2010 10th IEEE International Conference on Computer and Information Technology (CIT 2010)

Incorporating Trust in Network Virtualization

Loubna Mekouar
University of Waterloo
Waterloo, Canada
Email: lmekouar@bbcr.uwaterloo.ca

Youssef Iraqi
Khalifa University
Sharjah, UAE
Email: Youssef.Iraqi@kustar.ac.ae

Raouf Boutaba
University of Waterloo
Waterloo, Canada
Email: rboutaba@bbcr.uwaterloo.ca

***Abstract*—In this paper, we propose a trust management framework for network virtualization environments. The proposed framework helps distinguish among the infrastructure providers based on their past experiences and feedbacks from service providers. We describe the main components involved in gathering the feedbacks and managing the reputation data. We detail the proposed trust system and we define the concept of *Degree of Involvement* of an infrastructure provider in terms of nodes and links. We also perform a mathematical analysis of the proposed system. Simulation results confirm the effectiveness of the proposed trust management system in increasing the service providers' satisfaction and reducing the selection of infrastructure providers that do not fulfill their Service Level Agreement (SLA).**

***Keywords*-Network virtualization; Trust; Reputation; Degree of Involvement;**

I. Introduction

Network virtualization is the technology that allows the simultaneous operation of multiple logical networks on a single common physical platform. By using this technology, distributed participants are able to create their own network with application-specific naming, routing, and resource management mechanisms.

A virtual network (VN) in the network virtualization environment is managed by one service provider that may require physical resources from several infrastructure providers. A VN is composed of virtual nodes connected by a set of virtual links. Distinct VNs coexist within a common physical network.

The followings are the main players in the network virtualization model:

- Infrastructure Providers (InP): manage the underlying physical infrastructure. Infrastructure providers offer their resources to service providers through programmable interfaces.
- Service Providers (SP): create and manage virtual networks using physical resources of multiple infrastructure providers. They provide end-to-end services deployed on virtual networks to the end users.
- End Users: can choose a variety of services from service providers and can deal at the same time with different service providers.

Recently, network virtualization has received tremendous attention. Since infrastructure virtualization is a key concept in the future Internet, different projects have been developed recently all over the world that are related to network virtualization such as: 4WARD, CABO, GENI, UCLP, Clean Slate, Trilogy, VINI, AKARI and PlanetLab. A list of all these projects can be found in [1].

In this paper, we propose a novel trust management framework in the context of network virtualization. We will describe the main components involved in the reputation management process and we present a mathematical analysis of the proposed reputation system.

The paper is organized as follows. Section II highlights related works and discusses the relation between trust and security in the context of network virtualization. Section III describes the proposed trust framework and the different components involved in managing the reputation data. Section IV presents the system model and introduces the concept of *Degree of Involvement*. Section V shows how the reputation of the infrastructure provider is computed. Section VI presents a mathematical analysis of the proposed reputation system. Section VII describes the performance evaluation conducted and presents the results that confirm the good performance of the proposed trust management system. Finally, Section VIII concludes the paper.

II. Related Works

PlanetLab is an open, global platform for developing, deploying, and accessing planetary-scale services [2]. PlanetLab takes advantage of nodes contributed by research organizations. These nodes host services on behalf of users (i.e. researchers and service developers). PlanetLab supports distributed virtualization. Each service runs in a slice of PlanetLab's global resources. Multiple slices run concurrently on PlanetLab.

PlanetLab is one of the projects that have tackled the problem of trust by defining trust relationships. In PlanetLab, the trusted PlanetLab Consortium plays the role of a trusted intermediary entity between the node owner and the service developer. Each owner is released from having to negotiate a hosting agreement with each user.

The Global Environment for Network Innovations (GENI) is a virtual laboratory for exploring future Internet at scale [3]. The designers have based GENI on the concept of slices which means that resources can be divided among different

978-0-7695-4108-2/10 $26.00 © 2010 IEEE
DOI 10.1109/CIT.2010.174

researchers in order to allow each researcher to run his experiment. In GENI, the basic role for a clearinghouse is to organize and manage trust relationships between the clearinghouse and research organizations, and between the clearinghouse and the aggregates operators.

In both PlanetLab and GENI, trust is addressed from the point of security only. In PlanetLab, the nodes should trust that the received code will not execute harmful programs and in GENI, the components should trust the identities (through authentication) of each other.

In the different network virtualization projects, trust if addressed, is always addressed from the security and privacy point of view only. Authentication, authorization, access control, ensuring integrity of information and protecting the source of information are used to provide a secure virtual network. However, there are other trust aspects that need to be taken into consideration. For example, we should be able to trust that an infrastructure provider will fulfill its part of the SLA by providing the agreed Quality of Service (QoS).

III. The proposed Trust Management

In service-oriented environments, Chang et al. [4], define trust as *"The belief the trusting agent has in the trusted agent's willingness and capability to deliver a mutually agreed service in a given context and in a given time slot"*.

To measure trust, reputation is used. The survey of different reputation systems reveals the important mechanisms used to achieve good reputation management [5]. We propose to address the reputation of an InP in terms of fulfillment of the required service as agreed (e.g., according to an SLA). We adapt our trust management framework [6] that was initially designed for P2P systems to the context of network virtualization.

In this section, we will describe the components involved in rating the infrastructure providers.

A. Gathering Trust Information

We assume that a SP can assess the quality of service of an infrastructure provider involved in a virtual network in terms of availability of resources, reliability, confidentiality and integrity, and adaptability to network conditions. The update of the InP reputation takes into consideration its *Degree of Involvement*. The *Degree of Involvement* represents the contribution of an infrastructure provider in the mapping of the virtual network. A more formal definition will be provided later in Section IV.

The feedback sent by a service provider could be a:

1) Binary value (1 or 0) indicating if the SP is satisfied from the transaction or not.
2) Discrete value (e.g. Excellent, Good, Fair, and Poor).
3) Real value on a continuous scale (e.g., [0,1]).

While a binary value does not allow partial trust, a continuous value expresses better how much trust is given. However, from the SP point of view, it is simpler to assign a binary value than a real one, hence, in this trust management framework, we adopt a binary value feedback.

The feedbacks sent by different service providers are gathered and stored. A *Trust Management Service* is used to keep track of trust data of infrastructure providers. In mapping a virtual network, the SP will take into consideration the reputation of the infrastructure providers.

B. Reputation Computation

A service provider may keep track of all the records of every infrastructure provider who was involved in a virtual network. Or, only one record that summarizes all the transactions is kept, which will reduce the storage cost.

After each transaction, we can consider the following scenarios:

- update the reputation of each infrastructure provider involved with an equal value.
- update the reputation of each infrastructure provider involved according to its *Degree of Involvement* (e.g., the number of nodes and links owned by this InP) in the virtual network. We adopt this approach in this trust management framework.

C. Using Reputation

Mapping a virtual network request requires the selection of specific nodes and links according to the requirement of a service providers in terms of resources (e.g., location and CPU of the nodes, and the bandwidth of the links) and cost [7], [8], [9], [10]. If service providers consider only the cost in the VN embedding, the infrastructure providers may be tempted to reduce the price by minimizing the quality of the physical underlying network. To make the right decisions, we propose to incorporate trust by taking into account the reputation of the infrastructure providers in a VN mapping. Avoiding untrusted physical network providers where failure of nodes and links could easily happen, will improve the service provided to the users. Service providers may reward reputable infrastructure providers by higher priority/probability of involvement in future VN mapping requests.

IV. The System Model

In this section, we consider similar notations as used in [10]. We model a substrate network i as a weighted undirected graph and denote it by $G_i^S = (N_i^S, E_i^S)$ where N_i^S is the set of substrate nodes and E_i^S is the set of substrate links. Each substrate node n_i^S is associated with the CPU capacity weight value $c(n_i^S)$. Each substrate link e_i^S between two substrate nodes is associated with the bandwidth capacity weight value $b(e_i^S)$ denoting the total amount of bandwidth. Similarly, VN requests are modeled by $G^V = (N^V, E^V)$ where each node from N^V is hosted by a substrate node and a virtual link from E^V can be assigned to a set of substrate links.

In this work, each G_i^S belongs to an InP_i and a service provider deals with several InPs.

The *Degree of Involvement* of an InP_i in a VN in terms of the nodes involved can be defined as follows:

$$I_{node}(i, N^V) = \frac{\left|\{n^V, n^V \uparrow n_i^S\}\right|}{|N^V|} \tag{1}$$

Where $n^V \uparrow n_i^S$ means that the virtual node n^V is assigned to the physical node n_i^S that belongs to InP_i.

For fairness issues, we propose an alternative way in computing the *Degree of Involvement* in terms of the nodes involved by considering the amount of CPU of these nodes:

$$I_{node}(i, N^V) = \frac{\sum_{\{n^V, n^V \uparrow n_i^S\}} c(n^V)}{\sum_{\{n^V \in N^V\}} c(n^V)} \tag{2}$$

Note that

$$\sum_i I_{node}(i, N^V) = 1 \tag{3}$$

Similarly, we define the *Degree of Involvement* of an InP_i in a VN in terms of the physical links involved as follows:

$$I_{link}(i, E^V) = \frac{\left|\{e_i^S, \exists e^V e^V \uparrow e_i^S\}\right|}{\left|\{e^S, \exists e^V e^V \uparrow e^S\}\right|} \tag{4}$$

Where $e^V \uparrow e_i^S$ means that the physical link e_i^S that belongs to InP_i is part of the virtual link e^V.

By considering the bandwidth of the physical links involved in the VN assignment, we obtain the following:

$$I_{link}(i, E^V) = \frac{\sum_{\{e_i^S, \exists e^V e^V \uparrow e_i^S\}} (b(e^V), e^V \uparrow e_i^S)}{\sum_{\{e^S, \exists e^V e^V \uparrow e^S\}} (b(e^V), e^V \uparrow e^S)} \tag{5}$$

Similarly, we have

$$\sum_i I_{link}(i, E^V) = 1 \tag{6}$$

Finally, we define the *Degree of Involvement* of an infrastructure provider as follows:

$$I(i, G^V) = \alpha I_{node}(i, N^V) + (1 - \alpha) I_{link}(i, E^V) \tag{7}$$

Where α represents the weight given to the *Degree of Involvement* in terms of nodes $I_{node}(i, N^V)$ and the *Degree of Involvement* in terms of links $I_{link}(i, E^V)$ of an infrastructure provider InP_i such that $0 \leq \alpha \leq 1$. An in-depth analysis can be realized for parameter α settings to achieve the best performance.

The objective is to maximize the *Degree of Involvement* of highly reputable InPs in embedding the VN k:

Maximize

$$\sum_{i/\exists n^V, n^V \uparrow n_i^S} R_i I_{node}(i, N^V) + \sum_{i/\exists e^V, e^V \uparrow e_i^S} R_i I_{link}(i, E^V) \tag{8}$$

Where R_i represents the reputation of InP_i.

In general, the cost of the physical infrastructure will be considered while mapping a VN request. Minimizing the cost will allow the service providers to increase their revenues. In this paper, we focus on the trust management framework. Incorporating the cost factor in Eq. 8 is left as future work.

V. Infrastructure Providers' Reputation

A. Notations and Assumptions

The following notations will be used:

- Let A_k be the rating of a virtual network VN_k by a service provider.
- Let $I(i, G^k)$ be the *Degree of Involvement* of the infrastructure provider InP_i in the virtual network VN_k.
- Let T_k be the lifetime of the VN_k.

B. The Reputation Management Scheme

At the end of a virtual network VN_k, the requesting service provider will evaluate the quality of service of the VN_k. If satisfied then a positive feedback is sent to the trust manager entity (i.e. the centralized server) otherwise a negative feedback is sent. We set $A_k = 1$ if the transaction is considered successful. If not, we set $A_k = -1$. In this case, the quality of service was not acceptable.

Each infrastructure provider InP_i in the system has the following *reputation data* (REP_{InP_i}), stored by the trust manager entity:

1) D_i^+: Satisfactory VN involvement,
2) D_i^-: Unsatisfactory VN involvement,

If we use the number of times, an infrastructure provider has been involved in the VN mapping, we get the following operation:

If $A_k = 1$ then $D_i^+ ++$, else $D_i^- ++$.

This scheme allows to rate infrastructure providers according to the number of times they participated in the VNs mapping requets. However, it does not take into consideration the lifetime of the VN and the *Degree of Involvement* since many infrastructure providers could be part of a single VN and the lifetime of each VN is variable.

We propose to take the lifetime of a VN and the *Degree of Involvement* of each InP_i into consideration. The reputation data of InP_i is updated according to the following operation:

$$\begin{array}{ll} if A_k = 1 & \text{then } D_i^+ = D_i^+ + I(i, G^k) T_k \\ & \text{else } D_i^- = D_i^- + I(i, G^k) T_k \end{array} \tag{9}$$

To compute the reputation of an InP_i, we propose to take into consideration the difference between D_i^+ and D_i^- and also the sum of these values as follows:

$$\begin{array}{ll} R_i = \frac{D_i^+ - D_i^-}{D_i^+ + D_i^-} & \text{if } (D_i^+ + D_i^-) \neq 0 \\ R_i = 0 & \text{otherwise} \end{array} \qquad (10)$$

Note that the reputation as defined in equation 10 is a real number between -1 (if $D_i^+ = 0$) and 1 (if $D_i^- = 0$).

When using this reputation scheme, a service provider can do one of the following in a new VN request mapping:

1) Choose the InP_i with the maximum value of R_i, or
2) Choose the set of InPs such that $R_i \geq R_{threshold}$, where $R_{threshold}$ is a parameter set according to the SP requirements (e.g. the cost of the infrastructure). If highly reputable InPs require a higher cost, reducing the required reputation value will give the SP the opportunity to satisfy the cost constraints.

VI. Mathematical Analysis

Let's assume that the infrastructure provider InP_i will provide a bad quality of service with a probability p_i. The goal is to show that the proposed reputation system is able to deduce this probability from the received feedbacks.

- Let X_n^i be the value of D_i^+ after contributing to the n^{th} virtual network.
- Let Y_n^i be the value of D_i^- after contributing to the n^{th} virtual network.
- Let $I(i, G^n)$ be the degree of involvement of the InP_i in the n^{th} virtual network.
- Let T_n be the lifetime of the n^{th} virtual network.

According to Eq. 10 we have $R_i = \frac{X_n^i - Y_n^i}{X_n^i + Y_n^i}$

Since InP_i does not fulfill the SLA with a probability p_i. This means that the value of Y_n^i will increase by $I(i, G^{n+1})T_{n+1}$ with probability p_i and the value of X_n^j will increase by $I(i, G^{n+1})T_{n+1}$ with probability $(1 - p_i)$. In other words, the new values of X_n^i and Y_n^i are:

$X_{n+1}^i = X_n^i + (1 - p_i)I(i, G^{n+1})T_{n+1}$
$Y_{n+1}^i = Y_n^i + p_i I(i, G^{n+1})T_{n+1}$

Let's find a closed formula for X_n^i and Y_n^i.
We have $X_n^i = X_{n-1}^i + (1 - p_i)I(i, G^n)T_n$
and $X_{n-1}^i = X_{n-2}^i + (1 - p_i)I(i, G^{n-1})T_{n-1}$
Similarly $X_2^i = X_1^i + (1 - p_i)I(i, G^2)T_2$
and $X_1^i = X_0^i + (1 - p_i)I(i, G^1)T_1$
$X_0^i = 0$
Summing up will lead to:
$X_n^i = (1 - p_i)\sum_{k=1}^{k=n} I(i, G^k)T_k$
Using the same approach we have:
$Y_n^i = p_i \sum_{k=1}^{k=n} I(i, G^k)T_k$
This means that the *reputation* value of InP_i is:

$$\begin{aligned} R_i &= \frac{X_n^i - Y_n^i}{X_n^i + Y_n^i} \\ &= \frac{[(1-p_i)\sum_{k=1}^{k=n} I(i,G^k)T_k] - [p_i \sum_{k=1}^{k=n} I(i,G^k)T_k]}{[(1-p_i)\sum_{k=1}^{k=n} I(i,G^k)T_k] + [p_i \sum_{k=1}^{k=n} I(i,G^k)T_k]} \\ &= \frac{(1-p_i-p_i)\sum_{k=1}^{k=n} I(i,G^k)T_k}{(1-p_i+p_i)\sum_{k=1}^{k=n} I(i,G^k)T_k} \\ &= 1 - 2p_i \end{aligned} \qquad (11)$$

If $p_i = 0$, which means that the InP_i fulfills the SLA, its reputation will be equal to 1. Using the same approach, a probability p_i equals to 1 will lead to the worst reputation value (i.e. -1). This shows that the reputation assigned by the trust management scheme reflects the behavior of the infrastructure provider.

This reputation computing technique is more general and can capture more elaborated InP behaviors. The results in Eq. 11 can be explained by the fact that we considered that the behavior of InP_i is totally captured by the probability p_i independently from other factors. However, an InP may have different probabilities of providing the agreed QoS depending on the VN lifetime and/or its involvement. In this case, we can consider the behavior of the InP to be captured by a probability distribution. Consequently, the involvement of the InP and the lifetime of the VN will affect the terms in X_n^i and Y_n^i.

VII. Performance Evaluation

We simulate the proposed trust management system and compare it with the Random Way algorithm (*RW*). Since no reputation management has been proposed previously for virtual networks environments, the selection of infrastructure providers by the service providers may be done in a random way.

A. Simulation Parameters

We use the following simulation parameters:

- The number of infrastructure providers is 1000 and the number of service providers is 10,000.
- VN lifetimes follow an exponential distribution with a mean of 100 hours.
- At the beginning of the simulation, each infrastructure provider can provide some of the resources (nodes and links) for a VN mapping.
- InP behavior distribution is as depicted in Table I.
- We simulate 40,000 requests. The simulations were repeated several times over which the results are averaged.

B. Performance Metrics

In this trust management framework, we consider that the feedbacks are given by the SP for each InP involved in its VN. The feedback is based on the quality of service as perceived by the SP for each InP. A successful transaction with an InP is when the SP is satisfied from the service

Category	Percentage	Probability of providing the agreed QoS
$InP1$	40%	0.95
$InP2$	20%	0.6
$InP3$	20%	0.5
$InP4$	20%	0.2

Table I
INP BEHAVIOR DISTRIBUTION

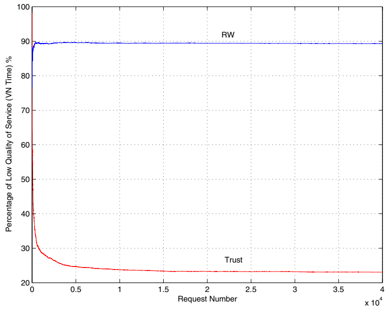

Figure 1. Percentage of low Quality of Service

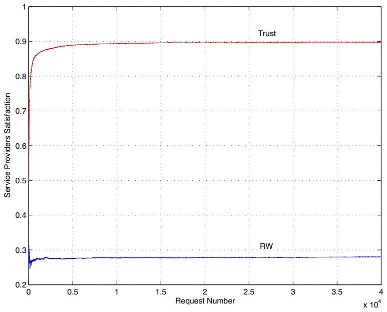

Figure 2. Service Providers' Satisfaction

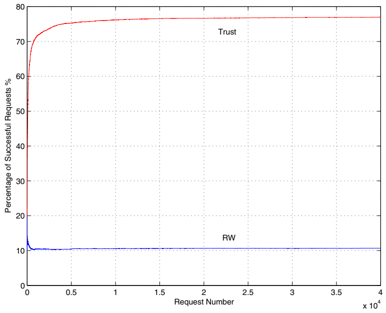

Figure 3. The Percentage of Successful Requests

provided. A successful transaction with all InPs is when the SP is satisfied from all of them. In these simulations, we focus on the following performance metrics:

- The percentage of low QoS: computed as the sum of the lifetime of all unsuccessful transactions over the total time of all VN.
- The service providers satisfaction: computed as the difference of successful and unsuccessful VNs in terms of time over the total lifetime of all the VNs.
- The percentage of successful requests: we consider a transaction to be successful only when all the InPs involved have provided the QoS as agreed.
- InP mapping share: to investigate the impact of the proposed trust management on the mapping distribution among InPs. The InP mapping share is computed as the sum of lifetime of all the VN mapped by this InP over the lifetime of all the VNs.

C. Simulation Results

Figure 1 depicts the percentage of low QoS (not as agreed) achieved by the two considered schemes. The X axis represents the number of requests while the Y axis represents the percentage of low QoS in terms of VN time. According to the figure, it is clear that the proposed trust management scheme (*Trust*) outperforms the *RW* scheme in terms of QoS provided to service providers. Without any reputation management scheme, we get 90% of VN with QoS not as agreed. The proposed trust management reduces this value to only 23%. The bad performance of *RW* can be explained by the fact that it does not distinguish between InPs that fulfill their SLA and those that do not. An InP is chosen randomly regardless of its behavior. The proposed trust management scheme is able to make the distinction and does not choose an InP if it is detected as an InP that does not provide a QoS as agreed. As a result, this technique controls the provided QoS for VN mapping requests and reduces the percentage of VN mapping transactions with a low QoS value. This will definitely increase the service providers' satisfaction.

Figure 2 depicts the difference of successful and unsuccessful VNs in terms of time over over the total lifetime of all the VNs. In this figure, the X axis represents the number of requests while the Y axis represents the SP satisfaction value. The maximum SP satisfaction that can be achieved is 1 while the minimum value is -1. SP satisfaction can be negative in case that the unsuccessful VN transactions surpass the successful ones. According to the figure, the proposed trust management scheme achieves a 0.9 value compared to the *RW* that achieves only a 0.28 value. This means that almost all the VN transactions were successful and the SP were satisfied from the service provided. Selecting highly reputable InP leads to increasing service providers satisfaction.

Figure 3 shows the percentage of successful transactions for both schemes. In *RW*, InPs are chosen randomly, and InPs can be selected from the ones that do not fulfill the agreed SLA, leading to a lower percentage of successful VN requests (11%) compared to the proposed trust management scheme (77%). The proposed scheme can quickly detect those InPs and, hence avoid choosing them for future VN

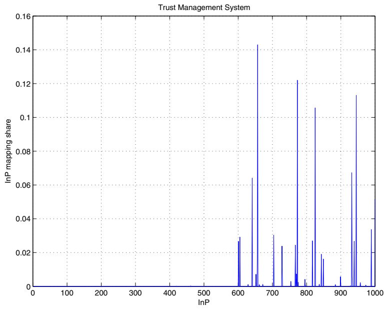

Figure 4. InP Mapping Share for the Proposed Trust Management

mapping requests. This achieves a high percentage of successful VN transactions and also, higher service providers' satisfaction as shown in figure 2.

To investigate the distribution of VN mapping share among InPs, figure 4 depicts the VN mapping share for each InP for one simulation run for the proposed trust management scheme.

Taking into consideration the InP behavior distribution and for clarity reasons, InPs with index from 1 to 200 belong to category $InP4$, InPs with index from 201 to 400 belong to category $InP3$, InPs with index from 401 to 600 belong to category $InP2$ while InPs with indix from 601 to 1000 belong to $InP1$.

The *RW* scheme distributes the load uniformly among the InPs regardless of the quality of service they are providing. It could easily happen that InPs that do not satisfy the agreed SLA are more often chosen to map VN requests which lead to a low service providers' satisfaction and a low percentage of successful transactions.

In figure 4, we can see that InPs that fail to satisfy their SLA are isolated and are not requested to perform any VN mappings. This is why the VN mapping share of these InPs (index from 1 to 600) is very small. On the other hand, the VN mappings supported by InPs providing a QoS as agreed (index from 601 to 1000) is higher than the ones in the *RW*. Almost all the VN mapping requests are performed by the InPs that provide the agreed QoS since the ones that do not satisfy the SLA are quickly detected and isolated.

VIII. Conclusion

Trust is of paramount importance in network virtualization. In this paper, we presented a novel trust management framework for network virtualization environments. We described the different components involved and we presented a mathematical analysis of the proposed reputation system. The performance evaluation results show that the proposed trust management system is able to identify the infrastructure providers that do not fulfill the agreed Qos and avoid selecting them to map future VN requests. This way, increasing service providers' satisfaction and the ratio of successful transactions.

In the proposed trust management system, service providers are motivated to deal only with highly reputable infrastructure providers for VN establishment. Another important factor to consider is the cost of the physical infrastructure. Combining the cost constraints with the proposed trust management is under investigation and is left for future work.

References

[1] "Future-internet," http://www.future-internet.eu/activities/fp7-projects.html/.

[2] "PlanetLab: An Open Platform for Developing, Deploying, and Accessing Planetary-scale Services," http://www.planet-lab.org/.

[3] "Geni," http://www.geni.net/.

[4] E. Chang, T. Dillon, and F. K. Hussain, *Trust and Reputation for Service-Oriented Environments*. Wiley, 2006.

[5] L. Mekouar, Y. Iraqi, and R. Boutaba, *Handbook of Peer-to-Peer Networking*. Springer, 2009, ch. Reputation Management in Peer-to-Peer Systems: Taxonomy and Anatomy.

[6] ——, "Peer-to-Peer Most Wanted: Malicious Peers," *Computer Networks Journal, Special Issue on Management in Peer-to-Peer Systems: Trust, Reputation and Security*, vol. 50, no. 4, pp. 545–562, 2006.

[7] Y. Zhu and M. Ammar, "Algorithms for Assigning Substrate Network Resources to Virtual Network Components," in *IEEE Conference on Computer Communications*, 2006, pp. 1–12.

[8] J. Lu and J. Turner, "Efficient Mapping of Virtual Networks onto a Shared Substrate," Washington University, USA, Tech. Rep., 2006.

[9] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking Virtual Network Embedding: Substrate support for Path Splitting and Migration," *ACM SIGCOMM Computer Communications Review*, vol. 38, no. 2, pp. 17–29, 2008.

[10] M. Chowdhury, M. Rahman, and R. Boutaba, "Virtual Network Embedding with Coordinated Node and Link Mapping," in *IEEE Conference on Computer Communications*, 2009, pp. 783–791.