traffic and provides host independent communication. The compatibility of NetInf with the existing Internet architecture is an issue which still remains to be addressed.

The Autonomic Network Architecture (ANA) [15] is a meta-architecture which introduces a 'clean-slate' approach in order to promote the evolution, adaptation and federation of novel networking schemes and protocols for the Internet. Its API uses a CONTEXT/SERVICE communication model to fit most networking applications and protocols.

All of these architectures are new, and very little work has been done to compare them and identify the best way forward for the future Internet. From the perspective of networking community, a backward compatible architecture would definitely be a better choice. The implementation of NetInf over the existing IP-based Internet infrastructure, as mentioned earlier, remains an open research problem. ANA presents a clean-slate approach i.e., to re-architecture the Internet from scratch. This does not appear to be a feasible option. CCN is the only architecture which is fully compliant with the existing infrastructures (hardware and protocols) among the three. It has already been incrementally implemented on the existing IP network for a Voice over CCN application [16].

Some previous notable projects and programs which contributed in the development of information centric networking were DONA, PSIRP, SAIL [17–19]. Significant effort has been devoted recently by the future Internet programs, such as GENI, FIND [20, 21] (in USA) and FIA [22] (in Europe) towards developing a network of information where content becomes a first class network entity.

2.2 Content centric networking (CCN)

Project PARC CCNx [11] was initiated to develop, promote, and evaluate a new communication architecture called content centric networking. This new architecture is the focus of a long-term research and development program. There are interesting problems in many areas of CCNx that remain to be solved. Jacobson et al. in [5] provide the full architectural details of CCN. They argue that today's Internet is predominantly a distribution system where users are more interested in the distribution of the content rather than point to point communication between hosts. The existing Internet architecture was designed to share resources, not data. CCN is a networking architecture built on similar principles as the existing IP architecture, but instead of using location based (host to host) communication to retrieve content CCN uses named data. Content can be of various nature, including data generated "on the fly" for conversational and interactive services or the dissemination of live events. Every content is split into chunks whose unique names follow a hierarchical structure inspired by that of webs URIs. The hierarchical naming structure enables aggregation as in IP routing tables and more generally allows system scalability. There are two types of packets used: *Interest* and *Data*. During the content publication phase, nodes with content that may be interesting to others announce content names or the prefixes of the name, similar to IP routing announcements. Once the content is published a consumer node broadcasts *Interest* packets having the names of the requested chunks. Any node within its communication radius with the required content sends back the *Data* packet. The flow balance of CCN's *Interest* and *Data* packets is similar to TCP's *Data* and *ACK* packets but the CCN model provides many-to-many multi-point delivery which is not possible with TCP.

CCN routers are composed of three main blocks: a content store, a pending *Interest* table and a forwarding information base. These routers are responsible for processing incoming *Interests* at the network layer. For every incoming *Interest*, the router first checks the content store for the requested chunk. If the chunk is present it is directly returned to the requesting node. Otherwise, it forwards the *Interest* to a next hop determined via the forwarding information base. It also creates an entry in the pending *Interest* table, so that when the chunk is received, the router will know where to send it. A 'strategy layer' defines the forwarding strategies. Due to these forwarding strategies the *Interest* packets propagate along multiple paths toward the potential locations of the data and pull the data down to the requesting nodes, with no need for explicit routing. This data forwarding is different from IP forwarding as intermediate nodes may transparently cache data packets in the content store for later transmission e.g., in response to an *Interest* for the same chunk. If several users request the same chunk, only one *Interest* is forwarded upstream and a single pending *Interest* table entry records the address of all the interfaces over which the chunk must be returned. This feature of CCN realizes multi-cast. The content store plays the role of cache by storing the received chunks, which helps to eliminate the need to repeatedly fetch popular content. Within a cache, content chunks are stored and replaced according to a specific policy e.g., Least Frequently Used (LFU) or Least Recently Used (LRU).

Mobility and rich connectivity are also among some distinguishing features of CCN. IP-based protocols cannot take advantage of more than one interface on a mobile device while CCN can. "CCN talks *about* data, not *to* nodes" [5], so it does not need to bind an IP address to a MAC address, hence rich connectivity and mobility can be achieved by exploiting multiple interfaces. Each node gets per-prefix, per-interface performance information, allowing it to adaptively choose the "best" interfaces for forwarding *Interests* matching some prefix. Although the main idea behind CCN is to replace the existing IP-based networking, the current implementation is deployed as overlay architecture for the IP network.

The two main projects other than PARC CCNx which are contributing to CCNx are Content-oriented networking: a new experience for content transfer (CONNECT) [23] and named data networking (NDN) [24]. CONNECT aims to complement the existing work on CCN with original proposals in three technical areas: (i) Rethinking of traffic control (TCP cannot be used so queue management strategies require new name-based criteria in order to achieve fairness); (ii) Designing algorithms suitable for full-scale implementation of PARC's proposal of naming, routing and forwarding; (iii) Defining replication and caching strategies for CCN-based network deployment. PARC's CCN project provides a functional feasibility demonstration of content centric networking but feasibility alone is not sufficient to claim that this architecture is the way forward for the future Internet. The NDN project aims to use the architectural framework of CCN to solve real problems, particularly in application areas that are poorly served by today's Internet.

2.3 Modeling content centric networking

Our aim in this thesis is to compare CCN with the existing IP-based Internet architecture. The first step towards this goal is to develop a working model of CCN. The basic guidelines to implement a CCN architecture were provided by Jacobson et al. in [5]. Since then several models have been proposed to capture the main functionalities of CCN. In [25] an analytical model based on continuous time Markov-chains is derived to determine the percentage of time a particular piece of content is cached in any node along the path to the server. The analytical model is developed using strong assumptions of Poisson behaviour but subsequent analysis shows that many of these assumptions can be weakened and the model remains valid. In [26] Muscariello et al. focus on the performance evaluation of content centric networks. They propose an analytical model of bandwidth and storage sharing under limited resources. It is assumed that each user implements a receiver-driven flow control protocol so that the bandwidth is utilized in a fair and efficient way. The content storage is managed by the least-recently used (LRU) policy. Under these assumptions, a closed form expression for average content delivery time is presented. In [27], Carofiglio et al. instead focused on a set of different storage management techniques and analyzed the relation between transport and caching. They extend the CCNx prototype to support a generic function-based chunk replacement strategy, and to provide per-application storage management.

Tortelli et al. in [28] proposed a closed form expression that models CCN fairness in cache usage. Their model takes into account several network parameters e.g., network size (i.e., number of nodes), cache size, number of content items and the content popularity. In [29] Carofiglio et al. provided a detailed model which captures all the basic functionalities of CCN. They implemented their model for the single cache miss probability and miss rate under a two-level Markov modulated rate process of requests where the content popularity is assumed to follow a Zipf distribution. They also modeled a network of caches with and without request aggregation.

In summary, the work in [28, 29] captures most of the functionalities (e.g., routing, caching, fetching) of CCN and provides a complete and general model for future analysis. As for the other related work, it mainly focuses on identifying better caching strategies for CCN.

2.4 Green Networking

The research on energy-efficient communication networks, especially with an emphasis on the environmental concerns is referred to as "green networking". The most closely related work in terms of green networking for CCN is that of Lee and Guna et al. [30, 31]. In [30] Lee et al. propose an architecture for an energy-efficient CCN router and claim that energy benefits can be reaped by deploying these routers incrementally throughout the Internet. The energy analysis in [30] only considers the energy cost when clients are downloading content and all the network devices are operating at their peak power ratings. There is no consideration of the manufacturing energy cost nor of the extra energy required to power the additional cache memory of the CCN routers. Lee et al. suggest that energy savings can be obtained by switching the CCN routers on or off according to the network traffic requirements.

Guna et al. in [31] performed an energy efficiency analysis of content delivery architectures. They studied energy consumption in disseminating content through three different networking architectures: conventional CDN, centralized CDN with dynamic optical bypass (where the content is delivered by a server that is attached to a core router and transported across the core network mostly via transparent optical connections) and CCN. To perform the analysis they assumed that the content routers in CCN (those with caching capability) carry out the tasks of content servers of CDNs. Guna et al. used four different network topologies to conduct their analysis. The results reveal that CCNs are more energy efficient in delivering popular content while optical bypass CDNs are better in terms of energy for delivering infrequently requested content. The relative energy benefits of CCNs and conventional CDNs are not very clear in this study.

Both of the studies discussed above only consider the energy aspect of green networking. To the best of our knowledge, there is no study available in the literature which provides a detailed study of a complete green analysis (i.e., energy and carbon footprint analysis) of CCN. Other work that conducts a complete green analysis does not involve CCNs. Seetharam et al. in [32] provide a green analysis comparing an IP-network movie streaming application with a postal-service movie distribution system. They provide a detailed discussion about methods for estimating the operational and manufacturing energy consumption of network devices. Raghavan et al. in [7] provide an approximate analysis of the Internet's energy consumption. They calculate both the operating energy consumption and the emergy (embodied energy), the energy required to construct the Internet. This emergy is equivalent to the manufacturing energy consumption discussed by Seetharam et al. in [32].

2.5 Dynamic Adaptation

In previous work [8, 9], it has been suggested that the current Internet backbone networks are over-provisioned, primarily to allow for the significant daily variation in the traffic load and to provide redundant capacity to cater for problems such as link failures. However, the energy consumption of current network elements (e.g., routers and switches) appears to be largely constant.

Nedevschi et al. in [8] argue that energy can be saved by periodically switching network devices into sleep mode. They present the design and evaluation of simple power management algorithms that slow down link rates and put components to sleep to reduce the energy consumption of the network.

Antonakopoulos et al. in [9] present the notion of power-aware routing with rate-adaptive network devices. They demonstrate that the combination can save a significant fraction of network power, for a wide variety of network topologies, traffic loads, and device startup energy requirements. The rate-adaptivity involves a fairly simple model of network power use, where power consumption is attributed to links. A link can be turned off, in which case power use is zero. A fixed startup power is required to turn on each link then the connected device consumes energy at its full power rating.

Soteriou and Peh in [33] proposed a dynamic power management policy, where network links are switched off and on in a distributed fashion, depending on network utilization. They argue that a fully random approach to switch a link on or off is damaging to the network performance, because it can lead to intermittent network connectivity and result in a core network router being disconnected from the network. To handle these issues the authors devised a systematic approach based on a connectivity graph which is able to balance power and performance in a 2D mesh topology. In addition to that a proactive routing algorithm has been proposed which exploits the advance knowledge of link on or off state and guarantees packet delivery throughout the network.

Chiaraviglio et al. in [34] use a similar idea as used in [8, 9, 33]. They consider a wide-area network scenario and evaluate the possibility of turning off some routers and links under connectivity and Quality of Service (QoS) constraints by following different simple strategies. The authors use integer linear programming formulation to show that, defining the minimum set of routers and links that have to be used in order to support a given traffic demand is an NP-complete problem. They proposed some greedy heuristic algorithms to approximately solve the problem.

Chiaraviglio et al. in [35] applied the approach introduced in [34] to a realistic IP network topology and evaluated the amount of energy that can be potentially saved when nodes and links in the network are turned off during off-peak periods. The power consumption for the routers and links is estimated using realistic figures. A simple algorithm has been proposed to identify network equipment that must be powered on to guarantee packet delivery in the network.

Another important aspect in regard to saving energy consumption in the network is that the network traffic fluctuates considerably during day and night times. Energy can be saved by automatically changing link speeds according to the network traffic utilization. The notion of adaptive link rate (ALR) was first introduced by Chamara et al. in [36, 37]. As a continuation of their work Chamara et al. presented a system design of ALR in [38]. To correctly model the behaviour of a packet switched network a Markov model for a state-dependent service rate was developed. The service rate transition can only occur at the completion of a service period. The mean number of arrivals in the system is increased due to rate transition at service completion. They developed and evaluated a dual utilization-threshold policy and a time-out-threshold policy to eliminate oscillations and it was shown that these policies are successful in eliminating rate oscillations for smooth traffic. This work shows that significant energy savings within LAN switches and Ethernet-connected devices can be achieved by using ALR as it allows internal components of the network devices to operate at a lower clock rate when the Ethernet link data rate is reduced. The proposed ALR policy can support only two data rates.

Gupta and Singh in [39, 40] go one step further and argue that the network interfaces of LAN switches can be switched into standby mode during packet inter-arrival times. If the inter-arrival time is greater than a predetermined value then the LAN interface can be powered down altogether to achieve more significant energy savings. Gupta and Singh in [40] expand this idea by designing and evaluating a dynamic Ethernet link shutdown algorithm that utilizes current technology. The proposed algorithm uses the programmable capability of the sleep timer to dynamically change the operating power modes. This method allows the interface to either operate at full speed, referred to as the *'On'* state, or in the extreme low power mode, referred to as the *'Off'* state. The state change occurs with the exchange of beacons between the transmitting and the receiving side. The algorithm checks the feasibility of turning the link on or off and then computes the length of sleep.

These collective works aim to reduce the overall network power consumption by turning nodes and links off, and re-routing traffic to save energy consumption. However, these approaches can introduce disconnections and unwanted delays. We consider that this is not a feasible option for a streaming network application. Therefore we require solutions that do not involve turning nodes off.

Lombardo et al. in [10, 41] suggest that capacity and switching control mechanisms [37–40] should be used together. They name this capability "active capacity scaling" (ACS). The existing routers and switches do not offer support for capacity scaling but the technology to implement variable capacity electronic devices is readily available. Mobile devices and modern PCs in general are a good example. Moreover, novel devices capable of entering different energy states according to the system requirements are expected to be developed in the near future [42]. In [10, 41], a combined congestion control and rate-adaptation scheme for the Internet access node with ACS capability has been proposed. To accomplish this Lombardo et al. use a congestion control technique named active window management (AWM) [43, 44] coupled with an energy aware service rate tuner handling (EARTH) mechanism. EARTH invokes power management primitives at the hardware level to increase or decrease the rate at which a node is operating. This improves network performance by enabling rate adaptation. The authors have modified EARTH such that it can detect the minimum transmission resources an access router needs to allocate for providing best performance at the cost of minimum energy usage. The proposed mechanisms would work only if the access node is able to estimate both the minimum capacity available in the network path between destination and source, and the maximum capacity the user traffic requires. Such a router is named a G-Router (Green Router) in this study.

2.6 Summary

This chapter strives to summarize the core ideas of content centric networking, green network analysis, and rate adaptation. We presented a brief survey of the existing projects working to develop content based networking architectures and also argued why CCNx appears to be one of the most promising. We discussed the basic functionalities and modeling techniques for CCN. We explained the importance of green networking and reviewed some relevant literature in this field. In the last section we introduced the notion of dynamic rate adaptation for network routers and switches as this concept can be used to achieve substantial energy savings in a CCN architecture.

Chapter 3

Network and Energy Model

In the first section of this chapter we discuss the network topologies to be used in our analysis and explain the differences between them. The second section consists of a detailed energy model and provides explanation of the two types of energy consumptions considered for the analysis. The notion of rate adaptation is explained in the third section. Section 4 consists of a brief discussion on carbon footprint.

3.1 Network Model

Consider a network composed of multiple edge routers and a single server running a video streaming service. Assume we have two similar networks with the above characteristics, one designed according to a CCN architecture and the other according to an IP-based architecture. The routers involved in these two networks have one important difference: CCN routers have an additional cache that allows them to store data that passes through them. We perform a steady-state analysis on three different network topologies.

- General tree topology.
- Tree topology with Content Delivery Networks (CDNs).

- Partial mesh topology.

In the general tree topology with N levels (see figure 3.1) each level represents a tier in the network. Each node in this tree structure depicts a router on a given level. The server is on the N+1th level. In a tree based topology, the routers at a given level are not able to establish communication with one another. They can only interact with their child and parent routers at different levels.

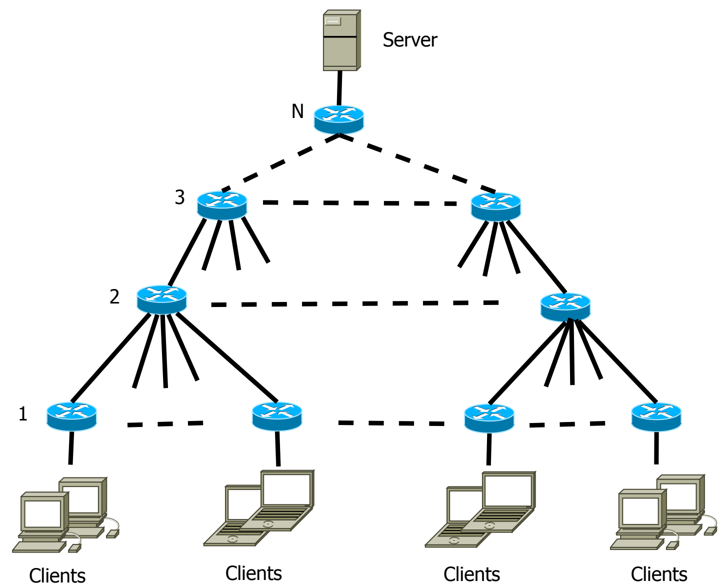

Fig. 3.1 General tree topology; the server at N+1th level. Routers at different levels can have more than one children but can only have a single parent.

As an extension to our analysis we consider a conventional CDN based tree topology for comparing CCN and IP-based networks in a video streaming scenario (see figure 3.2). A CDN is host-oriented setup consisting of an interconnected system of computers on the Internet that provides Web content rapidly to numerous users by duplicating the content on multiple servers and directing the content to users based on proximity. We use the same general tree topology as explained in figure 3.1 but replace some of the routers on some levels with CDN servers. This emulates a distributed system of servers.

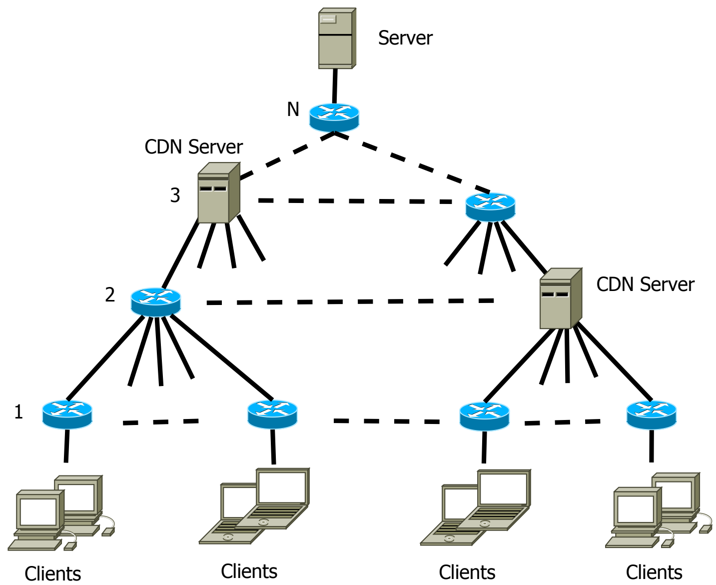

Fig. 3.2 Tree topology for CDNs; content is delivered to the clients through a host server with the help of centrally managed CDN servers.

The tree topologies described above are very useful for developing an understanding of the behaviour and performance of the different network architectures, but they do not accurately reflect real world network deployments. Although many networks do display the hierarchical characteristics of the tree topologies, there are frequently links between routers on the same level. We therefore also conduct an analysis for a partial mesh topology (see figure 3.3), where routers on the same level may communicate. To limit the complexity of the analysis to a manageable level, we only consider the case where a fraction of the routers on each level may communicate with other routers on that or an upper level in addition to their parent and children routers.

3.2 Energy Model

We consider two main aspects of energy consumption in a network: emergy and the operating energy consumption. It is common to use the power-rating when analysing energy

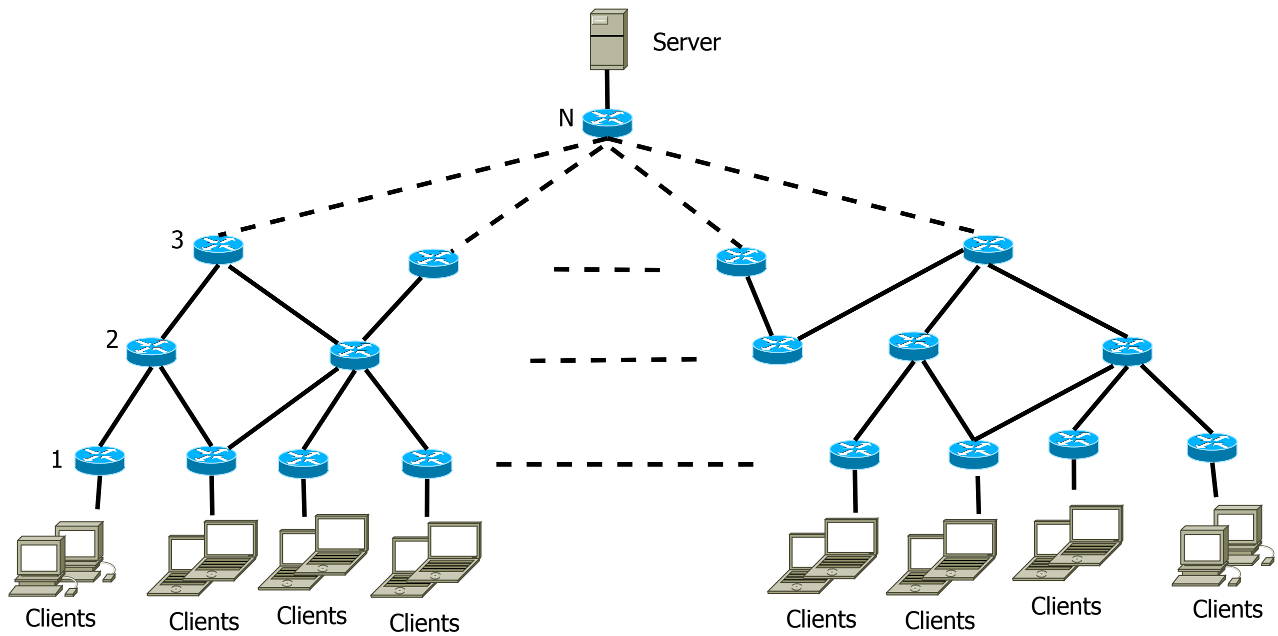

Fig. 3.3 Partial mesh topology; Routers on a given level may communicate with each other or with the routers on an upper level in addition to their parent and children nodes

consumption during the operation of a network device [7]. However, emergy which represents the manufacturing energy is expressed in joules so it ignores the time dimension. To be consistent in terms of units we define *embodied power* which is calculated by dividing emergy by the expected life cycle of a device. Table 3.1 summarizes the parameters employed in our analysis. An important component of our analysis is the impact of link rate on power consumption. In our study we assume that the power consumption of a router varies in proportion to the link rate [8,9]. We define the power reduction factor ϕ_k, $k = 1, 2, ..K$ as:

$$\phi_k = \frac{R_h(k)}{R_{IP}(k)} \tag{3.1}$$

where $R_h(k)$ and $R_{IP}(k)$ are defined in table 3.1.

Parameter	Description
N	Total number of levels in a general tree topology
K	Total number of routers
$K(i)$	Total number of routers at ith level
M_{IP}	IP router embodied power consumption
M_{CCN}	CCN router embodied power consumption
S_M	Server embodied power consumption
S_O	Server operating power consumption
S_{M-CDN}	CDN server embodied power consumption
S_{O-CDN}	CDN server operating power consumption
n	Number of CDN servers in the network
C_S	Power required to operate the server storage device
C_{CDN}	Power required to operate the CDN-server storage device
C_M	Power required to operate the router cache
O_{IP}	IP router operating power consumption
O_{CCN}	CCN router operating power consumption
$R_{IP}(k)$	Link rate at IP router k
$R_h(k)$	Link rate threshold at CCN router k

Table 3.1 Parameters used in the energy model

3.2.1 Power consumption

Although embodied power consumption is a one time cost, it is important to include this value in our analysis in order to have a fair comparison between the energy consumptions of IP and CCN-based networks for VoD streaming applications. Equations (3.2) and (3.3) describe the embodied power cost of IP and CCN networks respectively.

$$P_{IP}^{M} = KM_{IP} + S_M + nS_{M-CDN} \tag{3.2}$$

$$P_{CCN}^{M} = KM_{CCN} + S_M + nS_{M-CDN} \tag{3.3}$$

We assume that the manufacturing energy needed to produce a server is not related to the network architecture.

The power consumption relations for both IP and CCN-based networks are identified as follows:

$$P^{O}_{IP} = C_S + nC_{CDN} + S_O + nS_{O-CDN} + KO_{IP} \tag{3.4}$$

$$P^{O}_{CCN} = C_S + nC_{CDN} + S_O + KC_M + nS_{O-CDN} + O_{CCN} \sum_{k=1}^{K} \phi_k \tag{3.5}$$

Equation (3.5) introduces the use of the power reduction factor ϕ_k (see equation (3.1)). A discussion on ϕ_k is presented in subsection 3.3. To estimate the power consumption on a particular level, K can be replaced with $K(i)$ where i is 1,2.., N and the terms related to the server become zero.

The total power consumption is the sum of embodied and operational power consumption of the network. Equations (3.6) and (3.7) describe the total power consumption of the IP and CCN networks respectively. It should be noted here that these are the general equations which can be applied to any network topology. For example, in the case of a simple IP or the case of the CCN-based deployment the quantities related to CDN-server become zero.

$$P_{IP} = P^{M}_{IP} + P^{O}_{IP} \tag{3.6}$$

$$P_{CCN} = P^{M}_{CCN} + P^{O}_{CCN} \tag{3.7}$$

3.3 Rate adaptation (RA)

In a CCN-based network deployment, routers present on different tiers of the network are able to cache the content. This feature of CCN enables routers at higher levels to operate with smaller link rates than routers at lower levels. Figure 3.4 illustrates the impact on the fetching of content. The figure depicts a CCN-based network with a single server and multiple clients. Client 1 requests a particular content from the server, and the server then responds to the client by sending the requested content. This content is also cached at cache-capable intermediate routers. If client 2 requests the same content from the server, it is able to get a response from router E instead of going all the way to server. As a result fewer requests reach the upper levels and server, hence decreasing the traffic on links at upper levels.

We may exploit this feature of CCN by using routers which are capable of changing their operating rate according to the traffic rate on a given link. Many different techniques for rate adaptability have already been proposed. Section 2.5 provides a brief overview. We have mentioned earlier that the rate adaptation strategies of [8, 9, 33, 39, 40] are not well-suited to the VoD network application we study, which has strict delay requirements. The idea of "active capacity scaling (ACS)" introduced by Lombardo et al. in [10, 41] appears to be promising in achieving impressive rate adaptation, where routers can adapt their rate according to the load demand with minimum delays. But the routers and switches which populate the existing Internet are not capable of implementing ACS. Chamara et al. [37, 38] provide a more practical approach for rate adaptation in Ethernet networks. According to their proposal each switch monitors the traffic flowing through it. This information is then used to identify traffic thresholds. When the traffic exceeds or falls below these thresholds, the link rate is adjusted to a higher or lower value accordingly, which in turn impacts the consumed power of the device. Adjusting the link rates using traffic thresholds is accomplished using two policies: i) ALR utilization-threshold policy, ii) ALR time-out-threshold policy. In the utilization-threshold policy the decision to transition between different link rates is made by counting the number of bytes transmitted during a time period, which is referred to as utilization monitoring. Counting the number of bytes transmitted within a given time requires additional accumulators and registers that could increase the complexity of network equipment. To handle this issue a heuristic approach has been followed in the time-out-threshold policy, where a link is maintained at a high data rate for a predetermined period of time following a switch from a low to a high link data rate and vice-versa. This is achieved using timers with heuristic settings. These policies, if adopted, can provide substantial savings in energy consumption for different network deployments.

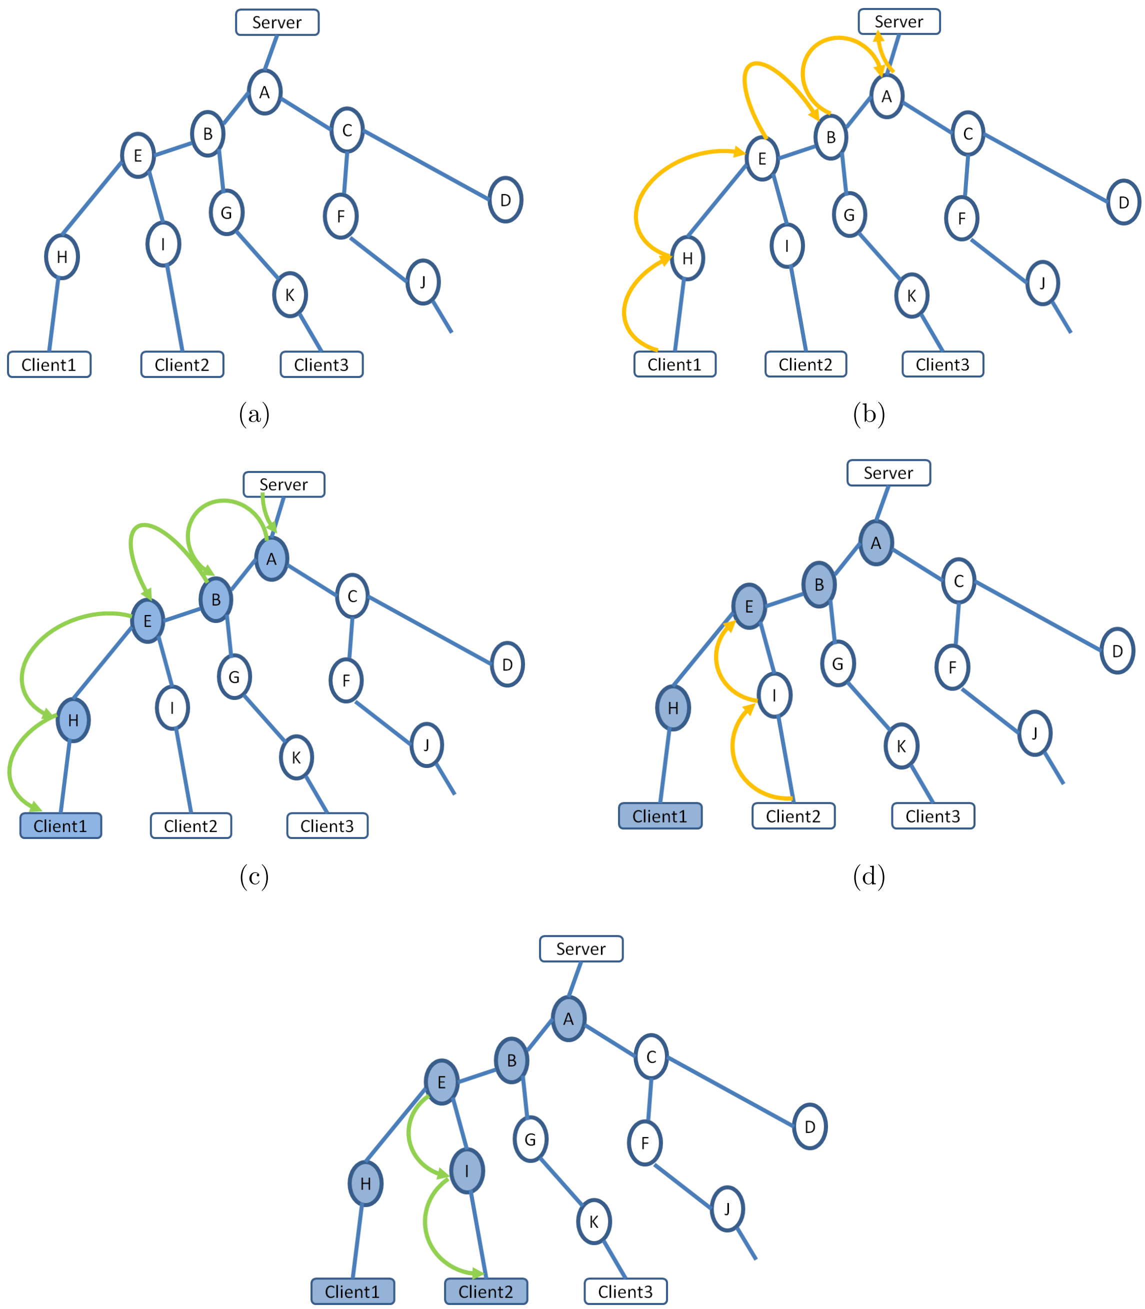

(a) (b) (c) (d) (e)

Fig. 3.4 Fetching of content in a CCN-based network deployment.(a) Network with a single server and cache capable routers with empty caches.(b) Client 1 requests for a content; request goes all the way up to the server. (c) Server responds with the required content which is cached in the intermediate routers before reaching to the client. (d) Client 2 requests the same content.(e) Client 2 gets the response from the nearest router which already had cached the content.

Our study therefore considers three scenarios for rate adaptation. In the first scenario, the IP and CCN network links are provisioned with the same capacities and the network devices operate at full power, i.e., there is no rate adaptation ($\phi_k = 1$). We label it as "Without RA". The second scenario is labeled as "Practical RA". Here we try to incorporate the strategy similar to [37, 38]; the IP and CCN networks are provisioned with links rates that provide sufficient capacity to support the generated traffic demand. We define the thresholds based on the generated traffic demand and provision the links according to the actual real world implemented values (1Mbps, 10 Mbps, 100 Mbps, etc.) of the available Ethernet.

The third scenario is inspired by the work in [10, 41] and provides a close-to-optimal approach, with link rates being adjusted seamlessly in response to changes in the traffic demand. We refer to this as "with RA", and we assume that the link rates are able to follow the traffic load instantaneously. In this instance ϕ is defined as the ratio of the exact values of link rates calculated using our simulators. As already introduced in equation (3.1), the rate factor ϕ_k allows us to quantify the relation between link rate and power consumption.

3.4 Carbon Footprint

We have discussed only the energy aspect of our green analysis so far. To complete the analysis we need to estimate the carbon footprint for the devices used in the different topologies we have mentioned in the network model. Seetharam et al. in [32] provide a carbon footprint analysis for a DVD shipping service compared to an IP video streaming service. We take motivation from their work and estimate the carbon footprint values for our analysis. Carbon footprint of CCN and IP-based networks can be determined by calculating the amount of carbon dioxide emitted due to manufacturing and operation of network equipment used in both kind of networks. The carbon footprint is the product of the carbon dioxide emission coefficient and the energy consumed. The mean value of the carbon coefficient for electricity is 0.585kg/kWh as given in [45]. We use this value in our calculations. The method of electricity generation assumed in [45] while calculating the carbon coefficient is hydro-electricity generation. In situations where electricity generation is done using nuclear plants in addition to hydro-electric, the carbon footprint of electricity generation is much lower [46]. We note that the reductions in carbon loading would be experienced by both networks so this does not affect our analysis. Recycling of IT equipment also helps to recover the carbon cost incurred due to manufacturing. In [47], the authors suggest that the carbon cost recovered due to recycling for a laptop computer is approximately 15%. We apply the same value to all network equipment (servers and routers) due to unavailability of data.

3.5 Summary

This chapter summarizes the basic theme of our analysis. We have presented the details of three network topologies (i.e., tree, Content Distribution Network tree and partial mesh) which will be used in our later analysis. The chapter also sheds light on the concepts of emergy (manufacturing energy) and operating energy of network devices. We estimated the total power consumption with the help of some simple equations. In the later part of this chapter we argued that rate adaptation can prove to be beneficial in terms of reducing energy consumption, if used in a CCN. The discussion on how to estimate the carbon footprint for the network devices concluded the chapter.

Chapter 4

Performance Evaluation

This chapter provides a performance evaluation of CCN compared to IP-based networks on the basis of power consumption and carbon footprint. We consider a Video on Demand (VoD) scenario where a single server is providing services to the clients. This scenario is simulated using all the topologies described in section 3.1.

In a VoD scenario, various kinds of traffic demand can occur during different instances of time. For example, the demand for the last five video lectures for "Computer Networks" before the examination would generate a uniform traffic demand on the network as the students would be equally interested in all of them. On the other hand, the demand of a famous music video would be much higher than any other content available in the content library of a VoD server. In our analysis the network traffic is generated using two different distributions—Uniform to capture the Uniformly distributed traffic demand and Zipf to model demand for content of varying popularity (see figure 4.1). We assume that our network is deployed using Verivue VoD and CDN servers [48] as the main and distributed servers respectively. The storage capacity assumed for the VoD server is 10 TB and that for the CDN server is 3 TB. The routers which populate the network are M10i Juniper routers [49]. An additional cache is added to the routers when simulating a CCN-based network.

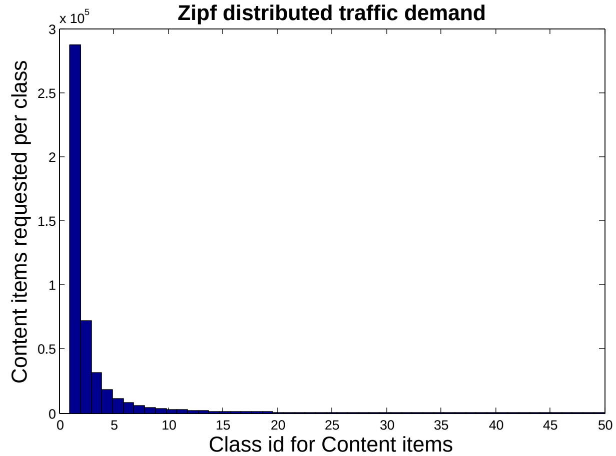

(a)

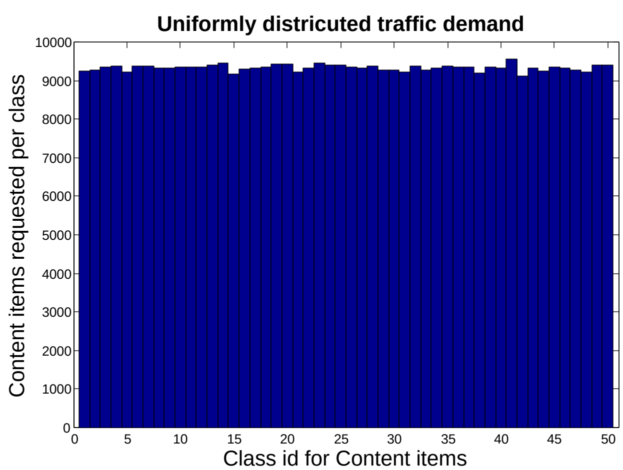

(b)

Fig. 4.1 Traffic patterns using Zipf and Uniform distributions. Here content items are distributed into 50 classes. In the Uniform distribution, the content items are requested from any class with equal probability. For the case of Zipf distribution, the content items belonging to the most popular class (i.e., with lowest class id) get the highest number of requests

In the first four sections of this chapter, we calculate the approximate embodied powers, operating powers and carbon-footprints of the network devices used in our analysis. The fifth section describes the details of our simulation model. We present the results of our performance analysis comparing CCN with IP-based networks in Section Six. Section Seven provides the summary of this chapter .

4.1 Embodied power consumption

To the best of our knowledge, there is no study available in the literature which provides an exact analysis of the energy or power consumed in manufacturing storage devices, servers and routers. Therefore we estimate these costs from data given in [7,32,50,51] and make use of the available data-sheets for different devices which we are considering in our analysis.

In the year 2000 the emergy of a disk drive of 30 GB was estimated to be 2926 MJ [50]. According to Kryder's law [52], the storage capacity of storage devices doubles approximately every 18 months for the same cost (see figure 4.2). Here we are assuming that the cost and manufacturing energy are directly proportional, so by applying Kryder's law we can calculate the amount of storage which can be manufactured in 2011 using 2926 MJ of energy as follows:

$$2^{\left(\frac{\text{number of months since year 2000}}{\text{18 months}}\right)} \times \text{Disk capacity in year 2000}$$

Inserting the values for the "number of months since year 2000" and "Disk capacity in year 2000" in the above expression we get:

$$DiskStorage \approx 2^{\left(\frac{12\times 11}{18}\right)} \times 30\text{GB} \approx 5\text{TB} \tag{4.1}$$

The emergy of a 10 TB storage device at the server can then be approximated as $2926 \times 2 \approx 5850$ MJ. The size of the CCN router cache memories should depend on the scale of the network. We consider seven candidate sizes (96 GB—1.5 TB) to assess the difference the choice makes on the overall power consumption of the CCN network. The

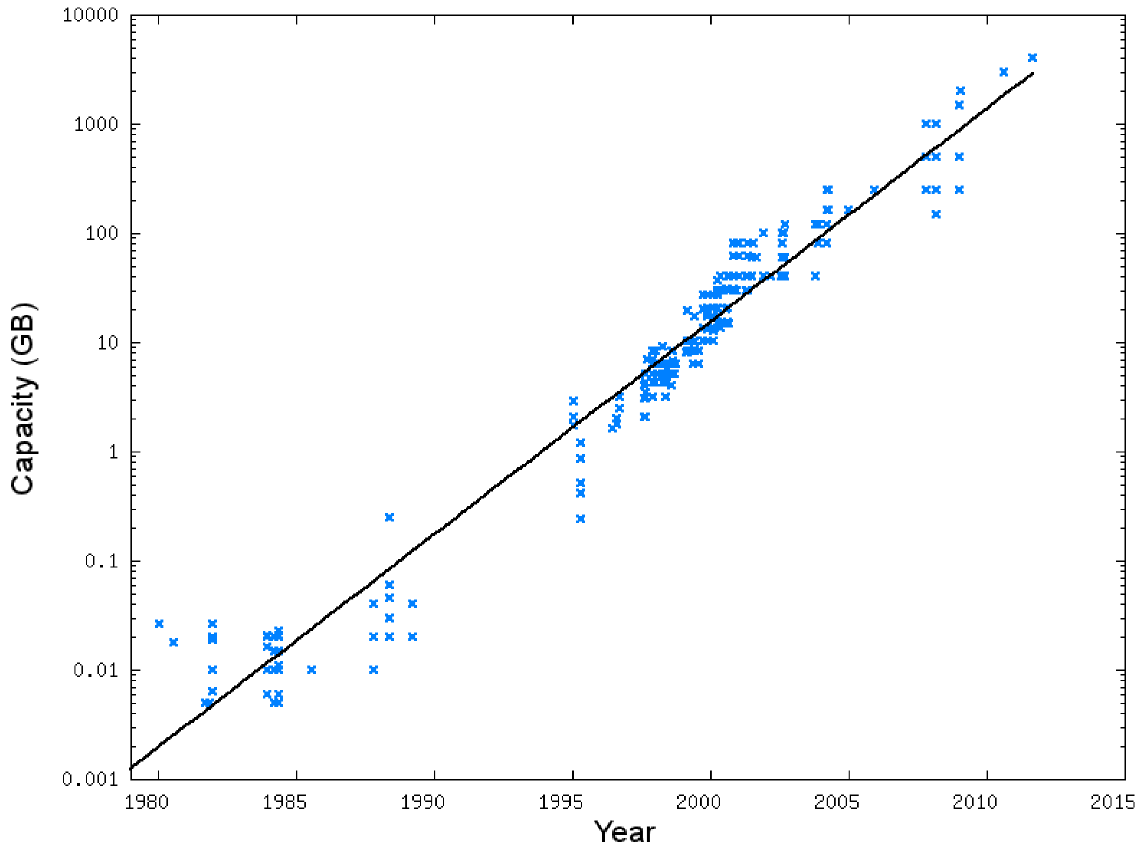

Fig. 4.2 Increase in hard disk capacity between 1980 and 2011, following Kryder's law. Graph is reproduced from the data available at [1]. The graph is logarithmic, so the fitted line corresponds to exponential growth.

emergy values for a subset of them are provided here. Hence the emergies for 1 TB, 512, 256 and 128 GB caches are 585, 290, 150 and 75 MJ respectively. We assume that the life cycle for all the devices we are using in our analysis is 3 years [51]. Hence the embodied power for these different size caches can be calculated by dividing the emergy over the life cycle.

The server emergy estimated by Seetharam et al. in [32] using the study of Williams in [50] is 550 MJ. We use the same estimate for our analysis. We are considering a simplistic dedicated video streaming scenario so we can assume that the network is populated with the same type of routers throughout (Juniper M10i routers in this case). Seetharam et al.

estimated the emergy of a router in [32] by scaling the weight of the router relative to the weight of the PC. The weight of a desktop PC varies between 13 to 35lbs [53]. For this analysis we assume that it is approximately 25lbs. The weight of a M10i router is 79lbs [49] and therefore the emergy of an edge router becomes

$$\frac{79}{25} \times 550 \times 10^6 \approx 1200\text{MJ}, \tag{4.2}$$

which is equal to 13 J/s (embodied power). The total embodied power for server and network routers is summarized in table 4.1. We are considering the same M10i routers for both CCN and IP-based networks. The only difference is that the routers in CCN have an additional cache memory.

Device	**Embodied power (J/s)**
Server (S_M)	68
CDN server (S_{M-CDN})	25
IP-router (M_{IP})	13
CCN-router (M_{CCN}) 1 TB, 512, 256, 128 GB	$20, 17, 15, 14$

Table 4.1 Embodied power consumption for the network devices used in this analysis

The difference between the values estimated in table 4.1 for server and router's embodied power is of the order of 4 to 6 times. These values are consistent with the values estimated by Raghavan et al. in [7].

4.2 Carbon footprint(CFP) of network devices during manufacture

The energy required to manufacture a 10TB memory is 5850 MJ (calculated in the previous section). The carbon footprint of this memory device can be calculated by using the definition provided in section 3.4, $0.585kg/3600kJ \times 5850MJ$. The final value is calculated after subtracting the 15% carbon cost reduction due to recycling. The same method is applied to estimate the carbon footprint for all sizes of cache memory used in the network. The carbon footprint for a network router's or server's manufacture is estimated from [45]. Carbon footprint values for the network devices used in our analysis are provided in table 4.2. For the case of the VoD-server and CCN-routers these values are estimated by taking into account the effect of storage and cache memories respectively.

Device	**Carbon footprint (Kg)**
Server	900
CDN-server	350
IP-router	175
CCN-router 1 TB, 512, 256, 128 GB	$240, 200, 185, 180$

Table 4.2 Carbon footprint for devices due to manufacture

In the next section we estimate the power consumed in transmitting video content (e.g., movie) through the IP and CCN-based network.

4.3 Operating power consumption

We assume that the library size for the server is $M = 10,000$ videos. Each video is compressed using DiVX codec and the average size is 700MB. The server we are using has a storage capacity of 10 TB consisting of an array of ten 1 TB devices (e.g., RAID configuration). The maximum downstream rate the server can provide is bottlenecked by the read capacity of memory storage (1385 Mbps, or approximately 175 MB/s, in this case). The power consumed by 10 TB storage is 480 J/s [54] during streaming operation. The operating power consumption for 1 TB, 512, 256 and 128 GB cache memory is 6.8, 4.6, 4.2 and 3.4 J/s respectively [55].

The total energy spent by a server's chip-set during video streaming can be calculated by following the model adopted in [32]. For a typical server the operating power of a multimedia streaming application is 251 J/s. The chip-set of edge routers we are using for this analysis (i.e., M10i) operate at 116 J/s. For the IP-based network the router operating power is the only power consumed in operating the M10i routers but for CCN the router operating power also includes power consumed by the cache memory during its operation (see table 4.3). We also consider the power required to keep alive the router cache and

Device	**Operating Power (J/s)**
Server (O_S)	730
CDN-server (S_{O-CDN})	260
IP-router (O_{IP})	116
CCN-router (O_{CCN}) 1 TB, 512, 256, 128 GB	$123, 121, 120, 118$

Table 4.3 Operating Power of the network devices used in this analysis

server's memory storage for the duration of the network deployment. These values can be extracted from the available data-sheets of the memory devices [54, 55] and are presented in table 4.4.

Power required to power up memory	Value (J/s)
Router cache memory (C_M) 1 TB, 512, 256, 128 GB	0.053
Server memory (C_S)	20
CDN-server memory (C_{CDN})	5

Table 4.4 Power required to power up the storage and cache memory

4.4 Carbon footprint of network devices during operation

The carbon footprint of network devices during operation is estimated by taking the product of the operating energy of all the devices with the carbon coefficient defined in section 3.4, over the duration of the network deployment. Table 4.5 presents the estimated values of carbon footprint due to the operation of the network devices.

Device	Carbon Footprint (Kg)
Server	9525
CDN-server	3510
IP-router	1550
CCN-router 1 TB, 512, 256, 128 GB	$1850, 1700, 1650, 1600$
CCN-router-cache 1 TB, 512, 256, 128 GB	700
Server-cache	260

Table 4.5 Carbon footprint for devices due to operation

4.5 Simulation Model

We implemented the basic relevant functionalities (caching, routing, fetching) of content centric networking in MATLAB and C++. We developed two simulators: *CCNSim* and *CCNSim++*. The description of both simulators is presented in the following subsections.

4.5.1 CCNSim

CCNSim, implemented in MATLAB, can generate a general tree topology with N levels (see figure 3.1), and extends the binary tree model developed in [29]. The simulator works as follows: when a request is generated, it is sent to the node in the first level. The node checks if the data is in its cache. If it is, the request is counted as being successfully sent to the user. If it is not, the request is sent to the next level. This process is repeated until the request is successfully served either by one of the nodes or because it reaches the server which can ultimately serve the request.

The simulator provides three main services: request generation, router caching and routing. The content items of various sizes are segmented into equal-size chunks. The request arrival process is modeled through a Markov Modulated Rate Process (MMRP) that captures the behaviour of the system at both the content level and the chunk level. Chunk level simulations allow us to have a better understanding of the temporal dynamics of the system. The cache process in each router employs a Least Recently Used (LRU) replacement policy. The cache maintains a table with a list of chunks which have already been cached. If a chunk arrives which is not in the table, we update the table by adding the recent chunk and discarding the least recently used entry. The system assumes that there is a fixed routing path between children and their parent (general tree structure). We consider a set of M different contents equally partitioned into D classes of popularity. In our simulations, we assume a Zipf and Uniform popularity distribution (see figure 4.1 for pictorial representation), given for class d as:

$$q_d = \begin{cases} c/d^{\gamma} & \text{if Zipf} \\ 1/D & \text{if Uniform} \end{cases}$$

Here γ is the value of exponent characterizing the Zipf distribution and it is greater than 1. c is the normalization constant. Each content request in class d is generated according to a Poisson process with rate $\lambda_d = \lambda q_d$, where λ is the total content request rate at the first node. This generation coincides with the request of the first chunk, and it is important to note that contents are uniformly chosen among the $m = M/D$ different content items in class d. The inter-arrival times between consecutive chunks are deterministic, and the number of chunks in a given content item is assumed to have a geometric distribution with average size σ. The superposition of different content requests defines the MMRP process. The underlying Markov chain is presented in [56]. Each node in the network has a cache of size X contents.

4.5.2 CCNSim++

I developed *CCNSim++* which is an improved version of *CCNSim*. Although the basic theme behind *CCNSim++* is the same as that of *CCNSim* there are some important improvements.

- The use of a low level language, i.e., C++, makes *CCNSim++* computationally more efficient than *CCNSim*.

- We can implement any general network topology (e.g., mesh, ring) in *CCNSim++* rather than just trees.
- The cache process in *CCNSim++* employs Least Frequently Used (LFU) replacement policy instead of LRU. In a VoD environment (e.g., Netflix [57]) it would be important to cache the content which remains popular irrespective of time (e.g., 250 top rated movies in IMDB [58]). LFU can better capture this notion of popularity of content.
- *CCNSim++* can generate a network topology having network routers with different cache sizes at different tiers of the network. This feature was not available in *CCNSim*.

The source code and executables for both the simulators are available on the project page of the McGill Computer Networks Lab [59].

4.6 Simulation Results

In our simulations, we have chosen to model a library of $10,000$ videos of size 700 MB on average, organized in $D = 50$ classes of popularity with parameter $\gamma = 2$, each one with $m = 200$ content items. Chunks are of size 1 MB and the average number of chunks is $\sigma = 700$. The system generates $\lambda = 1$ content request per second at each entry node. This is considered to be the peak demand. We analyze the network under the simplifying assumption that it is operating at peak load for the duration of its lifetime. The transmission rate for each content item is 1 Mbps and we consider that it is only possible for the system to request one chunk at a time (transmission window size $W = 1$). The cache size of routers on a given level is uniform but may vary along different levels per simulation requirements. The values of interest for the cache sizes of routers in our analysis ranges from 96 GB—1.5 TB. These settings provide a reasonably good platform to simulate a small scale VoD service (e.g., a scaled down version of Netflix [57] which supports a city of a few million users).

All the parameters defined in the paragraph above except for the cache size remain constant throughout our analysis as they have relatively little impact on the power ratio. The equations 3.6 and 3.7 presented in section 3.2 are used to define the power ratio (P_{IP}/P_{CCN}) which is a measure that compares CCN and IP-based networks in terms of power consumption.

Figure 4.3(a) and 4.3(b) present the variation in the overall power ratio corresponding to the change in the values of parameters γ and λ. Figure 4.3(a) illustrates that increasing

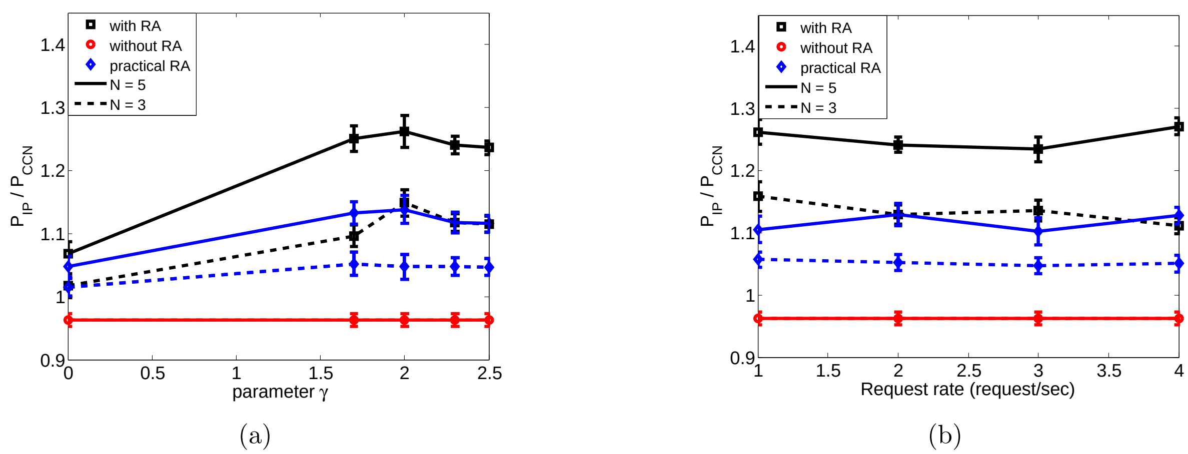

Fig. 4.3 Sensitivity of power ratio with respect to different network parameters; binary tree topology. The power ratio values in each of these figures are estimated by running 50 iterations of the simulation using randomly generated traffic load.(a) Variation with respect to the popularity parameter of the Zipf distribution, γ with $X = 128$ GB, $\lambda = 1$ request/sec. (b) Variation with respect to the content request rate λ with $X = 128$ GB, $\gamma = 2$.

the value of popularity distribution parameter γ impacts the power ratio. We observe a gain in power ratio even in the case when $\gamma = 0$ (traffic demand is uniformly distributed). Increasing the value of gamma translates into increasing demand for popular content. We get the optimal value of power ratio at $\gamma = 2$ for our particular implementation. Further increase in the value of gamma means that probability of demanding a particular class of content is getting higher. This in turn results in extensive caching of that particular content class. The client which requests a content with lower popularity has to travel more hops to get the content which degrades the performance gain (power ratio). In figure 4.3(b) we present the impact of request rates on power ratio. Assuming that the network is capable of responding fast enough to the request demand, having a higher request rate only results in quicker caching of content within the network. Hence increasing request rates does not produce a major effect on our steady state analysis of power ratio.

We conduct a simulation using *CCNSim* for a period of 1 hour to assess the traffic load on each link in the network. The time duration of 1 hour in this instance is enough to clearly see the difference in link rates at different levels of the network.

Figure 4.4 shows a comparison of CCN and IP link utilization. The CCN link rates are less than IP link rates due to the caching process at each router. Routers at lower levels store popular content and as a consequence the link rates at higher levels are low (see figure 4.4(c)). One important aspect to consider in a CCN network is the average number of hops that a request has to travel in order to reach content (video in our case). Fig. 4.5 shows the average number of hops as a function of the popularity class for a binary tree network with 3, 5 and 7 levels. The request generation process for this case follows the Zipf distribution. It is clear that content with the lowest class ID (very popular content) has a smaller average number of hops. The most popular content has an average of 1.5 hops, indicating that request of such content is served mainly by nodes at the first level. For the least popular content, the client has to reach the server to satisfy its request.

To evaluate the energy performance, we introduced the notion of rate adaptation (see section 3.3). In the light of the results presented in figures 4.4, 4.4(c) and 4.5 CCN unlike IP, can greatly benefit from rate adaptation. We have defined three scenarios for our evaluation. The first case labeled "without RA" involves no change in link rates, irrespective of the traffic load. Hence the power reduction factor ϕ defined in equation 3.1 remains unity. In the second scenario, links are provisioned according to the actual implemented values (1Mbps, 10Mbps, 10Gbps etc.) of the available Ethernet so ϕ becomes a ratio of realistic

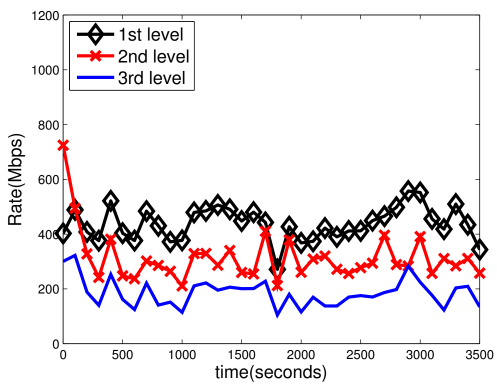

(a)

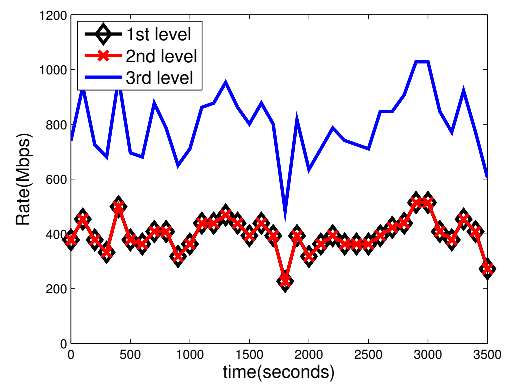

(b)

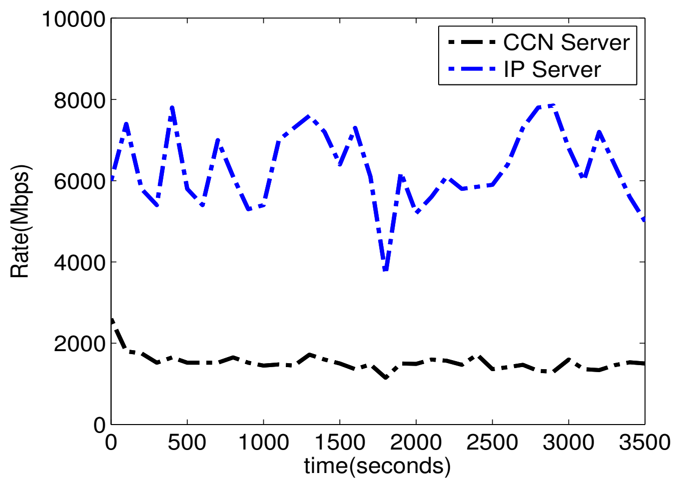

(c)

Fig. 4.4 Instantaneous link rates (a) CCN with $N = 5$, $X = 128$ GB, $\gamma = 2$, $\lambda = 1$ request/sec. (b) IP with $N = 5$, $\gamma = 2$, $\lambda = 1$ request/sec. (c) Comparison of CCN and IP servers levels in term of link rates.

link values. We label this as “practical RA”. The final scenario “with RA” corresponds to seamless rate adaptation in which we assume that the link rates are able to follow the traffic load instantaneously. Here ϕ becomes the ratio of the exact values of link rates calculated using our simulators.

CCNSim and *CCNSim++* produce similar results for the general tree topology but *CCNSim* cannot operate with complex topologies. Therefore, we use only *CCNSim++* for

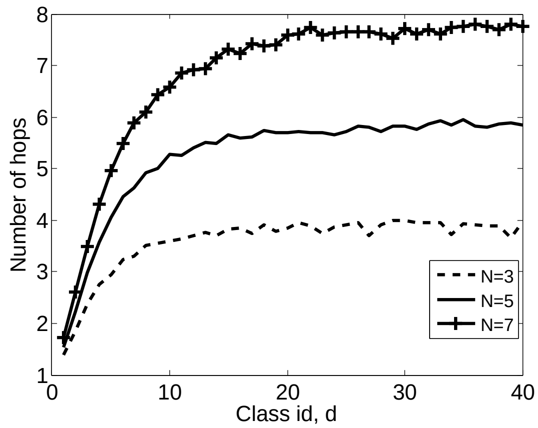

Fig. 4.5 CCN average number of hops per content request as a function of the popularity distribution class ID ($d = 1$ is the most popular content). $X = 128$ GB, $\gamma = 2$, $\lambda = 1$ request/sec

further analysis.

4.6.1 Green Analysis

Power Analysis for the Tree Topology

The power analysis is performed by comparing CCN and IP-based networks in terms of power ratio ($\frac{P_{IP}}{P_{CCN}}$). The power ratio is calculated by using the power consumption values for different network devices presented in sections 4.1 and 4.3. To cater for the imprecisions in our estimated values of power consumption for the network devices we introduce some error in our estimates. We consider a $\pm 10\%$ error in each component value and estimate the upper and lower bounds as: $\frac{P_{IP}-10\%P_{IP}}{P_{CCN}+10\%P_{CCN}}$ and $\frac{P_{IP}+10\%P_{IP}}{P_{CCN}-10\%P_{CCN}}$ respectively. While estimating the error bounds we are dealing with the power ratio of the network at each level, hence the power consumption values related to the server are set to zero (see section 3.1). The introduction of $\pm 10\%$ error in individual component values result in the overall error bound of up to $\pm 25\%$.

An important aspect in the estimation of the error bounds is to analyze error sensitivity of the individual components present in equations 3.6, 3.7 and their impact on the overall error bound. To perform the sensitivity analysis we apply $\pm 10\%$ *and* $\pm 20\%$ error to every parameter in the power ratio ($\frac{P_{IP}}{P_{CCN}}$), taking one parameter at a time and keeping others error free.

Table 4.6 presents the use case where we consider a 5 tiered binary tree topology (with $X = 1$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.) and observe the effect of introducing errors in different components of the power ratio at level 5. The power ratio without the error is 6.4459 in this instance. In table 4.6 we can see that the error in operating power

Power Ratio with different parameters in error	**Power Ratio**	**Error in Power Ratio**
Power Ratio with O_{IP} in $+10\%$ error	7.0255	0.5796
Power Ratio with O_{IP} in -10% error	5.8663	-0.5796
Power Ratio with O_{IP} in $+20\%$ error	7.6052	1.1593
Power Ratio with O_{IP} in -20% error	5.2866	-1.1593
Power Ratio with M_{IP} in $+10\%$ error	6.5109	0.0650
Power Ratio with M_{IP} in -10% error	6.3809	-0.0650
Power Ratio with M_{IP} in $+20\%$ error	6.5758	0.1299
Power Ratio with M_{IP} in -20% error	6.3160	-0.1299
Power Ratio with O_{CCN} in $+10\%$ error	6.6130	0.1638
Power Ratio with O_{CCN} in -10% error	6.2900	-0.1559
Power Ratio with O_{CCN} in $+20\%$ error	6.9221	0.4762
Power Ratio with O_{CCN} in -20% error	6.1415	-0.3044
Power Ratio with M_{CCN} in $+10\%$ error	6.9682	0.5223
Power Ratio with M_{CCN} in -10% error	5.9965	-0.4494
Power Ratio with M_{CCN} in $+20\%$ error	7.5826	1.1367
Power Ratio with M_{CCN} in -20% error	5.6056	-0.8403

Table 4.6 Error sensitivity analysis

consumption for IP routers i.e., O_{IP} is most influential in affecting the power ratio. This is a reasonable outcome as O_{IP} has a higher value compared to M_{IP} and M_{CCN}. Although O_{CCN} has a larger value than O_{IP}, the rate adaptation factor ϕ present in equation 3.5 lessens its effect.

Figure 4.6 presents the power ratio for the IP and CCN-based network deployment at different levels under the three rate-adaptation scenarios. It confirms our intuition that at higher levels in a tree topology, CCN with RA is energy efficient as IP is unable to take advantage from RA due to the absence of cache capable routers. Figure 4.6(a) shows the power ratio for the case where each router has two children. The request generation follows the Zipf distribution. In the absence of RA P_{IP}/P_{CCN} is very close to 1, indicating that the energy cost associated with manufacturing and powering the additional cache memory is relatively small. The practical RA implementation, based on provisioning according to demand, ensures reduced power consumption for CCN at higher levels in the network but it has poorer behavior at lower levels. In figure 4.6(b) we report the results for Uniformly distributed traffic demand while all other parameters remain the same. It is evident from the figures 4.6(a), 4.6(b) that although CCN outperforms IP when RA is used, the performance gain (power ratio) is lower than the case where traffic demand was Zipf distributed. This is due to the fact that all the content items are equally popular (Uniform demand) which prevents CCN from taking advantage of the frequent requests for popular content items.

Figures 4.6(c) and 4.6(d) present similar trends as shown in figures 4.6(a) and 4.6(b). One noticeable difference in this case occurs when the number of nodes in the network is increased from 2 children per node to 3 children (see figure 4.6(c)), the power ratio at different levels of the network also increase compared to the power ratio in figure 4.6(a). The increase in the number of routers on a given level results in an increase in the caching capability of the network which makes it easier for the popular content to cache. This phenomenon further decreases the load on the network and consequently the power ratio for CCN improves.

Another interesting observation that can be made here is about the power ratio in the case of uniform traffic demand. The power ratio at the lower levels in figure 4.6(d) is degraded compared to the ratio in figure 4.6(b). This stems from the fact that in the presence of uniform traffic demand, the caching property of a CCN network cannot be utilized properly and the rate adaptation becomes of little use. Hence the increased number of routers in the network increases the overall power requirements and the power

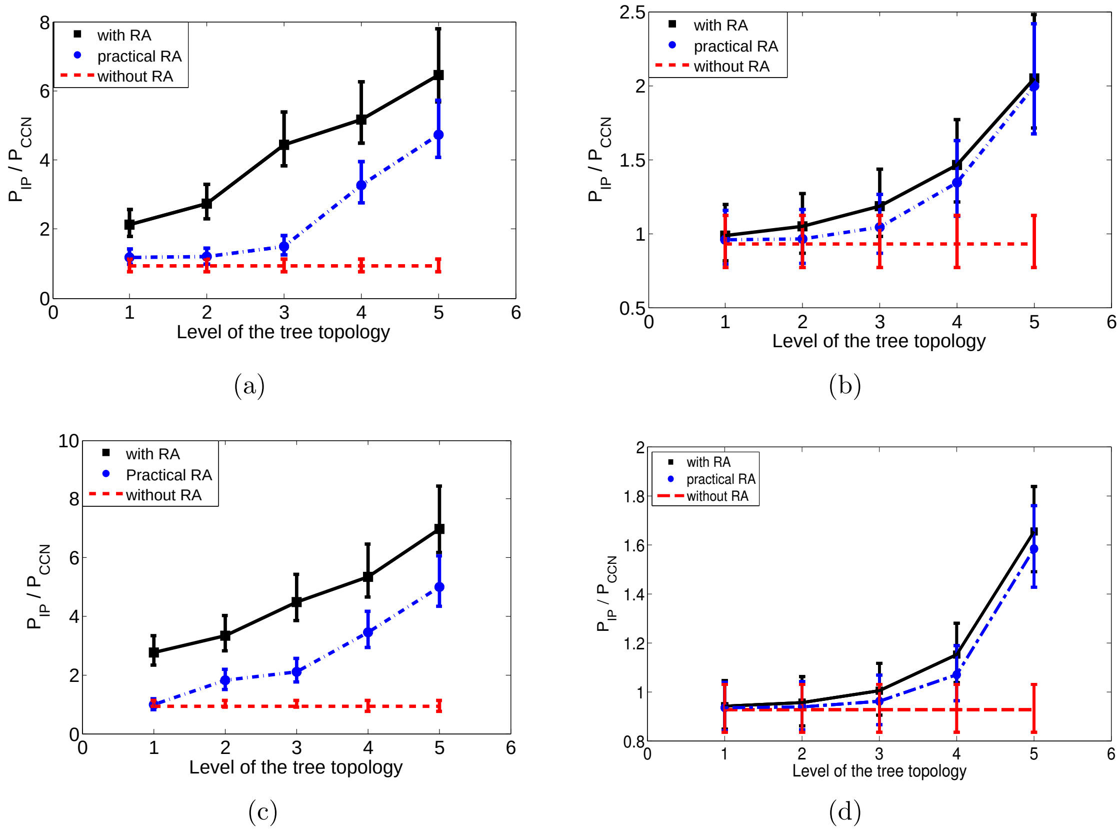

Fig. 4.6 Power ratio at each level of a general tree topology; $N = 5$, $X = 1$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.(a) 2 children per node and Zipf traffic load. (b) 2 children per node and Uniform traffic load. (c) 3 children per node and Zipf traffic load. (d) 3 children per node and Uniform traffic load.

ratio for CCN-based network degrades.

The curves in these figures are showing non-linear behavior. The explanation for this is twofold: 1) The number of nodes present on a given level is exponentially distributed, hence we have exponentially higher capability to cache the content at lower levels. 2) In case of zipf traffic demand the network takes full advantage of caching capability of routers and hence we get higher performance gain (power ratio) at higher levels (see figures 4.6(a) and 4.6(c)). But when we have uniformly distributed traffic demand (as in figures 4.6(b) and 4.6(d)) the caching feature of the network is not utilized at the fullest because of the frequent cache update as a result of uniformly generated traffic. However the content still gets cached in the network. The effect of this cached content on the power ratio, although less significant compared to the zipf demand, can be seen at the levels very close to the server. The arguments presented here remain valid for the similar observations we make in our further analysis.

To further elaborate our point and substantiate our findings, we provide results using larger networks in terms of number of nodes and also the number of levels. We have repeated the same simulations as discussed in the above paragraph for a larger network i.e., with 7 levels (see figure 4.7). The power ratio in all the three scenarios (i.e., with RA, Practical RA, without RA) follows the same trends as followed by the network with 5 levels but there is also a subtle difference in overall power ratio. The global gain (average power ratio of the entire network) of the network with 7 levels has slightly decreased. This fact is explained further in figure 4.8. In figure 4.8, we have presented the global power ratio for networks with 4, 5 and 7 levels respectively. Here we can clearly see a slight decrease in the average power ratio as the scale of the network increases. The reason behind this observation is twofold. In a large scale network, the content has to traverse larger number of hops on average (see figure 4.5) to reach the clients. This fact slows down the caching process which results in an increased network load and degraded global power ratio of the network.

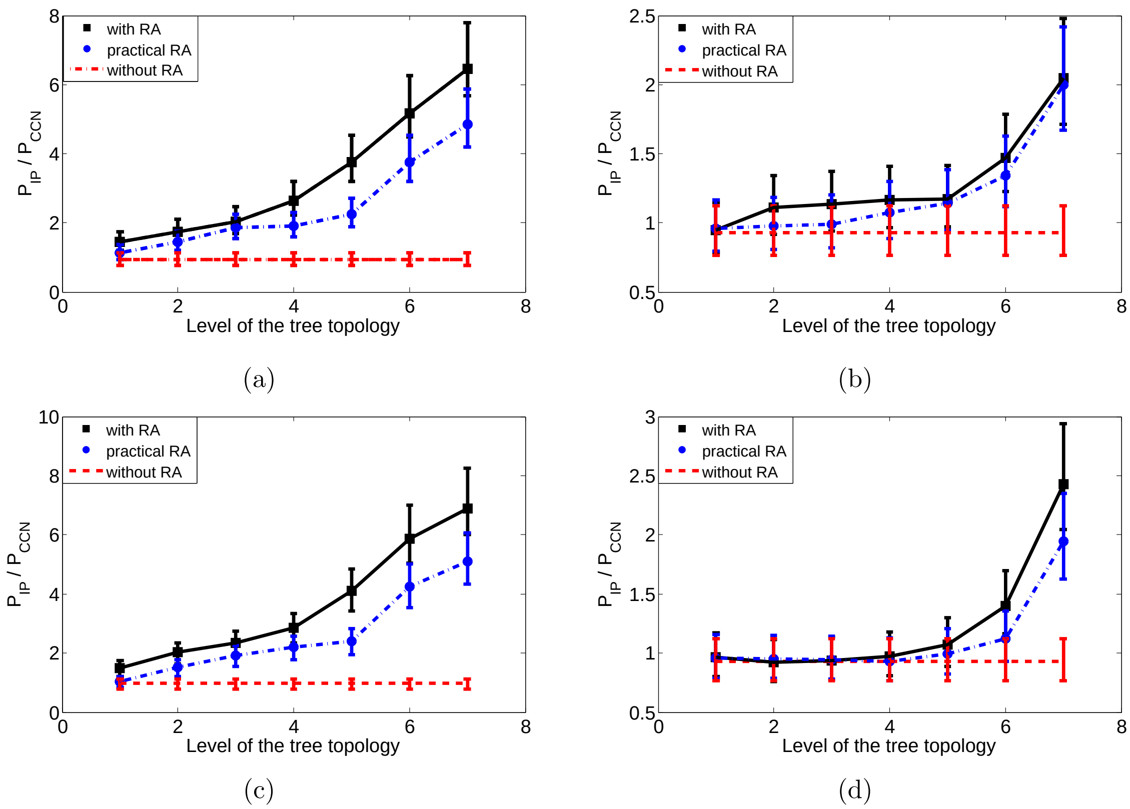

Fig. 4.7 Power ratio at each level of a binary tree topology; $N = 7$, $X = 1$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.(a) 2 children per node and Zipf traffic load. (b) 2 children per node and Uniform traffic load. (c) 3 children per node and Zipf traffic load. (d) 3 children per node and Uniform traffic load.

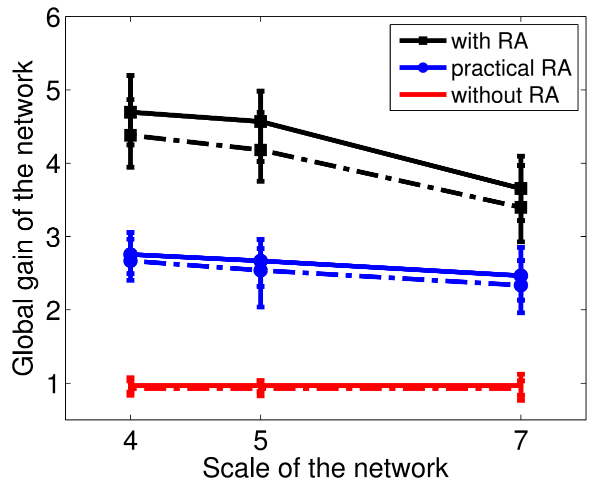

Fig. 4.8 The global gain (average power ratio) of networks with different scales; $X = 1$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.

So far in our analysis, we have considered a uniform cache size for the routers in a CCN network which is a somewhat simplistic assumption. To make our analysis closer to the real world implementations, we performed the following simulations by using different cache sizes at different network levels. We have designed our network in such a way that the levels closer to the server have larger cache sizes compared to the levels closer to the clients. Figure 4.9 shows the power ratio comparison of networks with uniform and incremental cache sizes. The thicker lines show the power ratio of network with incremental caches. It is evident from figures 4.9(a) and 4.9(b) that networks with cache sizes incrementing per level from client to server perform better than the networks with uniform cache sizes throughout the network.

One of the important aspects in implementing a CCN is the selection of appropriate cache sizes for the network deployment. If there are no bounds on the cache size in a CCN network, its performance can degrade in terms of energy consumption. We have so far performed simulations for our analysis by assuming reasonable cache sizes based on the results in figures 4.10, 4.11. Figure 4.10 presents the average power ratio of a 5-level general tree topology for different cache sizes. It can be seen in each sub-figure that the

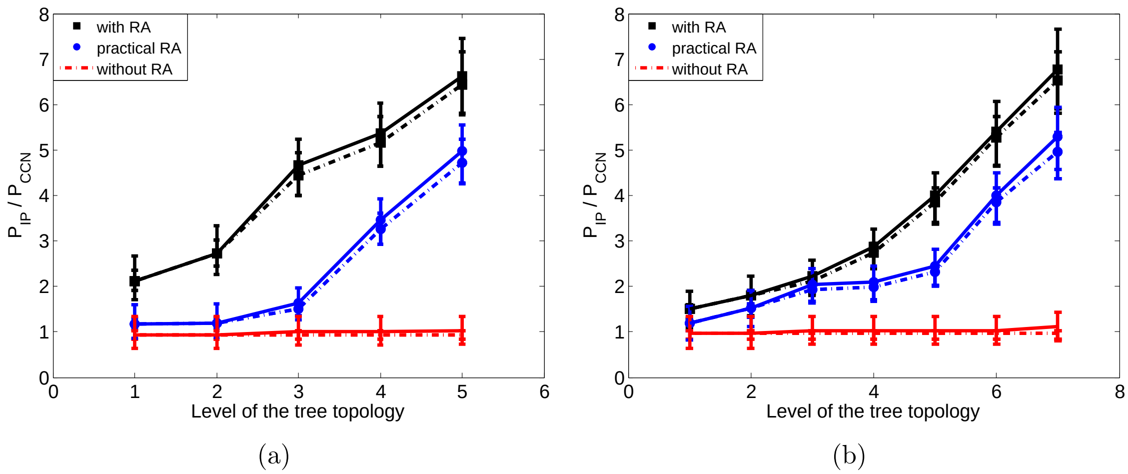

Fig. 4.9 Power ratio at each level of a binary tree topology with Zipf traffic load; Smaller caches are used on the lower levels of the network. The solid lines represent the network with uniform cache and the dotted lines depict the network with different cache sizes at different levels; The size of the uniform cache in both instances is 1TB(a) $N = 5$, $\gamma = 2$, $\lambda = 1$ request/sec, $X = 128, 256, 512$ GB, and 1 TB. (b) $N = 7$, $\gamma = 2$, $\lambda = 1$ request/sec $X = 128, 256, 512, 768, 1024$ GB, and 1.5 TB.

CCN network's power ratio is optimized at a particular cache size and starts to degrade if the size increases further. A careful analysis of figures 4.10(a), 4.10(b) and 4.10(c), 4.10(d) provides further evidence for our observation that increasing the number of routers at a given level of the general tree topology decreases the average traffic load on each link and hence reduces the cache size requirements throughout the network. This fact is more visible in a 7 tiered network shown in figure 4.11.

In figure 4.11(a) and 4.11(b) when the tree topology changes from binary to tertiary tree the cache size required to optimize the power ratio for CCN-based network decreases from 1 TB to 768 GB. These results provide some valuable insight for designing the cache size of a CCN network. A thorough study of this topic is beyond the scope of this thesis. It is evident from the results discussed so far that IP never outperforms CCN in the presence of any type of rate adaptation under both types of traffic demands. It should also be noted that we can better benefit from the merits of CCN if the traffic demand follows Zipf

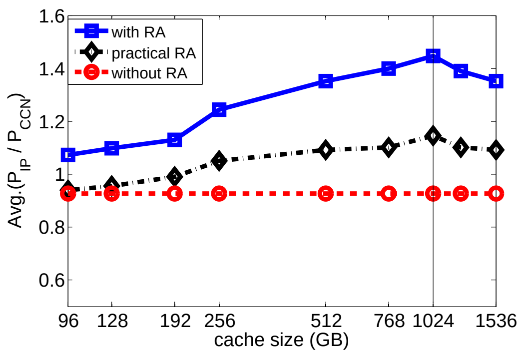

(a)

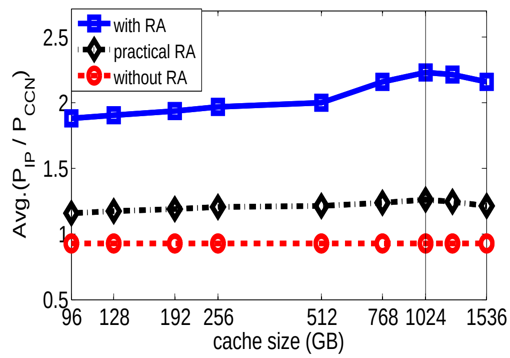

(b)

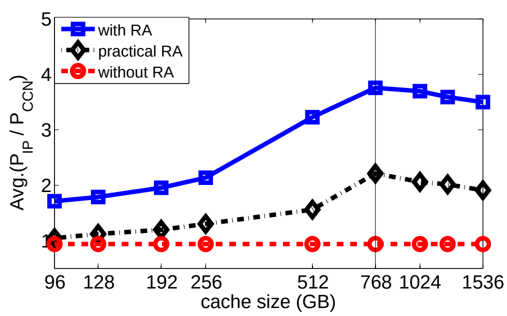

(c)

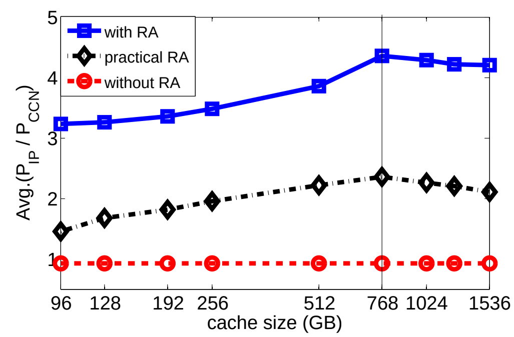

(d)

Fig. 4.10 Average power ratio for different cache sizes in general tree topology. The power ratio values in each of these figures are estimated by running 50 iterations of the simulation using randomly generated traffic load. The standard deviation of these estimates are between $0.004 - 0.005$, hence we are just presenting the mean values here; $N = 5$, $\gamma = 2$, $\lambda = 1$ request/sec.(a) Tree with 2 children per node, with uniform traffic demand. (b) Tree with 3 children per node, with uniform traffic demand (c) Tree with 2 children per node, with Zipf traffic demand. (d) Tree with 3 children per node, with Zipf traffic demand.

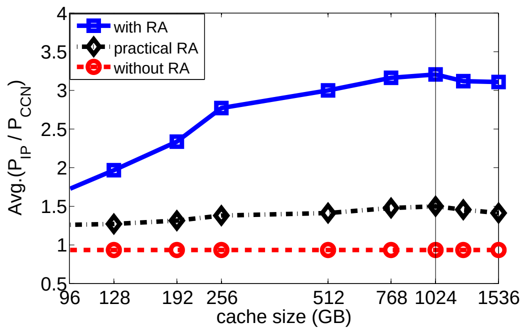

(a)

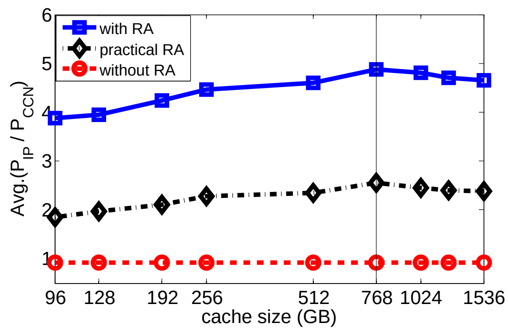

(b)

Fig. 4.11 Average power ratio for different cache sizes in general tree topology with Zipf traffic demand; $N = 7$, $\gamma = 2$, $\lambda = 1$ request/sec.(a) Tree with 2 children per node. (b) Tree with 3 children per node.

popularity distribution compared to the Uniform demand. Henceforth, we perform our further analysis using Zipf traffic demand only.

Comparison with CDN

In subsection 3.1 we have provided the definition of the CDN-based tree topology used in our analysis. Here we provide the simulation results for the comparison of CDN and CCN-based network topologies. Figure 4.12 shows the power ratio comparison between CDN and CCN-based networks. The CDN is designed in such a way that the CDN-servers are placed on the 2nd, 3rd and 4th levels of the 5 tiered network topology. Although the presence of CDN-servers decrease the load in the network, the absence of any rate adaptation incurs additional operation and manufacture energy costs to the network. Hence in this case CDN performs even worse than IP in terms of power consumption. CDN outperforms CCN at higher levels of the tree topology if rate adaptation is used. This is an expected outcome because using CDN-servers in a content centric networking environment would further decrease the traffic load in the network.

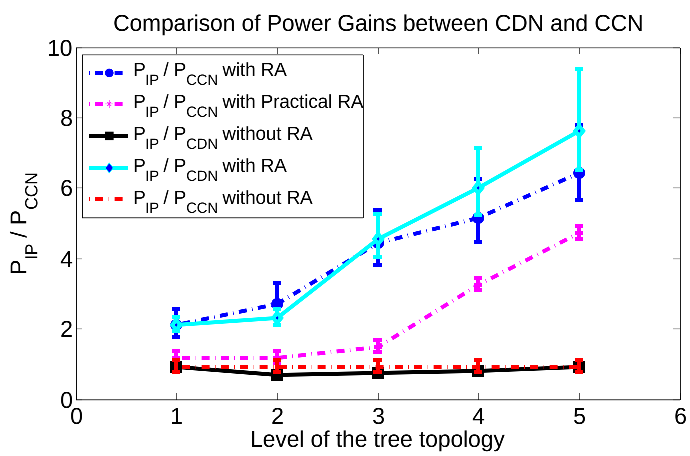

(a)

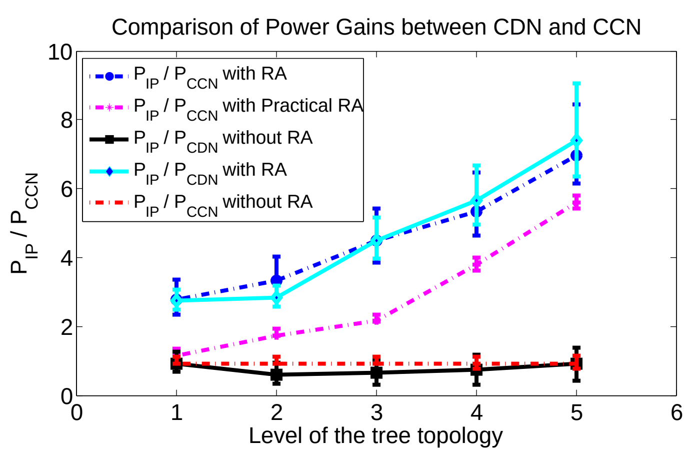

(b)

Fig. 4.12 Power ratio comparison between CCN and CDN with Zipf traffic demand; $N = 5$, $X = 1$ TB, $X_{CDN} = 3$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.(a) Tree topology with 2 children per node. (b) Tree topology with 3 children per node.

Power Analysis Using Mesh Topology

The general tree topology is a reasonably good choice to simulate a computer network but in a real world network deployment connections do not always follow the tree hierarchy. To address this issue, we present the power ratio results for CCN and IP-based network in a partially connected mesh topology. We have chosen two networks with 5 and 7 tiers respectively. The network is populated such that the number of nodes at each level is the same as binary tree but the nodes can have more than one connection on their own level and multiple connections to nodes in the layers directly above and below. Figure 4.13 provides the snapshot of a five tiered network showing connections of 2nd and 3rd level. Figures 4.13(a), 4.13(b) depict approximately 35 and 70 percent nodes have more than 3 neighbors. Figures 4.14(a) and 4.14(b) show the power ratio comparison for CCN and IP-based network in a partially connected mesh topology. As expected the CCN outperforms IP in the presence of a rate adaptation scheme. It is however important to notice the decrease in power ratio at the lower levels of the mesh topology compared to the power ratios of a similar tree based topology (see figures 4.6 and 4.7). This effect is more prominent in figures 4.14(c) and 4.14(d). Here we can see that performance of the network is reduced if the number of connections between the routers increases. The simulations are designed in such a way that when a router receives a request for a particular content, it broadcasts the request to all its parents in case of CCN. Upon reception of multiple copies of the same content, only one copy is stored and others are discarded. This phenomenon of broadcast puts added burden on the network on one hand (especially on the lower tiers of the network where there are more routers and connections) but on the other hand speeds up the caching of popular content and results in providing substantial power ratio at higher levels of the network.

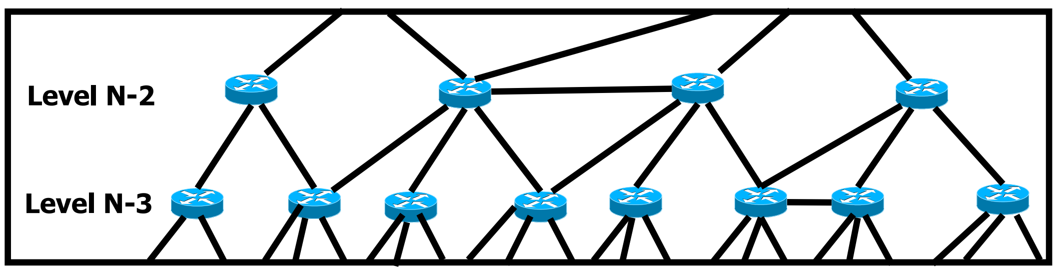

(a)

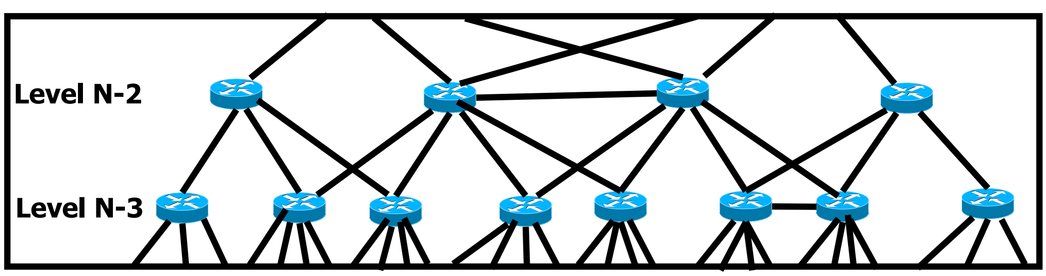

(b)

Fig. 4.13 Snapshot of the 5 tiered mesh topology showing the connections between routers at the same and different tiers. (a) When 30% nodes have more than 3 connections. (b) When 70% nodes have more than 3 connections

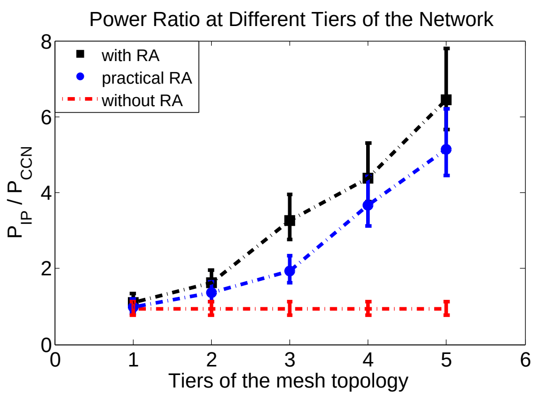

(a)

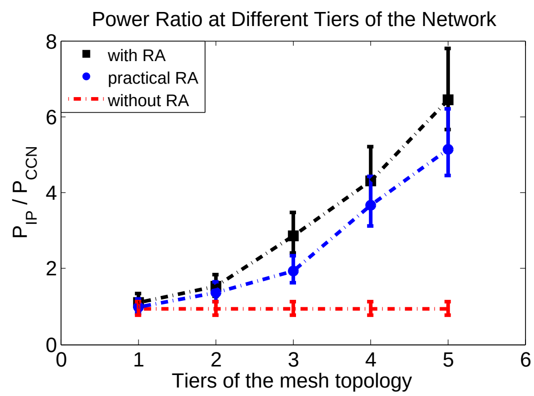

(b)

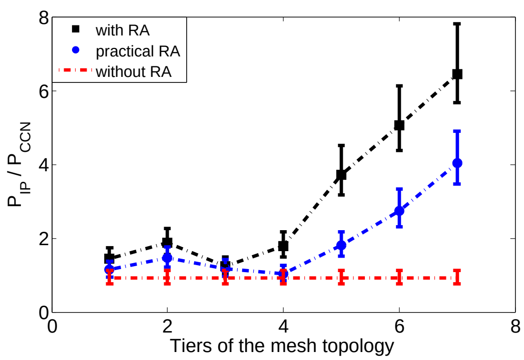

(c)

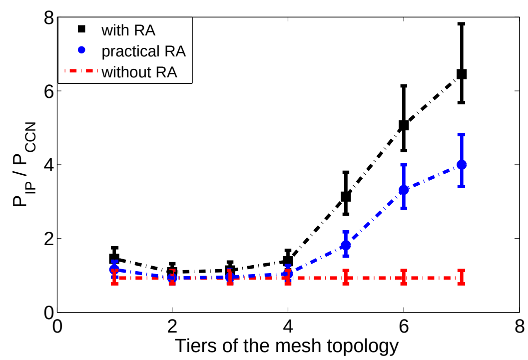

(d)

Fig. 4.14 Power ratio at each tier of partially connected mesh, Zipf traffic demand; $X = 1$ TB, $\gamma = 2$, $\lambda = 1$ request/sec.(a) $N = 5$, when 30% nodes have more than 3 connections. (b)$N = 5$, when 70% nodes have more than 3 connections.(c) $N = 7$, when 30% nodes have more than 3 connections. (d)$N = 7$, when 70% nodes have more than 3 connections.

4.6.2 Carbon footprint

We have used the values presented in tables 4.2 and 4.5 to calculate the carbon-footprint for the three scenarios we have considered for our analysis. We consider a 5 level binary tree where $X = 1$ TB, $\gamma = 2$ and $\lambda = 1$ request/sec. Figure 4.15 presents the comparison of carbon dioxide emission in grams for IP, CDN and CCN-based networks. The carbon footprint due to manufacture is higher in case of CDN as compared to IP and CCN. This is due to the presence of CDN-servers in the network in addition to the network routers. Also, the cost to operate a CDN-server is much higher than a network router when operated at peak power rating. In case of rate adaptation, the presence of CDN-servers considerably decrease the traffic load on the upper levels of CDN, reducing the carbon dioxide emission due to the operation compared to a CCN-based deployment (see figure 4.15). But CCN-based network substantially outperforms IP and marginally outperforms CDN in terms of lesser overall carbon dioxide emissions. In case of no rate adaptation, IP produce less carbon dioxide than CDN and CCN-based networks. CDN has the highest carbon footprint in this case due to the extra emissions incurred by the CDN-servers.

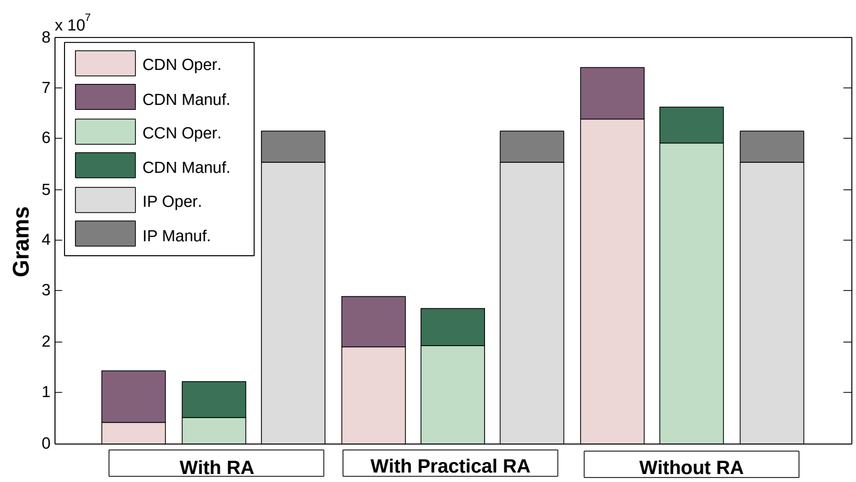

Fig. 4.15 The carbon footprint values for CCN, CDN and IP-based networks; The labels on the x-axis are representing the three scenarios for CCN-based networks. Due to the absence of cache capable routers IP-based networks cannot take advantage of rate adaptation.

4.7 Summary

This chapter describes the main contribution of our work. The first four sections discuss the methodology by which we estimated the power consumption and carbon-footprint values for the network devices used in our analysis. The research on the topic of CCNs is still in its early stages so a real world implementation of these kind of networks is not available at the moment. We developed *CCNSim and CCNSim++* to simulate CCN. The details of our simulators have been provided in the fifth section of this chapter. The penultimate section presented results and detailed discussion on the comparison of CCN with IP-based networks in terms of power consumption and carbon-footprint.

Chapter 5

Conclusion

5.1 Summary and Discussion

In this work, we conducted a comparative study of CCN and IP-based networks in terms of energy consumption and carbon footprint. The aim of our study was to identify which one of these networking architectures is greener for the environment. The analysis was performed using a dedicated video streaming application with a variety of traffic demands for CCN and IP-based networks. In our analysis we estimated the energy consumption during both manufacturing and operation of the network devices and calculated the carbon footprint using these estimates. Content centric networking is a relatively recent research topic and the existing literature on the green aspect of CCN is quite limited. To make our analysis conservative we introduced an error of $\pm 10\%$ in all the estimates related to the manufacturing and operation of network devices.

An initial assessment of the energy aspect of the CCN architecture suggested that the cache capable routers of CCN would incur an extra energy expenditure compared to the IP-based network due to the manufacturing and operation of the additional cache memory. The benefit of using a network with caching capability is the decrease in traffic load on many links of the network. We used this property of CCN and proposed the use of rate adaptation [9, 38] for CCN routers. This lead to a decreased energy consumption [8–10] in the case of CCN-based deployments. For the case of IP routers, the feature of caching was not present, hence the use of rate adaptation proved to be of little effect. We described two different mechanisms for rate adaptation: (a) On-line, where routers adjusted their operating rates seamlessly according to the traffic demand. (b) Static, where operating rates of the routers were provisioned according to the anticipated load.

In chapter 3, we presented our network model. We discussed the details of three network topologies (i.e., general tree, CDN tree and partial mesh) which were used in our analysis and also presented the general equations to estimate the total power consumption for the CCN and IP-based networks. We observed that rate adaptation provided a major reduction in the energy requirements for a CCN based network implementation. We concluded the chapter with the methodology for estimating the carbon footprint of the network devices.

The first half of chapter 4 provided the procedure by which we estimated the power consumption and carbon-footprint values for the network devices used in our analysis. Section 4.5 provided a detailed discussion on the simulation model and the two simulators i.e., *CCNSim++* and *CCNSim* developed to perform the analysis. The ability of *CCNSim++* to implement general network topologies for CCN and IP-based networks helped us to analyze more practical network deployments e.g., partial mesh, CDN tree.

The rest of the chapter discussed the results obtained using the above mentioned simulators. The findings from these results suggest that CCN's performance in terms of power consumption and carbon footprint is comparable to IP-based network deployments in the absence of any rate adaptation. The presence of any rate adaptation mechanism makes CCN a much greener networking architecture compared to IP.

5.2 Future Work

Our current study is based on results achieved using the simulators I developed with some help from colleagues in the McGill Computer Networks Lab. Although they provide a good platform to conduct our analysis, it would be useful to emulate a CCN using real network devices and then compare the emulated results with the findings obtained using simulations.

We discussed the effect of the cache size in a CCN network deployment and provided a procedure based on simulations to find the optimal cache size for the network. More insight can be obtained if a more rigorous analysis for optimal cache size is performed by taking into account the effect of different caching polices and location of a router in the real world network.

References

[1] H.-K. Nienhuys, "Hard drive capacity over time," [Online] http://commons.wikimedia.org/wiki/File_talk:Hard_drive_capacity_over_time.svg.

[2] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of peer-to-peer overlay network schemes," *IEEE, Comm. Surveys Tutorials*, vol. 7, pp. 72 – 93, May. 2005.

[3] M. Pathan and R. Buyya, "A taxonomy of CDNs," in *Lec. Notes Elect. Engg.*, 2008, vol. 9, pp. 33–77.

[4] A. Ghodsi, T. Koponen, B. Raghavan, S. Shenker, A. Singla, and J. Wilcox, "Information-centric networking: seeing the forest for the trees," in *Proc. ACM Workshop Hot Topics in Networking*, Cambridge, MA, USA, Nov. 2011.

[5] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," in *Proc. Int. Conf. Emerging Networking Experiments and Tech. (CoNEXT)*, Rome, Italy, Dec. 2009.

[6] A. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier, "A survey of green networking research," *IEEE, Comm. Surveys Tutorials*, vol. 14, pp. 3 –20, May. 2012.

[7] B. Raghavan and J. Ma, "The energy and emergy of the Internet," in *Proc. ACM Workshop Hot Topics in Networks*, Cambridge, MA, USA, Nov. 2011.

[8] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proc. Symp. Networked Systems Design and Implementation*, Berkeley, CA, USA, Apr. 2008.

[9] S. Antonakopoulos, S. Fortune, and L. Zhang, “Power-aware routing with rate-adaptive network elements,” in *Proc. Globecom Workshops*, Miami, Florida, USA, Dec. 2010.

[10] A. Lombardo, C. Panarello, and G. Schembra, “Achieving energy savings and QoS in Internet access routers,” *SIGMETRICS Perform. Eval. Rev.*, vol. 38, pp. 76–80, Jan. 2011.

[11] PARC, “Project CCNx,” [Online] http://www.ccnx.org/content/docs.

[12] R. Valadas and P. Salvador, *Traffic Management and Traffic Engineering for the Future Internet.* Springer, 2009, vol. 5464.

[13] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, “Understanding user behavior in large-scale video-on-demand systems,” in *Proc. ACM European Conf. Comp. Sys.*, Leuven, Belgium, Apr. 2006.

[14] B. Ahlgren, M. D’Ambrosio, M. Marchisio, I. Marsh, C. Dannewitz, B. Ohlman, K. Pentikousis, O. Strandberg, R. Rembarz, and V. Vercellone, “Design considerations for a Network of Information (NetInf),” in *Proc. Int. Conf. Emerging Networking Experiments and Tech. (CoNEXT)*, Madrid, Spain, Dec. 2008.

[15] G. Bouabene, C. Jelger, C. Tschudin, S. Schmid, A. Keller, and M. May, “The Autonomic Network Architecture (ANA),” *IEEE J. Selected Areas in Comm.*, vol. 28, pp. 4 –14, Jan. 2010.

[16] V. Jacobson, D. Smetters, J. Thornton, M. Plass, P. Stewart, N. Briggs, and R. Braynard, “VoCCN: voice-over content-centric networks,” in *Proc. Workshop Re-architecting the Internet*, Rome, Italy, Dec. 2009.

[17] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, “A data-oriented (and beyond) network architecture,” *SIGCOMM Comput. Commun. Rev.*, vol. 37, pp. 181–192, Aug. 2007.

[18] EFIP, “PubSub Internet Routing Paradigm,” [Online] http://www.psirp.org.

[19] SAIL, “Network of Information,” [Online] http://www.sail-project.eu.

[20] GENI, "Exploring networks of the future," [Online] http://www.geni.net.

[21] NSF, "FIND (Future Internet Design)," [Online] http://www.nets-find.net/.

[22] EFIP, "The Future Internet Assembly (FIA)," [Online] http://www.future-internet.eu.

[23] ANR Connect, "Content-Oriented Networking: a New Experience for Content Transfer," [Online] http://www.anr-connect.org.

[24] NDN, "Named Data Networking," [Online] http://www.named-data.net.

[25] I. Psaras, R. G. Clegg, R. Landa, W. K. Chai, and G. Pavlou, "Modelling and evaluation of CCN-caching trees," in *Proc. Int. IFIP TC 6 conf. on Networking*, Valencia, Spain, May. 2011.

[26] L. Muscariello, G. Carofiglio, and M. Gallo, "Bandwidth and storage sharing performance in information centric networking," in *Proc. ACM workshop Information-centric networking, (Sigcomm)*, Toronto, Ontario, Canada, Aug. 2011.

[27] G. Carofiglio, V. Gehlen, and D. Perino, "Experimental evaluation of memory management in content-centric networking," in *Proc. IEEE Int. Conf. Communications (ICC)*, Kyoto, Japan, Jun. 2011.

[28] M. Tortelli, I. Cianci, L. Grieco, G. Boggia, and P. Camarda, "A fairness analysis of content centric networks," in *Proc. Int. Conf. Network of the Future*, Paris, France, Nov. 2011.

[29] G. Carofiglio, M. Gallo, L. Muscariello, and D. Perino, "Modeling data transfer in content-centric networking," in *Proc. Int. Teletraffic Congress*, San Francisco, California,USA, Sept. 2011.

[30] U. Lee, I. Rimac, and V. Hilt, "Greening the Internet with content-centric networking," in *Proc. Int. Conf. Energy-Efficient Computing and Networking*, Passau, Germany, Apr. 2010.

[31] K. Guan, G. Atkinson, D. Kilper, and E. Gulsen, "On the energy efficiency of content delivery architectures," in *Proc. IEEE Int. Conf. on Communications*, Kyoto, Japan, Jun. 2011.

[32] A. Seetharam, M. Somasundaram, D. Towsley, J. Kurose, and P. Shenoy, "Shipping to streaming: is this shift green?" in *Proc. ACM SIGCOMM Workshop Green Networking*, New Delhi, India, Aug. 2010.

[33] V. Soteriou and L.-S. Peh, "Dynamic power management for power optimization of interconnection networks using on/off links," in *Proc. Symp. on High Performance Interconnects*, Stanford, CA, USA, Aug. 2003.

[34] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *Proc. Int. IEEE Conf. Communications(ICC)*, Dresden, Germany, Jun. 2009.

[35] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proc. INFOCOM*, Phoenix, USA, Apr. 2008.

[36] C. Gunaratne, K. Christensen, and B. Nordman, "Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections, and scaling of link speed," *Int. J. Netw. Manag.*, vol. 15, pp. 297–310, Sep. 2005.

[37] C. Gunaratne, K. Christensen, and S. Suen, "Ethernet adaptive link rate (ALR): Analysis of a buffer threshold policy," in *Proc. IEEE Global Telecommunications Conf.(Globcom)*, San Francisco, California, USA, Dec. 2006.

[38] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption of ethernet with Adaptive Link Rate (ALR)," *IEEE Trans. Comput.*, vol. 57, pp. 448–461, 2008.

[39] M. Gupta and S. Singh, "Using low-power modes for energy conservation in ethernet LANs," in *Proc. Int. IEEE Conf. on Computer Communications (Infocom)*, May 2007.

[40] ——, "Dynamic ethernet link shutdown for energy conservation on ethernet links," in *Proc. Int. IEEE Conf. Communications (ICC)*, Glasgow, Scotland, Jun. 2007.

[41] C. Panarello, A. Lombardo, G. Schembra, L. Chiaraviglio, and M. Mellia, "Energy saving and network performance: a trade-off approach," in *Proc. Int. Conf. Energy-Efficient Computing Networking*, Passau, Germany, Apr. 2010.

[42] R. S. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. V. Sorin, "Energy consumption in IP networks," in *Proc. European Conf. Optical Communication*, Brussels, Belgium, Sept. 2008.

[43] M. Barbera, A. Lombardo, C. Panarello, and G. Schembra, "Active Window Management: An efficient gateway mechanism for TCP traffic control," in *Proc. IEEE Int. Conf. Communications (ICC)*, Glasgow, UK, Jun. 2007.

[44] M. Barbera, M. Gerla, A. Lombardo, C. Panarello, M. Sanadidi, and G. Schembra, "Active window management: performance assessment through an extensive comparison with XCP," in *Proc. Int. IFIP-TC6 Networking Conf. AdHoc Sensor Networks, Wireless Networks, Next Generation Internet*, Singapore, May 2008.

[45] E. Carlson, "Carbon Fund." [Online]. Available: http://www.carbonfund.org

[46] The Holyoke Green High Performance Computing Centre (GPHCC), "Greening of computing." [Online]. Available: http://www.hged.com/2008-Annual-Report-FINAL-WEB.pdf

[47] C. Herrmann, "Environmental footprint of ict equipment in manufacture, use and end of life," in *European Conf. and Exhibition on Optical Comm.*, Brussels, Belgium, Sept. 2008.

[48] Verivue, "OneVantage VoD Streamer content delivery solution," [Online] http://www.verivue.com/pdf/verivue-vodstreamer.pdf.

[49] Juniper Networks, "M7i and m10i multiservice edge routers," [Online] http://www.juniper.net/us/en/local/pdf/datasheets/1100032-en.pdf.

[50] E. Williams, "Energy intensity of computer manufacturing: Hybrid assessment combining process and economic input-output methods," *Int. J. Envr. Sci. and Tech.*, vol. 38, pp. 66–74, 2004.

[51] L. Deng, E. Williams, and C. Babbitt, "Hybrid life cycle assessment of energy use in laptop computer manufacturing," in *Proc. Int. IEEE Symp. Sustainable Systems and Tech.*, Tempe, AZ, USA, May. 2009.

[52] C. Walter, "Kryder's Law," *Scientific American*, pp. 32–33, Aug. 2005.

[53] H. Duan, M. Eugster, R. Hischier, M. Streicher-Porte, and J. Li, “Life cycle assessment study of a chinese desktop personal computer,” *Science of the Total Environment*, vol. 407, no. 5, pp. 1755–1764, 2009.

[54] Seagate, “Data-sheet Barracuda 7200.11,” [Online] http://www. seagate.com/docs/pdf/datasheets/disc/ds_7200_12.pdf.

[55] ——, “Data-sheet Momentus 5400.4,” [Online] http://www.seagate.com /docs/pdf/datasheet/disc/ds_momentus_5400_4.pdf.

[56] G. Carofiglio, M. Gallo, L. Muscariello, and D. Perino, “Modeling data transfer in content-centric networking,” in *Research report*, Sept. 2011, [Online] http://perso.rd.francetelecom.fr/muscariello/report-itc-transport.pdf.

[57] Netflix, “Online movie streaming,” [Online] http://www.netflix.com.

[58] IMDB, “The internet movie database,” [Online] http://www.imdb.com/chart/top.

[59] M. Butt and O. Delgado, “The source codes for CCNSim++ and CCNSim,” [Online] http://networks.ece.mcgill.ca/node/167.