INSTITUTO DE COMPUTAÇÃO
UNIVERSIDADE ESTADUAL DE CAMPINAS

Live Migration in Green Virtualized Networks

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

Technical Report - IC-13-31 - Relatório Técnico

October - 2013 - Outubro

Live Migration in Green Virtualized Networks

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. An open issue in virtualization is the management of network resources in a way that energy savings are achieved without compromising the Quality of Service (QoS) requirements of the virtual networks. The dynamic allocation and deallocation of virtual networks can lead the state of the substrate to a less than optimum energy consumption. This paper introduces two algorithms for the migration of virtual routers and/or links which aims to allocate resources so that energy consumption is minimized. The efficacy of the migration of virtual routers and/or links and its impact on energy consumption are analyzed based on results derived via simulations.

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.

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 composed of virtual

*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. Alkmim, D. Batista, and N. da Fonseca, "Live migration in green virtualized networks" in 2013 IEEE International Conference on Communications (ICC 2013), 2013, pp. 1–6. [1]

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 virtual networks onto the substrate network [5] [3]. 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 [4].

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. [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 advance 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 use physical devices in a more efficient way, reducing the need of acquisition of new devices.

In [8], 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 optimum, leading to greater consumption of energy. This deviation from the minimum consumption of energy can be counteracted by reallocating resources to the virtual networks which demands live migration of virtual routers [9]. Figure 1 illustrates the benefit of resource reallocation. In Figure 1(a), virtual networks VN1, VN2 and VN3 use, the pair of nodes (R2,R3), (R1,R2) and (R1,R3), respectively. Figure 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 1(b). Figure 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 anymore, leading to a reduction of power consumption.

This paper 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

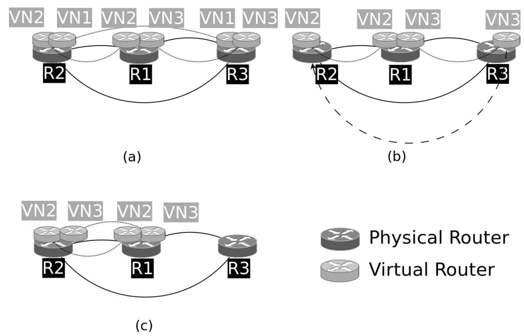

Figure 1: Example of reallocation of resources

that the remapping of virtual resources has a significant impact on energy consumption.

This paper is organized as follows: Section II summarizes related work. Section III introduces the proposed algorithms. Section IV discusses results produced by the algorithms and Section V draws conclusions.

2 Related Work

The work presented in [9] 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 [10] 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. This article shows that some of the transfer times involved in the live migration across large geographical distances can be negligible.

The work in [11] 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 [12] 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 [13] 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 [14] and [15] presented a virtual network mapping solution, considering a set of main characteristics of real networks. They extend their work in [8] by the introduction of the objective of minimizing the energy consumption. However, none of these works considered live migration.

The work in [16] 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.

3 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 network, finalize the transferring of the old virtual network and interrupt the services of the old virtual network is also negligible [10].

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. Before explaining the NRU algorithm, 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;

- $G^k = \{VN_m | a_{ijk}(t_k^-) = 1$ and $a_{ijk}(t_k^+) = 0_i^d \in D_k\}$ - 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 1: Algorithm 1: 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** $E(S(VN', t_z^+)) \leq E(S(VN, t_z^t))$ **then**
10 | return $S(VN', t_z^+)$;
11 **end if**
12 **else**
13 | return $S(VN, t_z^+)$;
14 **end if**

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.

Algorithm 2, called Live Migration-ALL (LM_ALL) aims to reduce power 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 2: Algorithm 2: 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**

4 Performance Evaluation

To assess the effectiveness of the proposed algorithms, the simulator introduced in [8] was extended to realize migration of virtual routers. The simulator receives a description of the substrate network as input and requests to instantiate virtual networks over the substrate. The algorithm used to re-map virtual networks employed in the experiment was introduced by the authors in [8]. It is based on an integer linear programming formulation and it has two main steps: one to map virtual routers and links and the second to define the routers on the substrate. The performance of the proposed algorithms was compared with a scheme that performs no migration. Both the topology of the substrate networks and the topology of virtual networks were generated using the BRITE [17] generator, with the BA-2 [18] method. Confidence intervals with 95% confidence level were derived using the independent replication method. The simulations experiments were executed on a computer running the operating system Debian GNU/Linux Squeeze, which was equipped with two Intel Xeon 2.13GHz processors, with 4 cores each, and 8GB of RAM. The values of the parameters used in the simulations are shown in Table 1.

Figure 2 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 request when

Table 1: Values of the parameters used in the simulation

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 × 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

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 smaller overhead.

Besides reducing the energy consumption per request, the proposed algorithms also consume less bandwidth than when no migration occurs (figure 3). The reduction can be at most of 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 difference on the physical length of paths associated to virtual links.

Figure 4 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 and 6 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

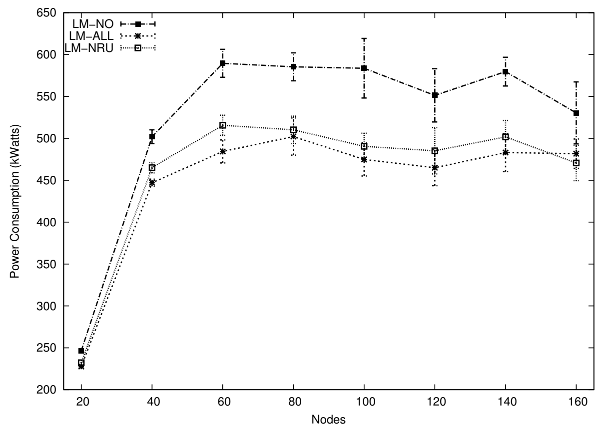

Figure 2: Average Power Consumption per Request

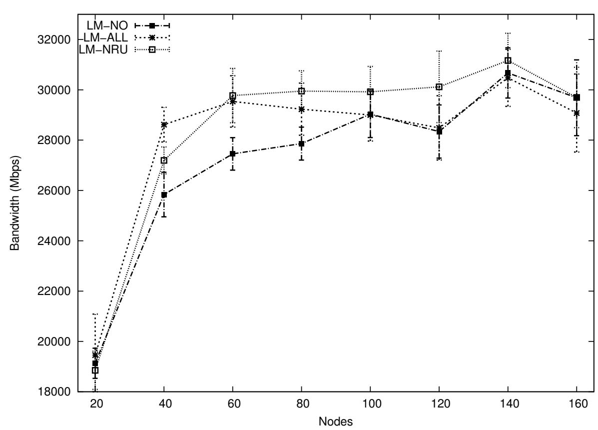

Figure 3: Average Allocated Bandwidth per Request

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 7 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.

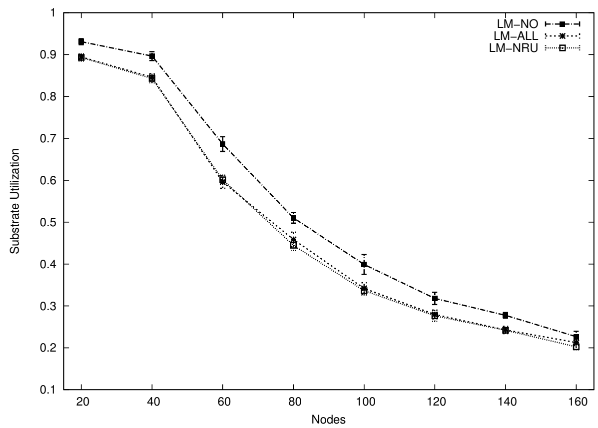

Figure 4: Substrate Utilization as a Function of the Substrate Size

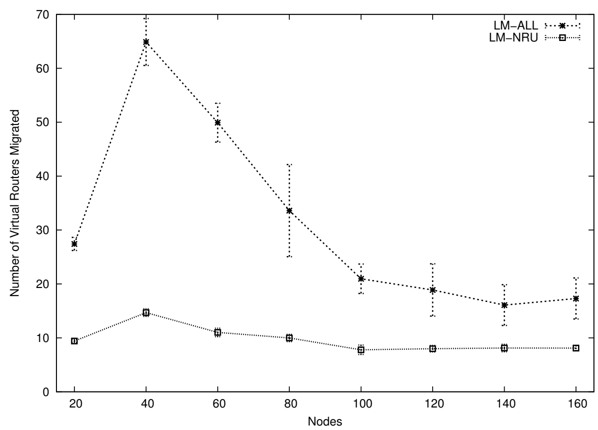

Figure 5: Average Virtual Routers Migrated per Event of Virtual Network Termination

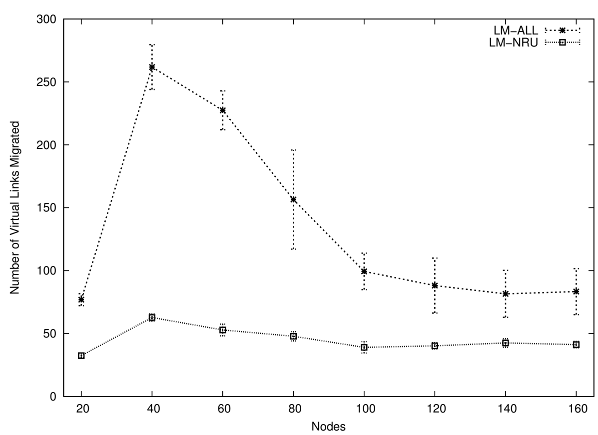

Figure 6: Average Virtual Links Migrated per Event of Virtual Network Termination

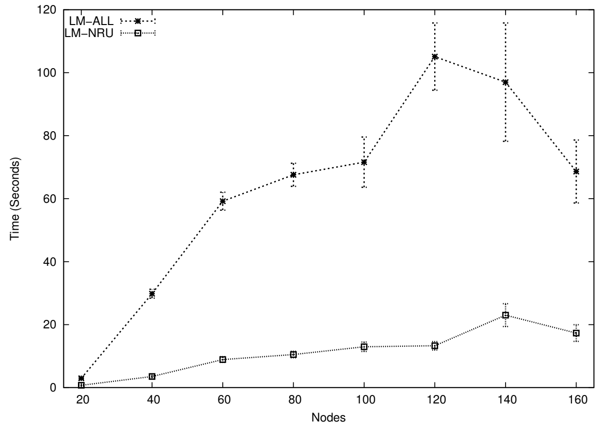

Figure 7: Execution Time per Event of Virtual Network Termination

5 Conclusion

This paper investigated the possibility of reduction in network energy consumption by remapping virtual networks on substrate networks. When there is a change in the state of the substrate occupancy, the remapping of virtual networks after the termination of one of them yields to significant energy savings. This paper presented two algorithms that use live migration for improving energy consumption of networks. 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.

References

[1] E. Rodriguez, G. Alkmim, D. Batista, and N. da Fonseca, "Live migration in green virtualized networks," in *2013 IEEE International Conference on Communications (ICC 2013)*, 2013, pp. 1–6.

[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 *In Proceedings of the 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] 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.

[9] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, ser. NSDI'05. Berkeley, CA, USA: USENIX Association, 2005, pp. 273–286.

[10] F. Travostino, P. Daspit, L. Gommans, C. Jog, C. de Laat, J. Mambretti, I. Monga, B. van Oudenaarde, S. Raghunath, and P. Y. Wang, "Seamless live migration of virtual machines over the man/wan," *Future Gener. Comput. Syst.*, vol. 22, no. 8, pp. 901–907, Oct. 2006.

[11] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, "Virtual routers on the move: live router migration as a network-management primitive," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 231–242, Aug. 2008.

[12] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation," *CoRR*, vol. abs/1109.4974, 2011.

[13] P. S. Pisa, N. C. Fern, H. E. T. Carvalho, M. D. D, M. Elias, M. Campista, L. H. M. K. Costa, and O. C. M. B. Duarte, "Openflow and xen-based virtual network migration," The World Computer Congress 2010 - Network of the Future Conf., Brisbane, Australia, Sept 2010.

[14] 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.

[15] 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.

[16] 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*, june 2012, pp. 261 –268.

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

[18] 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.