

CLIO – A Cross-Layer Information Service for Overlay Network Optimization

Dirk Haage, Ralph Holz, Heiko Niedermayer, and Pavel Laskov

Wilhelm-Schickard-Institut für Informatik - Universität Tübingen, Germany

Abstract New overlay-based services aim to provide properties like resilience, availability or QoS. To achieve this, automatic organization and optimization is required, which again demands accurate information on the network. Collecting and exchanging this data has a significant impact on the network, especially if several overlays are used on the same host. In this paper, we provide a survey of the current state of the art and identify challenges which must be addressed in order for new overlay-based services to be successful. We present our own solution CLIO, a cross-layer information service for overlays. CLIO provides information for the automatic creation and optimization of overlays. The service supports multiple overlays on the same node, the provided information is overlay-independent, and collected information is reused.

1 Introduction

The Internet’s underlying infrastructure connects heterogeneous networks ranging from classic Ethernet to WiFi, or 3G mobile communication networks. However, the functionality of IP is often considered insufficient, as it lacks desired features such as group communication, path diversity, or quick adaptation to changes. Hence new services often deploy their own overlays on top of the Internet. A typical example is peer-to-peer based Voice-over-IP telephony, e.g. P2PSIP [1] and Skype [2].

Several proposals for the ‘Future Internet’ are based on self-organizing overlays. The Autonomic Network Architecture (ANA) [3], Ambient Networks [4] or our own work, SpoVNet [5], use overlays to streamline complex networks. SpoVNet provides e.g. an underlay abstraction and generic services for applications to simplify the utilization of overlays.

An essential problem to be addressed by all self-organizing overlay network architectures is overlay optimization under massive usage. In this paper, we identify the requirements and challenges of overlay optimization and present our concept of a cross-layer information service (CLIO) that is aimed at efficient overlay management.

2 Challenges and State of the Art

Overlays have various advantages: Thanks to independent routing and redundancy mechanisms, they can provide higher resilience and availability [6,7]. More

complex services, e. g. end system multicast/anycast or distributed storage, are also possible. Overlays hide the details of the underlying network layers, leveraging simpler APIs for applications. Furthermore, it is possible to find alternative paths in the Internet providing better properties such as QoS, low latency, and resource availability[8].

However, these properties are hard to achieve without interaction between an overlay and the underlying layers. The following challenges faced by modern overlay-based applications can be identified:

Challenge 1: Application-Specific Overlay Optimization. Overlays must be optimized with regard to the requirements of the application that uses them. These requirements can differ widely. Video streaming in a multicast tree, e. g. requires participating nodes to offer certain bandwidths with respect to their upstream and downstream links. Massively multiplayer online games, on the other hand, require certain latencies between nodes. The requirements of applications that run on the same overlay can thus be in conflict, making it difficult to ‘hardcode’ an overlay that would satisfy all of them.

Challenge 2: Collecting Information About the Underlay. Knowledge about underlying layers allows better optimization of overlays. The challenge is to obtain accurate and reliable information about a dynamically changing underlay infrastructure. Some recent overlay implementations realize a limited measurement functionality, e. g. round-trip time measurements [9]. However, there are more possibilities to use network measurement results in overlays. Besides accuracy, a key requirement to such a measurement system is scalability with respect to the number of nodes in an overlay and the number of overlays on a host.

Challenge 3: Low Impact on the Network. Currently, the impact of overlay measurements on the network is relatively low, as most overlay networks make only little use of it, if at all. Measurement data intensively collected in many overlays can be quite voluminous. Hence, the applied mechanisms should minimize the measurement impact on the network by enabling the reuse and optimization of measurements across several overlays.

These challenges suggest a *consolidation* of the overlay measurement functionality in a dedicated component, denoted here as a ‘cross-layer information service’. This component must clearly be distributed, as a centralized service contradicts all paradigms of self-organizing overlays. It should deliver measurements to applications via a well-defined and overlay-independent interface while minimizing the measurement effort by capitalizing on the ‘economy of scale’. The following requirements are imposed on such a service.

Accuracy: The service must provide the most accurate information possible to multiple overlays running in parallel.

Efficiency: Resources must be used efficiently. The service must be scalable with respect to the number of overlays on a host, the number of nodes in each overlay and the number of overlays in a network. The measurement overhead (e. g. active measurements and inter-node communication) must be minimized.

Anonymity and Privacy: Privacy issues become important when data is shared between different entities. It is essential that the data acquired does not leak overlay-specific information between overlays (e. g. identity of nodes, etc.).

Simplicity: The complexity of information gathering must be hidden inside the service. Simple interfaces are essential for sharing information across heterogeneous overlays.

Conventional network management techniques do not meet these requirements. Most network management systems gather data at a central point of operation. The Management Information Base (MIB) concept of the Simple Network Management Protocol (SNMP) offers a good representation of data, but neither does SNMP provide a measurement infrastructure with a generic interface nor does it address overlays. Ganglia [10] is a tool for managing large distributed systems. However, it focuses on node properties rather than link properties, which are crucial for overlay optimization.

Most techniques for overlay optimization have been developed for one particular overlay, and often with respect to only one metric [11,9]. A more comprehensive approach is presented in [12]. The described routing underlay allows optimization of the overlay via primitives. The functions return the best metric-dependent next hop or a path between two hosts without revealing information to the requestor. Applications, however, can only use available metrics. Multi-dimensional metrics are not addressed.

Furthermore, several approaches are currently discussed to provide ISP-specific network information to overlay applications [13,14,15]. These concepts introduce a central point in a network sharing insight about the networks internal structure. In contrast to CLIO, these solutions do not involve end hosts but run entirely on a provider's infrastructure. Thus, CLIO can be seen as a concept that is complementary to these approaches.

3 CLIO - a Cross-Layer Information Service for Overlays

The design of CLIO is intended to address all the requirements specified in Section 2. As part of the SpoVNet architecture [5], CLIO is to be deployed in overlays in which applications can request connectivity with special requirements (e.g. security, latency, QoS) in an abstract way, and free from any network-specific concerns. Each application instantiates its own overlay with its own identifier space. An overlay can be optimized according to an application-defined metric. The goal of CLIO is to provide measurement information for optimization of application-specific overlays. Additional SpoVNet services, such as multicast or event notification, may instantiate their own overlays, therefore CLIO has to deal with multiple overlays, each with different optimization metrics. CLIO gathers a multitude of information and can remote-query other hosts for status information. A generic interface makes this data available to overlays on the host.

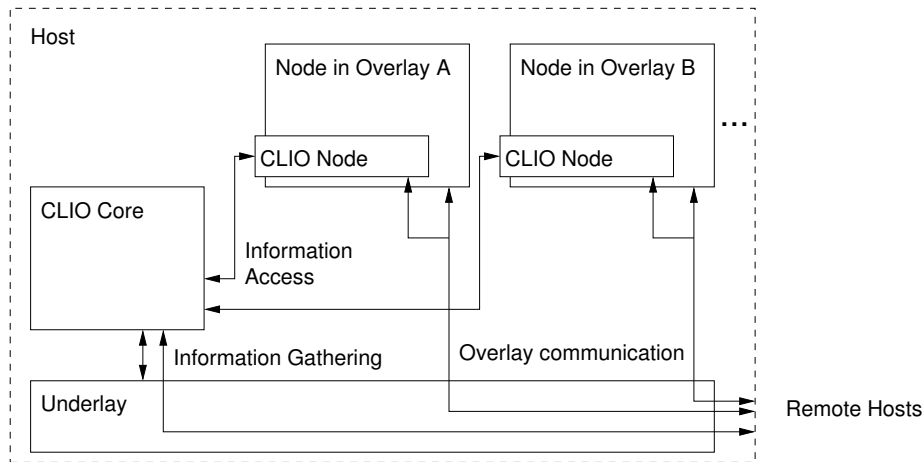


Figure 1. The CLIO architecture and information flow.

3.1 Inner Workings

CLIO consists of two parts. This reflects the need for a component that operates on the semantics of the underlying networks and for a further component that can interact with the overlays on the host. The *Core Component* is responsible for most of the functionality: taking measurements and monitoring, pre-processing of data and data storage, and finally request management and aggregation. The *Node Component* is the link to an overlay node that uses CLIO. It provides the Interface for the applications to request data items. Furthermore, it gives the core access to selected overlays function, e. g. it resolves the overlay identifiers to underlay identifiers (i.e. IP addresses). The architecture and interaction of the components are shown in Figure 1. While the core component works in underlay notation (e. g. IP in the current Internet), the node component for each overlay uses the overlay-specific identifiers to deliver measurement results.

3.2 Data Representation and Generic Interface

Irrespective of its internal data structure, CLIO presents its information via key:value pairs. Data items can only be queried in overlay-specific notation. Each overlay sees the information based on its own namespace. Available data items have unique keys, e.g. RTT, or RTTDeviation. Data items can be node-specific information (e.g. network type, load) and link-specific (e.g. bandwidth, MTU, RTT), while a link is an underlay connection between to nodes.

Overlays access the data through the generic interface. This interface can remain stable even if the methods for the information gathering change. This enables later optimization or exchange of the measurement and prediction algorithms.

Queries can be made as a one-time request, a periodic update or a triggered notification when a value has changed. This allows flexible and asynchronous data access for the applications on the one hand, and efficient measurement scheduling for the CLIO service on the other. If two requests are made for the same data, CLIO attempts to execute the measurements only once. This also helps to reduce the number of measurements and the traffic for information exchange between nodes.

4 Discussion and Conclusion

The proposed cross-layer information service offers several features for automatic creation and maintenance of multiple overlay networks. These features can be classified into the following three categories.

Access to Information: CLIO offers a generic interface for gathering up-to-date network information. It enables overlay-based applications to optimize their communication without collecting this information themselves. This eases application development and allows efficient usage of resources. For security reasons, data is collected in the underlay notation and is translated for each overlay into its own overlay notation. Hence, it is not possible for an application to gain knowledge about other overlays running on the same host.

Efficient Data Maintenance: CLIO implements several mechanisms aimed at minimizing the impact on a network. Measurements are only initiated by request, consequently, unless CLIO is used by any overlay, resource usage is near zero. If a host is involved in several overlays with little or no overlap, CLIO has the same impact on a network as each overlay would have without using CLIO. For overlapping overlays, resource usage is reduced considerably by CLIO through measurement order aggregation and result caching. Furthermore, CLIO offers an important advantage for bootstrapping of overlays: Previously collected information is available immediately. This can speed up the creation and initial optimization of overlays.

Enhancing Services: The data collected by CLIO can be used for additional services. Even though the overlays work on top of a best-effort network, more complex requirements like ‘probabilistic’ QoS support can be addressed by the means of network measurements. Furthermore, applications can adapt to varying QoS levels by utilizing the up-to-date measurements. For example, if the desired quality of service cannot be met anymore, an application may try to obtain a different link from the base overlay software. A further attractive feature offered by our approach is the possibility of ‘forward diagnostics’. Dynamically changing underlay infrastructures bear the risk of performance degradation as well as security incidents. The information provided by our service can be used to identify anomalous conditions *before* they affect the functionality. Thus, countermeasures can be initiated early enough, to maintain efficiency and resilience of both the overlay and CLIO.

We have implemented a first version of CLIO for use in a demonstrator including a number of measurement modules. We are currently continuing the

development with the goal of carrying out performance measurements and evaluating its utility in the SpoVNet architecture.

References

1. P2PSIP Status Pages. <http://tools.ietf.org/wg/p2psip/> (April 2008)
2. Baset, S.A., Schulzrinne, H.: An analysis of the Skype Peer-to-Peer Internet telephony protocol. arXiv:cs/0412017v1, <http://arxiv.org/abs/cs/0412017v1> (2004)
3. ANA: Autonomic Network Architecture. <http://www.ana-project.org/> (Feb 2008)
4. L. Cheng et al.: Service-aware overlay adaptation in Ambient Networks. In: Proc. of the International Multi-Conference on Computing in the Global Information Technology (ICCGI '06). (2006)
5. The SpoVNet Consortium: SpoVNet: An Architecture for Supporting Future Internet Applications. In: Proc. 7th Würzburg Workshop on IP: Joint EuroFGI and ITG Workshop on “Visions of Future Generation Networks”. (July 2007)
6. D. Andersen et al.: Resilient overlay networks. In: Proc. of the 18th ACM symposium on Operating systems principles (SOSP 2001). (2001)
7. Gummadi, K., Gummadi, R., Gribble, S., Ratnasamy, S., Stoica, I.: The impact of dht routing geometry on resilience and proximity. In: Proc. 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications. (2003) 381–394
8. Stefan Savage et al.: The end-to-end effects of internet path selection. SIGCOMM Comput. Commun. Rev. (1999)
9. A. Rowstron, P. Druschel: Pastry: Scalable, distributed object location and routing for large-scale Peer-to-Peer systems. In: Proc. of the IFIP/ACM International Conference on Distributed Systems Platforms, Heidelberg, Germany (Middleware 2001). (2001)
10. Ganglia Monitoring System. <http://ganglia.info> (April 2008)
11. J. Han et al.: Topology aware overlay networks. In: Proc. 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2005). (2005)
12. A. Nakao et al.: A routing underlay for overlay networks. In: Proc. of the 2003 conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '03). (2003)
13. Aggarwal, V., Feldmann, A.: Locality-aware p2p query search with isp collaboration. *Networks and Heterogeneous Media* **3**(2) (2008)
14. Xie, H., Yang, Y.R., Krishnamurthy, A., Liu, Y.G., Silberschatz, A.: P4p: provider portal for applications. SIGCOMM Comput. Commun. Rev. **38**(4) (2008) 351–362
15. Hilt, V., Rimal, I., Tomsu, M., Gurbani, V., Marocco, E.: A Survey on Research on the Application-Layer Traffic Optimization (ALTO) Problem. Internet-Draft draft-hilt-alto-survey-00, Internet Engineering Task Force (July 2008) Work in progress.