

OverSim: A scalable and flexible overlay framework for simulation and real network applications

(Invited Talk)

Ingmar Baumgart, Bernhard Heep, Stephan Krause
Institute of Telematics, Universität Karlsruhe (TH)
Zirkel 2, D-76128 Karlsruhe, Germany
Email: {baumgart, heep, stkrause}@tm.uka.de

I. INTRODUCTION

A fundamental problem in studying peer-to-peer networks is the evaluation of new protocols. This commonly involves both the simulation of the protocol in a large-scale network as well as the testing of the protocol in connection with real applications in networks like PlanetLab. To facilitate these tasks we have developed the overlay simulation framework *OverSim*. It is designed to fulfill a number of requirements that have been partially neglected by existing simulation frameworks.

II. OVERVIEW

OverSim[1] is a flexible overlay network simulation framework for Linux, Windows, Mac OS X and Nokia Internet Tablets (Maemo) based on OMNeT++. The framework includes many structured peer-to-peer protocols like Chord, Kademlia, Pastry, Bamboo, Koorde, and Broose. Additionally, *OverSim* contains several event distribution protocols like Vast and QuON and the unstructured overlay GIA. Several applications like i3, Scribe, P2PNS, and SimMUD based on these protocols are available. All these implementations can be used for both simulation as well as real world networks. *OverSim* provides several common functions for structured peer-to-peer networks to facilitate the implementation of additional protocols and to make them more comparable.

OverSim utilizes the graphical interface of OMNeT++ to display overlay and underlay topology and all network packets in detail (see Fig. 1), thus allows for intuitive debugging. Additionally, the framework provides a central module for gathering and processing statistics. *OverSim*'s post-processing scripts facilitate the generation of publication quality plots.

OverSim's modular architecture (see Fig. 2) allows the modeling of all components of a P2P network. Modules can be easily exchanged or extended. The resulting flexibility facilitates code reuse. Several exchangeable underlay network models allow to simulate complex heterogeneous underlay networks as well as simplified networks for large-scale simulations. With *OverSim* simulations of overlay networks with up to 100,000 nodes are feasible. *OverSim* is well documented and published under the *GPL*.

III. UNDERLAY ABSTRACTION

OverSim provides different models for underlay abstraction which differ in complexity and accuracy:

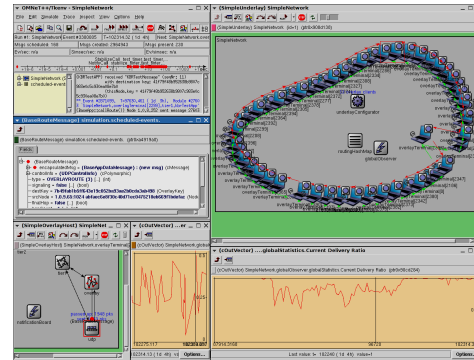


Fig. 1. Screenshot showing overlay topology under churn

- *Simple*: The *Simple Underlay* is the default underlay model for *OverSim*. It combines a low computational overhead with high accuracy, making it a good model for simulating large overlay networks. In this model, nodes are placed into a n -dimensional Euclidean space. The delay between two nodes is then calculated by the distance between those nodes. The positions of the nodes are chosen to match the Internet latency measurements from the *CAIDA/Skitter* project. Additionally, each node is assigned to a logical access network characterized by bandwidth, access delay, jitter and packet loss parameters to allow the simulation of heterogeneous access networks. Node mobility can be achieved by changing coordinates, access network characteristics and the IP address of a node. To model bandwidth effects, each node contains a logical sending queue. The *Simple Underlay* allows the simulation of underlay network partitioning and merging.
- *INET*: For simulations of heterogeneous access networks, backbone routers and terminal mobility, *OverSim* provides an underlay model based on the *INET* framework. Here, the IP stack is completely modeled and even routers can be part of the simulated overlay. *INET* also contains several MAC protocol implementations, which e.g. allow to model wireless IEEE 802.11 interfaces.
- *SingleHost*: The *SingleHost Underlay* provides real network support for *OverSim*. It acts as a middleware to support deploying overlay protocols developed for *OverSim* on real networks. In Sect. IV, this feature will be discussed in greater detail.

All of the underlying network models have a consistent UDP/IP interface to the overlay protocols. Hence, using a different underlying network model is fully transparent to the overlay layer.

IV. SUPPORT FOR REAL NETWORKS

All OverSim protocol implementations can be used without code modifications in real networks. This can be done in two different ways. With the *SingleHost Underlay* an OverSim instance emulates a single overlay host. This host can be connected to other protocol instances over a real network like the Internet. With the *INET Underlay* OverSim can simulate an arbitrary number of overlay hosts. For real network support, OverSim’s simulation time can be synchronized with the real time. Using the Linux TUN interface, the hosts in the INET Underlay can communicate with an external network. This can be used to demonstrate overlay protocols and applications with a limited number of physical devices by connecting them to a large number of emulated overlay nodes.

V. OVERLAY PROTOCOLS

To facilitate the implementation of overlay protocols in OverSim, several functions that many overlay protocol implementations have in common were identified and integrated into the simulation framework:

- Overlay message handling (e.g. RPCs)
- Generic lookup with support for different routing modes
- Node failure discovery and routing table recovery
- *Common API*[2] support
- Bootstrapping support
- Proximity awareness (e.g. Vivaldi, GNP)

The overlay message handler provides a common message format as well as an RPC interface to facilitate dealing with timeouts and packet retransmission due to packet losses. The RPCs can be either routed over the overlay to a specific overlay key or send directly to their destination using UDP/IP. This feature can be used by all components of a node and collects all message related statistical data. The lookup function separates the common functionality of lookup mechanisms in structured overlay networks and provides a generic iterative and recursive lookup interface. Thus, each overlay protocol only has to implement a method to return the closest node in the overlay topology from the local routing table. All of these features allow for rapid overlay protocol prototyping and make protocol implementations comparable and less error-prone.

VI. P2P APPLICATIONS AND USER MODELS

For overlay applications either a layered or a component-based architecture is feasible (see Fig. 2). Applications can use the *Common API* interface, *Application Layer Multicast* interface or *Virtual World* interface provided by the overlay. Additionally OverSim makes use of an XML-RPC interface to provide overlay services (e.g. distributed data storage) to external applications similar to the interface provided by the OpenDHT service. User behavior can be scripted by using the *trace manager* which parses scenario or trace files containing application events. We have successfully used OverSim for

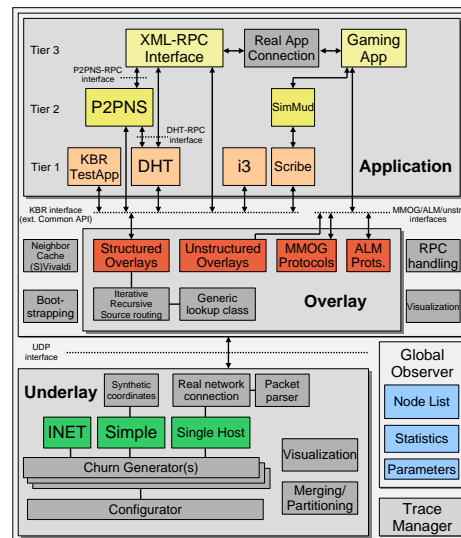


Fig. 2. OverSim’s modular architecture

rapid prototyping of novel peer-to-peer applications like the distributed name service P2PNS (<http://www.p2pns.org/>) and a distributed virtual world, both deployed in PlanetLab.

VII. MODELING OF CHURN

OverSim provides several models for generating churn. This includes a lifetime based churn model supporting different distribution functions (e.g. Weibull, Pareto or Exponential). Alternatively, a scenario or trace file containing join and leave events can be used to model churn behavior. It is possible to use more than one *churn generator* at the same time to simulate groups of nodes with different churn behaviors. For each churn generator different node configurations and overlay parameters can be specified, which allows to easily generate complex scenarios with heterogeneous node behavior.

VIII. CONCLUSION

OverSim facilitates rapid prototyping of new overlay protocols by providing functions common to most overlay protocols. The framework allows to use the same implementations for scalable simulations and in real networks. OverSim benefits from features like an efficient event scheduler and strong GUI support. The large number of implemented overlay protocols and the availability to collect various statistical data make OverSim a powerful tool and reference platform for the peer-to-peer research community. OverSim is well-documented, actively developed as an open source project on <http://www.oversim.org/>, has an active mailing list with a strong user base and is open to contributions.

REFERENCES

- [1] I. Baumgart, B. Heep, and S. Krause, “OverSim: A flexible overlay network simulation framework,” in *Proceedings of 10th IEEE Global Internet Symposium (GI '07) in conjunction with IEEE INFOCOM 2007, Anchorage, AK, USA, May 2007*, pp. 79–84.
- [2] F. Dabek, B. Zhao, P. Druschel, J. Kubiatowicz, and I. Stoica, “Towards a Common API for Structured Peer-to-Peer Overlays,” in *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, vol. 2735/2003, 2003, pp. 33–44.