SpoVNet: An Architecture for Supporting Future Internet Applications

The SpoVNet Consortium
http://www.spovnet.de

The current development towards user generated content ("Web 2.0") indicates that future Internet applications will be highly interactive and cooperative. Similar to current applications, they will require support for quality-of-service, group communication like multicast, mobile users, and will be run over a variety of heterogeneous networks. Currently, it is not obvious whether next generation networks will offer native support for these functions. In today's Internet, these functions are only partially supported, e.g., QoS in UMTS- or IEEE 802.11e-based access networks or multicast in the core networks of certain internet service providers (ISPs). However, innovative applications have to implement end-to-end support for these functions at application level, which is either inefficient (e.g, multicast) or impossible (e.g, for QoS).

This talk presents an approach for providing *Spontaneous Virtual Networks (SpoVNets)* that enable flexible, adaptive, and spontaneous provisioning of application-oriented and network-oriented services on top of heterogeneous networks. SpoVNets supply new and uniform communication abstractions for future Internet applications so applications can make use of advanced services not supported by today's Internet. We expect that many functions, which are currently provided by SpoVNet on the application layer will become an integral part of future networks. Thus, SpoVNet will transparently use advanced services from the underlying network infrastructure as they become available (e.g., QoS-support in access networks or multicast in certain ISPs), enabling a seamless transition from current to future generation networks without modifying the applications.

A SpoVNet allows spontaneous creation of a common communication context based on application-specific requirements. For example, a SpoVNet instance can comprise all nodes cooperating in a specific global application, or there can be several SpovNet instances running the same application within different regional or logical scopes, thus, building different groups. Since virtual networks constitute overlay networks on top of the physical network infrastructure, we use overlay network techniques as shown in Figure 1 and exploit their self-organization, scalability and robustness features. Furthermore, the distributed nature of overlay networks allows spontaneous creation of such networks without support of dedicated infrastructure such as servers. The SpoVNet approach is, however, different from existing pure peer-to-peer overlay approaches, because it is aware of the underlying network infrastructure (underlay-awareness) in two aspects. First, it allows using cross-layer information for performance optimizations. Second, we consider the use of optional SpoVNet supporting nodes (*SpoVNet Booster Nodes*) in the infrastructure in order to increase the efficiency and performance of the communication. This is different from approaches like Cabo[1] that try to apply virtualization to network nodes in the underlay. Though Service-aware Adaptive Transport Overlays[2] (SATO) developed by the Ambient Networks project have similar objectives, SpoVNet's focus lies more on application-independent communications abstractions like group communication whereas SATO provides application-specific transport overlays.

The SpoVNet architecture is implemented on top of Layer 3 in the network stack, i.e., it is independent from lower layers that could easily be exchanged (e.g. IPv4, IPv6, HIP). It facilitates the deployment of future Internet applications by providing new communication abstractions and services such as stable identifiers, group communications, QoS support, as well as event observation and notification services. Further communication abstractions can easily be added. All communication abstractions are implemented using overlay-based *communication services*. These services provide the best possible QoS guarantees, preferentially by using native QoS support of the underlay if available. To offer best possible QoS where native support is not available, overlay connections are optimized based upon cross-layer information provided by the Cross-Layer Information Overlay (CLIO). CLIO provides abstract cross-layer information to other SpoVNet components, for example bandwidth capabilities, link latencies or the nature of the transmission medium (e.g., broadcast or wireline). This information is collected from standardized network functions, such as IEEE 802.21 where available, or by passive or active measurements. Furthermore, CLIO enables the communication services to adapt to changing underlay conditions and allows controlling underlay properties where possible. This includes, e.g., QoS- and mobility-support.

To make mobility and multi-homing transparent to higher layers where appropriate, SpoVNet includes two components providing basic transport connectivity based on a SpoVNet-specific addressing scheme. As upper component, a *base overlay* connects all nodes that participate in an instance of a SpoVNet application. It maps SpoVNet identifiers to actual network locators. Thus, applications use SpoVNet identifiers only, whereas several different addresses (e.g., locators) in the underlay may actually be used. Nevertheless, applications may optionally use cross-layer information to implement mobility-awareness. Data may be exchanged either via the overlay or via direct transport connections between SpoVNet nodes, which are managed by the *base communication layer* as lower component. It supports different kinds of transport connections, i.e. reliable and unreliable as well as secured and unsecured connections. Connections can be established over heterogeneous underlay networks, combining, e.g., IPv4 and IPv6 networks.

SpoVNet provides application interfaces, which allow a quick adaptation to changing network conditions. We use the interfaces to implement two examples for future Internet applications as proof of concept: First, a Video-streaming application that adapts to changing network conditions by reducing or increasing the video coding rate. Second, a real-time-multiplayer game that is able to adapt to the given conditions by influencing its network traffic. Other applications can be easily deployed upon the SpoVNet architecture, which will facilitate their transition towards future network technologies due to its abstraction layers and its standardized cross-layer interfaces.

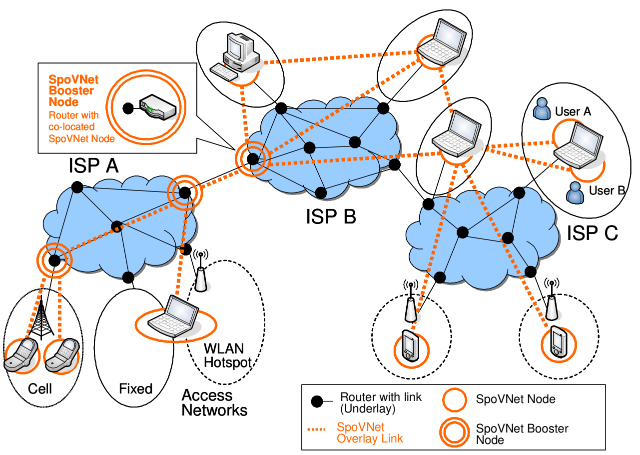

Fig. 1. A SpoVNet deployed in a heterogeneous network environment.

[1] N. Feamster, L. Gao, J. Rexford; How to lease the Internet in your spare time; ACM SIGCOMM Computer Communications Review, Vol. 37, No.1, January 2007

[2] Martin Stiemerling et al, System Design of SATO and ASI, Ambient Networks Deliverable D12-F.1, 2006.