Parameters and Challenges for Virtual Network Embedding in the Future Internet

David Stezenbach, Matthias Hartmann, Kurt Tutschku
University of Vienna, Chair of Future Communication (Endowed by A1 Telekom Austria AG)
{david.stezenbach, matthias.hartmann, kurt.tuschku}@univie.ac.at

***Abstract*—Recent developments in network science will facilitate the Internet as we know it today to integrate existing and upcoming technologies into a heterogeneous and highly dynamic resource pool. This enables the design of new applications and services which will form the Future Internet (FI). From this, many interesting prospects and unknown flexibility in terms of resource usage arise. In order to achieve this flexibility and make it usable for Service Providers (SPs) a key concept of the FI will be Virtual Networks (VNs) embedded into this resource pool. Furthermore, *Federation* which includes a closer interaction, resource sharing, and information exchange between providers will be an enabler for this freedom of design. This allows the operation of heterogeneous and divergent services on the same physical infrastructure substrate. It is clear to see that all these upcoming possibilities offer new areas of research. In this work we structure the parameters required to describe resources and services accurately. Furthermore, we summarize research questions that need to be resolved in order to make the FI a story of success.**

***Index Terms*—Future Internet; Virtual Network; Federation**

I. INTRODUCTION

The Internet as we know it today has evolved from a plain data transmission network to a highly dynamic and heterogeneous system. Recently developed and future applications will have quality requirements which go far beyond basic connectivity with a specific bandwidth. As applications are globally distributed a single provider which has only a regional scope is not able to deal with these requirements. It is necessary to have a closer interaction between providers. Federation enables the providers to share resources across their borders which is a deeper cooperation than Service Level Agreements (SLAs) or peering contracts.

One possibility to establish Federation is to deploy VNs across multiple providers. These VNs are highly flexible and can be compared to the resource usage in cloud computing environments where resources are dynamically added when needed. This fulfills the requirement of variable infrastructure regarding daytime or even minutes which will be required for future applications and services.

From this approach many questions and research topics come up and will be discussed in this paper. The physical resource description and the application requirements characterization are fundamental questions as both are required to solve the Virtual Network Embedding (VNE) problem. Thus, we will structure parameters and metrics for both and outline which are usable at different points in the hierarchy from the Physical Infrastructure Providers (PIPs) up to the SPs. We agree, that a brief description of resources and services offered by the different providers is essential. In addition, we argue that basic descriptions like connectivity or bandwidth will not be enough. The current approach of resource description inludes topology information and mean value analysis. For a highly dynamic scenario as the FI these are not adequate and metrics like standard deviation or distribution functions need to be taken into account. But on the other hand it is important to have only a small subset of parameters in order to keep the VNEs problem small.

Furthermore, we discuss how certain parameters could be measured and how these information are exchanged between different federating providers. This includes questions on how resources are broadcast and who deploys and operates the different VNs. We argue, that at least on a logical level a new “player” that handles the shared and federated resources is required between the network and the applications.

The remainder of this paper is structure as follows. In Section II we present the main concepts of the FI. We continue with Section III that presents an overview of metrics and parameters for resource, service, and application description. We discuss possible design and research challenges of the FI in Section IV. Finally, Section V wraps up the paper.

II. BUILDING BLOCKS OF THE FUTURE INTERNET

The Internet as we know it today does not have any significant problems in terms of performance or expandability. From the perspective of basic connectivity and the increased demand of bandwidth, overprovisioning as a strategy of topology dimensioning, is currently used by Internet Service Providers (ISPs) and does its job. But, regarding efficient and flexible resource utilization in the case of dynamically changing demands it performs fairly poorly. Without the investment in physical equipment topology changes are not possible. Also the ossification [1] respecting the TCP/IP protocol stack prohibits the development and deployment of new protocols and applications in the current Internet.

Optimization Challenges in the Operation of the Future, Federated Internet (OptFI) has been funded by the Vienna Science and Technology Fund (WWTF) through project ICT10-027.

This work was funded by EuroNF through the SJRP JRA.S.45 Algorithms and Operating Procedures for Federated Virtualized Networks (FEDNET) and the WP IA.4.1 on ”Mobility of Ph.D. Students

A. Future Internet

In the early years of its development the Internet applications just used the given capabilities. These mainly consists of basic connectivity and were used for data transmission only. Subsequently, new use cases were build on top of these capabilities and new services evolved. This approach changed significantly in the last years. Today the application is defined first and the network is then adapted in order to fit the requirements of the desired service. A good example for that is the usage of HTTP based on TCP for video streaming [2]. There are many protocols specially designed for video streaming with more capabilities. But the deployment of new protocols is difficult. Thus we use HTTP for this purpose and stay with less potential as the Internet was not designed for this evolution strategy. One strategy could be the modification of HTTP which blows up the complexity. A very straight forward approach is the direct support of the application by the network with means for traffic handling and the dynamic integration of new protocols and applications within an isolated VN.

Several research projects aim to find solutions for the FI. The GENI program [3] offering a virtual laboratory at the frontiers of network science, the FIND initiative [4] working out the requirements of a global network 15 years from now, the Autonomic Internet project [5] evaluating the possibilities of an intelligent autonomic network supporting end-to-end applications, and the FIRE initiative as an experimentally driven search for revolutionary ideas for new networking and service paradigms.

A hot discussion in the research community is whether to instantiate a new clean-slate version of the Internet [6] or to extend the current one with required features [7], [8]. As fundamental changes in the Internet architecture are hardly deployable, we argue that an approach based on the current structure will lead to a faster deployment. Furthermore, the Internet has proven to be an amazingly scalable distributed system. Thus, we recommend an adaptation of the provided capabilities based on the experiences made with a system supporting billions of users, while retaining the existing basic structure.

Taking both of these arguments into account, the design of the FI should be a step towards an approach that offers as much support for application requirements as possible. As the FI is supposed to be a service and content centric [8] network, one should consider connectivity as a service itself. This should be a main objective of the FI since services have proven themselves being more flexible than physical infrastructure. This is a first step towards an ubiquitous network which is one prerequisite for applications that do not bother its users with insufficient connectivity. A good example of an application that handles its connectivity transparent for the user is Amazon Kindle.

A further goal of the FI will be the support of a heterogeneous pool of applications with rather divergent requirements. Thus, it is inevitable to have a closer interaction between the applications and the network in order to adapt proactivly to changed situations. As typical network applications and services operate across ISP borders it is necessary to operate the network adaptation in a borderless fashion too.

B. Federation

A further question in the FI is how the application, service, and infrastructure provider landscape will be structured. Some of theses questions cannot be answered by computer scientists since they require economic, legal, and even social decisions. But we can provide options pointing out design pattern for the FI.

In this work we will have a closer look at two patterns. The first we refer to as *vertical interaction* is how infrastructure, service provider, and application providers are going to interact. Second the *horizontal interaction* which describes the cooperation and information exchange of providers on the same hierarchy level. In recent publications the term Federation is common regarding horizontal cooperation. As providers could operate on multiple hierarchies and thus both a horizontal and a vertical resource leasing is possible, we use the term Federation in both scenarios in order to highlight the close collaboration of the different providers.Virtual Network Providers (VNPs)

The current trend leads to ISPs which provide their customers with a whole bunch of services. This has two drawbacks in our opinion. First, the users get services which they do not need, but pay for them as they are bundled in a package. Second, the providers are loosing their specialization which is their core competence.

As proposed in the literature [9]–[11] connectivity should be offered by specialized providers. In this work we follow the naming of [9] which depicts the PIPs as the physical substrate owner, the VNPs bundle these resources to Virtual Topologies (VTs) and offer plain connectivity. On top of that the Virtual Network Operators (VNOs) offer full service connectivity. This could include routing, traffic management, accounting, security, and monitoring. These full featured topologies are named VNs. For clarification, the different hierarchy levels do not need to be strictly separated as they describe only logical administrative responsibilities. In the following, we describe which possibilities and challenges arise from an approach that is based on the outlined breakdown of responsibilities.

The hierarchy in connectivity provisioning offers the SPs a topology based on heterogeneous and tailored resources. Each topology can be adapted to the respective resource requirements of the application that runs within this topology only. This enables the SP in cooperation with the VNO to operate the VT with optimizations on routing, security, and general traffic management. At the very top of this hierarchy pyramid the user is able to either select single services as needed or packages offered by a third-party reseller.

One reason for ISPs to act as a representative for each stage of the pyramid is that this is currently their only way to guarantee certain service quality. In a multi provider infrastructure and service landscape resources and services operated by others are used and redistributed. The benefits for the providers are that they gain a high grade of flexibility and diversity [12] as they do not need to deploy new hardware at different sites if they want to modify or extend their service. This resource pooling across multiple providers we refer to as *Federation*. This clearly requires the FI to include new types of SLAs and their validation [13] in order to enable each provider to guarantee a service quality for its offered service (e.g., connectivity, topology, or application).

C. Network as a Service

The concept of Federation requires higher flexibility, faster resource acquisition, and closer reliable coupling of providers. Furthermore, transparent resource booking with less configuration effort and administrative interventions is required in order to allow an fully automated leasing process. These are at a first glance mutually contradictory goals that need to be consolidated.

In order to accomplish these challenges we need to detach from the assumption that connectivity is predefined and hardly changeable. The FI needs to treat connectivity as a variable and dynamically modifiable service, we refer to as *Network as a Service (NaaS)*. In the FI this will be enabled with network virtualization [1]. In combination with others services this provides a heterogeneous resource pool containing the building blocks for future applications and services.

Currently there are many projects running investigations on the performance and the design of virtualized networks. The *Trellis* [14] platform is a container based approach using common virtualization technologies like VServer. The links are virtualized with packet tunneling. Each virtual network consists of isolated links and nodes which are mapped on the physical hardware. An experimental platform for new network protocols based on virtualization is VINI [15]. It enables the experimenters to evaluate their protocols and architectures in a realistic network environment. The project ExpoNet [16] presents a software and hardware based router virtualization. The customization of the control plane and the data plane is possible with respecting experiment requirements. The previous mentioned examples based on virtualization of a hardware system. The ADVisor [17] project uses OpenFlow to establish virtual networks based on OpenFlow enabled switches.

The VNs established by one of the different technologies are offered by the VNOs as a basic connectivity service and are made available to the SPs.These VNs need not only be described in terms of topology and traffic management patterns, but also the underlying infrastructure description is required. Sufficient substrate description is required to enable the VNPs and VNOs to the select adequate infrastructure from the PIPs. Fine grained resource descriptions will facilitate the development of autonomous resource leasing mechanisms. These mechanisms are a central element of the FI, as only this way a rapid resource selection and adaptation is possible. To be able to guarantee a certain service quality the autonomous validation of SLAs and adaptation of the VNs is mandatory. The proper selection of resources from an underlying substrate is the *VNE problem.*

In the following, we summarize possible metrics and important parameters for the VNE problem that need to be discussed as applications and services in the FI will have fairly divergent resource requirements. These need to be described accurately by the SPs in order to generate tailored VNs.

III. APPLICATION AND RESOURCE DESCRIPTION

SPs that want to use a virtual network, and the different parties (VNOs, VNPs, and PIPs) that together provide the VT need a common language to describe their resources and requirements. Most previous work considers only link-bandwidth and router CPU requirements. Some papers also include geo-location, either as fixed coordinates, or in terms of delay. In this section, we provide an overview over parameters, performance metrics, and requirements that can be important in the VNE problem.

A. Substrate Network Parameters

A network can be considered as a number of network nodes which are interconnected by network links. For the description of the substrate network, the important parameters describe the properties of these links and nodes. It is essential that metrics such as delay, delay variation and loss are defined consistently [18] in order to support QoS meaningfully across multiple providers.

1) Resource Description: The first group of parameters consists of the physical descriptions of network entities. A link is described by the maximum transmittable bandwidth (i.e., its capacity) and its propagation delay. But there are also other parameters that could or must be taken into account:

- capacity
- propagation delay
- bit error rate (BER)
- technology
- geographical link location

A router in its simplest form is described by its total forwarding capacity and the forwarding delay. A more adequate description of a router could include all of the following:

- computation capacity (not only for routing, but also for general purpose tasks in virtualized environments)
- network I/O
- storage I/O (e.g., for routing tables, VLAN tags)
- memory (RAM)
- disk capacity
- buffer size
- forwarding delay
- forwarding capacity
- processing delay
- geographical location

Each link and router can fail and failures can affect many different virtual networks. Thus, information about the failure behavior of network components is also important. A PIP must also provide information about shared risk groups and also about the resilience mechanisms that are in place to alleviate the effects of failures. The following list summarizes these parameters:

- packet loss probability
- loss patterns
- mean time between failures (MTBF)
- mean time to repair (MTTR)
- shared risk group
- resilience mechanisms

When network components are not observed individually, but as a whole topology, more parameters are becoming important. Several links and nodes form a path between ingress and egress routers with additional parameters:

- jitter
- packet reordering
- path bandwidth
- delay
- delay/jitter under changing load
- current Workload
- capacity

Another important parameter of network elements is energy consumption. A PIP could either hide this information in its general cost calculation, or transparently pass it to the VNP. This information could be important to save costs by creating energy efficient virtual networks.

2) Virtualization Techniques: In order to lease only parts of physical entities these need to be virtualized. Depending on the virtualization technology used the behavior is different. Thus, it is necessary to describe how a virtual entity is realized on the physical substrate. Therefore, it is required to provide the following information.

Layer of link virtualization:

- physical layer (multiplexing: WDM, TDM, FDM, etc.)
- data link layer (VLAN tagging)
- network layer (IP, tunnels (IP in IP, GRE, IPsec))
- transport layer (e.g., VPN)

Router/Node virtualization:

- flow-based (e.g., OpenFlow)
- hardware isolated (e.g., line cards, network processing engines, fast path processing)
- software isolated (e.g., software router, operating system)

3) Parameter Properties: Many of the previously described parameters cannot simply be described by a fix scalar number, but might require the specification of a parameter range or even a distribution. If, for example, a minimum bandwidth for certain links is leased by a VNP, the question arises, whether and how much of the additional (available) bandwidth can be granted to this VNP. If too much spare capacity is provided by the PIP, the VNP will never pay for a higher minimum capacity. On the other side, the unused capacity of a PIP does not bring profit, and could / should be used to improve performance and happiness of already paying customers. The specification of a parameter range could be an alternative or better way compared with minimum description only. For example, describe a link with a minimum bandwidth of 1MBit that must be available under all circumstances, up to a maximum (best-effort) bandwidth of 10MBit. Maybe even the specification of an additional guaranteed long-term average bandwidth of 3Mbit could be required.

The degree of isolation between different VT can also be specified. Most networks are not designed and deployed exclusively for virtual networks on top, but themselves carry different types of traffic. If best-effort traffic is using capacity, it must be handled or limited in such a way that it does not interfere with traffic in the virtual networks.

4) Time Scales: A physical network cannot be changed or adapted quickly to changing requirements. Providers usually prefer long-scale agreements and contracts for planning dependability. On the other hand, VNPs want to react quickly to changing demands and often want to be able to lease additional resources quickly.

When a VN is embedded into the physical substrate and demands change over time, it might be wise to re-embed the network in a different way. As this might cause some service interruption it should happen rarely.

Thus, time constraints must be provided that specify

- minimum lease time
- setup time requirements for new virtual networks
- minimum / maximum time between re-embedding

5) Federation Parameters: Running virtual networks in a federated environment that consists of several different PIPs requires even more parameters for an adequate description. Which parameters are of interest in this context? Federation metrics (between PIPs and VNPs) are probably agreed upon on a much larger timescale, compared to the more dynamic behavior of application metrics.

- resource ownership
- restrictions (e.g., technical, legal, social, or business)
- offered services (e.g., configuration, embedding, management, or monitoring)

6) Obscuring Parameters: Another problem is that PIPs are usually not willing to reveal their complete internal network structure for competitive reasons. Thus, the exact geo-location and exact parameters of link-placement are most likely not known to higher levels in the VT creation hierarchy. This is especially bad for resilience and security requirements, when it is unknown, whether several links from different providers are using the same cable ducts (e.g., across a river or sea, or through a tunnel), or whether several routers from different providers are actually located in the same area or even the same building.

B. *Virtual Network Requirements*

Network virtualization requires a combination of node and link virtualization. A virtual node is usually deployed inside a single physical node and a virtual link is often represented by a path inside the substrate network. The VNPs, VNOs, and SPs can have the following requirements for the links, paths and nodes in their virtual network:

- bandwidth demands per link / bandwidth as traffic matrix
- minimum availability
- maximum end-to-end packet loss
- maximum hop-by-hop packet loss
- security (e.g., encryption, authentication)
- reliability
- configuration costs
- maximum delay
- jitter

The VNE tries to fulfill these requirement parameters by mapping the VN elements onto adequate physical substrate entities.

In each level of virtualization, it is important to be able to find resources provided by the level below. So, either some form of centralized resource broker is needed that collects resource information and advertises them to the next higher level, or some form of decentralized resource discovery protocol is needed.

The arrival rate of VN requests must somehow be determined in order to scale resource discovery and embedding mechanisms. Also the dynamic behavior of virtual networks and network services must be considered and described. Especially, the description of requirements that change over time. Here, different types of dynamics can be observed. For example, the more or less predictable demand differences between day and night, or weekday and weekends, or less predictable behavior, like customer behavior changes or application popularity.

The delay of a virtual link is usually the delay sum of each physical link on the path, combined with the total delay in all hidden-hop-routers. The router delay is not fixed, but may depend on the current load situation. When the router is under high load, a) the processing delay could increase due to high CPU load, and b) the total delay may increase due to extensive packet-buffering.

C. *Parameter Selection*

In the end, an optimal solution for the VNE problem, considering all possible parameters and dynamic behavior, is impossible to achieve. The large number of parameters does not allow efficient algorithms. So, meaningful subsets of parameters must be selected that represent the important properties.

An embedding can be evaluated with different metrics: Simple metrics are the *cost*, i.e. the resources spent, and the *revenue* metric, i.e. the economic benefit of accepting VN requests. A list of such metrics can for example be found in [19].

IV. DESIGN AND RESEARCH CHALLENGES

In this paper we presented many parameters and additional descriptions regarding networks, applications, and services. Beyond that, there are many other questions that need to be answered in order to provide the FI with the right means for being a well designed successor of todays Internet. This section presents questions and challenges that the research community should think about.

A. *Physical Substrate Description*

- *How to identify, broadcast, and lease federated resources?*
  The FIs will have a very high number of resources with tremendous diversity. Thus, one need to find feasible ways to advertise these resources. This needs to take into account where resource information is stored (centralized or decentralized), which information is exchanged between whom, and how new resources are integrated in the system.
- *Which of the parameters are measurable?*
  Certainly, it is always good to know as much as possible about the physical substrate or the underlying hierarchy level. But, one need to question if the benefit gained from a parameter is higher than the cost for its measurement. Thus, it is necessary to evaluate which parameters are measurable within a running system and if its possible to express one as the combination of others with the same information received.
- *Which is the right subset of parameters for the VNE problem?*
  This question cannot be answered generally. The selection of parameters highly depends on the current use case. But, one goal of further research has to be the appropriate selection of parameters for given services. This is required in order to keep the balance between a small and solvable VNE problem and a sufficient embedding quality.

B. *Virtual Network Embedding*

- *Do we need independent VNs for each application?*
  The answer to this question has two sides of a coin. One VN for each application allows to parametrize it very close to the requirements. But, with an increasing number of VNs the system is getting more complex and we run the risk of getting into the same problems as we currently do with large routing tables. Additionally, virtualization technologies often exhibits a non-linear scaling which necessitates an accumulation of similar traffic in order to keep the number of VNs small.
- *How to describe the resource usage of the VNs?*
  The description of the resources needed by a VN goes beyond simple parameters like maximum bandwidth usage. As VNs are changing dynamically they have to be described like this. Average value, coefficient of variation, or distribution function could be the right choice for that and needs to be respected by the VNE problem.

- *How to map the request parameters of the VNs to the underlying substrate network?*
  The resource requirements of the VNs needs to be somehow mapped into the physical substrate in order to solve the VNE problem. Since, the behavior of virtualized resources often does not scale linear and no full isolation is given, one need to take into account the interaction between the different VNs and the physical substrate. Furthermore, some limitations of the physical substrate (e.g., computation capacity on a router) can not be directly mapped to usage parameters of a VN (e.g., bandwidth usage).
- *How to run the embedding process of the VNs?*
  The embedding could run based on fine grain resource description. This would allow an appropriate selection of the entities of a VN. As this increases the complexity of the VNE problem one need to think about an hierarchical embedding. Could it be implemented in a way analog to the routing in IP networks, which are Autonomous System (AS) based on a higher level, without forfeiting the benefits? This would require the ISPs to announce their resources based on connectivity through their own network without spreading out the full topology.

C. Virtual Network Operation

- *How to operate the VNs?*
  This question covers multiple topics. First, one needs to think about at which point the traffic of the users enters the VNs. Should this be done directly on the users machine, like a tunnel from the application to its first node in the core network? Should the provider take care of the classification and aggregate the traffic on the access nodes? Furthermore, this question includes how static the VNs are operated regarding time scales of modifications and setup times.
- *Which virtualization technology to use for establishing the VTs?*
  In a heterogeneous environment with multiple sometimes competing players a uniform solution is not possible and sometimes not even desirable. Thus, further research needs to be done in order to find the right technology for each use case. At this juncture, one should focus on how a single virtualization influences the quality of a service running in it and how different virtualization technologies interfere each other.
- *Are VNs with full isolation always the right choice?*
  A nearly full isolation clearly has its advantages. The application can be handled disregarding the requirements of others. Fairness can be achieved on VN level and applications do not have to deal with it. On the other hand, sometimes the effort of creating and maintaining VNs exceeds the benefits gained from it. Lightweight virtualization systems build on top of connection based forwarding could be an alternative here.
- *How to administer resources in a foreign network?*
  If a VNP uses resources from multiple PIPs an administration interface needs to be set up in order to allow the VNP to configure the foreign resources in an appropriate way. This also needs to take into account questions on security, monitoring and the surveillance of SLAs.
- *Is a controller necessary for each VN or is a self-organizing approach sufficient?*
  In current testbeds the resources are managed by a central instance, mostly known as a clearinghouse. Is it feasible to transfer this approach to federated productive resources? One can imagine a control instance for each VN, a central clearinghouse within each provider network, or even a self-organizing mechanism like a peer-to-peer system as a way to organize the network.

D. Virtual Network Management

- *Do we need new SLAs and monitoring approaches?*
  Probably these new upcoming technologies will require flexible and adaptable methods for the validation of SLAs. Should SLAs be validated from within the VNs or be based on measurements of the underlying technology.
- *How to account the usage of these new services types?*
  Accounting based on bandwidth or throughput will not adequately fit the requirements of FI services. The upcoming of highly dynamic resource usage will require new accounting strategies taking into account the variability of resources over time. Thus, new methods of resource evaluation and usage quantification are necessary.

This section gave an overview on possible research fields in the FI. We focused on topics related to network scientists. Nevertheless, it is also worth to think about questions regarding economic (e.g., business models), social (e.g., user acceptance), and legal (e.g., privacy) issues.

V. CONCLUSION

The FI will have VNs and Federation, a concept of resource sharing, as its main building blocks in order to allow a coexisting operation of different technologies on the same physical substrate. In this paper, we provided an overview over possible parameters and metrics for the description of future applications, services, and the physical resource substrate. It was detailed that a mean value analysis of the parameters will sometimes fail to describe the VNE problem appropriately. Thus, one need to think about alternative ways to express parameters and metrics in order to reflect the highly dynamic characteristics of future applications and services in an adequate way. We provided a structuring and pointed out to which level of the hierarchy they relate to. We recommend that these parameters should be covered by future researches in this area. Resource descriptions currently in use will not cope the requirements of heterogeneous resources and highly dynamic services in the FI. Thus, we summarized further research challenges in the field of Federation and the VNE problem. We encourage the research community to think about these questions and to come up with solutions that enables the FI to be a successor of the current Internet that offers a broad range of design freedom to its users.

VI. Acknowledgment

The authors would like to thank Günther Raidl and Johannes Inführ from the Vienna University of Technology for their highly stimulating collaboration in this joint research project. Furthermore, we would like to acknowledge the competent input of Katharina Pfeffer and the fruitful discussions we had with her.

References

[1] T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the internet impasse through virtualization," *Computer*, vol. 38, no. 4, pp. 34– 41, Apr. 2005.

[2] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "Youtube traffic characterization: a view from the edge," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, Oct. 2007, p. 15–28. [Online]. Available: http://doi.acm.org/10.1145/1298306.1298310

[3] "Global environment for network innovation (GENI) program," http://www.geni.net/. [Online]. Available: http://www.geni.net/

[4] "Future internet design (FIND) program," http://www.nets-find.net/. [Online]. Available: http://www.nets-find.net/

[5] "EU-IST autonomic internet (AutoI) project," http://ist-autoi.eu/autoi/. [Online]. Available: http://ist-autoi.eu/autoi/

[6] T. Koponena, S. Shenkerb, H. Balakrishnanc, N. Feamsterd, I. Ganicheve, A. Ghodsif, P. Godfreyg, N. McKeownh, G. Parulkari, B. Raghavanj *et al.*, "Architecting for innovation," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 3, pp. 24–36, Jul. 2011.

[7] A. Feldmann, "Internet clean-slate design: what and why?" *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 3, p. 59–64, Jul. 2007.

[8] K. Tutschku, P. Tran-Gia, and F. Andersen, "Trends in network and service operation for the emerging future internet," *AEU-International Journal of Electronics and Communications*, vol. 62, no. 9, p. 705–714, Nov. 2008.

[9] G. Schaffrath, C. Werle, P. Papadimitriou, A. Feldmann, R. Bless, A. Greenhalgh, A. Wundsam, M. Kind, O. Maennel, and L. Mathy, "Network virtualization architecture: proposal and initial prototype," in *Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures (VISA)*, Aug. 2009, p. 63–72.

[10] N. Chowdhury and R. Boutaba, "Network virtualization: state of the art and research challenges," *IEEE Communications Magazine*, vol. 47, no. 7, p. 20–26, Jul. 2009.

[11] J. Rubio-Loyola, J. Serrat, A. Astorga, A. Fischer, A. Berl, H. de Meer, and G. Koumoutsos, "A viewpoint of the network management paradigm for future internet networks," in *Proceedings of the 1st IFIP/IEEE International Workshop on Management of the Future Internet (ManFI)*. IEEE, Jun. 2009, pp. 93–100.

[12] P. Antoniadis, S. Fdida, T. Friedman, and V. Misra, "Federation of virtualized infrastructures: sharing the value of diversity," in *Proceedings of the 6th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Dec. 2010.

[13] J. Schonwalder, M. Fouquet, G. Rodosek, and I. Hochstatter, "Future internet = content + services + management," *IEEE Communications Magazine*, vol. 47, no. 7, pp. 27–33, Jul. 2009.

[14] S. Bhatia, M. Motiwala, W. Muhlbauer, Y. Mundada, V. Valancius, A. Bavier, N. Feamster, L. Peterson, and J. Rexford, "Trellis: a platform for building flexible, fast virtual networks on commodity hardware," in *Proceedings of the 2008 ACM CoNEXT Conference*. ACM Press, 2008, pp. 1–6. [Online]. Available: http://dl.acm.org/citation.cfm?id=1544084

[15] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford, "In VINI veritas: realistic and controlled network experimentation," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, p. 3–14, 2006.

[16] Y. Li, L. Su, D. Jin, and L. Zeng, "ExpoNET: a flexible platform for concurrent experiments on programmable infrastructure," in *2011 IEEE Global Telecommunications Conference (GLOBECOM 2011)*. IEEE, Dec. 2011, pp. 1–5.

[17] E. Salvadori, R. D. Corin, A. Broglio, and M. Gerola, "Generalizing virtual network topologies in OpenFlow-Based networks," in *2011 IEEE Global Telecommunications Conference (GLOBECOM 2011)*. IEEE, Dec. 2011, pp. 1–6.

[18] MIT Communications Futures Program (CFP), "Inter-provider Quality of Service White paper draft 1.1," Nov. 2006.

[19] A. Fischer, J. Botero, M. Duelli, D. Schlosser, X. Hesselbach, and H. De Meer, "ALEVIN-A framework to develop, compare, and analyze virtual network embedding algorithms," *Electronic Communications of the EASST*, vol. 37, 2011.