A P2P Network Overlay for Distributed QoS Management and Routing

Lisandro Z. Granville, Luci Pirmez, Elias P. Duarte Jr., Jose Neuman de Souza, Rossana C. Andrade, Liane R. Tarouco, Reinaldo B. Correia, Alexandre Lages

Abstract— The management of long distance backbones based on high speed optical networks requires new solutions for challenging tasks. For instance, operators and users located at different administrative domains must communicate with each other in order to configure and monitor agreed quality of service levels. This paper proposes GigaManP2P: a novel peer-topeer (P2P) management architecture for optical networks, which was originally developed for the Brazilian RNP Giga backbone. In the architecture, peers provide management information in a ubiquitous fashion to modules that interface with both the optical infrastructure and network users. The architecture has a specific focus on QoS monitoring and routing. After QoS constraint violations are detected, a proactive rerouting strategy is triggered based on redundant virtual circuits, allowing both full and partial rerouting. GigaManP2P has been implemented, and experimental results are presented, showing the overhead the P2P infrastructure poses on raw SNMP. In order to show meaningful experimental rerouting results a simulation environment was constructed, and results were obtained for static and dynamic networks.

Index Terms: Network Management, P2P, QoS, Rerouting.

I. INTRODUCTION

THE proper provisioning of quality of service (QoS) in computer networks depends, among other factors, on efficient and accurate network management. QoS management has become one of the key challenges in high-speed networks because only properly managed QoS support can enable the deployment of critical QoS sensitive applications such as transmission of high definition TV (HDTV), telemedicine, real-time remote control, and videoconferencing.

In the recent past, QoS management was primarily concerned with the proper use of the available (and frequently scarce) bandwidth. Nowadays, with the increasingly fast deployment of optical communication technologies, bandwidthrelated issues have become less critical because the large amount of available bandwidth is often sufficient for the majority of critical applications. In this scenario, where bandwidth is not the primary concern, other QoS related issues requiring proper attention arise.

Networks with QoS support must be able to sustain the quality of transmissions for the whole lifetime of critical flows. Delivering traffic with QoS degradation is obviously unacceptable because it breaks the Service Level Agreement (SLA) established with the network customer. In optical networks, despite the usually large bandwidth available, QoS degradation may be the result of several events (e.g., congestion, faults). While some portions of the optical network may be experiencing QoS problems due to the lack of resources,

other portions may have unused resources available. When critical flows are routed through virtual circuits that cannot sustain the corresponding QoS requirements, rerouting is a practical action that may be employed in order to prevent QoS degradation and keep the applications running according to their needs.

In this article we present GigaManP2P, a peer-to-peer (P2P) management platform that features a proactive rerouting strategy used to sustain QoS in optical networks. The strategy is proactive (not reactive) in the sense that it anticipates the violation of QoS requirements. Reactive rerouting is only triggered after a given flow has suffered QoS degradation. We employ a set of rerouting-aware management agents operating above the optical infrastructure where rerouting takes place. Network devices, including optical devices, usually run simpler management software, such as popular daemons, SNMP (Simple Network Management Protocol) agents, and lightweight Web servers. Thus, more sophisticated entities, like our rerouting-aware agents, must be placed on a tier above that of the optical devices. In this context, we propose a management overlay based on the peer-to-peer (P2P) paradigm which allows services to be dynamically configured and in which the rerouting-aware agents execute and communicate with each other. The GigaManP2P management overlay is responsible for monitoring flows and triggering the rerouting process.

We evaluate our proposed solution by simulating the rerouting process as well as observe the performance aspects of each management peer when they need to access SNMP (Simple Network Management Protocol) agents at the optical devices. Since GigaManP2P has been designed for the management of the Brazilian RNP Giga backbone, our evaluations has been performed considering actual optical networks from that backbone.

The rest of the article is organized as follows. First we review traditional management technologies currently available for QoS management and routing. The GigaManP2P management overlay is then introduced for flow monitoring and rerouting. The proactive rerouting strategy is then presented in the following section. Experimental results are also given, obtained from simulation, in which we evaluated both the overhead and latency of the proposed approach. Finally we conclude the article presenting closing remarks and future work.

II. GIGAMANP2P: A QOS MANAGEMENT OVERLAY

P2P systems [1] are logical overlays deployed on top of physical networks. The nodes (peers) of a P2P system typically run at inexpensive end-user computers and have logical connections with other nodes forming a P2P network. Peers join and leave the P2P network as time goes by, which turn the network topology very dynamic. P2P network models can be characterized in several ways, using parameters such as the employment of flat organizations where every peer has the same capabilities, or having super-peers with enhanced roles like resource indexing and application-layer routing.

In order to manage the optical network of the Brazilian RNP Giga backone, we designed a P2P-based management solution (GigaManP2P) to bridge the needs of end-users and their applications with the services exposed by the optical infrastructure. In this section we focus on the P2P-based management infrastructure provided by the GigaManP2P management overlay.

A. Architecture and Management Peer Clients

GigaManP2P offers management services considering three different types of clients: network operators, end-users, and end-user applications. Figure 1 presents the general environment where one can observe the optical infrastructure, the GigaManP2P management overlay, and the management clients located along the managed network.

Peers are placed across the managed optical network to monitor and control network devices found at different administrative domains. Each peer locally offers management services to local clients (i.e., local network operators, local users, and local user applications). In addition, each peer provides additional services to other remote peers in order to form the management overlay.

Each peer, in order to accomplish its management tasks, employ a set of elements that forms its internal architecture, as presented in Figure 2.



Fig. 2. Architecture of the GigaManP2P peers

At the top of the architecture the basic communication support is found, i.e., the HTTP and JXTA elements required to enable the communication between peers and clients, and between peers themselves. JXTA is also need to maintain the P2P overlay, which is a tasks coordinated, at the core of the architecture, by the peer daemon.

B. Communication between Clients, Peers, and Management Services

Clients access GigaManP2P to request the execution of a management service. The set of available services are placed in a middle layer inside the peer architecture. If the set of management services needs to be expanded, new services can be installed using the P2P overlay itself. Such new services are stored in the "other services" repository. New services can also be deployed as management scripts, stored in the script repository. The difference between scripts and the services from the "other services" repository is that scripts implement lighter services that tend to have a schedule of execution and, after completion, are removed from the repository. Regular services, in turn, are more complex services that can be seen as extensions to the basic peer architecture.

Network operators are GigaManP2P clients that access the management services available to them through either a local GigaManP2P peer or via dynamic Web pages exposed by remote peers using HTTP. On of the key items operators are responsible for is the definition of management policies that are stored in the distributed database formed by the collection of local policy repositories of each peer.

End-users are clients that also access management services through dynamic Web pages. The set of services available to end-users, however, is restricted if compared with the set of services available to network operators. For example, endusers cannot define management policies. End-user applications are clients able to access the same services available to end-users. The difference resides in the fact that enduser applications uses Web services interface via SOAP/HTTP requests, while the human user access the same services via conventional Web pages.

C. Peers and Optical Infrastructure Communication

The modules for the communication of peers with the optical infrastructure are located at the bottom of the peer architecture. They provide an interface to access optical devices in a transparent way, regardless the actual management protocol used to access the managed equipments. These modules in fact implement an adaptation layer used by the remaining internal elements when a communication with optical devices is required. The current implementation of the GigaManP2P peers supports SNMP (for monitoring) and SSH/CLI (for configuration) in the optical communication layer.

III. THE PROACTIVE REROUTING SERVICE

GigaManP2P employs a distributed approach for QoS management. Agents are employed for setting and monitoring virtual circuits, as well as rerouting. A proactive rerouting strategy selects alternative paths for critical flows before users perceive a QoS degradation. The proposed strategy allows both full and partial rerouting, depending on whether the broken virtual circuit is completely or partially replaced by the new



Fig. 1. P2P-based optical network management overlay

one. The strategy is proactive (not reactive) in the sense that it anticipates the violation of QoS requirements. Rerouting starts with the discovery of a critical path, a subset of the virtual circuit to be replaced, and the configuration of a new alternative. The critical path is determined taking into account QoS metrics relevant to the application.

A. Rerouting Agents

Three types of agents support the proactive rerouting strategy: *InputNodeAgent*, *IntermediateNodeAgent*, and *AlternativeRouteAgent*.

The *InputNodeAgent* operates at the input (or first) device of the virtual circuit. This agent triggers the rerouting of flows belonging to the same virtual circuit and is also responsible for interacting with external modules (e. g. the module used to create virtual circuits), offering the interface through which the rerouting is accessed by the rest of the system.

The *IntermediateNodeAgent* has two main goals: (i) to monitor the optical switches that belong to the virtual circuit, and (ii) to feed the *InputNodeAgent* with performance information of the QoS metrics relevant to the flow.

Finally, the *AlternativeRouteAgent* is responsible for discovering alternative paths. The discovery mechanism operates in a limited area around the critical path called *search area*. The size of a search area is defined by a parameter called the *search radius*, which refers to the maximum number of links starting from the *InputNodeAgent*. The value of this parameter is a pre-defined for each virtual circuit. In addition, the *AlternativeRouteAgent* has two other goals: (i) to select the best alternative path, and (ii) to reconfigure the devices in order to establish the new route.

B. Rerouting Phases

The complete proactive rerouting process consists of five phases in which the operations required to reroute a flow take place:

- 1) Agent activation;
- 2) Virtual circuit monitoring;
- 3) Discovery of alternative routes;
- 4) Alternative route monitoring;
- 5) Route change configuration.

These phases, except the route change configuration phase, are executed sequentially after a critical flow is started. The route change configuration is executed when needed. The phases implement the proactive rerouting strategy, in the sense that they are executed before any QoS failure occurrs or even before detecting any QoS failure trend. The the route change configuration phase is triggered by an event explicitly requesting rerouting, as described below. This approach aims at decreasing the rerouting latency.

In the Agent activation phase, GigaManP2P activates a *InputNodeAgent* at the peer responsible for the virtual circuit's input (the first) routing device (usually an optical switch). The *InputNodeAgent* then activates an *IntermediateNodeAgent* at the next peer responsible for controlling the optical switch in the route to the destination. This *IntermediateNodeAgent* then activates another *IntermediateNodeAgent* at the next peer to the destination, and so on, step by step until all devices have an associated *IntermediateNodeAgent* activated.

The next phase, *Virtual circuit monitoring*, starts immediately after the last *IntermediateNodeAgent* is activated at the peer responsible for the output device of the virtual circuit. This phase is concluded only when the flow finishes. During this phase the operations required to obtain relevant information from the virtual circuit are executed, and subsequently sent to the *InputNodeAgent*. Only one message is employed by all *IntermediateNodeAgents* to communicate monitoring information – QoS parameters which may be configured for each flow depending on its requirements. Periodically the *IntermediateNodeAgent* at the last node creates a message with its monitored information and sends the message backwards to the previous *IntermediateNodeAgent* which updates the information, and send the message in turn to the previous agent, and so on, until the message reaches the *InputNodeAgent*.

The sequential monitoring strategy allows the discovery of the *critical paths*, a critical path is a part of the route which is the bottleneck for the set of monitored parameters in the whole route. After a critical path is detected, an *AlternativeRouteAgent* is activated by the *InputNodeAgent* at the first node of the critical path, i.e. the peer of the critical path that is closest to the *InputNodeAgent*.

The Discovery of alternative routes phase starts after the AlternativeRouteAgent is activated at the peer responsible for the first device of the critical path. The phase concludes when the last AlternativeRouteAgent is activated at the peer of the last device of the critical path, indicating the discovery of the last alternative route. A simple algorithm is employed

for the discovery of alternative routes: the limited diffusion of *AlternativeRouteAgents* in a search area with a predefined search radius.

In the *Alternative route monitoring* phase, information about the alternative routes are obtained and sent periodically to the *InputNodeAgent*. Messages flow through the *AlternativeRouteAgents* until the agent at the last peer of the alternative path, each agent updates the message with local information. The destination of these messages is the *AlternativeRouteAgent* of the first peer node of the critical path. These messages contain information that allows the *AlternativeRouteAgent* to choose the best path to employ given the resources available and flow and the requirements of the QoS flow to be rerouted.

The Route change configuration phase consists of the set of operations executed after a rerouting request is issue by an IntermediateNodeAgent. This phase is responsible for redirecting the flow to the alternative path. The redirection is executed by the AlternativeRouteAgent at the peer of the first node of the critical path. The operations involved in this phase depend on the approach adopted in the generation and association of the local identifiers of the new virtual circuit, which can be anticipated or on demand. In the anticipated approach, the local identifiers (i.e., labels) are generated in DAP phase (proactive phase) without impact in the latency of RC phase. On the other hand, the on demand approach incurs into an increase of the latency of RC phase, since the operations related to the creation and association of the local identificatiers must be done in this phase. Moreover, the latency of RC phase becomes dependent of the length of the alternative path, what does not occur in the case of the anticipated scheme.

IV. EXPERIMENTAL RESULTS

In this section we present two sets of experimental results obtained from the evaluation of GigaManP2P. A peer was implemented with the JXTA [?] toolkit and we measured the overhead of using this peer to access management information, instead of using the raw SNMP agent. A simulator was implemented in order to test the rerouting strategy, and several rerouting experiments are described.

A. Evaluation of the Peer Overhead

In this section, experimental results from the evaluation of the overhead of the P2P infrastructure on the SNMP agent are presented. This refers to the additional time required by a management peer to access an optical device in comparison to a raw SNMP agent. Please note that the SNMP agent is also employed by GigaManP2P, so the objective of the test is to quantify the latency introduced by the extra processing layers to access the optical device's information.

For accomplishing these tests, two Pentium-based machines on a Gigabit Ethernet LAN were employed. The first machine is was responsible for the execution of P2P management architecture. The second machine mimics the communication device, and runs the SNMP agent.

Results are shown in table 1, incluir uma tabela em que medidas obtidas para agentes SNMP sao comparadas com medidas obtidas para acesso com JXTA SNMP Peers

B. Evaluation of the Rerouting Strategy

A simulator was implemented with NS-2 [] for testing the rerouting strategy. The main metric of interest is the rerouting delay, i.e. the time interval the system takes between a rerouting action starts and completes. We measured the latency of three phases of the rerouting process: *agent activation*, *discovery of alternative routes*, and *reroute configuration*.

The latency of the *agent activation* phase consists of the time interval from the time instant the *InputNodeAgent* is installed at the input device to the time instant the *Alterna-tiveRouteAgent* is activated at the first node of the critical path. This latency was measured for radomly chosen virtual circuits with sizes varying from 2 to 8.

The latency of the *discovery of alternative routes* phase corresponds to the time interval from the discovery of the first alternative route, until the last *AlternativeRouteAgent* is activated at the last node of this route. This latency was measured for radomly chosen virtual circuits with sizes varying from 2 to 4.

The latency of the *reroute configuration* phase corresponds to the time instant the reroute request is received until a new route is configured, including the selection of the best alternative route to replace the critical path. This latency was measured for radomly chosen virtual circuits with sizes varying from 2 to 4.

The experiments were run on the topology shown in figure **??**. Each link was configured with a bandwidth of 10Mbps, and a delay of 2ms. Two types of traffic were configured: CBR (Constant Bit Rate) and Web. 80% of the total traffic corresponds to CBR, while 20% corresponds to Web traffic. We present results considering that the network traffic consumes from 80% to 100% of the available bandwidth.

Results are shown in figure ??.

V. RELATED WORK

VI. CONCLUSION

In this article we introduced GigaManP2P, a novel P2Pbased management architecture for QoS management and routing of large backbones that span several administrative domains. The P2P paradigm has facilities that allow the cooperation of management entities located at different domains. Traditional management solutions, based only on SNMP, do not provide this functionality.

GigaManP2P allows the negotiation of QoS requirements among three different types of clients: network operators, end-users and applications. The proposed architecture offers management services that act as a bridge between the user requirements and the optical infrastructure.

GigaManP2P is able to guarantee QoS requirements by rerouting flows before users perceive any QoS degradation. A proactive rerouting strategy was presented, in which agents monitor both the employed routes and redundant alternative routes which are selected for flows depending on user requirements and network conditions.

A peer was implemented, and an evaluation of the impact of the proposed solution on a raw SNMP agent was presented. A second set of experimental results were presented for rerouting. A simulator was implemented in order to test the proposed strategy on several network topologies under different traffic loads. Both sets of results confirm the feasibility of the proposed solution.

Future work includes evaluating the proposed approach in comparison with reactive rerouting.

REFERENCES

 E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A Survey and Comparison of Peer-to-Peer Overlay Network Schemes," *IEEE Commun. Surveys & Tutorials*, vol. 7, no. 2, pp. 72–93, second quarter 2005.