Electronic Communications of the EASST
Volume X (2008)

Proceedings of the Workshop on Overlay and Network Virtualization at KIVS 2009

Network Virtualization: Implementation Steps Towards the Future Internet

K. Tutschku, T. Zinner, A. Nakao and P. Tran-Gia

14 pages

Guest Editors: Martin Stiemerling, Xiaoming Fu, Roland Bless
Managing Editors: Tiziana Margaria, Julia Padberg, Gabriele Taentzer
ECEASST Home Page: http://www.easst.org/eceasst/ ISSN 1863-2122

Network Virtualization: Implementation Steps Towards the Future Internet

K. Tutschku[1], T. Zinner[2], A. Nakao[3] and P. Tran-Gia[2]

[1] kurt.tutschku@univie.ac.at http://www.cs.univie.ac.at
Universität Wien, Professur "Future Communication",
Universitätsstrasse 10, 1090 Wien, Austria

[2] [zinner|trangia]@informatik.uni-wuerzburg.de http://www3.informatik.uni-wuerzburg.de
Universität Würzburg, Lehrstuhl für Informatik III,
Am Hubland, 97074 Würzburg, Germany.

[3] nakao@iii.u-tokyo.ac.jp http://nakao.iii.u-tokyo.ac.jp/
University of Tokyo, Graduate School of Interdisciplinary Information Studies,
7-3-1, Hongo, Bunkyo-ku, Tokyo 113-0033, Japan,
also with the NICT (National Institute for Communication and Information Technology),
1-33-16, Hakusan, Bunkyo-ku, Tokyo, 113-0001, Japan.

Abstract: In this paper we will investigate why and how *Network Virtualization (NV)* can overcome the shortfalls of the current system and how it paves the way for the Future Internet. Therefore, we will first discuss some major deficiencies and achievements of today's Internet. Afterwards, we identify three major building blocks of NV: a) the use of application-specific routing overlays, b) the safe consolidation of resources by OS virtualization on a generic infrastructure, and c) the exploitation of the network diversity for performance enhancements and for new business models, such as the provisioning of intermediate nodes or path oracles. Subsequently, we discuss an implementation scheme for network virtualization or routing overlays based on *one-hop source routers (OSRs)*. The capabilities of the combination of NV and OSRs are demonstrated by a *concurrent multipath transmission (CMP) mechanism* (also known as stripping) for obtaining high throughput transmission pipes. The suggested stripping mechanism constitutes a first instance of a refinement of the concept of NV, the idea of *transport system virtualization*.

Keywords: Network Virtualization, Overlays, Transport System Virtualization, Future Internet Architecture

1 Introduction

While today's Internet and its protocols are apparently reaching their limits, a new technology is emerging which promises to overcome many of the deficiencies of the current system. This technology is denoted as *Network Virtualization (NV)*.

NV is the technology that allows the simultaneous operation of multiple logical networks (also known as overlays) on a single physical platform. Similar to server and operating system virtualization [1], NV might be able to reduce significantly the operational expenditures (OPEX) of networks since it can consolidate multiple virtual structures efficiently into only a single topology requiring small configuration efforts and potentially lesser hardware than before. In addition, NV has appealing features for its use in the Future Internet. It permits distributed participants to create almost instantly their own network with application-specific naming, topology, routing, and resource management mechanisms such as server virtualization, and enables users to use even a whole computing center arbitrarily as their own personal computer. Based on these features, NV technologies received recently tremendous attention and is expected to be one of the major paradigms for the Future Internet as proposed by numerous international initiatives on future networks, e.g. PlanetLab (USA, International) [2], GENI (USA) [3, 4], AKARI (Japan) [5], and G-Lab (Germany) [6].

In this paper we will investigate why and how Network Virtualization can overcome major shortfalls of the current system and how it paves the way for the Future Internet. We will first discuss in Section 2 significant major deficiencies and achievements of today's Internet. Afterwards, we identify in Section 3 three major building blocks of NV: a) the use of application-specific routing overlays, b) the safe consolidation of resources by OS virtualization on a generic infrastructure, and c) the exploitation of the network diversity. Subsequently, Section 4 discusses an implementation scheme for network virtualization of routing overlays based on *one-hop source routers (OSRs)*. We suggest a *concurrent multipath transmission (CMP) mechanism* (also known as stripping) which demonstrates the capabilities of the combination of NV and OSRs. In addition, we outline the parallels of CMP transmission to the successful multiple source download mechanisms of P2P content distribution applications. The suggested CMP mechanism constitutes a first instance of a refinement of the concept of NV, the idea of *transport system virtualization*. Finally, the paper will be concluded in Section 5 with a short summary.

2 Some Deficiencies and Achievements of Today's Internet

When addressing the shortfalls of the current Internet, the discussion usually focuses quickly on architectural and operational issues such as the anticipated lack of IP addresses [7], the complexity of today's management [8], or the insufficient extensibility of today's IP protocol family (denoted as *protocol ossification*)[9]. However, this discussion neglects often the requirements of the future applications and users. Since it is particularly hard to foresee the future, we restrict this discussion to accepted requirements of current applications and usages which are not solved, even until today. This approach provides a benchmark whether a Future Internet architecture will be at least superior to today's system by solving current problems. After the discourse of the deficiencies, we acknowledge that the current Internet is still a success story. We will outline selected achievements of the current system and its applications and investigate what we can learn from these successes for the future system.

2.1 Deficiencies

A major deficiency of today's Internet is still the *missing control of the end-to-end quality of service (QoS)*. Many solutions such as IntServ or DiffServ have been developed and certain

QoS islands have been formed depending on the technology and the capabilities of the providers applying these mechanisms. As a result, a user may ask ”Why can't I take advantage of these islands?”.

Although the protocols of the current Internet haven been designed for catastrophic failures, the *reliability* of the current system and its application is very poor. However, the sophisticated resilience concepts exist, e.g. for MPLS, and are available at experienced Internet Services Providers (ISPs). Again, this fact raises the question why the reliability islands can't be exploited for better system or service reliability.

Finally, a major deficiency is the lock-in of users to their ISPs which suppresses competition among ISPs. John Crowcroft expressed this shortfall precisely in a posting to the End2End-Interest Mailing on April 26^{th} 2008: ” ... i can go on the web and get my gas, electricity, ... changed, why is it not possible to get a SPOT price for broadband internet?”. This feature is similar to the ”call-by-call” provider selection scheme in some deregulated telephone service markets such as Germany. After passing the access system, the traffic is forwarded to the appropriate transport network. Currently, a user may ask: ”Why is the data traffic not being relayed after passing the access network to the most cost-efficient ISP selected by me?”. By the way, this would help even travelers in hotels to reduce their roaming costs.

2.2 Achievements

Despite all its deficiencies, the current Internet has facilitated never expected ways of using and operating efficiently the networks.

2.2.1 P2P-based Content Distribution

One of the fastest revolutions in Internet usage was the development of *Peer-to-Peer (P2P) content distribution applications*. P2P systems are a specific type of distributed systems, which consist of equal entities, denoted as *peers*, that share and exploit resources in a cooperative way by direct end-to-end exchanges on application layer.

P2P content distribution systems are used to distribute very large video and audio files like DVDs or CDs. The first major P2P content distribution application was Gnutella [10], released in 1999. After only four years, P2P contribution applications have become the major source of Internet traffic. Table 1 shows the shares of the different traffic types at a residential access system [11] .

Type	P2P	not identified	Web	eMail	FTP
Percent	67.3 %	23.3 %	7.9%	1.2%	0.3%

Table 1: Typical Traffic Distribution in Residential Access Systems, after [11]

Traditional P2P content distribution applications consider a loose notion for quality, i.e. a file will eventually be downloaded after some time. P2P-based IP-TV applications are even capable to support strict quality constraints for video playback. They are capable to relay sufficient video data to an end user peer such that the peer is able to play out continuously a moving image with sound. The popularity of P2P-based IP-TV was revealed in recent studies. Table 2 depicts observed and estimated traffic volumes of different IP-TV applications reported in [12]. Again, P2P-based IP-TV has gained a significant market share in very short time. It can even compete with conventional Content Distribution Networks (CDNs) as used by YouTube [13].

Traffic Type	Terrabytes per month
YouTube – worldwide (Cisco est., May 2008)	100.000
P2P Video Streaming in China (Jan. 2008)	33.000
YouTube – United States (May 2008)	30.500
US Internet backbone at year end 2000	25.000
US Internet backbone at year end 1998	6.000

Table 2: Amount of IP-TV Traffic, after [12]

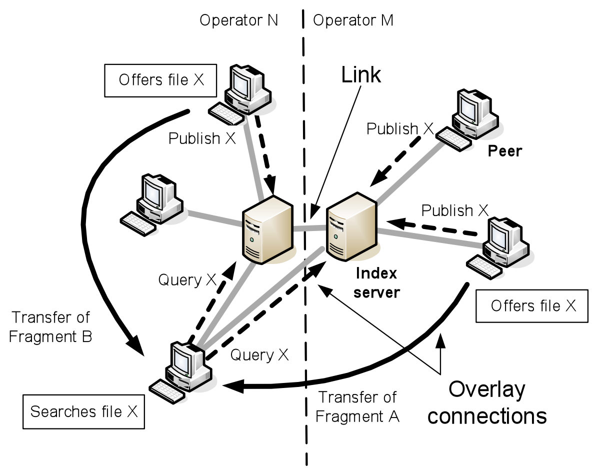

Figure 1: Hybrid P2P Content Distribution Application

In order to understand the success of P2P-based content distribution, we will investigate now briefly the highly popular eDonkey system [14, 15]. eDonkey is a typical representative for P2P content distribution applications. The eDonkey architecture is depicted in Figure 1. eDonkey is denoted as a hybrid-P2P system since it consists of two kinds: a) end-user peers (for short denoted as peers) providing and downloading files, and b) index servers providing the information on the locations of a file or parts of it. When a peer wants to download a file, it queries the index servers and then asks the providing peers for data transmission. The data transmission can be accelerated by using the *multiple source download principle*. Here, two or more different pieces of a file are downloaded in parallel from different providing peers. Due to the availability of order information, the pieces can be reassembled appropriately. Since peers can both, downloading and providing information, the boundary between consumer and provider vanishes in P2P systems.

A closer look reveals, that P2P content distribution systems form two different overlays. One overlay is dedicated to the distribution of query information, while the other one is dedicated for user data exchange, i.e. for transmitting video or audio information. It becomes also evident that the two overlays may have different topologies, addresses, and routing principles. In addition, a downloading peer remains in command where to download the data from. If numerous peers provide the same information, the downloading peer can choose the best peers to download from. This characteristic facilitates also the feature of P2P overlays to be more reliable than conventional client/server systems since they don't rely on a single source. Another feature of P2P systems is the use of their own addressing schemes. In this way, they are able to circumvent the problems of Internet hosts being behind NAT (Network Address Translation). Moreover, P2P overlays enable the integration of networks of different technologies and of different administrative domains into a single virtual structure. Thus, they facilitate the notion of *multi-network services* [16].

2.2.2 Diversity in Connectivity and Quality

Another achievement of today's Internet is its diversity in connectivity and quality. The Internet is not a homogenous network with a flat topology. Figure 2 depicts the topologies of three North-american Tier 1 network operators (AS3967, AS3356, AS6467) on Point-of-Presence (POP) level [17]. The figure reveals that a very large number of locations have many different routes to an arbitrary destination. These routes are often spread among different operators. Hence, a user would have theoretically the possibility to choose among the multiple providers and even within a provider among multiple routes. This characteristic would not only facilitate better performance but might also increased the competition among providers. A user can chose the most cost-efficient provider. Additionally, this picture shows that a significant redundancy is present in the networks. A better exploitation of this characteristic might enhance the reliability of the system.

The current Internet is not only diverse in its topology. Accompanying this feature is its *diversity in quality*. Theoretically, the current Internet protocols should find the ”shortest” route to a destination. This feature means that theoretically the *triangle inequality (TI)* holds for the packet delay, cf. Figure 3. However, recent measurements within PlanetLab have demonstrated that this inequality is violated more often than one has expected so far. The violation might be as high as 25% [18].

This results shows a) that the current Internet routing is far from being optimal, b) better routes exist and sufficient capacity is often available in the networks and c) it can potentially be exploited and offered. Unfortunately, current IP transport protocols are not readily capable for multi-homing.

2.2.3 Operating System (OS) Virtualization

The virtualization of operating systems has become very popular recently due to its capability to consolidate multiple virtual servers into a single physical machine [19]. The application of OS virtualization reduces directly the operational expenditures (OPEX) of multiple servers. Elaborated virtualization techniques like ”Hypervisor-based” or ”Host-based” virtual machine control

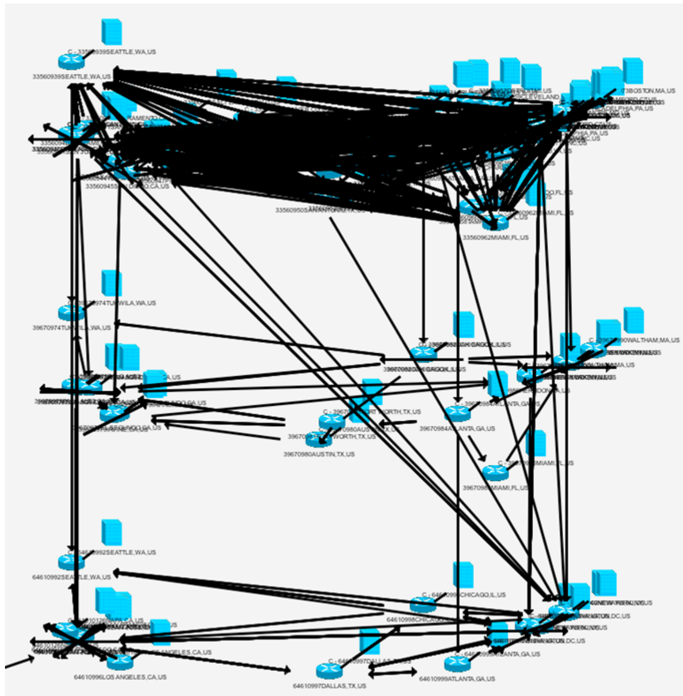

Figure 2: Selected North-american Tier 1 Provider Networks

(cf. Figure 4) permit a fair and reliable *resource isolation* among virtual machines. In this way, virtualization allows a safe testing of server configurations without harming the other virtual machines. Furthermore, a personal machine configuration running on large servers is permitted for individual users. In this way, the users can use even a complete computer center as a PC which is located next to their desks.

Another advantage is that virtualization enables applications to be moved arbitrarily within the memory. This *memory invariance* can be exploited. Applications and systems can easily be moved and relocated to arbitrary physical locations.

In oder to speed up the relocation of an instance of a virtual computer, efficient compression technologies for complete machine states such as SBUML (Scrap-Book User Mode Linux) [20] have been developed. SBUML can compress a state down to 10% of the real memory size. This compression ration enables a fast relocation of even large router operating system images within a network.

3 Network Virtualization: Solving the Puzzle

The puzzle how Network Virtualization can overcome the shortfalls of today's Internet and paving the way for the future Internet resolves rapidly when the above outlined achievements are combined.

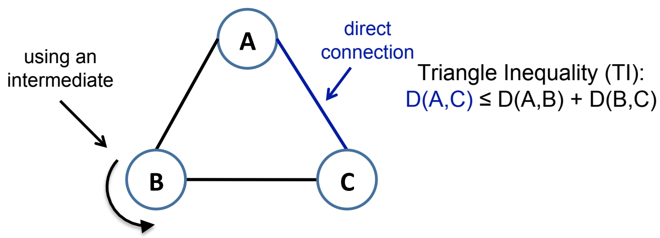

Figure 3: Triangle Inequality Violation

3.1 Building Blockings

The concept of *virtual network structures*, such as P2P overlays, form the first major building block for Network Virtualization, cf. Figure 5. Due to their ability to form arbitrary application-specific network structures, overlays can achieve higher performance and are more reliable than other network architectures. In addition, the specific ability of P2P overlays for defining their own address scheme prevents a look-in of users into a specific provider. In contrast to current standard IP technology, an overlay address can be maintained while the actual physical address (e.g an IP address) can change. The capability of overlays for bridging between various network architectures facilitates services across multiple technical and operational domains.

The second building block is *the diversity of networked resources*, e.g. with respect to connectivity and quality. The diversity of transport resources in today's Internet will even be increased in the Future Internet due to new physical transport systems for core networks, such as 100Gb Ethernet, and more network providers. As a result, it can be assumed that high amounts of data transmission capacity will be available in the future networks. If one is able to locate these resources, they can be utilized for achieving high performance and reliability of the system.

A *transport resource* can be defined as a resource which is able to forward data to a given destination, e.g. a link, a route, or a path. This definition allows to outline the parallels of the concept of a sliver in PlanetLab, which is a specific computing resource on an PlanetLab PC, to a transport resource. Transport resources can be combined to a routing overlay (see Section 3.2) as slivers can be combined to a slice.

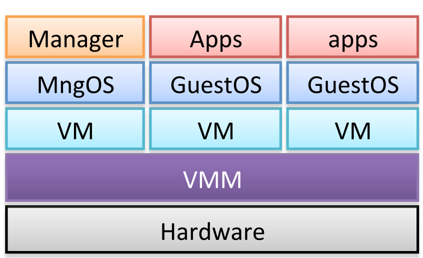

(a) Hypervisor-based

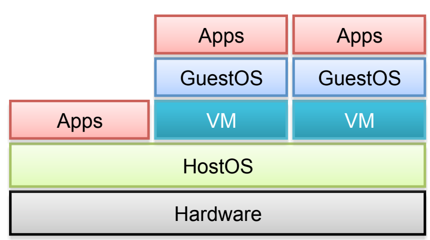

(b) Host-based

Figure 4: Virtualization Control Options

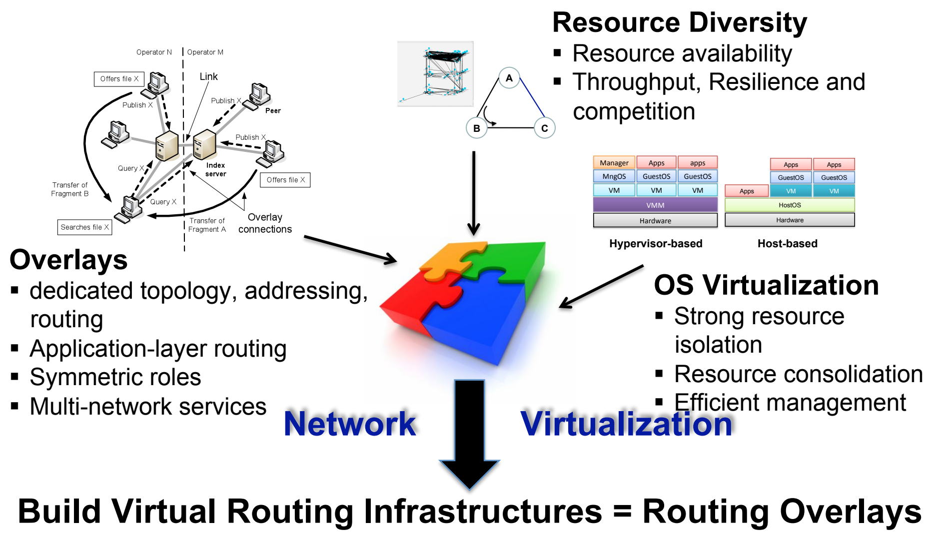

Figure 5: Components of Network Virtualization

The finding of transport resources can be performed in pure P2P manner (e.g. by flooding protocols as in unstructured P2P protocols such as Gnutella or by DHTs-based mechanisms as in structured P2P protocols such as Chord) or by a hybrid or infrastructure-based P2P architecture. In particular the use of additional infrastructure for locating these resources draws currently a lot of attention, for example in the newly formed IETF working group "Application-Layer Traffic Optimization (ALTO)". In addition, operator-supported solutions [21, 22] might enable operators to enhance the efficiency of their physical networks.

Finally, *OS virtualization* constitutes the third building block. It provides the opportunity to consolidate safely multiple networks in one physical platform. In addition, it may simplify the management of the system due to the reduction of physical entities.

3.2 Routing Overlays: The Basis for Evolving Today's Network and Operating the Future Internet

The combination of the above outlined building blocks of Network Virtualization found the basis for *routing overlays*. Routing overlays are virtual network structures dedicated to forwarding arbitrary data. Routing overlays can apply their own address and routing scheme and may have arbitrary topology. The actual topology and routing may be defined by available transport resources. An infrastructure which attempts to implement these routing overlays should a) enable its re-use on small scale, b) provide services invariant from the location of the service provider, and c) permit the use of application-layer mechanisms safely in lower layers of the stack. A generic infrastructure which facilitates this requirements can be based on OS virtualization on routers and the virtualization of interfaces [19]. A single physical router can execute multiple virtual router images. The physical interfaces are shared by virtualization techniques among all virtual routers. Since independent and full access to every virtual router is available, an application can set, for example, the routing tables arbitrarily. Hence, the application can adapt the network to its needs. In this way, routing overlays can build the basis for evolving today's networks and operating the Future Internet.

Only a few generic commands on a physical entity supporting virtual routers is necessary. The commands are: instantiate a virtual router image at a certain location, start a virtual router, suspend a virtual router, resume a virtual router, stop a virtual router and destroy a virtual router. In addition, mechanisms for establishing virtual connections between the virtual routers are necessary.

4 Implementing Advanced Routing Overlays

Recently, various architectures of routing overlays have been proposed [23, 24]. A highly promising approach is the concept of *one-hop source routing*. Hereby, the user data is forwarded one-time only to a specific intermediate node which then relays the traffic to its destination using ordinary IP routing. The dedicated forwarding can be easily achieved by establishing a tunnel to the intermediate node. The advantage of one-hop source routing is the simple control of performance by selecting an appropriate intermediate node while still being scalable.

4.1 An Efficient One-hop Source Routing Architecture

An efficient one-hop source routing architecture capable of NV was suggested in [25, 26]. This architecture is depicted in Figure 6. The architecture applies edge-based NV-boxes which can execute safely virtual router software. These software routers can accept incoming traffic from tunnels and forward the traffic to the destination using conventional IP routing protocols. When a source wants to send data with controlled performance, cf. Step 1 in Figure 6, then a virtual routing overlay is instanciated. Therefore, the source sends a signal to an NV-box running the *One-hop Source Router (OSR)* software. When an OSR router receives such a signal it asks a *Path Oracle* to provide him with the addresses of one or more intermediate nodes which can forward this data in the required way, cf. Step 2 in Figure 6. Subsequently, the ingress OSR router establishes tunnels to the selected intermediate OSR routers, cf. Step 3 in Figure 6. Finally, the intermediate OSR router will insert the traffic into the conventional IP routing process. After finishing the instantiation of the routing overlay, the actual traffic can be inserted into the virtual structure which will forward it towards the destination.

This architecture shows a separation of the former monolithic IP system into two virtual overlays, one for signaling and one for data forwarding. This separation can be seen in parallel to the two overlays in P2P content distribution applications. The two overlays can be structured and equipped with routing mechanisms according to their specific function.

However, it also has to be mentioned here that the edge-based nature of this architecture reduces the efficiency of the one-hop source routing concept. This disadvantage can be overcome by deploying OSR systems in the network core.

Due to the use of NV-boxes and their placement at arbitrary locations, virtual OSR systems can be also instantiated at arbitrary locations. An application-specific compressed OSR image can be

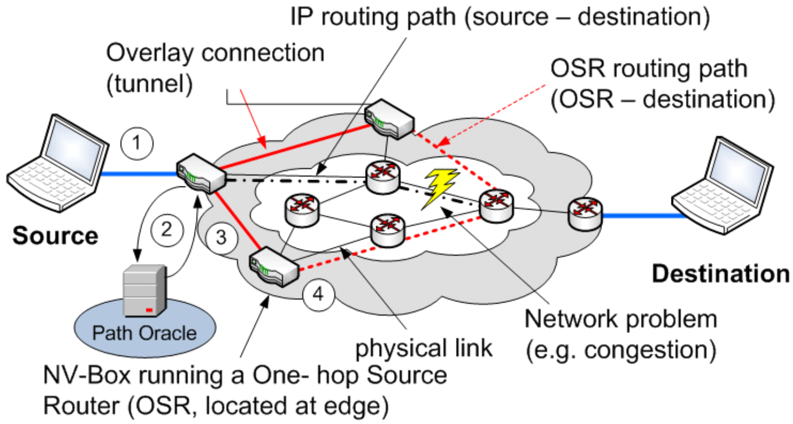

Figure 6: Routing Overlay using One-hop Source Routing

uploaded rapidly to a NV-box. Thus, an application can easily re-use the generic infrastructure.

4.2 Concurrent Multipath Transmission

The capability of the above introduced one-hop source architecture can be demonstrated readily by the problem of achieving very high throughput data transmissions. A well-known solution to the problem of high capacity transport is the combination of the multiple overlay paths into one large overall transport pipe by using concurrent multipath transfer. This concept is known in packet-switched networks also as *stripping*.

4.2.1 Overall Architecture

The considered CMP architecture sends data packets concurrently on different overlay paths from the source to the destination, cf. Figure 7. The paths can be chosen from different overlays which can span across different physical networks. The access to these paths for a specific application can be achieved by instantiating virtual OSRs dedicated to the specific overlay. Traffic is relayed through such an OSR when it seems to be of advantage for an application to use this specific forward node.

The combination of multiple paths achieves a direct increase of throughput and a higher reliability since the system does not rely on a single path anymore. In addition, this architecture facilitates interdomain traffic management and edge-based performance control due to the selection of appropriate intermediate nodes. Furthermore, the application of the path oracle can permits a rapid discovery of the best available resources in the network. Such a path oracle can be provided by the network operator or by other institutions [21].

The use of parallel transmissions is similar to the multi-source download principle in P2P content distribution systems. A path can be compared to a peer. In CMP, the peer which requests the transmission service selects the best peers, i.e. paths, which promise the highest combined throughput. In multi-source download, the peer selects the best peers which provide him the

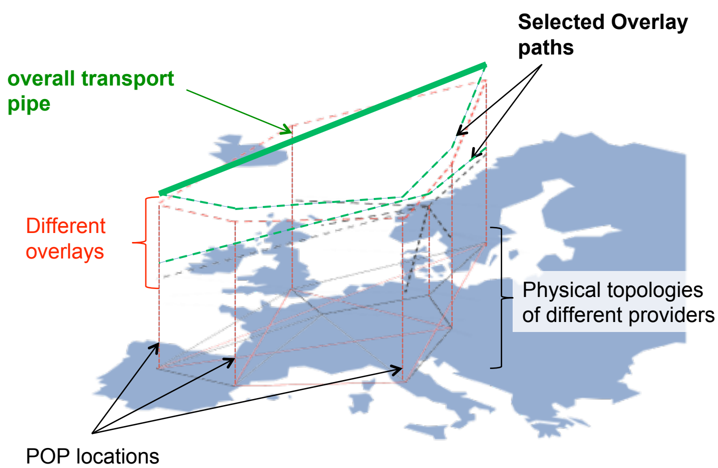

Figure 7: Providing a High-capacity Pipe by Combination of Multiple Overlay Paths

highest throughput for the download.

4.2.2 Transport System Virtualization

The combination and parallel use of multiple transport resources can be viewed as *transport system virtualization*. This view is in parallel to the PlanetLab slice concept for distributed applications. In PlanetLab, local resources are denoted as slivers and are combined together into a slice. In *transport system virtualization* local transport resources (cf. the definition of transport resource in Section 3.1) are combined together and form a larger, more powerful transport pipe.

4.2.3 Transmission Mechanism

Figure 8 shows a detailed model of the stripping mechanism. The data stream is divided at the OSR router into segments which are split into k smaller parts. The k parts are transmitted in parallel on k different overlay paths. The receiving OSR router reassembles these parts again into segments. The parts can arrive at the receiving router at different time instances since they are transmitted on paths with different varying delays. Therefore, it is possible that they arrive ”out of order”.

In order to avoid having this behavior impact on the application performance, the receiving OSR router maintains a finite re-sequencing buffer. However, when the re-sequencing buffer is filled and the receiving router is still waiting for parts, part loss can occur. This loss of parts is again harmful for the application and should be minimized. This objective can be achieved by an appropriate selection of the re-sequencing buffer size. An investigation of this specific problem of stripping mechanisms in the context of network virtualization is presented in an accompanying paper [27].

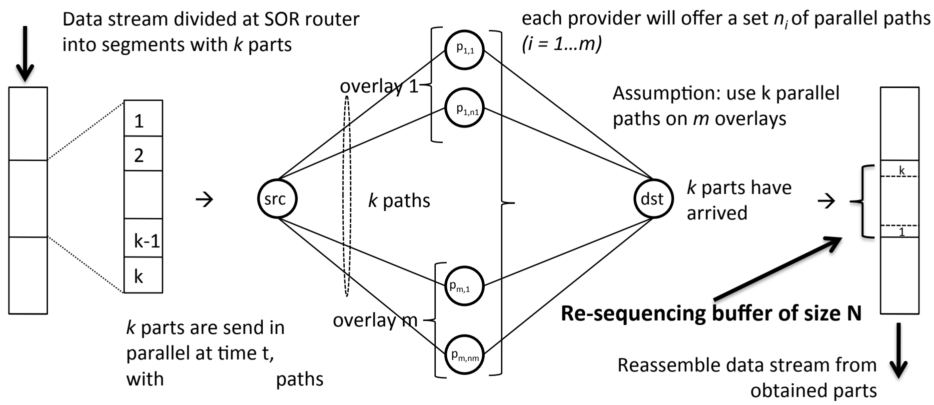

Figure 8: Transmission Mechanism

5 Conclusion

Network Virtualization is a set of technologies which has a very high potential for becoming the major paradigm for the Future Internet. Originally, NV allows mainly the simultaneous operation of multiple logical networks or overlays on a single physical platform. The application of the three major building blocks of NV (which are a) use of application-specific routing overlays, b) the safe consolidation of resources by OS virtualization on a generic infrastructure, and c) the exploitation of the diversity of networked resources), overcomes also the deficiencies of today's Internet. This can be verified by the following examples. The use of routing overlays can bridge between QoS and reliability islands. They also can facilitate the split between the identifier and locator function of an address. Thus, they overcome the lock-in of users to a specific provider. The safe consolidation of multiple virtual routers into one physical entity permits an easy re-use of a generic infrastructure. Finally, the sophisticated exploitation of the diversity of networked resources permits higher quality and reliability, and even more competition.

The capabilities of the combination of Network Virtualization and one-hop source routers (OSRs) have been discussed by a concurrent multipath transmission (CMP) mechanism for obtaining high throughput transmission pipes. The suggested stripping mechanism constitutes a first instance of a refinement of the concept of NV, the idea of *transport system virtualization*. The main idea of it is to combine local transport resources to form a larger, more powerful transport service.

Future investigations will be focused on two areas: a) The discussion of a stronger separation of the future network architecture into a generic data transport plane and a generic control layer should be intensified. Especially, the architectures and mechanisms of the future control should be refined. b) A detailed performance investigation of the mechanisms for transport system virtualization has to be performed, e.g. how can one find the required networked resources in an efficient way and what is the optimal combination of these resource.

6 Literature

Bibliography

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the Art of Virtualization," *ACM SIGOPS Operating Systems Review*, vol. 37, no. 5, pp. 164–177, Dec. 2003.

[2] T. Roscoe, *Peer-to-Peer Systems and Applications*, chapter 33. The PlanetLab Platform, Springer, Berlin, 2005.

[3] GENI Consortium, "GENI - Global Environment for Network Innovations," 2006, Information available at http://www.geni.net/.

[4] GENI Planning Group, "GENI Design Principles," *IEEE Computer*, vol. 39(9), Sep. 2006.

[5] NICT, "AKARI Architecture Design Project for New Generation Network," Information available at: http://akari-project.nict.go.jp/eng/index2.htm, 2007.

[6] P. Tran-Gia, "G-Lab: A Future Generation Internet Research Platform," Information available at: www.future-internet.eu, 2008.

[7] CNet News, "Net Number System at a Crossroads," Information available at: http://news.cnet.com/Net-number-system-at-a-crossroads/2009-1023_3-225712.html, 1999.

[8] D. Clark, C. Partridge, J. Ramming, and J. Wroclawski, "A Knowledge Plane for the Internet," in *Proc. of the ACM Sigcomm 2003 Conference*, Karlsruhe, Germany, Aug. 2003.

[9] M. Handley, "Why The Internet Only Just Works," *BT Technology Journal*, vol. 24, no. 3, Jul. 2006.

[10] Clip2, "The Gnutella Protocol Specification v0.4 (Document Revision 1.2)," Information available at http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf, 2001.

[11] J. Gabeiras, "P2P-Traffic," Presentation at the COST 279 Midterm Seminar; Rome, Italy, Jan. 2004.

[12] Cisco Inc., "Approaching the Zettabyte Era," Information available at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns82 7/white_paper_c11-481374.pdf, Jun. 2008.

[13] P. Gill, M. Arlittz, Z.Li, and A. Mahantix, "YouTube Traffic Characterization: A View From the Edge," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement (IMC 07)*, San Diego, CA., Oct. 2007.

[14] K. Tutschku, "A Measurement-based Traffic Profile of the eDonkey Filesharing Service," in *Proc. of the 5th Passive and Active Measurement Workshop (PAM2004)*, Antibes Juan-les-Pins, France, Apr. 2004.

[15] G. Wearden, "eDonkey Pulls Ahead in Europe P2P Race," http://business2-cnet.com.com/2100-1025_3-5091230.html.

[16] K. Tutschku, P. Tran-Gia, and F.-U. Andersen, "Trends in Network and Service Operation for the Emerging Future Internet," *International Journal of Electronics and Communication*, 2008.

[17] M. Liljenstam, J. Liu, and D. Nicol, "An Internet Topology for Simulation," Information available at http://www.crhc.uiuc.edu/ jasonliu/projects/topo/, 2003.

[18] S. Banerjee, T. Griffin, and M. Pias, “The Interdomain Connectivity of PlanetLab Nodes,” in *Proc. of the 5th Passive and Active Measurement Workshop (PAM2004)*, Antibes Juan-les-Pins, France., Apr. 2004.

[19] S. Rixner, “Network Virtualization: Breaking the Performance Barrier,” *ACM Queue*, Jan./Feb. 2008.

[20] O. Sato, R. Potter, M. Yamamoto, and M. Hagiya, “UML Scrapbook and Realization of Snapshot Programming Environment,” in *Proceedings of the Second Mext-NSF-JSPSInternational Symposium on Software Security (ISSS 2003)*, Tokyo, Japan., 2003.

[21] V. Aggarwal, A. Feldmann, and C. Scheideler, “Can ISPs and P2P Systems Co-operate for Improved Performance?,” *ACM SIGCOMM Computer Communications Review (CCR)*, vol. 37, no. 3, Jul. 2007.

[22] H. Xie, Y. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, “P4P: Provider Portal for Applications,” *ACM SIGCOMM Computer Communications Review (CCR)*, vol. 38, no. 4, pp. 351 – 362, Oct. 2008.

[23] A. Nakao, L. Peterson, and A. Bavier, “A Routing Underlay for Overlay Networks,” in *Proc. of the ACM Sigcomm 2003 Conference*, Karlsruhe, Germany, Aug. 2003.

[24] K. Gummadi, H. Madhyastha, S. Gribble, H. Levy, and D. Wetherall, “Improving the Reliability of Internet Paths with One-Hop Source Routing,” in *Proceedings of 6th conference on Symposium on Opearting Systems Design & Implementation (OSDI’04)*, San Francisco, Ca., USA, Dec. 2004.

[25] J. Lane and A. Nakao, “SORA: A Shared Overlay Routing Architecture,” in *Proceedings of the 2nd International Workshop on Real Overlays And Distributed Systems (ROADS)*, Warsaw, Poland., Jul. 2007.

[26] S. Khor and A. Nakao, “AI-RON-E: Prophecy of One-hop Source Routers,” in *Proc. of the 2008 IEEE Global Telecommunications Conference (Globecom08)*, New Orleans, LA., Nov./Dec. 2008.

[27] K. Tutschku, T. Zinner, A. Nakao, and P. Tran-Gia, “Re-sequencing Buffer Occupancy of a Concurrent Multipath Transport Mechanism for Network Virtualization,” in *Proceedings of the 16. ITG/GI - Fachtagung Kommunikation in Verteilten Systemen 2009 - KiVS 2009*, Kasel, Germany, Mar. 2009.