


Esteban de Jesus Rodriguez Brljevich

“Energy-aware Virtual Network Mapping”

*“Mapeamento de Redes Virtuais Ciente do
Consumo de Energia”*

CAMPINAS
2013



University of Campinas
Institute of Computing

*Universidade Estadual de Campinas
Instituto de Computação*

Esteban de Jesus Rodriguez Brljevich

“Energy-aware Virtual Network Mapping”

Supervisor: Prof. Dr. Nelson Luis Saldanha da Fonseca
Orientador(a):

***“Mapeamento de Redes Virtuais Ciente do
Consumo de Energia”***

MSc Dissertation presented to the Post Graduate Program of the Institute of Computing of the University of Campinas to obtain a Mestre degree in Computer Science.

Dissertação de Mestrado apresentada ao Programa de Pós-Graduação em Ciência da Computação do Instituto de Computação da Universidade Estadual de Campinas para obtenção do título de Mestre em Ciência da Computação.

THIS VOLUME CORRESPONDS TO THE FINAL VERSION OF THE DISSERTATION DEFENDED BY ESTEBAN DE JESUS RODRIGUEZ BRljeVICH, UNDER THE SUPERVISION OF PROF. DR. NELSON LUIS SALDANHA DA FONSECA.

ESTE EXEMPLAR CORRESPONDE À VERSÃO FINAL DA DISSERTAÇÃO DEFENDIDA POR ESTEBAN DE JESUS RODRIGUEZ BRljeVICH, SOB ORIENTAÇÃO DE PROF. DR. NELSON LUIS SALDANHA DA FONSECA.

A handwritten signature in blue ink, reading "Nelson Luis Saldanha da Fonseca".

Supervisor's signature / *Assinatura do Orientador(a)*

CAMPINAS
2013

Ficha catalográfica
Universidade Estadual de Campinas
Biblioteca do Instituto de Matemática, Estatística e Computação Científica
Ana Regina Machado - CRB 8/5467

R618e Rodriguez Brljevich, Esteban, 1984-
Energy-aware virtual network mapping / Esteban Brljevich Rodriguez. –
Campinas, SP : [s.n.], 2013.

Orientador: Nelson Luis Saldanha da Fonseca.
Dissertação (mestrado) – Universidade Estadual de Campinas, Instituto de
Computação.

1. Redes verdes. 2. Virtualização de redes. 3. Redes de computadores -
Conservação de energia. 4. Redes de computadores - Aspectos ambientais. I.
Fonseca, Nelson Luis Saldanha da, 1961-. II. Universidade Estadual de Campinas.
Instituto de Computação. III. Título.

Informações para Biblioteca Digital

Título em outro idioma: Mapeamento de redes virtuais ciente do consumo de energia

Palavras-chave em inglês:

Green networking

Network virtualization

Computer networks - Energy conservation

Computer networks - Environmental aspects

Área de concentração: Ciência da Computação

Titulação: Mestre em Ciência da Computação

Banca examinadora:

Nelson Luis Saldanha da Fonseca [Orientador]

Otto Carlos Muniz Bandeira Duarte


Luiz Fernando Bittencourt

Data de defesa: 29-11-2013


Programa de Pós-Graduação: Ciência da Computação

TERMO DE APROVAÇÃO


Dissertação Defendida e Aprovada em 29 de novembro de 2013, pela Banca examinadora composta pelos Professores Doutores:



Prof. Dr. Nelson Luis Saldanha da Fonseca
IC / UNICAMP



Prof. Dr. Otto Carlos Muniz Bandeira Duarte
GTA / UFRJ



Prof. Dr. Luiz Fernando Bittencourt
IC / UNICAMP

Energy-aware Virtual Network Mapping

Esteban de Jesus Rodriguez Brljevich

November 29, 2013

Examiner Board / *Banca Examinadora*:

- Prof. Dr. Nelson Luis Saldanha da Fonseca (Supervisor / *Orientador*)
- Prof. Otto Carlos Muniz Bandeira Duarte
Grupo de Teleinformática e Automação - UFRJ
- Prof. Dr. Luiz Fernando Bittencourt
Institute of Computing - UNICAMP
- Prof. Dr. Christian Esteve Rothenberg
FEEC - UNICAMP
- Juliana Freitag Borin
Institute of Computing - UNICAMP

Abstract

Network virtualization is a promising technology for the Internet of the Future since it facilitates the deployment of new protocols and applications without the need of changing the core of the network. A key step to instantiate virtual networks is the allocation of physical resources to virtual elements (routers and links). In order to contribute to the global effort of saving energy, choice of physical resources to instantiate a virtual network needs to minimize the network energy consumption. However, this is not a trivial task, since the QoS of the application requirements has to be supported. Indeed, the search for the optimal solution of this problem is NP-hard.

The mapping of virtual networks on network substrates at the arrival time of requests to the establishment of virtual networks may not lead to a global minimum energy consumption of energy due to the dynamic allocations and deallocations of virtual networks. Actually, such dynamics can lead to the underutilization of the network substrate. To mitigate the negative effect of this dynamics, techniques such as live migration can be employed to rearrange already mapped virtual networks to achieve energy savings.

This dissertation presents a set of new algorithms for the mapping of virtual networks on network substrates aiming to reduce energy consumption. Additionally, two new algorithms are proposed for the migration of virtual routers and links to reduce the number of powered routers and optical amplifiers. Results derived by simulation show the efficacy of the proposed algorithms.

Resumo

A virtualização de redes é uma tecnologia promissora para a Internet do futuro, já que facilita a implementação de novos protocolos e aplicações sem a necessidade de alterar o núcleo da rede. Um passo chave para instanciar redes virtuais é a alocação de recursos físicos para elementos virtuais (roteadores e enlaces). A fim de contribuir para o esforço global de poupança de energia, a escolha de recursos físicos para instanciar uma rede virtual deveria minimizar o consumo de energia rede. No entanto, esta não é uma tarefa trivial, já que requerimentos de QoS devem ser atingidos. A busca da solução ótima deste problema é NP-difícil.

O mapeamento de redes virtuais em substratos de rede físicos em cenários de alocação e desalocação de redes virtuais pode não levar a um consumo mínimo de energia devido à dinâmica das atribuições dos elementos virtuais previamente alocados. Tal dinâmica pode levar à subutilização da rede substrato. Para reduzir os efeitos negativos desta dinâmica, técnicas tais como a migração de redes virtuais em tempo real podem ser empregadas para rearranjar as redes virtuais previamente mapeadas para poupar energia.

Esta dissertação apresenta um conjunto de novos algoritmos para o mapeamento de redes virtuais em substratos de rede com o objetivo de reduzir o consumo de energia. Além disso, dois novos algoritmos são propostos para a migração dos roteadores e enlaces virtuais para reduzir o número de roteadores e amplificadores ópticos requeridos. Os resultados obtidos por simulação mostram a eficácia dos algoritmos propostos.

Acknowledgements

Muito obrigado Deus, por iluminar meu caminho quando estava mais escuro. Obrigado por sempre cuidar de mim e dos meus entes queridos, e dar sinais para me encaminhar quando mais precisei.

Ao meu pai, que sempre é um exemplo a seguir. Muito obrigado pelas sábias palavras! À minha mãe, por sempre estar presente para me escutar e dar suporte emocional e amor, dia após dia. Carlos, obrigado por cuidar das coisas na CR e também por ser um exemplo a seguir sempre. Muito obrigado a todos por tolerarem tanto tempo sem mim.

A Cristina, obrigado pelo suporte constante. Não existe distância entre dois corações, os nossos se atraem como imãs, e a cada dia me sinto mais perto de você!

Este trabalho não seria possível sem a força, motivação e paciência infinita do Prof. Nelson Fonseca. Prof. Nelson, muito obrigado por me aceitar como o seu aluno e por todos os ensinamentos. Sempre lembrarei os conselhos e a alegria de trabalhar junto ao senhor. Ao Prof. Daniel Batista, muito obrigado pelos comentários e todas as dicas durante o curso do trabalho, além, da paciência. Eu sei que não é fácil trabalhar com aluno estrangeiro, muito obrigado por não desistir de mim, e boas corridas!

A todas as minhas amigadas: todos vocês fizeram que minha experiência no Brasil fosse inesquecível! Vocês valem ouro e espero encontrar vocês na CR! Muito obrigado Aloísio, Eduardo, Pedro e Pri (tropaaas!), Wagner, por sempre acreditarem em mim, e me ajudar a acreditar em mim mesmo! Muito obrigado Arthur, Davi, Laura e Moacy!

Ao LRC, muito obrigado por sempre achar um jeito de botar um sorriso na minha cara! Pela força de manter uma estrutura tão eficiente e valiosa para a pesquisa! Acreditem que chegarão longe! Muito obrigado Alisson, Carlitos, Carlos Senna, César, Daniel L, Flávio, Gê, Gustavo, Helder, Juliana, Luciano, Luis Guilherme, Mariana, Pedro H, Pedro M, Rafael L, Rafael S, Raphael (Rose!), Rodolfo, Rodrigo, Takeo, Thiago (Whey), Tiago (Tiaguito), Walisson. Ao IC, muito obrigado por abrir a suas maravilhosas portas para mim. Muito obrigado Andrei, Erick, Junior, Ramon (Jamon), Ricardo, Rosana. Muito obrigado, Daniel Capeleto, por correr atrás de mim quando precisava!

Gostaria de agradecer ao CNPq, The Horizon Project, MICIT, CONICIT e SAMOSOL S.A. pelo apoio financeiro. E finalmente, muito obrigado Brasil!

Contents

Abstract	ix
Resumo	xi
Acknowledgements	xiii
1 Introduction	1
2 Key concepts	4
2.1 Virtual Networks	4
2.2 Virtual Network Mapping Problem	6
2.3 Green Networking	6
2.4 The Green Virtual Network Mapping Problem	8
2.5 Live Migration	9
3 Related Work	11
3.1 Green Virtual Networks	11
3.2 Live Migration	12
3.3 Comparison of related work	13
4 Green Virtualized Networks	15
4.1 Motivation	15
4.2 Energy Consumption Model	16
4.2.1 Physical Router	17
4.2.2 Physical Links	18
4.3 Proposed Mathematical Model	19
4.4 Proposed Algorithms	25
4.4.1 Virtual Network Mapping Algorithms	25
4.5 Configuration of the Experiments	27
4.6 Results	27

4.6.1	GREEN Experiment 1	28
4.6.2	GREEN Experiment 2	29
4.6.3	GREEN Experiment 3	31
5	Live Migration of Virtual Networks	37
5.1	Introduction	37
5.2	Proposed Algorithms	38
5.2.1	Terminology	40
5.2.2	Nodes Recently Used (NRU)	41
5.2.3	Live Migration-ALL	42
5.3	Configuration of the Experiments	44
5.4	Results	44
5.4.1	Live Migration Experiment 1	44
5.4.2	Live Migration Experiment 2	47
6	Conclusions	54
	Bibliography	56
A	BAND Formulation	62
B	Original GREEN Formulation	65

List of Tables

3.1	Comparison between mapping algorithms	14
4.1	Values of the parameters used in GREEN Experiment 1	28
4.2	Values of the parameters used in GREEN Experiment 2	30
4.3	Values of the parameters used in GREEN Experiment 3	34
5.1	Values of the parameters used in Live Migration Experiment 1	45
5.2	Values of the parameters used in Live Migration Experiment 2	49

List of Figures

2.1	Network Architecture: Virtual Network and physical network with software repository	7
4.1	simple simulated scenario	16
4.2	Energy consumption per alternating link power	16
4.3	Energy consumption components	18
4.4	Average Energy Consumption per request as a function of the number of nodes of the substrate	29
4.5	Average Allocated Bandwidth per request as a function of the number of nodes of the substrate (GREEN Experiment 1)	30
4.6	Average Energy Consumption (Second set of experiments of a duration of the inter arrival time of request for a substrate with 30 nodes)	31
4.7	Average Energy Consumption (Second set of experiments of a duration of the inter arrival time of request for a substrate with 150 nodes)	32
4.8	Average Energy Consumption per request as a function of the number of nodes of the substrate	33
4.9	Average Allocated Bandwidth per request as a function of the number of nodes of the substrate	33
4.10	Substrate Utilization (Physical Nodes)	35
4.11	Substrate Utilization (Physical Links)	35
4.12	Blocking Rate	36
4.13	Execution Time	36
5.1	Example of reallocation of resources	38
5.2	Step by step of Live Migration to reduce energy consumption	39
5.3	Step by Step of LM NRU Algorithm	43
5.4	Average Energy Consumption per Request	46
5.5	Average Allocated Bandwidth per Request	46
5.6	Substrate Utilization as a Function of the Substrate Size	47
5.7	Average Virtual Routers Migrated per Event of Virtual Network Termination	48

5.8	Average Virtual Links Migrated per Event of Virtual Network Termination	48
5.9	Execution Time per Event of Virtual Network Termination	49
5.10	Average Energy Consumption per Request	50
5.11	Average Allocated Bandwidth per Request	51
5.12	Substrate Utilization as a Function of the Substrate Size	52
5.13	Average Virtual Routers Migrated per Event of Virtual Network Termination	52
5.14	Average Virtual Links Migrated per Event of Virtual Network Termination	53
5.15	Execution Time per Event of Virtual Network Termination	53

Chapter 1

Introduction

The reduction of energy consumption has become crucial for the sustainability of the world. Besides this, there are also engineering and marketing reasons involved. Emphasis has been given to the reduction of energy consumption in the past years and it is still a very important issue to consider in the design and implementation of information systems today.

Several studies [26] [14] have pointed out that the massive greenhouse gas emissions have a profound impact on the environment and climate change. It is required, according to [34], to reduce the greenhouse gas emissions from 15% to 30% by 2020 in order to guarantee that the global temperature will increase less than 2°C.

In recent years, telecommunication companies and Internet Service Providers have faced an increase in energy consumption due to the growing spread of broadband access and the expansion of the services offered. According to Bolla et al. [5], the increase in the volume of network traffic follows Moore's law, doubling every 18 months; while silicon technologies improve their energy efficiency according to Dennard's law, by a factor of 1.65 every 18 months. Thus, there is a constant increase in energy consumption related to communication networks, which corresponds to 2% to 10% of the world current energy consumption and this is expected to increase in the coming years.

The minimalist approach and the independence of specific network technology at the link layer have enabled the global spread of the Internet. The core of the Internet was designed to be simple, using the TCP/IP stack operational over different types of link layer technologies. However, as a consequence of this simplicity, various attempts have been made to provide missing features in its original design.

One of the main issues in network virtualization is the efficient mapping of virtual networks onto the substrate network [8] [51]. This mapping determines the allocation of routers and links of a virtual network onto the routers and links of the substrate network. However, the search for the optimal mapping of virtual networks (VNs) is an NP-hard

problem [23].

To overcome these limitations, various new architectures and mechanisms have been proposed to promote the evolution of the Internet [17] [51] [23]. Several of these mechanisms are based on network virtualization which allows the definition of virtual networks composed of virtual routers and links; these are then hosted by routers and links in real networks called “network substrates”. Network virtualization allows the coexistence of different protocol stacks and architectures on the same substrate, without the need of modifying the physical network. Moreover, it imposes no restrictions on these protocols and architectures.

Advances in hardware have allowed the design of energy efficient network devices by the adoption of “power on demand” operation. Techniques employed at the physical layer have made transmission more energy efficient. However, the advance of the state of the art in energy efficient networking is expected to happen at the architectural level [13]. In this context, network virtualization plays a key role since it can use physical devices in a more efficient way, reducing the need of acquisition of new devices.

The use of virtualization and energy efficient network devices have allowed several companies to employ different strategies in order to reduce the energy consumption of their networks. The CANARIE (Canada’s advanced research and innovation network) has employed the use of virtualization in order to localize services by energy source availability [1]. Google has displaced their server farms on the banks of the Columbia River in order to take advantage of the energy offered by nearby hydroelectric power plants, while exploring the possibility of using the river flow for natural cooling [27]. Amazon also explores geographic localization of services in an attempt to lower energy costs. [39] The U.S. Postal Service has been able to eliminate 791 of its 895 physical servers, thus decreasing their energy consumption costs with the use of virtualization [38]. These decisions are important due to the under utilization of servers. The survey in [4] shows that in a sample of 5000 typically only 10-15% of a server’s total capacity is employed. The survey in [16] states that idle devices consume around 70% of total energy consumption of datacenters. All these survey point out clearly the need for energy-aware utilization of resources.

The virtual network mapping problem consists in allocating virtual networks onto physical substrates. This problem is approached in several works [3], which explore several solutions in order to achieve bandwidth efficient virtual network mappings. A bandwidth efficient virtual network mapping may come with a hidden cost of energy.

It is possible to extend the virtual network mapping problem into an energy efficient virtual network mapping problem. In order to do this, the physical network model must be modified in order to include the cost of energy of the network components. Furthermore, the minimization algorithm must consider that using more routers and links implies on a higher energy consumption.

In addition to using virtualization techniques to improve the energy efficiency of virtual networks, live migration has also proven to be a very effective technique in reducing energy consumption. Live migration allows resources to be reallocated in order to turn off existing devices, thus allowing further energy savings.

The objective of this dissertation is to evaluate the possibility of using virtual networks to reduce the energy consumption of core networks by proposing energy efficient algorithms and resources reallocation. The dissertation formulates and extends the virtual network mapping problem with the objective of reducing the energy consumption in core networks. The dissertation also introduces techniques by evaluating the possibility of using live migration in order to further increase energy savings in core networks.

The work in this dissertation has been reported in the following publications:

- E. Rodriguez, G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca. Green Virtualized Networks. In Proceedings of the IEEE International Conference on Communications (ICC), pages 1–5, 2012. [42]
- Rodriguez E., Alkmim G. P., D. M. Batista, and N. L. S. Fonseca. Mapeamento energeticamente eficiente de redes virtuais em substratos fisicos. In Proceedings of the XXX Simposio Brasileiro de Redes de Computadores, pages 1–14, Ouro Preto, Minas Gerais, Brasil, 2012. SBC. [15]
- E. Rodriguez, G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca. Live Migration in Green Virtualized Networks. In Proceedings of the IEEE International Conference on Communications (ICC), pages 1–5, 2013. [43]
- E. Rodriguez, G. Alkmim, D.M. Batista, and N.L.S. da Fonseca. Trade-off between bandwidth and energy consumption minimization in virtual network mapping. In Communications (LATINCOM), 2012 IEEE Latin-America Conference on, pages 1–6, 2012. [41]
- E. Rodriguez, G. Prado Alkmim, D. Macedo Batista, and N.L. Saldanha da Fonseca. Trade-off between bandwidth and energy consumption minimization in virtual network mapping. Latin America Transactions, IEEE (Revista IEEE America Latina), 11(3):983–988, 2013. [44]

This dissertation is organized as follows: Section 2 explains the key concepts of this work. Section 3 summarizes related work. Section 4 presents the energy consumption model as well as the ILP formulation for Green Virtual Network mappings and the results. The algorithms for energy savings using live migration and the results are presented in Section 5. Finally, in Section 6, conclusions are drawn.

Chapter 2

Key concepts

2.1 Virtual Networks

The concept of virtualization is not new, as it originated in the 60's. Virtualization was used in IBM mainframes [30] [35]. In [37], the concept of virtual machines was formalized, which is an efficient and isolated duplicate of real machines. A duplicate is an exact identical copy of a machine, meaning, that any program executed in a virtual machine must exhibit the same output as if it was executed in a physical machine. The term isolation means that a virtual machine must work as an independent device [37]. This aspect allows a user to perceive that he is working on a real machine rather than on a virtual machine.

The use of virtualization allows several interesting benefits [12]:

- Energy savings: the use of servers into virtual machines can be translated into low energy and cooling costs, which implies in a low carbon footprint. This holds true because the use of virtualization translates into allocating several services that would traditionally be allocated on several physical machines on fewer machines.
- Reduction of overall footprint: this translated to fewer servers, networking equipment, racks, energy outlets and anything related to equipment in general and less floor space required in order to have an operational network.
- Experimental environments: allows the creation of experimental, self contained networks operating in an isolated way. This prevents conflicts between existing implementations, thus allowing experimentation without compromising stable environments as well as to build environments that would otherwise be costly prohibitive based only on pure physical implementations.

- Application isolation: virtualization allows the “one application/one server” model to be more efficient in terms of server waste, as it fully utilizes physical server resources. It allows the provisioning of VMs with the exact amount of CPU, memory and storage requirements, which can resolve application compatibility issues.

A virtual network is a computer network that includes virtual network links and virtual network routers. A virtual network link is a link that does not consist of a physical (wired or wireless) connection between two computing devices but is implemented using methods of network virtualization. One virtual link can be the concatenation of several physical links. A virtual router is a virtual machine that does not consist of a physical router. In virtual networks, virtual routers are virtual machines, which when interconnected with other virtual machines by virtual links make up a virtual network. Each virtual router has its own protocol stack, independently from the protocol stack of the physical network. This allows the execution of different network architectures over the same physical network.

In the context of the Internet of the Future, the coexistence of several virtual networks with different protocol stacks allows solutions for several of the existing issues related to Internet rigidities [32].

The need to create virtual networks arises from the current limitations of the current architecture of the Internet. In virtual network based architectures, it is possible to define two different layers: the physical layer and the virtual layer. In [17], an architecture called Cabo is proposed, in which the *Infrastructure Providers* (IP) are responsible for the control of the physical layer of the network, and the *Service Provider* (SP) are responsible for providing services to the network. In this way, SPs make virtual network request to the IPs which are responsible for allocating the resources to them.

In [53], a three layer architecture called Cabernet is proposed, which reduces the limitations of the implementation of services in a WAN. The main idea behind this is the creation of a network of connectivity between the services and the infrastructure layers. This layer uses the acquired virtual links from the infrastructure service providers to execute the virtual networks with the necessary support to provide the desired services. In this architecture, the resources of a virtual network can be provided from several and interconnected infrastructure providers.

In [52], the UFO *Underlay Fused with Overlays* project is presented. In this architecture, the underlay layer of the networks notifies the overlay layer about the conditions of the network to help improve the efficiency and scalability of the routing of the proposed networks.

In [22], the DaVinci architecture is introduced. This architecture periodically reallocates the resources of the virtual machines. Another characteristic of this architecture is that it realizes packet routing. It is supposed that all the virtual networks are within

the same institution and there is no interest in acquiring more resources, consequently reducing the performance of other virtual networks.

In [45], the authors define interactions between the business functions of the networks, without defining its implementation, since each entity must define these functions independently. Besides that, it is emphasized that a virtual network control not located in the virtual network is required. The *VNet Control Plane Architecture* is presented, which provides management control for the involved entities.

2.2 Virtual Network Mapping Problem

The virtual network mapping problem consists basically in finding an optimal mapping of virtual networks onto physical networks. For a given virtual network request, a solution must be presented within a reasonable amount of time. The solution to this problem must guarantee the mapping of the greatest possible amount of virtual routers within a reasonable time.

In a real life scenario, virtual network requests are not previously known, which increases the complexity of the problem. Even by knowing the topology of each virtual network request, the problem is an NP-Hard problem [23].

In [51], issues which makes this a challenging problem are stated. The first is the combination of restrictions involving nodes and links, which make the problem a lot harder to solve. Besides that, the adoption of admission control is required. The third aspect is the uncertainty factor of each request, since the lifetime of each virtual request is usually unknown. Finally, the diversity of existing topologies adds complexity to this problem.

The work in [3] considers a unique characteristic of the Virtual Network Mapping Problem. It considers that virtual networks require image repositories that contain sets of software and protocol stacks in order to instantiate virtual networks. These images are used to instantiate the virtual routers onto physical routers, as shown in Figure 2.1. The virtual network mapping algorithm requires the images to be transferred from the repository to the physical router before initiating each virtual network and must choose the images and the paths required to accomplish this.

2.3 Green Networking

Green Networking is defined as the set of heuristics, approaches and proposals which aim to reduce energy consumption in computer networks. It responds to several global necessities, which can be approached from an environmental, engineering and economical




Figure 2.1: Network Architecture: Virtual Network and physical network with software repository

point of view [11].

From an environmental point of view, it responds to the issue of reducing global warming by trying to reduce the gas emissions that cause the greenhouse effect. This can be approached by using renewable energy as much as possible. Another way to guarantee this is to design for lower energy consumption components while guaranteeing the same level of performance. Architecture redesign by de-locating network equipment can also be employed, as it allows equipment to be placed closer to energy sites, thus reducing energetic transfer costs. It is also possible to place sites on natural cooling sites, such as rivers or lower temperature sites. Services can be reallocated at optimal places, according to daily load. For example, operations can be executed in a way that allows processing to follow seasonal or day/night patterns. It is noticeable that there is a challenge in performing such services migration without any service disruption, preserving fault-tolerance and data security.

Energy prices in the market tend to be volatile. From an economical point of view, computational displacement can be exploited on energy sources that, at a given time, can be low for any reason, such as energetic surpluses caused by weather.

From an engineering point of view, it should reduce the energy from networks required to carry out a given task while maintaining the same level of performance. This point of view alone is still relevant as system efficiency as a whole.

The problem with current network design is that over-positioning and redundancy are employed extensively. Over-positioning address the worst-case scenarios, typically in peak hours. These resources normally remain active even under low loads.

In order to minimize energy consumption, the following types of strategies are em-

ployed:

- Resources consolidation: consists in rearranging a set of resources in a network to reduce the global consumption due to devices underutilized at a given time. Traffic levels follow specific daily and weekly patterns, and this allows the adaptation of the networking settings, for example, turning off and rerouting traffic during idle times.
- Selective connectivity: consists on distributed mechanisms that allow single equipment to go idle for some time as transparently as possible for the rest of the network. Selective connectivity permits the turning on and off of resources at certain sections of the network. Nodes could go idle to prevent processing unnecessary network traffic (heartbeats, broadcasts, etc). These tasks can be temporarily assumed by other nodes by “faking” their identity for a short period of time.
- Virtualization: there are several reasons that virtualization should be used in order to save energy. The crucial point to take in consideration is that virtualization allows the utilization of existing powered on devices, which prevents the cost of powering a whole new device. Another important issue to consider is cooling. For any given resource that can be placed on a current turned on device, the cooling cost for that device will be automatically transferred as part of the current turned on device cooling cost.
- Proportional computing: proportional computing consists on efforts to make energy consumption of devices to work as a function of their workload. Ideally, it is desirable for a device energy consumption to be totally dependent on its workload. However, in many cases, the energy consumption of devices is energy-agnostic, that is, regardless of their workload, their energy consumption is always the worst. Between these two cases, there is an infinite amount of intermediate profiles, in which the energy consumption adapts to the device workload.

2.4 The Green Virtual Network Mapping Problem

The Green Virtual Network Mapping Problem can be stated as an extension of the Virtual Network Mapping Problem. Traditionally, the Virtual Network Mapping Problem has as objective the use of a set of routers and links that minimize the amount of bandwidth required to satisfy the virtual request. Rather than attempting to minimize bandwidth, the Green Virtual Network Mapping Problem will attempt to find the virtual network mapping that issues the lowest possible energy consumption.

The biggest determining factor in order to make an energy efficient virtual network mapping is the re utilization factor. An energy efficient virtual network mapping should initially take into consideration nodes and links that are already powered on, and try to utilize them as much as possible.

Other factors that could affect Green Virtual Network Mappings are cooling factors and alternate sources of energy. One could, for example, add an additional cost for cooling on each site. It is possible also to add a percentage reduction to determined sites that may depend partially on alternative sources of energy, such as wind and solar power.

The complexity of the Virtual Network Mapping problem applies to this problem as well. It is to be expected that the problem is even more complex than the virtual network mapping problem, due to the energy consumption constrains that must be placed. Furthermore, all the link and node restrictions that were mentioned in Section 2.2 must be present. It must be taken to consideration that no elements can be allocated on powered off devices. The energy consumption of several devices must be taken into consideration as well.

2.5 Live Migration

Live migration is defined as the process of transferring a virtual machine or application currently in execution to another physical machine, without requiring the interruption of the service itself. Memory, storage, and network connectivity of the virtual machine are transferred from the original host machine to its destination.

There are several reasons that make the migration of a service from one host to another desirable. Live Migration can alleviate the workload of several servers, thus ensuring load balancing. It also allows for separation of hardware and software considerations, and consolidating clustered hardware into a single coherent management domain. Live Migration can also be used as an approach to save energy.

It is possible to generalize the transferring of the state of the memory from one machine to another by summarizing the three phases [10]:

- Push phase: while the source VM is in execution, certain pages are sent across the network to the new destination. In order to guarantee consistency, pages modified during this process must be transferred again.
- Stop-and-copy phase: the source VM is halted. During this period, pages are transferred to the destination VM. The new VM is started once again.
- Pull phase: the new VM executes and if it accesses a page that has not yet been copied, this page is faulted in (pulled) across the network from the source VM.

In practice, most migration solutions will implement one or two of these phases.

Live Migration has been explored as a tentative way to achieve energy savings in virtual networks. It is possible to move services to sites that offer lower energy costs in a certain period of the day. It is also possible to rearrange a network configuration to fill in gaps created by virtual network terminated processes to allow great energy savings. This specific scenario will be approached with more detail in Chapter 6.

Chapter 3

Related Work

This chapter includes a literature review of the works related to green virtual network mapping. Section 3.1 presents the work related to greening of virtual networks and Section 3.2 includes the work related to live migration.

3.1 Green Virtual Networks

The seminal work in [21] opened avenues for several others in the literature. This work discusses the issue of energy consumption in networks. The work establishes that the energy consumption of the Internet is indeed something that must be taken to consideration and explores the possibility of putting devices to sleep in order to attain energy savings.

The idea of using virtualization as a technique for the reduction of energy consumption is mentioned in [11] as one of the strategies used for Green Networking. This work attempts to formalize the different Green Networking strategies into several categories.

A summary of energy consumption elements in IP over WDM is presented in [18]. The energy consumption of a core router chassis, line cards and processor are used as well as the values of a physical link amplifier. There is an extensive data collection of energy consumption elements in this work, which is used in the present work.

A Mixed Integer Linear Programming (MILP) model for the design of energy-efficient WDM network was introduced in [46]. This model attempts to minimize the number of physical components that are used on the network as well as the energy consumption of the network. It considers an array of physical elements, and traffic aggregation. The energy consumption value of linecards is computed as the consumption of a chassis distributed by the number of linecards, and this is used to calculate the energy consumption of the physical links in the present work.

There are several aspects in the work in [46] that are not taken into consideration in this dissertation. The cited work considers the IP router ports as a measure of the total

energy consumption of an IP router. This differs from the present work, that breaks down the energy consumption of the physical router into several elements.

The work in [7] considers a wide area network, and evaluates the possibility of turning off elements under QoS constraints. The objective of doing this is to minimize the energy consumption of the network. It considers the value of the physical links to be much lower than the power of the nodes, and thus discriminates these values. The work in [7] does not consider virtual networks explicitly.

The work in [3] approaches in detail the virtual network mapping problem. This work makes an extensive background collection of virtual networking technologies and algorithms, and proposes several algorithms. It is used as a basis for this dissertation.

The works in [6] and [19] consider the virtual network mapping problem in terms of energy efficiency. The work in [6] considers the trade-off between energy consumption and bandwidth consumption. There are several characteristics not considered on both works that are included in this dissertation, such as locality constraints, memory and operation system images.

3.2 Live Migration

The work in [10] was one of the first to introduce the concept of migrating operating systems under “liveness constraints”. Live Migration draws a clean separation between hardware and software facilitates, and facilitates fault management, load balancing, and low-level system maintenance. The work in [48] proposes the use of Xen for virtual machines migration when machines are distant apart. Such migration can be usually achieved within a downtime of 1-2 seconds. The article shows that some of the transfer times involved in the live migration across large geographical distances can be negligible. The work in [31] models the virtual machine transfer of virtual machines, and proposes an algorithm to improve the downtimes of virtual machines.

The work in [50] introduces a primitive, called VROOM, that “avoids unnecessary changes to the logical topology by allowing (virtual) routers to freely move from one physical node to another” and it indicates that live migration can be pursued to reduce energy consumption. However, the authors did not show any solution for that.

The authors in [49] quantify the effect of live migration in the Internet using a specific application. This work also indicates some other advantages of live migration, such as improved manageability, performance and fault tolerance. The authors suggest the adoption of load balancing and server maintenance after migrating their workload to other servers.

The work in [36] compares the Xen and OpenFlow live migration approaches. It shows that live migration can be employed to realize green networking strategy. However, no concrete proposal to reduce the energy consumption was presented.

The authors in [3] presented a virtual network mapping solution, considering a set of main characteristics of real networks. They extend their work in [42] by the introduction of the objective of minimizing the energy consumption. However, none of these works considered live migration.

The work in [28] explores the impact of live migration in data centers. Assumptions are made to reduce the energy consumption during migration, however, the work does not introduce any technique for such reduction.

The works in [33] and [40] presents algorithms to reduce the energy consumption in clouds using migration. [33] presents a solution, named SERCON, which also tries to minimize the amount of migrations done. The work in [40] presents a dynamic threshold based scheme, which rearranges virtual machines depending on their utilization. Both work shows that migrations can be performed to reduce energy consumption in networks, though it is focused on data centers.

3.3 Comparison of related work

Table 3.1 compares several works with that in the present dissertation which is unique since it considers: sets of images with different sizes, the time required to instantiate virtual routers, the locations of the repository in which images are stored and the available memory of the physical routers, reduction of energy consumption as well as live migration to further minimize energy consumption. This dissertation significantly improves the state of the art of green virtual network mappings onto substrate networks, since they provide a more realistic assessment of operational networks.

This dissertation does not impose any alignment constraints between virtual topologies and physical topologies. It is possible that the topology of physical routers and links allocated to a virtual network will be the same of that of the requested virtual network, but this happens only if the topology is the one which minimizes the bandwidth allocated. Moreover, such an alignment is not necessary to guarantee the QoS requirements of the applications, which are indeed assured by the constraints of the mapping problem.

Table 3.1: Comparison between mapping algorithms

Reference	Virtual Network Mappings	Green Networking	Live Migration	Joint bandwidth/power reduction
[47]	no	yes	no	no
[51]	yes	yes	no	no
[8]	yes	yes	yes	no
[23]	yes	yes	no	no
[29]	no	yes	no	no
[46]	yes	yes	no	no
[7]	no	yes	no	no
[6]	yes	yes	no	yes
[3]	yes	no	no	no
[19]	yes	yes	no	no
[33]	yes	yes	yes	no
[40]	no	yes	yes	no
This Dissertation	yes	yes	yes	yes

Referência	Processor Cores	Bandwidth	Locality Constrains	Images for Instantiating
[47]	no	no	no	no
[51]	no	no	no	no
[8]	yes	no	no	no
[23]	no	no	no	no
[29]	no	no	no	no
[46]	yes	yes	no	no
[7]	no	yes	no	no
[6]	no	yes	no	no
[3]	yes	yes	yes	yes
[19]	no	yes	no	no
[33]	no	yes	no	no
[40]	no	no	no	no
This Dissertation	yes	yes	yes	yes

Referência	Link Delay	Available memory and image size	Repository location of images	Mapping Finalization Time
[47]	no	no	no	no
[51]	no	no	no	no
[8]	yes	no	no	no
[23]	no	no	no	no
[29]	no	no	no	no
[46]	no	no	no	no
[7]	no	no	no	no
[6]	no	no	no	no
[3]	yes	yes	yes	yes
[19]	no	no	no	yes
[33]	no	no	no	no
[40]	no	no	no	no
This Dissertation	yes	yes	yes	yes

Chapter 4

Green Virtualized Networks

4.1 Motivation

Virtualization can be employed as a technique used to reduce energy consumption in networks [3] since it can replace a large amount of physical elements, by replacing them by virtual counterparts. In the design of networks, minimizing bandwidth exclusively does not necessarily provide acceptable energy savings [3]. To illustrate this, a simple example is provided. The physical layout of this experiment as well as the virtual network that will be mapped is shown in Figure 4.1. All links have the same energy consumption, except the link (0,2) that has a significant low energy cost. Its energy consumption is later increased by steps of 250W, to the point that its energy consumption exceeds the one of the neighboring links. Two algorithms are executed on this scenario: a bandwidth minimization heuristic (BAND) and the green minimization heuristic (GREEN) proposed in this dissertation. All physical routers have two cores available, and the virtual routers request two cores each.

Figure 4.2 plots the energy consumption of the network when the allocation derived by each heuristics is considered. The consumption is plotted as a function of the energy cost of the link connecting the real routers 0 and 2. The energy saving when using the GREEN algorithm shows that when the cost of the link increases, the algorithm will not choose this link, thus choosing the lowest-cost links. The original algorithm will always allocate the same physical link, regardless of its energy consumption. This result highlights the motivation of this work, which is to show that it is possible to reduce the energy consumption on virtual network mapping.

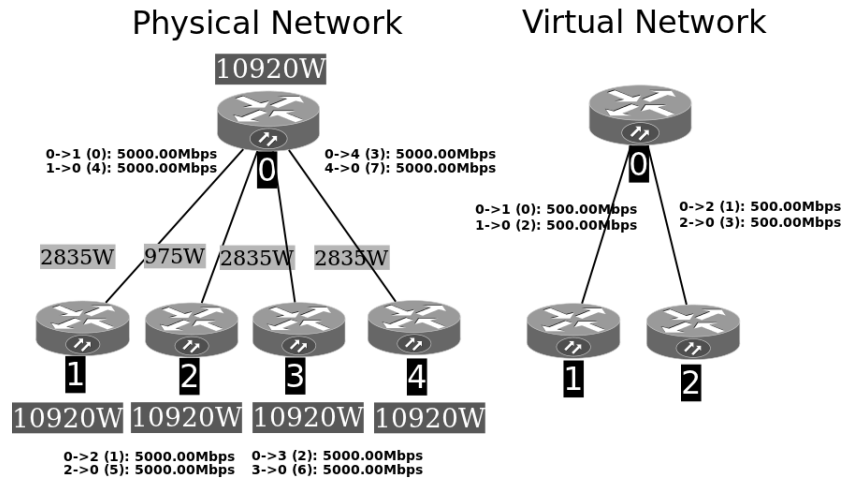


Figure 4.1: simple simulated scenario




Figure 4.2: Energy consumption per alternating link power

4.2 Energy Consumption Model

In order to build a virtual network mapping onto a physical substrate, an energy consumption model that should account the consumption of the physical components of the network is required. The energy consumption of the physical elements of the network can be divided into the components of the physical router and the components of the physical

links.

4.2.1 Physical Router

The energy consumption of a router can be divided into two components: traffic dependent and traffic independent components. The traffic independent portion typically represents around 90% of a router energy consumption. Three major components contribute to a router's energy consumption: the router chassis, the router processor and its linecards [9] [2] [18].

The router chassis corresponds to the case that is used to hold the rest of the components of the router. This box accounts for an overwhelming portion of the router energy consumption. All other components of the router require to be sheltered in the chassis, as it will provide cooling and protection. The energy consumption of the chassis is considered to be traffic-independent and constant. It is, thus, crucial to minimize the amount of chassis to be turned "on" in the network. The chassis should be turned off only when no processors are active. The cost of all the other electronic components of the router, such as the motherboard and its physical memory are taken into consideration jointly with the cost of the chassis.

The linecard is the slot required in order to link a physical interface to the router. The energy consumption of the linecard is considerably lower than that of the chassis, but this consumption adds up to the total energy consumption of the network. Also, the amount of linecards that will be used will determine how many physical links will be used in the network. The energy consumption cost of the linecards is considered to be traffic dependant, as more traffic demands will result in requiring more slots to satisfy the bandwidth demands of the network.

The router processors lie within one of the slots of the chassis. It is possible to have more than one processor array in a physical router. However, in this model, it is assumed that each router has one physical processor slot. Each slot is composed of several processors or cores, which are used to allocate different virtual routers in order to satisfy their processing demands. Regardless of the amount of processors that are active at a certain time, the chassis also needs to be active.

Figure 4.3 shows how each of the components in the physical router interacts with each other. The linecards and the processors lie within the router chassis, which encapsulates all the other components. The linecards serve as the components that connect the physical links of the network. Each processor slot, which contains several processors, lies within a chassis slot and serves as an extension of the router motherboard.

The incorporation of these three components into the energy model differentiates it from other models in the literature which assume a constant value to the energy con-




Figure 4.3: Energy consumption components

sumption of a router.

4.2.2 Physical Links

The model adopted in [46] [25] for the energy consumption of the physical link is employed in this paper. In this model, the energy cost of a physical link depends on the length of the link, that gives the amount of amplifiers required along the link and on the number of used wavelengths. A span distance \mathcal{P}^s determines that every \mathcal{P}^s kilometers (km) a new amplifier is required to properly propagate the signal. Equations 1 and 2 give the energy cost of a physical link.

$$\mathcal{P}_{u,v}^L = \mathcal{P}_{u,v}^{\#A} \times \mathcal{P}^{CA} \quad (4.1)$$

$$\mathcal{P}_{u,v}^{\#A} = \lceil \frac{J_{u,v}}{\mathcal{P}^s} - 1 \rceil + 2 \quad (4.2)$$

where:

- $\mathcal{P}_{u,v}^L$ gives the power consumed by the physical link (u, v) ;
- $\mathcal{P}_{u,v}^{\#A}$ gives the amount of amplifiers of link (u, v) ;

- \mathcal{P}^{CA} gives the energy cost of each amplifier;
- $J_{u,v}$ gives the length of link (u, v) ;

In this work, it is assumed that the links have a single fiber, and that two amplifiers are added at the endpoints of the link, since signal amplification is needed (hence the value of 2 added at equation 4.2).

4.3 Proposed Mathematical Model

The formulation in this paper models the dynamic arrival of requests for virtual network establishment on network substrates. Each request specifies the topology of the virtual network, the resource demanded by the virtual network elements, and the QoS requirements which include a bound on the time to instantiate the VNs and the location constrains to instantiate the nodes of the VN.

The model proposed takes the request for the establishment of the virtual network and tries to select which elements of the network substrate should be allocated to instantiate the virtual network. The selection criteria aim at minimizing the energy consumption and yet satisfying the requirements of the request.

The proposal is based on two 0-1 Integer Linear Programming (ILP) sub-models. The models consider realistic parameters, such as the existence of router images located in repositories. The formulation differs from that in [3] by the introduction of a two step formulation which reduces the computational complexity by reducing the memory consumption. The following notation is used in the formulation of the problem:

- N is the set of physical routers;
- F is the set of physical links, the physical link (n_1, n_2) connects the physical routers n_1 and $n_2 \in N$;
- M is the set of virtual routers;
- V is the set of virtual links, the virtual link (m_1, m_2) connects the virtual routers m_1 and $m_2 \in M$;
- I is the set of images stored in the repository. Each image corresponds to a file with an operating system and a specific set of software ready to be instantiated in a physical router;
- A is the number of available cores in the physical routers; $A(n)$ gives the number of cores of router n ;

- P is the set of the number of cores requested by the virtual routers; $P(m)$ gives the number of cores required by the virtual router m ;
- C is the set of values of available bandwidth in the physical links; $C(f)$, $f \in F$, gives the available bandwidth in the link f ;
- Q is the set of bandwidth values requested by the virtual links; $Q(v)$, $v \in V$, gives the bandwidth required by the virtual link v ;
- D is the set of values of delays in the physical links; $D(f)$, $f \in F$, gives the delay in the link f ;
- K is the set of values of maximum delay allowed on a virtual link; $K(v)$, $v \in V$, represents the maximum delay allowed in the virtual link v ;
- $L_{n,m}$ defines the restrictions related to the location of the physical routers. The value of the variable is 1 if the virtual router m can be mapped onto the physical router n . Otherwise, it is 0. This variable is useful for imposing policy restrictions related to the location of the physical routers. If a user does not want a virtual router m to be mapped onto a physical router, the variables of $L_{n,m}$ must be 0;
- $R_{n,i}$ provides details about the location where images are stored. If the image i is located in a repository with a direct link to the physical router n , the value of the variable is 1. Otherwise, it is 0;
- $E_{m,i}$ is related to software restrictions. If the image i contains all the software requirements (operating system, protocol stacks and kernel modules) required by the virtual router m , the value of the variable is 1. Otherwise, it is 0;
- B is the set of values of the available memory in the physical routers; $B(n)$ represents the memory available in the router n ;
- G is the set of image sizes; $G(i)$ represents the size of the image i ;
- S is the time threshold for instantiating the virtual network;
- $T_{n,i}$ represents the time the physical router n takes to boot the image i ;

The values of $\mathcal{P}^{chassis}$, \mathcal{P}^L , \mathcal{P}^{card} and \mathcal{P}^{core} are used in the constraints related to express the energy consumption of the network of chassis, physical link, linecards and cores, respectively.

The inclusion of D , K , $L_{n,m}$, $E_{m,i}$, B , G , S , $T_{n,i}$, I and $R_{n,i}$ in the formulation makes this work unique and realistic due to the consideration of diverse aspects of real networks.

The maximum delay allowed in the network links (D, K) affects the QoS furnished to applications sensitive to the delay, specially those involving video and audio. The specific image required by a virtual router should be defined and the content of each repository must be known ($R_{n,i}$) to determine from which repository the image should be downloaded ($I, E_{m,i}$). Locality restrictions and the size of the images should be known since routers have limited storage capacity (B, G) and the size of the image impacts the download time. Moreover, users can have policy issues that prevent the utilization of some physical routers ($L_{n,m}$) or can restrict the solution to employ energy efficient sites. Moreover, the maximum time acceptable to the instantiation of the virtual network is related to the urgency of virtual networks and service prioritization ($S, D, K, T_{n,i}$).

In the formulations, the two following variables define the state of occupancy of the network substrate:

- K_n denotes the number of cores allocated in the physical router n
- $O_{u,v}$ denotes the number of virtual links that use the physical link u, v

The values of K_n and $O_{u,v}$ are used in the computation of α_n and $\beta_{u,v}$, which simplify the objective function:

$$\alpha_n = \lceil \frac{K_n}{K_n + 1} \rceil \quad (4.3)$$

$$\beta_{u,v} = \lceil \frac{O_{u,v}}{O_{u,v} + 1} \rceil \quad (4.4)$$

The values of α_n and $\beta_{u,v}$ determine, respectively, whether or not a router and a physical link are already in use.

The solution of the problem is given by the binary variables:

- $X_{n,m,i}$: its value is 1 if the virtual router m is mapped onto the physical router n using the image i ; otherwise, it is 0;
- $Y_{u,v,w}$: its value is 1 if the physical path used by the virtual link w includes the physical link (u, v) ; otherwise, it is 0;
- $Z_{u,v,m}$: its value is 1 if the physical path (u, v) is used to transfer the image requested by the virtual router m ; otherwise, it is 0.
- U_n : its value is 1 if the physical router (n) is to be powered on; otherwise, it is 0.
- $W_{u,v}$: its value is 1 if the physical path (u, v) is to be powered on; otherwise, it is 0.

The mapping of the virtual networks is based on the sequential execution of two ILPs. The first (ILP-Green-Mapping) maps the virtual networks onto the substrate. The second (ILP-Green-Image) determines the path in the substrate used to transfer the images.

The ILP-Green-Mapping is formulated as follows:

$$\begin{aligned}
& \text{Minimize} \\
& \mathcal{P}^{chassis} \sum_{n \in N} (\alpha_n + (1 - \alpha_n)U_n) + \\
& \mathcal{P}^{core} \sum_{n \in N} \sum_{m \in M} \sum_{i \in I} (X_{n,m,i} \times P(m)) + \\
& (2\mathcal{P}_{u,v}^{card} + \mathcal{P}_{u,v}^L) \sum_{(u,v) \in F} (\beta_{u,v} + (1 - \beta_{u,v})W_{u,v})
\end{aligned}$$

subject to

$$\begin{aligned}
& \sum_{n \in N} \sum_{i \in I} X_{n,m,i} = 1 & (C1) \\
& \forall m \in M
\end{aligned}$$

$$\begin{aligned}
& \sum_{m \in M} \sum_{i \in I} X_{n,m,i} \leq 1 & (C2) \\
& \forall n \in N
\end{aligned}$$

$$\begin{aligned}
& \sum_{m \in M} \sum_{i \in I} P(m) \times X_{n,m,i} \leq A(n) & (C3) \\
& \forall n \in N
\end{aligned}$$

$$\begin{aligned}
& X_{n,m,i} = 0 & (C4) \\
& \forall n \in N, \forall m \in M, \forall i \in I | L_{n,m} = 0 \text{ or } E_{m,i} = 0
\end{aligned}$$

$$\begin{aligned}
& \sum_{w' \in V} Y_{u,v,w'} \times Q(w') \leq C(w) & (C5) \\
& \forall w = (u, v) \in F
\end{aligned}$$

$$\begin{aligned}
& \sum_{u \in N} \sum_{v \in N} Y_{u,v,w} \times D(u, v) \leq K(w) & (C6) \\
& \forall w \in V, (u, v) \in F
\end{aligned}$$

$$\sum_{m \in M} \sum_{i \in I} X_{n,m,i} \times G(i) \leq B(n) \quad (C7)$$

$$\forall n \in N$$

$$Y_{u,v,w} = 0 \quad (C8)$$

$$\forall u, \forall v, \forall w \in V | (u, v) \notin F$$

$$\sum_{f \in N} Y_{n,f,w} - \sum_{f \in N} Y_{f,n,w} = \quad (C9)$$

$$\sum_{i \in I} X_{n,a,i} - \sum_{i \in I} X_{n,b,i}$$

$$\forall w = (a, b) \in V, \forall n \in N$$

$$X_{n,m,i} \leq U_n \quad (C10)$$

$$\forall n \in N, \forall m \in M, \forall i \in I$$

$$U_n \leq \sum_{m \in M} \sum_{i \in I} X_{n,m,i} \quad (C11)$$

$$\forall n \in N$$

$$Y_{u,v,w} \leq W_{u,v} \quad (C12)$$

$$\forall v \in V, \forall (u, v) \in F$$

$$W_{u,v} \leq \sum_{v \in V} Y_{u,v,w} \quad (C13)$$

$$\forall (u, v) \in F$$

$$X_{n,m,i} \in \{0, 1\} \quad (C14)$$

$$\forall n \in N, \forall m \in M, \forall i \in I$$

$$Y_{u,v,w} \in \{0, 1\} \quad (C15)$$

$$\forall u, \forall v, \forall w \in V$$

$$U_n \in \{0, 1\} \quad (C16)$$

$$\forall n \in N$$

$$W_{u,v} \in \{0, 1\} \quad (C17)$$

$$\forall u, \forall v \in V$$

The objective function minimizes the energy consumed by the request.

Constraint (C1) establishes that a virtual router is assigned to a single physical router and that a single image is used to instantiate it. Constraint (C2) limits the amount of virtual routers that can be allocated to a physical router per request. Only one virtual router can be allocated to a physical router per request. Constraint (C9) ensures that the set of physical links that composes a virtual link is a valid path. It compares the in-degree and the out-degree of each physical router n .

Constraints (C3) and (C7) express the limitations of the physical routers. They ensure that each physical router will not allocate more than its maximum capacity of cores and memory, respectively.

Constraint (C4) guarantees that the virtual routers will be instantiated only using images that meet its software requirements as well as geographic location.

Constraints (C5) and (C6) express the limitations of the physical links. Constraint (C5) ensures that the bandwidth available in each physical link is greater than the bandwidth requirements of all virtual links using it. Constraint (C6) establishes that the total delay in the physical path allocated to a virtual link should not exceed the delay threshold allowed for that virtual link.

Constraint (C8) guarantees that if (u, v) does not correspond to a physical link, it will never be used in the mapping.

Constraints (C10) and (C11) express the energy constraints of the physical routers. Constraint (C10) ensures that no core can be assigned to a given router without turning on the device first. Constraint (C11) ensures that if the router is powered on, then at least one core needs to be assigned to that router.

Constraints (C12) and (C13) express the energy constraints of the physical links. Constraint (C12) ensures that a virtual link can be used on the physical link (u, v) only if the physical link is powered on. Constraint (C13) ensures that if the link is powered on, then at least one virtual link needs to be assigned to that physical link.

Constraints (C14), (C15), (C16) and (C17) define the domains of the binary variables.

After the solution of the ILP-Green-Mapping is found, the values of $X_{n,m,i}$ e $Y_{u,v,w}$ can be used as input to the second formulation, the ILP-Green-Image.

The ILP-Green-Image is formulated as follows:

$$\begin{aligned} & \text{Minimize } \sum_{m \in M} \sum_{u \in N} \sum_{v \in N | (u,v) \in F} Z_{u,v,m} \times D(u,v) \\ & + \frac{Z_{u,v,m} \times G(i | X_{n,m,i} = 1)}{C(u,v)} \text{ subject to} \end{aligned}$$

$$\begin{aligned} & \sum_{m \in M} Z_{u,v,m} = 0 \\ & \forall u, \forall v | (u,v) \notin F \end{aligned} \tag{C18}$$

$$\sum_{j \in N} Z_{u,j,m} - \sum_{j \in N} Z_{j,u,m} = \tag{C19}$$

$$\begin{aligned} & X_{n,m,i} \times R_{u,i} - X_{n,m,i} \times (1 - \lceil \frac{|u-n|}{\alpha} \rceil) \\ & \forall m \in M, \forall i \in I, \forall n, u \in N, \alpha = |N| \end{aligned}$$

$$\begin{aligned} & Z_{u,v,m} \in \{0, 1\} \\ & \forall u, \forall v, \forall m \in M \end{aligned} \tag{C20}$$

The objective function minimizes the time required to instantiate the virtual network. The time needed to instantiate each virtual router is the sum of the time required to transfer the image and to boot the operating system assuming that two or more images can be transferred simultaneously in the same physical link.

Constraint (C18) guarantees that (u, v) will not be used if it does not belong to the considered substrate. Constraint (C19) establishes that the set of physical links allocated to transfer an image consists of a valid path in the substrate network. Constraint (C20) defines the domains of the variables.

4.4 Proposed Algorithms

4.4.1 Virtual Network Mapping Algorithms

OPT

The formulation presented in 4.3 was implemented exactly as shown using the *IBM ILOG CPLEX Optimizer* [24]. The implementation is presented in Algorithm (ALG). The ILP is solved using *CPLEX*'s dynamic search algorithm. It is based on the *Branch & Cut* [20] algorithm.

The *ILP-GREEN-Mapping* searches through all the nodes of the search tree and returns a solution that minimizes the total bandwidth. The *ILP-GREEN-Image* also

searches thru all the nodes of the search tree and minimizes the instantiation time of the virtual network. The Opt algorithm to solve the ILP formulations is presented in 1.

Algorithm 1: Opt Algorithm

Data: Substrate γ with attributes α , virtual network δ com attributes β .

Result: Mapping of δ in γ with the paths of the physical network rede física θ used to transfer images.

```

1 Define  $\gamma$ ,  $\alpha$ ,  $\delta$  and  $\beta$  as parametros for ILP-GREEN-Mapping;
2 Travel the entire search tree in search of ILP-GREEN-Mapping and obtain the
  values of variables  $X_{n,m,i}$  and  $Y_{n,u,w}$  related to the best solution found;
3 if ILP-Mapping doesn't find a solution then
4   | Block the request;
5 end if
6 else
7   | Define  $\gamma$ ,  $\alpha$ ,  $\delta$ ,  $\beta$  and variable  $X_{n,m,i}$  as parameters for ILP-Image;
8   | Travel the entire search tree in search of ILP-Mapping and obtain the values of
  variables  $Z_{n,u,m}$  related to the best solution found;
9   if ILP-Image doesn't find a solution then
10  | Block the request;
11  end if
12  else
13  | Return the mapping of  $\delta$  in  $\gamma$  using the values of variables  $X_{n,m,i}$  e  $Y_{n,u,w}$ ;
14  | Return the paths  $\theta$  using the values of variable  $Z_{n,u,m}$ .
15  end if
16 end if

```

The Opt algorithm solves the ILP formulations (Lines 2 and 8), parameterizing the virtual network attributes as well as the physical substrate (Lines 1 and 7), and returns the virtual router and virtual link mappings (Lines 13 and 14) over the physical substrate. If no feasible solution is found, the virtual network request is rejected (Lines 3, 4, 9 and 10).

ROOT

Initial experiments with the algorithm show that the time it takes to find a solution for substrates over size 20 nodes is long. Experiments with substrate sizes of 25 physical nodes can take more than two days to execute. This motivated the implementation of a heuristic denominated *Root*. This algorithm stops the search at the root of the search tree, considerably lowering its execution time when compared to the optimal implementation of the problem. This criteria was derived by empiric observations that feasible solutions to the problems were attainable from the root of the search tree.

The *Root* algorithm differs from the implementation of the optimal implementation by Lines 2 and 8, which are respectively replaced by:

- **Line 2:** Stop the search of solutions for *ILP-GREEN-Mapping* at the root of the search tree and for the values of variables $X_{n,m,i}$ and $Y_{n,u,w}$;
- **Line 8:** Stop the search of solutions for *ILP-Image* at the root of the search tree and for the values of variable $Z_{n,u,m}$.

4.5 Configuration of the Experiments

To evaluate the performance of the proposed formulation, another optimization model was developed containing similar constraints, but a different objective function: the minimization of the allocated bandwidth. In the remainder of this dissertation, the proposed formulation is denoted as GREEN, while the formulation which minimizes the allocated bandwidth is denoted as BAND. Details of the BAND optimization model can be seen in [3]. The formulations were evaluated in dynamic scenarios, in which the availability of resources in the substrate network varies as a function of time.

The average energy consumption per request, the amount of bandwidth allocated per request and the substrate utilization for physical nodes and physical links were evaluated as a function of the substrate sizes (number of physical routers) and as a function of the inter arrival time of requests.

To assess the effectiveness of the proposed formulation, a simulator was implemented in C++. This simulator receives a description of the substrate network as input. Confidence intervals of 95% were derived using the independent replication method. The ILP formulations were implemented using the CPLEX optimization library version 12.0. The simulations were executed on a computer running the operating system Debian GNU/Linux Squeeze. The computer was equipped with two Intel Xeon 2.27GHz processors, with 6 cores each one, and 6GB of RAM.

4.6 Results

The results for the experiments are shown below. Two different types of experiments were performed: one that evaluated the growth of the network by altering the size of the network, and another that evaluated the load of the network by altering the average inter arrival time.

4.6.1 GREEN Experiment 1

Parameters

The parameters used for this experiment are shown in Table 4.1. For this experiment, the average arrival time between requests was fixed to 25 seconds. The substrate sizes were varied in steps of 20 nodes, from 20 to 300 nodes.

Table 4.1: Values of the parameters used in GREEN Experiment 1

Parameter	Value
Number of physical nodes	{20 40 60 ... 300}
Bandwidth of each physical link	10240Mbps
Number of images in the network	3
simulation time	5000s
Average arrival time per request	{ 50 }s
Average duration per request	1250s
Number of virtual nodes per request	4
Bandwidth of each virtual link	1024Mbps
Maximum time required to instantiate the network	100s
RAM memory	768MB
Image size	128MB
Cores per physical router	6
Cores per virtual router	6
Physical link delay	Defined by BRITE
Virtual link delay	15 × value defined by BRITE
Time required to process the image	10s

Results

Figure 4.4 shows the average energy consumption per request as a function of the number of nodes in the substrate for an average interval time of requests of 25 seconds. The average energy consumption increases for both formulations since a higher number of nodes in the substrate implies that the mapping tends to use a higher number of nodes. The energy consumption of the chassis is quite high compared to other elements, so the impact on the consumption per request increases since there is a higher number of active nodes. For substrates with more than 120 nodes, the GREEN model tends to allocate cores in routers already active (powered on) and the energy consumption per request stabilizes. However, for the BAND model which does not tend to utilize active routers, the consumption per request tends to increase due to a higher number of routers active. For substrates with 300 nodes, GREEN energy reduction was of the order of 50%. Figure 4.5 shows the average allocated bandwidth per request as a function of the number of nodes of the substrate. As expected, the bandwidth consumption per request is higher for the GREEN model since it tends to use routers already powered regardless of the location of these routers which can imply in long paths between nodes. The bandwidth consumption under the GREEN model is more than that consumed by BAND, but it is

important to state that the GREEN model does not violate the QoS requirements of the requests, since the constraints defined in the formulation guarantee the support of these requirements.




Figure 4.4: Average Energy Consumption per request as a function of the number of nodes of the substrate

4.6.2 GREEN Experiment 2

Parameters

The parameters used for this experiment are shown in Table 4.2. For this experiment, the physical substrate sizes were fixed to 30 and 150 nodes. The virtual network arrival times were varied from 25 to 300 seconds.

Results

Figures 4.6 and 4.7 show the average energy consumption per request as a function of the inter arrival time of requests for substrates with 30 and 150 nodes respectively. As the inter arrival time decreases, i.e. the arrival rate increases, the energy consumption tends to increase since a higher number of routers become active. However, the increase is much sharper to BAND than to GREEN. The difference is quite significant and the consumption of BAND can be almost twice for low loads and substrates of 30 nodes and




Figure 4.5: Average Allocated Bandwidth per request as a function of the number of nodes of the substrate (GREEN Experiment 1)

Table 4.2: Values of the parameters used in GREEN Experiment 2

Parameter	Value
Number of physical nodes	{30 150}
Bandwidth of each physical link	10240Mbps
Number of images in the network	3
simulation time	5000s
Average arrival time per request	{25 50 75 100 125 150 175 200 225 250 275 300}s
Average duration per request	1250s
Number of virtual nodes per request	4
Bandwidth of each virtual link	1024Mbps
Maximum time required to instantiate the network	100s
RAM memory	768MB
Image size	128MB
Cores per physical router	6
Cores per virtual router	6
Physical link delay	Defined by BRITE
Virtual link delay	15 × value defined by BRITE
Time required to process the image	10s

up to 1.5 more for substrates of 150 nodes. However, GREEN consumption of bandwidth (figured not shown due to space limitations) is quite longer. It consumes 3, 5 and 4 times more bandwidth for substrates with 30 and 150 nodes. Moreover, as mentioned before, the GREEN model supports the QoS requirements of the requests.




Figure 4.6: Average Energy Consumption (Second set of experiments of a duration of the inter arrival time of request for a substrate with 30 nodes)

4.6.3 GREEN Experiment 3

Parameters

The parameters used for this experiment are shown in Table 4.3. For this experiment, the average inter arrival time of requests was fixed to 50 seconds. The substrate sizes were varied in steps of 25 nodes, from 25 to 400 nodes.

Results

Figure 4.8 shows the average energy consumption per request as a function of the number of nodes in the substrate for an average interval time of requests of 25 seconds. The average energy consumption increases for both formulations since a higher number of nodes in the substrate implies that mapping tends to use a higher number of nodes. The energy consumption of the chassis is quite high compared to other elements, so




Figure 4.7: Average Energy Consumption (Second set of experiments of a duration of the inter arrival time of request for a substrate with 150 nodes)

the impact on the consumption per request increases since there is a higher number of active nodes. For substrates with more than 100 nodes, the GREEN model tends to allocate cores in routers already active (powered on) and the energy consumption per request stabilizes. However, for the BAND model which does not tend to utilize active routers, the consumption per request tends to increase due to a higher number of active routers. This energy consumption of BAND stabilizes at 250 nodes. The energy reduction produced by GREEN was of the order of 20% for substrates larger than 100.

Figure 4.9 shows the average allocated bandwidth per request as a function of the number of nodes of the substrate. The bandwidth consumption per request is higher for the GREEN model since it tends to use routers already powered regardless of the location of these routers which can imply in long paths between nodes. The bandwidth consumption under the GREEN model is more than that consumed by BAND, but it is important to state that the GREEN model does not violate the QoS requirements of the requests, since the constraints defined in the formulation guarantee the support of these requirements.

Figure 4.10 shows the physical node utilization, which is the ratio of physical routers occupied by the total number of physical routers, as a function of the substrate size. The utilization of the physical nodes in both scenarios decreases as the physical substrate sizes increases, since there are more nodes available to allocate the incoming requests.




Figure 4.8: Average Energy Consumption per request as a function of the number of nodes of the substrate




Figure 4.9: Average Allocated Bandwidth per request as a function of the number of nodes of the substrate

Table 4.3: Values of the parameters used in GREEN Experiment 3

Parameter	Value
Number of physical nodes	{25, 50, ..., 400}
Bandwidth of each physical link	10240Mbps
Number of images in the network	3
Simulation time	5000s
Average arrival time per request	{50 }s
Average duration per request	1250s
Number of virtual nodes per request	{4}
Bandwidth of each virtual link	1024Mbps
Maximum time required to instantiate the network	100s
RAM memory	768MB
Image size	128MB
Cores per physical router	6
Cores per virtual router	6
Physical link delay	Defined by BRITE
Virtual link delay	15 × value defined by BRITE
Time required to process the image	10s
Chassis Energy Consumption	10920W
Processor Energy Consumption	166W
Line Card Energy Consumption	450W
Amplifier Energy Consumption	15W

The GREEN formulation shows significant reduction in utilization, which is at most 14% when compared to the BAND formulation. Such difference is almost constant regardless of the substrate size, averaging at 10%. This happens because the GREEN formulation will tend to use powered on routers whenever possible, and thus will concentrate the virtual network requests on a smaller physical substrate.

Figure 4.11 shows the physical link utilization, which is the ratio of physical links occupied by the total number of physical links, as a function of the substrate size. Starting from substrate sizes 100 and onward, the GREEN formulation starts using more physical links than BAND. This is because it is less expensive in terms of energy to use physical links, as their energy costs are lower. The reason why this doesn't occur for substrate sizes 25, 50 and 75 is because the flow of requests saturates the physical nodes, and thus the network is not able to handle enough requests to show this behavior.




Figure 4.10: Substrate Utilization (Physical Nodes)




Figure 4.11: Substrate Utilization (Physical Links)




Figure 4.12: Blocking Rate

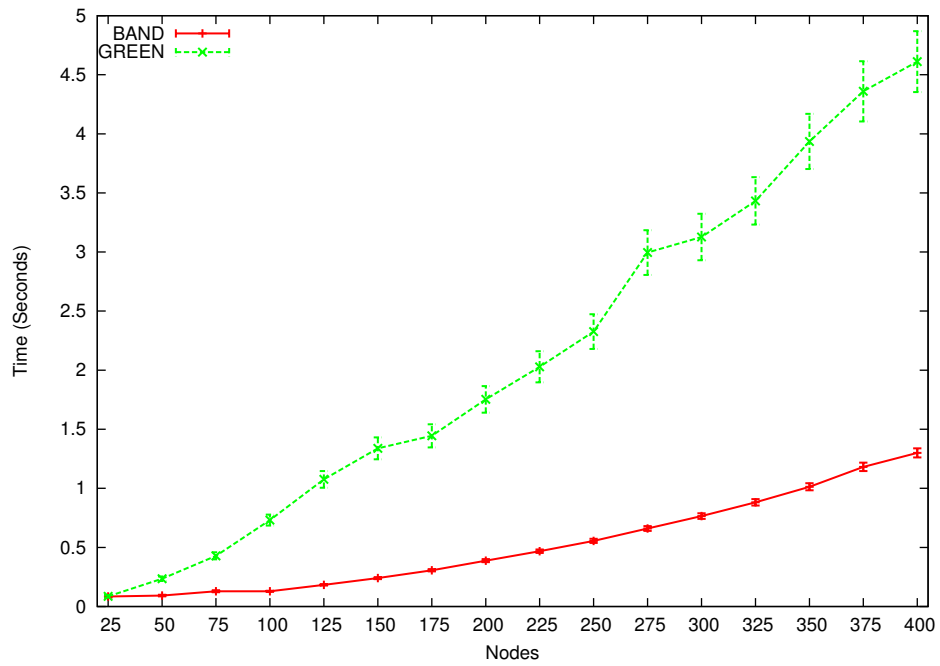


Figure 4.13: Execution Time

Chapter 5

Live Migration of Virtual Networks

5.1 Introduction

In [42], we proposed algorithms to map virtual networks on network substrates with the aim of minimizing the energy consumption of virtual network establishment. Although these algorithms produce optimal mappings, the dynamics of virtual network establishment and termination can make resource allocation less than optimal, leading to greater consumption of energy. This deviation from the minimum consumption of energy can be counteracted by reallocating resources to the virtual networks [10]. Figure 5.1 illustrates the benefit of resource reallocation. In Figure 5.1(a), virtual networks VN1, VN2 and VN3 use, the pair of nodes (R2,R3), (R1,R2) and (R1,R3), respectively. Figure 5.1(b) shows the state of the network after the termination of the virtual network VN1 which consumes more energy than the minimum to provide service to VN2 and VN3. It is, thus, possible to remap the virtual networks so the energy consumption is minimum. A solution for that is the migration of the virtual router VN3 from the physical router R3 to the physical router R2, as illustrated by the dashed line in Figure 5.1(b). Figure 5.1(c) shows the new mapping which allows the physical router R3 to be turned off, given that there is no virtual router mapped on it, leading to a reduction of energy consumption.

Figure 5.2 shows a more verbose example of the application of Live Migration in order to reduce energy consumption in green virtual networks. Figure 5.2(a) shows a physical network of 5 nodes and virtual networks 1 to 5. For simplicity, the virtual links are omitted from this demonstration. Figure 5.2(b) shows that VN1 being uninstantiated from the network and Figure 5.2(c) shows the network with VN1 uninstantiated of the network. Figures 5.2(d) and 5.2(e) show the same process described above, but for VN 3. In 5.2(e), it can be seen that there were “gaps” left by the virtual network requests that were unmapped from the substrate. Figure 5.2(f) shows a possible migration schema for routers from the virtual networks 4 and 5. The arrows show the places where the virtual

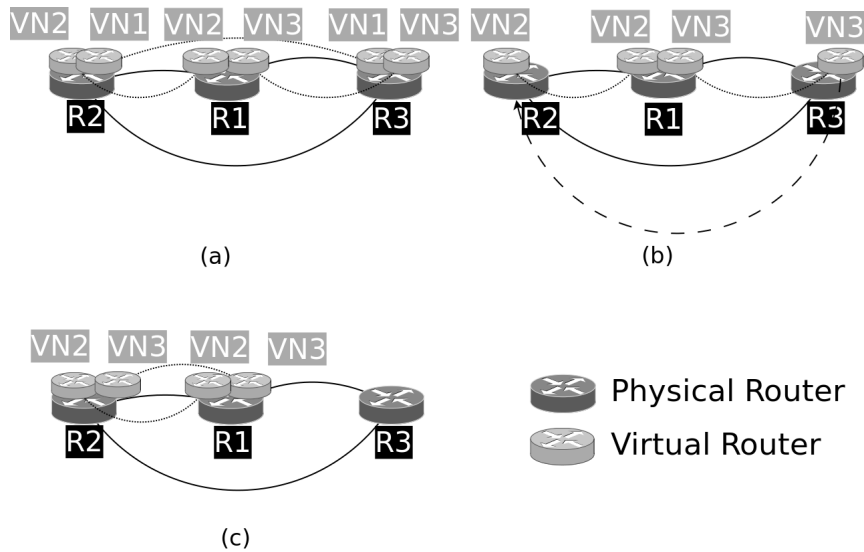


Figure 5.1: Example of reallocation of resources

routers will be allocated. Figure 5.2(g) shows the physical substrate after the migration takes place. The faded out colored elements show the physical routers and links that can be turned off after the migration takes place.

This chapter presents two algorithms for live migration of virtual networks with the objective of reducing the energy consumption of the network. The first algorithm remaps the virtual networks which virtual routers reside on the same physical router that hosted a virtual router of a terminated virtual network. The second algorithm remaps all the virtual networks after the termination of a virtual network. Results derived via simulation confirm that the remapping of virtual resources has a significant impact on energy consumption.

5.2 Proposed Algorithms

The central idea of remapping virtual networks is to restore a network configuration which consumes minimally energy. One of the events which offers most promising opportunities for remapping is the termination of a virtual network since virtual routers can be “packed” on a reduced number of physical routers, which can lead to deallocation of all virtual routers hosted on a physical router, allowing it to be turned off.

The algorithms presented in this paper are executed at the termination of virtual networks, when the active virtual networks can be re-arranged to obtain energy savings.

The algorithms assume that the transfer of information about the state of running processes is negligible given the high capacity of the physical links. Moreover, they assume that the time required to boot the virtual networks, redirect the traffic to the new virtual

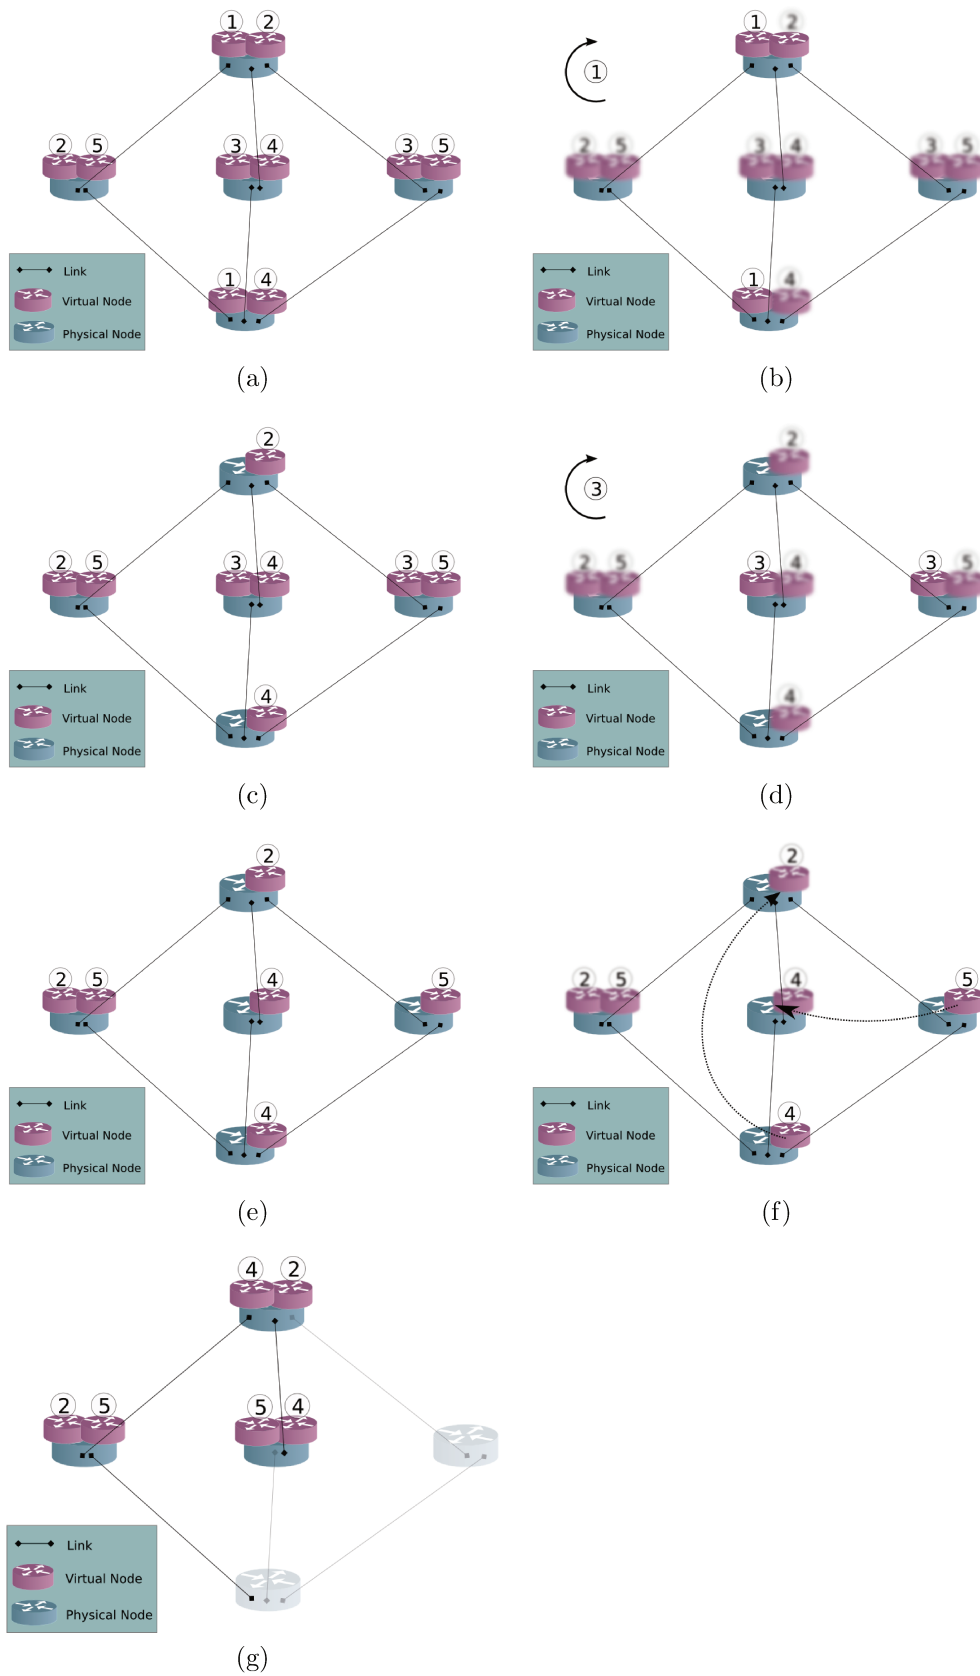


Figure 5.2: Step by step of Live Migration to reduce energy consumption

network, finalize the transferring of the old virtual network and interrupt the services of the old virtual network is also negligible [48].

5.2.1 Terminology

Before explaining any of the algorithms, some mathematical notations will be introduced.

- $P = \{p_i\}$ - set of physical routers;
- $C = \{c_{ij}\}$ - set of physical links (channels), C_{ij} connects physical routers (P_i, P_j) ;
- $V^k = \{V_i^k\}$ - set of virtual routers of virtual network k ;
- $L^k = \{l_{ij}^k\}$ - set of virtual links of virtual network k , l_{ij}^k connects virtual routers (V_i^k, V_j^k) ;
- $VN_k = (V^k, L^k)$ - virtual network k ;
- $VN = \{VN_k\}$ - set of active virtual networks;
- $A(t) = [a_{ijk}(t)]$ - matrix which elements are $a_{ijk} = 1$ if the virtual routers $V_j^k \in VN_k$ is allocated on the physical router $P_i \in P$ at time t ;
- $F(t) = [f_{ij}(t)]$ - matrix which elements $f_{ij}(t)$ gives the allocated bandwidth on the physical link $c_{ij} \in C$;
- $S(VN, t) = (A(t), F(t))$ - state of the substrate network hosting VM at time t ;
- $E(S(VN, t))$ - energy consumption of $S(VN, t)$;
- t_k - time at which the virtual network k terminated, $t_k \in R^+$
- $D_k = \{P_k \in P | a_{ijk}(t_k^-) = 1 \text{ and } a_{ijk}(t_k^+) = 0\}$ - set of physical routers on which VN_k had a virtual router allocated on;
- $G^k = \{VN_m | a_{ijk}(t_k^-) = 1 \text{ and } a_{ijk}(t_k^+) = 0 \text{ and } P_i \in D_k\}$ - set of virtual networks which have virtual routers allocated on the same physical router that hosted a virtual router of the terminated virtual network k ;

Algorithm 2: LM_NRU

Data: Z - terminating virtual network
Data: $S(VN, t_z)$ - state of the substrate when VN_z terminates
Result: S' - state of the substrate

```

1  $VN' = VN \cap G^z$ ;
2 for all  $VN_i \in G^z$ ;
3 if RE-MAP ( $VN_i$ ) then
4   |  $VN' = VN' \cup \{VN_i\}$ ;
5 end if
6 else
7   | return  $S(VN, t_z^+)$ ;
8 end if
9 if  $E(S(VN', t_z^+)) \leq E(S(VN, t_z^+))$  then
10  | return  $S(VN', t_z^+)$ ;
11 end if
12 else
13  | return  $S(VN, t_z^+)$ ;
14 end if
```

5.2.2 Nodes Recently Used (NRU)

Algorithm 1, called Nodes Recently Used (NRU), aims to reduce the energy consumption by reallocating virtual routers allocated on the same physical routers that hosted the routers of the terminated virtual networks.

The NRU algorithm receives as input the state of the network substrate at the time the virtual network z terminates and it returns the state of the network that minimizes energy consumption.

In Line 1, all virtual networks which have virtual routers allocated on the same physical router that hosted a virtual router of the terminating virtual network z are “un-mapped”.

In Lines 2-5, attempts are made to re-map all “un-mapped” virtual networks. The aim is to produce a new mapping that decreases the energy consumption in comparison to the consumption when the virtual network z terminated. RE-MAP is a binary function which returns true when the re-mapping of VN_i is feasible. In this case, a new mapping of virtual networks including VN_i is considered (Lines 6-8).

Otherwise, the process of re-mapping of all “un-mapped” virtual networks is interrupted and the substrate remains in the same state it was when the virtual network terminated since no virtual network can be terminated before its due time.

In case all “un-mapped” virtual networks are remapped and the new substrate state consumes less energy, the new state is returned as the solution given by the algorithm

and the virtual networks are migrated to the physical routers suggested by the solution.

If the new mapping does not lead to a reduction of energy consumption then no migration is pursued.

Figure 5.3 shows an example of the NRU algorithm. Figure 5.3(a) shows a physical substrate with several virtual routers mapped, corresponding to certain VNs. Figure 5.3(b) shows that VN 1 is going to be unmapped from the network. This unmapping occurs on figure 5.3(c). The physical routers related with VN 1 are physical routers 1 and 2. We consider those physical routers as part of set P . The virtual network associated with physical routers from the set P is VN 2, and therefore it is considered in set Q . The VNs in set Q are temporarily unmapped as shown on figure 5.3(d), and by executing a green virtual network mapping algorithm, it can be seen that it can be remapped into physical routers 3 and 4, allowing physical router 2 to be powered off.

5.2.3 Live Migration-ALL

Algorithm 2, called Live Migration-ALL (LM_ALL) aims to reduce the energy consumption by reallocating all virtual routers of the network. It “un-maps” all virtual networks (Line 1) and then tries to re-map all of them (Lines 2-8). The rest of Algorithm 2 is similar to Algorithm 1. It differs from Algorithm 1 since LM_ALL deallocates all virtual routers and LM_NRU deallocates only a portion of the network.

Algorithm 3: LM_ALL

Data: Z - terminating virtual network
Data: $S(VN, t_z)$ - state of the substrate when VN_z terminates
Result: S' - state of the substrate

```

1  $VN' = \{\}$ ;
2 for all  $VN_{i \neq z}$ ;
3 if RE-MAP ( $VN_i$ ) then
4   |  $VN' = VN' \cup \{VN_i\}$ ;
5 end if
6 else
7   | return  $S(VN, t_z^+)$ ;
8 end if
9 if  $E(S(VN', t_z^+)) \leq E(S(VN, t_z^t))$  then
10  | return  $S(VN', t_z^+)$ ;
11 end if
12 else
13  | return  $S(VN, t_z^+)$  ;
14 end if
```

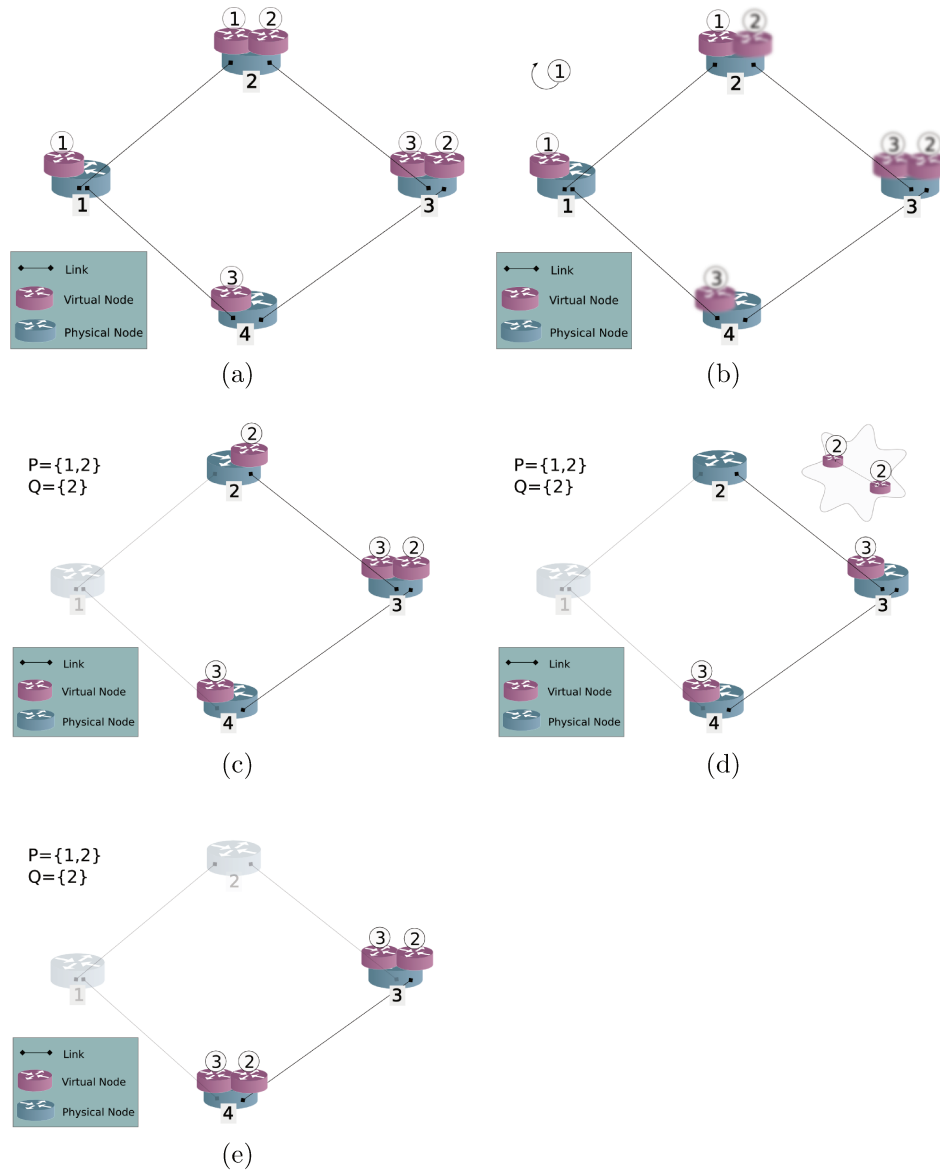


Figure 5.3: Step by Step of LM NRU Algorithm

5.3 Configuration of the Experiments

To evaluate the performance of the proposed formulation, another optimization model was developed containing similar constraints, but a different objective function that is the minimization of the allocated bandwidth. In the remainder of this paper, the proposed formulation is denoted as GREEN, while the formulation which minimizes the allocated bandwidth is denoted as BAND. Details on the BAND optimization model can be consulted in [3]. The formulations were evaluated in dynamic scenarios, in which the availability of resources in the substrate network varies as a function of time.

The average energy consumption per request, the amount of bandwidth allocated per request and the substrate utilization for physical nodes and physical links were evaluated as a function of the substrate sizes (number of physical routers) and as a function of the inter arrival time of requests.

To assess the effectiveness of the proposed formulation, a simulator was implemented in C++. This simulator receives a description of the substrate network as input and generates requests. Confidence intervals of 95% were derived using the independent replication method. The ILP formulations were implemented using the CPLEX optimization library version 12.0. The simulations were executed on a computer running the operating system Debian GNU/Linux Squeeze. The computer was equipped with two Intel Xeon 2.27GHz processors, with 6 cores each one, and 6GB of RAM.

5.4 Results

The results for the experiments are shown below. The experiments from this section evaluated the growth of the network by altering the size of the network.

5.4.1 Live Migration Experiment 1

Parameters

The parameters used for this experiment are shown in Table 5.1. For this experiment, the average inter arrival time of requests was fixed to 12 seconds. The substrate sizes were varied in steps of 20 nodes, from 20 to 160 nodes.

Results

Figure 5.4 shows the average energy consumption per request using the scheme with no live migration (LM_NO) and the two proposed algorithms, for network substrates of at most 160 nodes. Both LM_ALL and LM_NRU show significant energy savings per

Table 5.1: Values of the parameters used in Live Migration Experiment 1

Parameter	Value
Number of physical nodes	{20 40 60 80 100 120 140 160}
Bandwidth of each physical link	10240Mbps
Number of images in the network	3
Simulation time	3600s
Average arrival time per request	12s
Average duration per request	360s
Number of virtual nodes per request	{2 3 4 5 6}
Bandwidth of each virtual link	1024Mbps
Maximum time required to instantiate the network	100s
RAM memory	768MB
Image size	128MB
Cores per physical router	6
Cores per virtual router	6
Physical link delay	Defined by BRITE
Virtual link delay	$15 \times$ value defined by BRITE
Time required to process the image	10s
Chassis Energy Consumption	10920W
Processor Energy Consumption	166W
Line Card Energy Consumption	450W
Amplifier Energy Consumption	15W

request when compared to LM_NO, which was, on average, 15% and at most 18% for substrates with more than 40 physical routers. For small substrate (less than 20 routers) no significant difference was observed since there is not much room for optimized mapping of virtual networks given the reduced number of physical routers. Energy savings given by LM_NRU and LM_ALL present no significant difference which favors the adoption of the LM_NRU algorithm since it requires small overhead.

Besides reducing the energy consumption per request, the proposed algorithms also consume less bandwidth than when no migration occurs (Figure 5.5). The reduction can be of at most 8%. The LM_NRU algorithm reduces the bandwidth consumed per request only for substrates larger than 100 routers since for smaller substrates there was not much differences in the physical length of paths associated to virtual links.

Figure 5.6 shows the substrate utilization, which is the ratio of physical routers occupied by the total number of physical routers, as a function of the substrate size. The proposed algorithms show significant reduction in utilization and the difference between the utilization they produce is negligible. The reduction is around 13% when compared to a no migration scheme and such difference is almost constant regardless of the substrate size. This happens since the reallocation of virtual network resources yields to the concentration of larger number of virtual networks, on a smaller number of physical routers when compared to the no migration scheme given the savings obtained with a smaller number of physical router chassis powered.

Figures 5.7 and 5.8 display the average number of virtual router and virtual link migrations per event of virtual request termination, respectively. These figures show clearly




Figure 5.4: Average Energy Consumption per Request




Figure 5.5: Average Allocated Bandwidth per Request

that LM_ALL uses all available choices to re-map the virtual networks. As the number




Figure 5.6: Substrate Utilization as a Function of the Substrate Size

of physical nodes increases, the amount of virtual routers and virtual links migrated decreases since several virtual networks are mapped on the same physical nodes and links that they were originally mapped on. The amount of migrations of both virtual routers and links per event produced by LM_NRU is almost constant. LM_NRU migrates 41% less virtual routers than LM_ALL and 38% less virtual links than does LM_ALL. Considering the energy savings produced by these algorithms, the LM_NRU is more attractive than the LM_ALL since it produces a lower number of migrations.

Figure 5.9 shows the average execution time per event of virtual request termination. The LM_ALL demands increase sharply when compared to LM_NRU. LM_ALL is at most 750% slower than LM_NRU as a result of less deallocation and reallocation that needs to be done in LM_NRU. Considering the similarities in energy savings and the time complexity and overhead costs of the LM_ALL algorithm, the adoption of LM_NRU algorithm for live migration of virtual networks is recommended for adoption.

5.4.2 Live Migration Experiment 2

Parameters

The parameters used for this experiment are shown in Table 5.2. For this experiment, the average inter arrival time of requests was fixed to 12 seconds. The substrate sizes




Figure 5.7: Average Virtual Routers Migrated per Event of Virtual Network Termination




Figure 5.8: Average Virtual Links Migrated per Event of Virtual Network Termination

were varied in steps of 20 nodes, from 20 to 240 nodes. This experiment attempts to




Figure 5.9: Execution Time per Event of Virtual Network Termination

push the limits of the previous experiment even further, and bigger substrate sizes are considered. The number of virtual nodes per request is also increased.

Table 5.2: Values of the parameters used in Live Migration Experiment 2

Parameter	Value
Number of physical nodes	{20 40 60 80 100 120 140 160 180 200 220 240}
Bandwidth of each physical link	10240Mbps
Number of images in the network	3
Resolution time	3600s
Average arrival time per request	12s
Average duration per request	360s
Number of virtual nodes per request	{2 3 4 5 6 7}
Bandwidth of each virtual link	1024Mbps
Maximum time required to instantiate the network	100s
RAM memory	768MB
Image size	128MB
Cores per physical router	6
Cores per virtual router	6
Physical link delay	Defined by BRITE
Virtual link delay	15 × value defined by BRITE
Time required to process the image	10s
Chassis Energy Consumption	10920W
Processor Energy Consumption	166W
Line Card Energy Consumption	450W
Amplifier Energy Consumption	15W

Figure 5.10 shows the average energy consumption per request using the scheme with no live migration (LM_NO) and the two proposed algorithms, for network substrates of

at most 240 nodes. Both LM_ALL and LM_NRU show significant energy savings per request when compared to LM_NO, which was, on average, 20% and at most 25% for substrates with more than 60 physical routers. For small substrate (less than 40 routers) no significant difference was observed since there is not much room for optimized mapping of virtual networks given the reduced number of physical routers. Energy savings given by LM_NRU and LM_ALL present no significant difference which favors the adoption of the LM_NRU algorithm since it requires smaller overhead.




Figure 5.10: Average Energy Consumption per Request

The proposed algorithms also consumed a similar amount of bandwidth than when no migration occurs (Figure 5.11). The no migration scheme tends to consume the lowest amount of bandwidth. The reduction can be at most 8%. As the substrate size increases, this difference becomes negligible since for smaller substrates there was not much difference on the physical length of paths associated to virtual links.

Figure 5.12 shows the substrate utilization, which is the ratio of physical routers occupied by the total number of physical routers, as a function of the substrate size. The proposed algorithms show significant reduction in utilization and the difference between the utilization they produce is negligible. The reduction is around 13% when compared to a no migration scheme and such difference is almost constant regardless of the substrate size. This happens since the reallocation of virtual network resources yields to the concentration of larger number of virtual networks on a smaller number of physical routers




Figure 5.11: Average Allocated Bandwidth per Request

when compared to the no migration scheme due to the savings obtained with a smaller number of physical router chassis powered on.

Figures 5.13 and 5.14 display the average number of virtual router and virtual link migrations per event of virtual request termination, respectively. These figures show clearly that LM_ALL uses all available choices to re-map the virtual networks. As the number of physical nodes increases, the amount of virtual routers and virtual links migrated decreases since several virtual networks are mapped on the same physical nodes and links that they were originally mapped on. The amount of migrations of both virtual routers and links per event produced by LM_NRU is almost constant. LM_NRU migrates 41% less virtual routers than LM_ALL and 38% less virtual links than does LM_ALL. Considering the energy savings produced by these algorithms, the LM_NRU is more attractive than the LM_ALL since it produces a lower number of migrations.

Figure 5.15 shows the average execution time per event of virtual request termination. The LM_ALL demand increases sharply when compared to LM_NRU. LM_ALL is at most 750% slower than LM_NRU as a result of less deallocation and reallocation that needs to be done. Considering the similarities in energy savings and the time complexity and overhead costs of the LM_ALL algorithm, the adoption of LM_NRU algorithm for live migration of virtual networks is recommended for adoption.




Figure 5.12: Substrate Utilization as a Function of the Substrate Size




Figure 5.13: Average Virtual Routers Migrated per Event of Virtual Network Termination




Figure 5.14: Average Virtual Links Migrated per Event of Virtual Network Termination




Figure 5.15: Execution Time per Event of Virtual Network Termination

Chapter 6

Conclusions

Networking virtualization is a promising technique used to solve a lot of problems that have arisen in core networks. Virtualization has an important place in the Future Internet. Besides the benefits highlighted in the previous sections, virtualization is crucial for energy savings and for a greener Internet.

This dissertation presented novel algorithms for the reduction of energy consumption of core virtual networks. A new model to map virtual networks on network substrates with the objective of minimizing energy consumption was presented. This model is based on an ILP formulations and it considers several realistic parameters neglected by previous works in the literature. Also, this work proposes the use of live migration to further reduce the energy consumption, and presents two algorithms in order to achieve this.

The main contribution of this work is the consideration of real life parameters omitted in other works for the green virtual network mapping problem and the use of live migration for energy consumption reduction in core networks. The characteristics of the scenarios and the parameters are summarized below:

- Preconfigured images for each virtual router exists, allowing each virtual request to be instantiated swiftly. These images are located in repositories spread out in the physical substrate;
- Each virtual router specifies which image it will use;
- The flash memory of each physical routers is limited and must be used to save the image of the routers instantiated at a given time;
- Each virtual network request specifies a limit in the time needed to be mapped;

In order to minimize the energy consumption per request, a two step ILP was developed. Due to the complexity of the problem (NP-Hard), a heuristic was developed in

order to reduce the execution time of the solution. This proposal was compared to another one that minimizes the bandwidth consumption. Results show that acquired energy savings per request can be of the order of 50% for some scenarios. It is important to note that QoS requirements of the virtual networks demands were satisfied.

Results indicate that re-mapping of virtual networks can reduce the energy consumption in networks, of up to 18%, when compared to a no migration scheme without compromising the bandwidth demand. The LM_NRU algorithm leads to less overhead than that produced by the LM_ALL algorithm and to similar energy savings. Consequently, it is the best choice for adoption in virtual network re-mapping with live migration.

These greening techniques answer the call for an ever growing demand for decreasing the amount of energy consumption of communication networks. It can be used for the virtualized network of the future which definitely needs to refrain the consumption of energy in the face of the growth of network demand.

Bibliography

- [1] Canada's advanced research and innovation network (canarie). <http://canarie.ca/>.
- [2] *Product data sheet. T SERIES CORE ROUTERS.*
- [3] G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca. Optimal Mapping of Virtual Networks. In *Proc. of the IEEE GLOBECOM*, pages 1–5, 2011.
- [4] L.A. Barroso and U. Holzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, 2007.
- [5] R. Bolla, F. Davoli, R. Bruschi, K. Christensen, F. Cucchietti, and S. Singh. The potential impact of green technologies in next-generation wireline networks: Is there room for energy saving optimization? *IEEE COMMAG*, 49(8):80–86, august 2011.
- [6] J.F. Botero, X. Hesselbach, M. Duelli, D. Schlosser, A. Fischer, and H. De Meer. Energy efficient virtual network embedding. *Communications Letters, IEEE*, 16(5):756–759, 2012.
- [7] Luca Chiaraviglio, Marco Mellia, and Fabio Neri. Energy-Aware Networks: Reducing Power Consumption By Switching Off Network Elements. 2010.
- [8] N.M.M.K. Chowdhury, M.R. Rahman, and R. Boutaba. Virtual Network Embedding with Coordinated Node and Link Mapping. In *IEEE INFOCOM*, pages 783–791, Abril 2009.
- [9] Cisco. *Brochures for devices CRS-1 and components, CISCO ONS 15501 and others.* Available at <http://www.cisco.com/>, accessed 06/2011, 2011.
- [10] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul, Christian Limpach, Ian Pratt, and Andrew Warfield. Live migration of virtual machines. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 273–286, Berkeley, CA, USA, 2005. USENIX Association.

- [11] Rossi Dario, Bianzino Aruna Prem, Chaudet Claude, and Rougier Jean-Louis. A Survey of Green Networking Research. *Communications Surveys and Tutorials*, Nov 2009.
- [12] Infoworld Magazone David Marshall. <http://www.infoworld.com/d/virtualization/top-10-benefits-server-virtualization-177828>, accessed 03/2012, 2011.
- [13] C. Despins, F. Labeau, Tho Le Ngoc, R. Labelle, M. Cheriet, C. Thibeault, F. Gagnon, A. Leon-Garcia, O. Cherkaoui, B. St. Arnaud, J. Mcneill, Y. Lemieux, and M. Lemay. Leveraging green communications for carbon emission reductions: Techniques, testbeds, and emerging carbon footprint standards. *Communications Magazine, IEEE*, 49(8):101–109, Aug. 2011.
- [14] Jinxu Ding and A. Somani. Reduction of green house gas emission by clean power trading. In *Energy Conversion Congress and Exposition (ECCE), 2010 IEEE*, pages 4588–4592, 2010.
- [15] Rodriguez E., Alkmim G. P., D. M. Batista, and N. L. S. Fonseca. Mapeamento energeticamente eficiente de redes virtuais em substratos fisicos. In *Proceedings of the XXX Simposio Brasileiro de Redes de Computadores*, pages 1–14, Ouro Preto, Minas Gerais, Brasil, 2012. SBC.
- [16] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. Power provisioning for a warehouse-sized computer. *SIGARCH Comput. Archit. News*, 35(2):13–23, June 2007.
- [17] Nick Feamster, Lixin Gao, and Jennifer Rexford. How to Lease the Internet in Your Spare Time. *SIGCOMM Comput. Commun. Rev.*, 37(1):61–64, 2007.
- [18] F.Idzikowski. Power consumption of network elements in IP over WDM networks. TKN Technical Report Series TKN-09-006, Telecommunication Networks Group, Technical University Berlin, July 2009.
- [19] A. Fischer, M.T. Beck, and H. de Meer. An approach to energy-efficient virtual network embeddings. In *5th International Workshop on Management of the Future Internet (ManFI 2013)*, 2013.
- [20] R. E. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Society*, 64:275–278, 1958.

- [21] Maruti Gupta and Suresh Singh. Greening of the internet. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '03, pages 19–26, New York, NY, USA, 2003. ACM.
- [22] Jiayue He, Rui Zhang-Shen, Ying Li, Cheng-Yen Lee, Jennifer Rexford, and Mung Chiang. DaVinci: Dynamically Adaptive Virtual Networks for a Customized Internet. In *ACM CoNEXT '08*, pages 15:1–15:12, 2008.
- [23] Ines Houidi, Wajdi Louati, and Djamel Zeghlache. A Distributed and Autonomic Virtual Network Mapping Framework. In *In Proceedings of the ICAS '08*, pages 241–247, 2008.
- [24] IBM. IBM ILOG CPLEX Optimization Studio V12.2, 2010. <http://publib.boulder.ibm.com/infocenter/cosinfoc/v12r2/index.jsp?topic=/ilog.odms.cplex.help/Content/Optimization/Documentation/CPLEX/>. Accessed at 12/20/2010.
- [25] F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, and A. Wolisz. Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios. In *ONDM, 2010 14th Conference on*, pages 1–6, feb. 2010.
- [26] Xie Jie. Greenhouse effects on the world agriculture - based on computable general equilibrium model analysis. In *Mechanic Automation and Control Engineering (MACE), 2010 International Conference on*, pages 2082–2085, 2010.
- [27] R.H. Katz. Tech titans building boom. *Spectrum, IEEE*, 46(2):40–54, 2009.
- [28] S. Kikuchi and Y. Matsumoto. Impact of live migration on multi-tier application performance in clouds. In *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, pages 261 –268, June 2012.
- [29] Jing Lu and Jonathan Turner. Efficient Mapping of Virtual Networks onto a Shared Substrate. Technical Report WUCSE-2006-35, Washington University, 2006. <http://www.arl.wustl.edu/~jst/pubs/wucse2006-35.pdf>. Accessed at 12/20/2010.
- [30] R. A. MEYER and L. H. Seawright. A virtual machine time-sharing system. *IBM System Journal*, 9(3):199–218, 1970.
- [31] F.F. Moghaddam and M. Cheriet. Decreasing live virtual machine migration downtime using a memory page selection based on memory change pdf. In *Networking, Sensing and Control (ICNSC), 2010 International Conference on*, pages 355–359, 2010.

- [32] M. D. D. Moreira, N. C. Fernandes, L. H. M. K. Costa, , and O. C. M. B. Duarte. Internet do futuro: Um novo horizonte. In *Minicursos do Simpósio Brasileiro de Redes de Computadores - SBRC'2009*, pages 1–59. 2009.
- [33] Aziz. Murtazaev and Sangyoon. Oh. Sercon: Server Consolidation Algorithm using Live Migration of Virtual Machines for Green Computing. *IETE Technical Review*, 28(3):212–231, 2011.
- [34] D. Pamlin and K. Szomolanyi. Saving the climate @ the speed of light First roadmap for reduced CO2 emissions in the EU and beyond. *World Wildlife Fund and European Telecommunications Network Operators' Association*, April 2007.
- [35] R. P. PARMELEE, T. I. Peterson, C. C. Tillman, and D. J. Hatfiels. Virtual storage and machine concepts. *IBM System Journal*, 11(2):99–130, 1972.
- [36] Pedro S. Pisa, Natalia C. Fern, Hugo E. T. Carvalho, Marcelo D. D, Miguel Elias, M. Campista, Luís Henrique M. K. Costa, and Otto Carlos M. B. Duarte. OpenFlow and Xen-Based Virtual Network Migration. The World Computer Congress 2010 - Network of the Future Conf., Brisbane, Australia, Sept 2010.
- [37] Gerald J. Popek and Robert P. Goldberg. Formal requirements for virtualizable third generation architectures. *Commun. ACM*, 17(7):412–421, 1974.
- [38] U.S. Environmental Protection Agency – Energy Star Program. Report to congress on server and data center energy efficiency public law 109-431, August 2007.
- [39] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. Cutting the electric bill for internet-scale systems. *SIGCOMM Comput. Commun. Rev.*, 39(4):123–134, August 2009.
- [40] Sinha Richa, Purohit Nidhi, and Diwanji Hiteshi. Power Aware Live Migration for Data Centers in Cloud using Dynamic Threshold, 2011.
- [41] E. Rodriguez, G. Alkmim, D.M. Batista, and N.L.S. da Fonseca. Trade-off between bandwidth and energy consumption minimization in virtual network mapping. In *Communications (LATINCOM), 2012 IEEE Latin-America Conference on*, pages 1–6, 2012.
- [42] E. Rodriguez, G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca. Green Virtualized Networks. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 1–5, 2012.

- [43] E. Rodriguez, G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca. Live Migration in Green Virtualized Networks. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 1–5, 2013.
- [44] E. Rodriguez, G. Prado Alkmim, D. Macedo Batista, and N.L. Saldanha da Fonseca. Trade-off between bandwidth and energy consumption minimization in virtual network mapping. *Latin America Transactions, IEEE (Revista IEEE America Latina)*, 11(3):983–988, 2013.
- [45] Gregor Schaffrath, Christoph Werle, Panagiotis Papadimitriou, Anja Feldmann, Roland Bless, Adam Greenhalgh, Andreas Wundsam, Mario Kind, Olaf Maennel, and Laurent Mathy. Network virtualization architecture: proposal and initial prototype. In *VISA '09*, pages 63–72, New York, NY, USA, 2009. ACM.
- [46] Gangxiang Shen and Rodney S. Tucker. Energy-Minimized Design for IP Over WDM Networks. *J. Opt. Commun. Netw.*, 1(1):176–186, Jun 2009.
- [47] W. Szeto, Y. Iraqi, and R. Boutaba. A multi-commodity flow based approach to virtual network resource allocation. In *GLOBECOM '03. IEEE*, volume 6, pages 3004–3008, Dec. 2003.
- [48] Franco Travostino, Paul Daspit, Leon Gommans, Chetan Jog, Cees de Laat, Joe Mambretti, Inder Monga, Bas van Oudenaarde, Satish Raghunath, and Phil Yonghui Wang. Seamless live migration of virtual machines over the MAN/WAN. *Future Gener. Comput. Syst.*, 22(8):901–907, October 2006.
- [49] William Voorsluys, James Broberg, Srikumar Venugopal, and Rajkumar Buyya. Cost of virtual machine live migration in clouds: A performance evaluation. *CoRR*, abs/1109.4974, 2011.
- [50] Yi Wang, Eric Keller, Brian Biskeborn, Jacobus van der Merwe, and Jennifer Rexford. Virtual routers on the move: live router migration as a network-management primitive. *SIGCOMM Comput. Commun. Rev.*, 38(4):231–242, August 2008.
- [51] Minlan Yu, Yung Yi, Jennifer Rexford, and Mung Chiang. Rethinking Virtual Network Embedding: Substrate Support for Path Splitting and Migration. *SIGCOMM Comput. Commun. Rev.*, 38(2):17–29, 2008.
- [52] Yaping Zhu, Andy Bavier, Nick Feamster, Sampath Rangarajan, and Jennifer Rexford. UFO: a resilient layered routing architecture. *SIGCOMM Comput. Commun. Rev.*, 38(5):59–62, 2008.

- [53] Yaping Zhu, Rui Zhang-Shen, Sampath Rangarajan, and Jennifer Rexford. Cabernet: Connectivity Architecture for Better Network Services. In *ACM CoNEXT '08*, pages 64:1–64:6, 2008.

Appendix A

BAND Formulation

The following is the BAND formulation presented in [3]. The GREEN algorithm presented in this dissertation was compared to this formulation, which minimized bandwidth. The modeling of this problem is the same as the modeling of GREEN, with the exception that this formulation does not take in consideration the energy consumption restrictions and it does not minimize bandwidth.

$$\begin{aligned}
 & N \text{Minimize} \\
 & \mathcal{P}^{chassis} \sum_{n \in N} (\alpha_n + (1 - \alpha_n)U_n) + \\
 & \mathcal{P}^{core} \sum_{n \in N} \sum_{m \in M} \sum_{i \in I} (X_{n,m,i} \times P(m)) + \\
 & (2\mathcal{P}_{u,v}^{card} + \mathcal{P}_{u,v}^L) \sum_{(u,v) \in F} (\beta_{u,v} + (1 - \beta_{u,v})W_{u,v})
 \end{aligned}$$

subject to

$$\begin{aligned}
 & \sum_{n \in N} \sum_{i \in I} X_{n,m,i} = 1 & (C1) \\
 & \forall m \in M
 \end{aligned}$$

$$\begin{aligned}
 & \sum_{m \in M} \sum_{i \in I} X_{n,m,i} \leq 1 & (C2) \\
 & \forall n \in N
 \end{aligned}$$

$$\begin{aligned}
 & \sum_{m \in M} \sum_{i \in I} P(m) \times X_{n,m,i} \leq A(n) & (C3) \\
 & \forall n \in N
 \end{aligned}$$

$$\begin{aligned} X_{n,m,i} &= 0 & (C4) \\ \forall n \in N, \forall m \in M, \forall i \in I | L_{n,m} = 0 \text{ or } E_{m,i} = 0 \end{aligned}$$

$$\begin{aligned} \sum_{w' \in V} Y_{u,v,w'} \times Q(w') &\leq C(w) & (C5) \\ \forall w = (u, v) \in F \end{aligned}$$

$$\begin{aligned} \sum_{u \in N} \sum_{v \in N} Y_{u,v,w} \times D(u, v) &\leq K(w) & (C6) \\ \forall w \in V, (u, v) \in F \end{aligned}$$

$$\begin{aligned} \sum_{m \in M} \sum_{i \in I} X_{n,m,i} \times G(i) &\leq B(n) & (C7) \\ \forall n \in N \end{aligned}$$

$$\begin{aligned} Y_{u,v,w} &= 0 & (C8) \\ \forall u, \forall v, \forall w \in V | (u, v) \notin F \end{aligned}$$

$$\begin{aligned} \sum_{f \in N} Y_{n,f,w} - \sum_{f \in N} Y_{f,n,w} &= & (C9) \\ \sum_{i \in I} X_{n,a,i} - \sum_{i \in I} X_{n,b,i} \\ \forall w = (a, b) \in V, \forall n \in N \end{aligned}$$

$$\begin{aligned} X_{n,m,i} &\leq U_n & (C10) \\ \forall n \in N, \forall m \in M, \forall i \in I \end{aligned}$$

$$\begin{aligned} U_n &\leq \sum_{m \in M} \sum_{i \in I} X_{n,m,i} & (C11) \\ \forall n \in N \end{aligned}$$

$$\begin{aligned} Y_{u,v,w} &\leq W_{u,v} & (C12) \\ \forall v \in V, \forall (u, v) \in F \end{aligned}$$

$$\begin{aligned} W_{u,v} &\leq \sum_{v \in V} Y_{u,v,w} & (C13) \\ \forall (u, v) \in F \end{aligned}$$

$$\begin{aligned} X_{n,m,i} &\in \{0, 1\} & (C14) \\ \forall n \in N, \forall m \in M, \forall i \in I \end{aligned}$$

$$Y_{u,v,w} \in \{0, 1\} \quad (C15)$$
$$\forall u, \forall v, \forall w \in V$$

$$U_n \in \{0, 1\} \quad (C16)$$
$$\forall n \in N$$

$$W_{u,v} \in \{0, 1\} \quad (C17)$$
$$\forall u, \forall v \in V$$

Appendix B

Original GREEN Formulation

During the development of the formulation, an initial simulation was first derived. This model was slightly adapted, as improvements to the modeling of the energy consumption restrictions were considered. The old restrictions shown here have more “for all” statements, thus adding more dimensions of complexity to these definitions. This caused a direct impact in execution times and the quality of the solutions obtained by the executions. and is shown as follows:

$$\begin{aligned} & \text{Minimize} \\ & \mathcal{P}^{chassis} \sum_{n \in N} (\alpha_n + (1 - \alpha_n)U_n) + \\ & \mathcal{P}^{core} \sum_{n \in N} \sum_{m \in M} \sum_{i \in I} (X_{n,m,i} \times P(m)) + \\ & (\mathcal{P}_{u,v}^{card} + \mathcal{P}_{u,v}^L) \sum_{(u,v) \in F} (\beta_{u,v} + (1 - \beta_{u,v})W_{u,v}) \end{aligned}$$

subject to

$$\begin{aligned} \sum_{n \in N} \sum_{i \in I} X_{n,m,i} &= 1 & (C1) \\ \forall m \in M & \end{aligned}$$

$$\begin{aligned} \sum_{m \in M} \sum_{i \in I} X_{n,m,i} &\leq 1 & (C2) \\ \forall n \in N & \end{aligned}$$

$$\sum_{m \in M} \sum_{i \in I} P(m) \times X_{n,m,i} \leq A(n) \quad (C3)$$

$$\forall n \in N$$

$$X_{n,m,i} = 0 \quad (C4)$$

$$\forall n \in N, \forall m \in M, \forall i \in I | L_{n,m} = 0 \text{ or } E_{m,i} = 0$$

$$\sum_{w' \in V} Y_{u,v,w'} \times Q(w') \leq C(w) \quad (C5)$$

$$\forall w = (u, v) \in F$$

$$\sum_{u \in N} \sum_{v \in N} Y_{u,v,w} \times D(u, v) \leq K(w) \quad (C6)$$

$$\forall w \in V, (u, v) \in F$$

$$\sum_{m \in M} \sum_{i \in I} X_{n,m,i} \times G(i) \leq B(n) \quad (C7)$$

$$\forall n \in N$$

$$Y_{u,v,w} = 0 \quad (C8)$$

$$\forall u, \forall v, \forall w \in V | (u, v) \notin F$$

$$\sum_{f \in N} Y_{n,f,w} - \sum_{f \in N} Y_{f,n,w} = \quad (C9)$$

$$\sum_{i \in I} X_{n,a,i} - \sum_{i \in I} X_{n,b,i}$$

$$\forall w = (a, b) \in V, \forall n \in N$$

$$X_{n,m,i} \leq U_n \quad (C10)$$

$$\forall n \in N, \forall m \in M, \forall i \in I$$

$$U_n \leq \sum_{m \in M} \sum_{i \in I} X_{n,m,i} \quad (C11)$$

$$\forall n \in N$$

$$Y_{u,v,w} \leq W_{u,v} \quad (C12)$$

$$\forall v \in V, \forall (u, v) \in F$$

$$W_{u,v} \leq \sum_{v \in V} Y_{u,v,w} \quad (C13)$$

$$\forall (u, v) \in F$$

$$X_{n,m,i} \in \{0, 1\} \quad (C14)$$

$$\forall n \in N, \forall m \in M, \forall i \in I$$

$$Y_{u,v,w} \in \{0, 1\} \quad (C15)$$

$$\forall u, \forall v, \forall w \in V$$

$$U_n \in \{0, 1\} \quad (C16)$$

$$\forall n \in N$$

$$W_{u,v} \in \{0, 1\} \quad (C17)$$

$$\forall u, \forall v \in V$$