INSTITUTO DE COMPUTAÇÃO
UNIVERSIDADE ESTADUAL DE CAMPINAS

Trade-off Between Bandwidth and Energy Consumption Minimization in Virtual Network Mapping

E. Rodriguez *G. P. Alkmim* *D. M. Batista*
N. L. S. da Fonseca

Technical Report - IC-13-32 - Relatório Técnico

October - 2013 - Outubro

Trade-off Between Bandwidth and Energy Consumption Minimization in Virtual Network Mapping

Esteban Rodriguez [*] Gustavo P. Alkmim [†] Daniel M. Batista [‡]
Nelson L. S. da Fonseca[§]

10/29/2013

Abstract

Network virtualization is a promising technology for the Internet of the Future. Nevertheless, an open issue in virtualization is to satisfy the control of resources in a way that energy savings are achieved. This paper introduces a model for the mapping of virtual networks onto network substrates which aims to reduce the energy consumption as well as to reduce the bandwidth consumption. This model is based on an integer linear programming formulation and several parameters, corresponding to characteristic of real networks, are considered. The trade-off between energy and bandwidth consumption is analyzed based on results derived via simulation.

1 Introduction

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. The impossibility of inclusion of new features in the core of the Internet has prevented the development of several applications and services. This has often been labeled the “ossification of the Internet”.

---

[*] Institute of Computing, University of Campinas, Campinas, SP, Brazil, 13089-971.

[†] Institute of Computing, University of Campinas, Campinas, SP, Brazil, 13089-971.

[‡] Institute of Mathematics and Statistics, University of São Paulo, São Paulo, Brazil

[§] Institute of Computing, University of Campinas, Campinas, SP, Brazil, 13089-971. This research was partially sponsored by Fapesp, process number 2008/07753-6, and CNPq.

This paper was published in E. Rodriguez, G. Prado Alkmim, D. Macedo Batista, and N. Saldanha da Fonseca, “Trade-off between bandwidth and energy consumption minimization in virtual network mapping” 2012 IEEE Latin America Transactions (LATINCOM 2012), vol. 11, no. 3, pp. 983–988, 2012. [1]

To overcome these limitations, various new architectures and mechanisms have been proposed to promote the evolution of the Internet [2] [3] [4]. Several of these are based on network virtualization which allows the definition of virtual networks (VNs) 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.

One of the main issues in network virtualization is the efficient mapping of VNs onto the substrate network [5] [3]. This mapping determines the allocation of routers and links of a VN onto the routers and links of the substrate network and can be aimed to achieve different objectives. However, the search for the optimal mapping of VNs is an NP-hard problem [4].

In recent years, telecommunication companies (telcos) and Internet Service Providers (ISPs) 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. [6], the increase in the volume of the 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 power consumption related to communication networks, which corresponds to 2% to 10% of the world current power consumption and this is expected to increase in the coming years.

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 advancement of the state of the art in energy efficient networking is expected to happen at the architectural level [7]. In this context, network virtualization plays a key role since it can replace a great amount of physical elements. Besides that, protocol processing and virtualized elements can be placed at sites with renewable energy.

Depending on the objective assumed different mappings for a virtual network on the same substrate can be obtained. Figure 1 illustrates two different mappings for the same virtual network: one assuming bandwidth reduction and the other energy consumption. Indeed, reducing the energy consumption quite often compromises the quality of service provisioning due to lower transmission rates and longer paths adopted. Therefore, investigating the extension to which minimizing the energy consumption jeopardizes quality of service provisioning is essential for virtual network mapping.

This paper analyses such trade-off by jointly minimizing the energy and the bandwidth consumption. The model considers realistic assumptions of virtual and physical networks as well as it considers the existence of repositories of images of different software and protocol stacks, which are used to instantiate the virtual routers. The proposal is based on 0–1 integer linear programming (ILP) formulation. The weight value assigned to balance the minimization of energy and the minimization of bandwidth consumption is varied. Results derived via simulation indicate that an equal weight leads a small variation of bandwidth consumption and yet an energy consumption close to the minimum value obtained when only energy consumption is minimized.

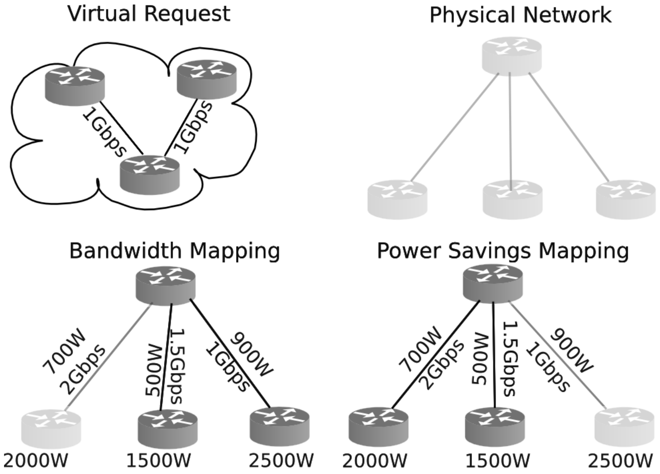

Figure 1: Mapping with different objectives

2 Proposed Formulation

The formulation in this paper models requests for virtual network establishment on network substrates that arrive dynamically. 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 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 energy and bandwidth consumption, and yet satisfying the requirements of the request.

The proposal is based on two 0-1 Integer Linear Programming (ILP) sub-models. The model considers realistic parameters, such as the existence of router images located in repositories. The formulation in this paper is a bi-criteria minimization, which considers both the objective function in [8] [9] and the objective function in [10]. 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; $C(u, v)$, $u, v \in N$, 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; $D(u, v)$, $u, v \in N$, 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 the power consumption of the network of chassis, physical link, linecards and cores, respectively.

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:

$$\alpha_n = \lceil \frac{K_n}{K_n + 1} \rceil \tag{1}$$

$$\beta_{u,v} = \lceil \frac{O_{u,v}}{O_{u,v} + 1} \rceil \tag{2}$$

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 formulation combines the minimization of bandwidth and energy consumption. The weight given to the bandwidth consumption is denoted by ϕ and the weight given to energy consumption is $1 - \phi$.

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+Band-Mapping) maps the virtual networks onto the substrate. The second (ILP-Green+Band-Image) determines the path in the substrate used to transfer the images.

The ILP-Green+Band-Mapping is formulated as follows:

Minimize

$$\phi[\sum_{u \in N}\sum_{v \in N}\sum_{w \in V} Y_{u,v,w} \times Q(w)]+$$
$$(1-\phi)[\mathcal{P}^{chassis}\sum_{n \in N}(\alpha_n + (1-\alpha_n)U_n)+$$
$$\mathcal{P}^{core}\sum_{n \in N}\sum_{m \in M}\sum_{i \in I}(X_{n,m,i} \times P(m))+$$
$$(2\mathcal{P}^{linecard}_{u,v} + \mathcal{P}^{E}_{u,v})\sum_{(u,v) \in F}(\beta_{u,v} + (1-\beta_{u,v})W_{u,v})]$$

subject to

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

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

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

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

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

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

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

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

$$\forall n \in N$$

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

$$\begin{aligned}
&\frac{\sum_{m \in M} \sum_{i \in I} X_{n,m,i}}{|M| \times |I|} \leq U_n && (C10)\\
&U_n \leq \sum_{m \in M} \sum_{i \in I} X_{n,m,i} && (C11)\\
&\forall n \in N &&
\end{aligned}$$

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

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

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

$$\begin{aligned}
&Y_{u,v,w} \in \{0,1\} && (C15)\\
&\forall u, \forall v, \forall w \in V &&
\end{aligned}$$

$$\begin{aligned}
&U_n \in \{0,1\} && (C16)\\
&\forall n \in N &&
\end{aligned}$$

$$\begin{aligned}
&W_{u,v} \in \{0,1\} && (C17)\\
&\forall u, \forall v \in V &&
\end{aligned}$$

The objective function minimizes the energy and bandwidth consumed by a request. The formulation can be customized to emphasize energy savings or bandwidth consumption. Such capacity is quite important for the management of such networks.

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+Band-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+Band-Image.

The ILP-Green+Band-Image is formulated as follows:

$$\text{Minimize} \sum_{m \in M} \sum_{u \in N} \sum_{v \in N|(u,v) \in F} Z_{u,v,m} \times D(u,v)$$

$$+ \frac{Z_{u,v,m} \times G(i|X_{n,m,i} = 1)}{C(u,v)} \text{ subject to}$$

$$\textstyle\sum_{m \in M} Z_{u,v,m} = 0 \qquad (C18)$$
$$\forall u, \forall v|(u,v) \notin F$$

$$\sum_{j \in N} Z_{u,j,m} - \sum_{j \in N} Z_{j,u,m} = \qquad (C19)$$

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

$$Z_{u,v,m} \in \{0, 1\} \qquad (C20)$$
$$\forall u, \forall v, \forall m \in M$$

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.

3 Performance Evaluation

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 with 95% confidence level were derived using the independent replication method. All 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 2.13GHz processors, with 4 cores each one, and 8GB of RAM.

To evaluate the performance of the proposed formulation, it was compared with two other formulations. The first one minimizes only the allocated bandwidth and was proposed in [8]. The second one minimizes only the power consumption and was proposed in [10]. In the remainder of this paper, the proposed formulation is denoted as BAND+GREEN, while the formulation which exclusively minimizes the allocated bandwidth is denoted as BAND($\phi = 1$) and the formulation which exclusively minimizes power consumption is denoted as GREEN($\phi = 0$). 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 and the amount of bandwidth allocated per request were evaluated as a function of ϕ.

Table 1 shows the value of the parameters used. Energy parameters in the simulations were obtained from several sources [11].

Due to the fact that the presented formulation is NP-Hard, an heuristic was developed. This heuristic finds solutions exclusively in the root node of the search tree, employing the *branch and cut* method [12]. The use of this heuristic was motivated by the fact that several solutions to the problem can be found at the root node of the search tree.

Both the topology of the substrate networks and that of the virtual networks were generated using the topology generator `BRITE` [13], with the BA-2 [14] method, a method that generates network topologies similar to the Internet. For the substrate network, the link delays were the values returned by `BRITE`.

Figure 2 shows the energy and bandwidth consumption for a network substrate with 200 nodes for different values of ϕ. The consumption is similar to that found in substrates with different size. As ϕ decreases the energy consumption also decreases. For $\phi = 0.5$ the energy consumption stabilizes at a values which is very close to that when $\phi = 0$. The bandwidth consumption, however, increases 10% when compared to the value when $\phi = 1.0$. The bandwidth consumption increases as ϕ decreases and when $\phi = 0$, it is extremely large.

Figures 3 and 4 compare, respectively, the energy and bandwidth consumption per

Table 1: Values of the parameters used in the simulation

Parameter	**Value**
Number of physical nodes	{140 200 300}
Bandwidth of each physical link	∼10240Mbps
Number of images in the network	3
Simulation time	5000s
Average arrival time per request	{25 }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 Power Consumption	10920W
Processor Power Consumption	166W
Line Card Power Consumption	450W
Amplifier Power Consumption	15W

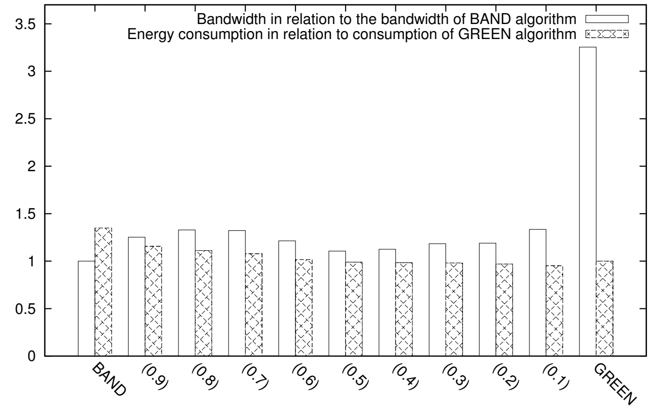

Figure 2: Results for scenarios with 200 nodes

request, obtained when $\phi = 1$ (BAND) and $\phi = 0$ (GREEN) and $\phi = 0.5$. It is clear that for different values of the substrate size the energy consumption per request when $\phi = 0.5$ is less than 10% higher than that produced by GREEN $\phi = 0$ and this difference decreases for larger substrates. Moreover the bandwidth consumption per request is less than 30% of that green by BAND.

Figure 5 shows the execution time of the algorithms per request. The GREEN algorithm execution time increases sharply and much higher than when no consideration of energy consumption is made (BAND). The execution time when $\phi = 0.5$ is closer to that of BAND for small networks but deviates from BAND when the substrate size increases.

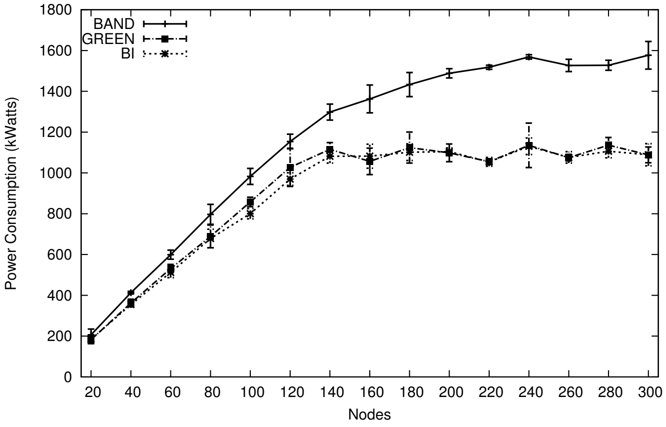

Figure 3: Average Power Consumption per Request

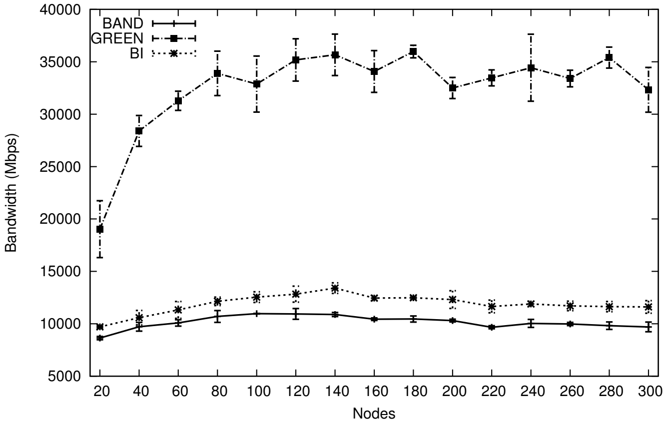

Figure 4: Average allocated Bandwidth

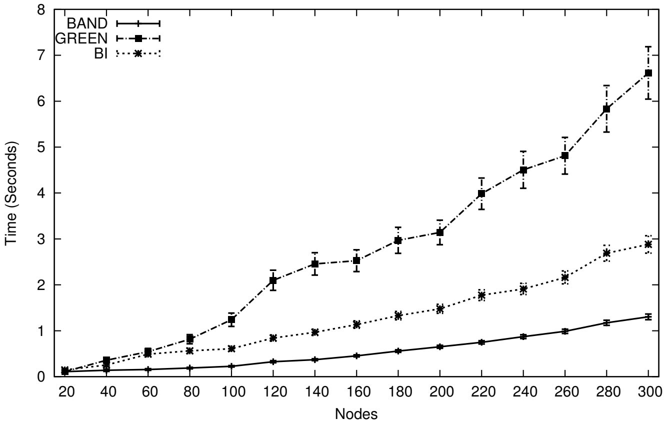

Figure 5: Execution Time

4 Conclusion

Minimization of energy consumption is currently a major concern in communications given its impact on the global energy consumption. However, having its reduction as a single objective can degrade the quality of service provisioning. This paper investigated the trade-off between the reduction of energy and bandwidth consumption. It was found out that to minimize only the energy consumption leads to extremely high bandwidth consumption. Moreover, giving equal weights to both consumptions yields an energy consumption per request very close to the minimum achieved and yet increases the bandwidth consumption by less than 30%.

References

[1] E. Rodriguez, G. Prado Alkmim, D. Macedo Batista, and N. Saldanha da Fonseca, "Trade-off between bandwidth and energy consumption minimization in virtual network mapping," *2012 IEEE Latin America Transactions (LATINCOM 2012)*, vol. 11, no. 3, pp. 983–988, 2012.

[2] N. Feamster, L. Gao, and J. Rexford, "How to Lease the Internet in Your Spare Time," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 61–64, 2007.

[3] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking Virtual Network Embedding: Substrate Support for Path Splitting and Migration," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 17–29, 2008.

[4] I. Houidi, W. Louati, and D. Zeghlache, "A Distributed and Autonomic Virtual Network Mapping Framework," in *ICAS '08*, 2008, pp. 241–247.

[5] N. Chowdhury, M. Rahman, and R. Boutaba, "Virtual Network Embedding with Coordinated Node and Link Mapping," in *IEEE INFOCOM*, Abril 2009, pp. 783–791.

[6] 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*, vol. 49, no. 8, pp. 80–86, august 2011.

[7] C. Despins, F. Labeau, T. L. 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*, vol. 49, no. 8, pp. 101–109, aug. 2011.

[8] G. Alkmim, D. Batista, and N. da Fonseca, "Optimal mapping of virtual networks," in *2011 IEEE Global Telecommunications Conference (GLOBECOM 2011)*, 2011, pp. 1–6.

[9] G. Alkmim, D. Batista, and N. da Fonseca, "Approximated algorithms for mapping virtual networks on network substrates," in *2012 IEEE International Conference on Communication (ICC 2012)*, 2012, pp. 1460–1465.

[10] E. Rodriguez, G. Alkmim, D. Batista, and N. da Fonseca, "Green virtualized networks," in *2012 IEEE International Conference on Communications (ICC 2012)*, 2012, pp. 1970–1975.

[11] Cisco, *Brochures for devices CRS-1 and components, CISCO ONS 15501 and others.*, Available at http://www.cisco.com/, acessed 06/2011, 2011.

[12] R. E. Gomory, "Outline of an algorithm for integer solutions to linear programs," *Bulletin of the American Society*, vol. 64, pp. 275–278, 1958.

[13] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE," 2010, http://www.cs.bu.edu/brite/. Accessed at 12/20/2010.

[14] R. Albert and A. L. Barabási, "Topology of Evolving Networks: Local Events and Universality," *Physical Review Letters*, vol. 85, no. 24, pp. 5234–5237, Dec 2000.