Future Generation Computer Systems 25 (2009) 137–141

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

# OptIPuter: Enabling advanced applications with novel optical control planes and backplanes

Joe Mambretti *

*International Center for Advanced Internet Research (iCAIR), Northwestern University, United States*

ARTICLE INFO

*Article history:*
Received 3 March 2008
Received in revised form
9 June 2008
Accepted 23 June 2008
Available online 5 July 2008

*Keywords:*
Optical networks
Control planes
Lightpaths
GMPLS
Dynamic provisioning

ABSTRACT

Many emerging and anticipated advanced applications have demanding requirements that cannot be supported using traditional communication services, architecture, and technologies. Such requirements include those related to high performance, large scale data volume transport, dynamic provisioning, path flexibility, and stringent determinism for all service parameters. The OptIPuter research initiative was established to investigate the potential for meeting these and other exceptional challenges by creating new types of communication services that could be implemented as complements to traditional approaches and, in some special cases, as alternatives. In particular, the OptIPuter initiative has been exploring advanced communication services based on individually addressable, dynamically allocated lightpaths. These approaches require novel control plane and backplane architecture and technology. As part of its core research agenda, the OptIPuter initiative designed and deployed such architecture and technology, which is described here along with a large scale testbed implementation, including national and international extensions. This testbed was used for a series of experiments, and the results are summarized.



## 1. Introduction

Several large classes of emerging and potential advanced applications have demanding requirements that exceed the capabilities of traditional communication services, such as Layer 3 services designed for extremely large numbers of fairly small information flows. Such services do not provide adequate support for some classes of traffic, for example, extremely large, long duration streams required by scientific research. Such applications frequently are the first to encounter problems with the technology limitations of existing infrastructure. Increasingly, global science projects require the sophisticated management, transport, and analysis of extremely large sets of data residing in multiple locations, including highly remote sites. Many such applications are at the terascale and petascale level. The challenging issues that must be addressed are not only those related to bandwidth and traffic volume. To some degree, if these were the only issues, they could be approached, although at a high cost, by adding more resources to existing infrastructure. However, merely adding resources does not sufficiently address requirements for guaranteed high performance, especially on shared infrastructure. Also, this approach would not meet the need for fast responses to changing application data flows through dynamic provisioning. General routed network environments are designed to service fairly predictable traffic behavior within set parameters and not to support many sudden changes among multiple large scale traffic flows.

The OptIPuter research initiative was funded by the National Science Foundation to investigate alternatives to traditional routed communication services by designing and implementing new services, especially those based on individually addressable, dynamically allocated lightpaths [1]. A research goal of the OptIPuter project is not simply to provide high performance, high capacity circuits, but to provide a totally new type of flexible communications environment extending to optimal path discovery and implementation while ensuring rigorous service parameters are met, for example, by supporting extremely stringent determinism by completely eliminating jitter and loss and providing for the lowest possible latency. This research requires innovative considerations of control plane and backplane architecture. Traditionally, backplane considerations relate only to infrastructure within a local computing environment. In contrast, this architecture allows multiple lightpath channels to be integrated into large scale, extended backplanes within highly distributed infrastructure, among multiple sites that can span a nation, a continent, or the globe. These backplanes become the foundation for large scale adaptable virtual environments that can include many types of resources, including computers, storage systems, instruments, and sensors. Traditional control planes are used to address fairly static resources. This architecture

* Tel.: +1 312 503 0735; fax: +1 312 503 0745.
*E-mail address:* j-mambretti@northwestern.edu.

0167-739X/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2008.06.008

allows for control planes to be integrated with other processes so that those resources can be used dynamically, with continuous interactions though specialized signaling. These new services will be implemented initially by universities, and a few corporations, which are creating regional and state-wide fiber networks, as well as national fiber infrastructure facilities, such as the National Lambda Rail (NLR) and international facilities, such as the Global Lambda Integrated Facility (GLIF) [2,3].

## 2. Problem definition

One key goal of OptIPuter research is to create a large scale distributed environment that is virtualized to such an extent that it has no dependencies on specific physical infrastructure and configurations, in part, to remove restrictive performance barriers of such infrastructure. Also, this concept envisions not creating an application that can merely utilize a particular infrastructure (whether or not it is virtual), but allowing applications to dynamically create their own environments, designing, implementing, and continually reconfiguring them. To accomplish these goals, all resources, including communications, must be addressable and configurable through high level processes. Traditional L3 services are external, non-addressable, non-configurable resources that require applications to conform to multiple preset architecture and technology parameters, for example, those that limit support for exceptionally large scale traffic flows, e.g., 1 Gbps, several 1 Gbps, 10 Gbps and higher. The OptIPuter was designed as a highly distributed infrastructure with processes that allow applications to have direct, controllable access to extremely large scale L1 and L2 resources, not as aggregated channels of smaller flows but as channels supporting a few dedicated individual flows or a single large flow. Such channels can be used to extend backplanes so they are no longer isolated resources within local computational environments but foundations for global infrastructure fabrics. For example, an integrated set of international lightpaths can be used as a global backplane to interlink high performance computers on several continents to allow all resources to be used as a single contiguous environment. Also, individual applications can continually reconfigure these backplanes as required by changing conditions. The architecture anticipates that high volume flows will be terascale in the near future, terabits per second serving terabyte data sets, soon to be followed by petascale communications. Although supporting such large scale flows is important, an equally key objective is to allow for the flexible dynamic provisioning of L1 and L2 paths. Currently, such paths are almost always statically provisioned in data networks. Using the OptIPuter control plane, applications can directly implement and manipulate backplane channels. Also, many next generation services and applications require non-traditional stream distributions within supporting infrastructure. For example, general network architecture anticipates small data flows at the edge and large scale aggregations at the core. The OptIPuter architecture anticipates large scale data flows both at the edge and within the core as well as flows that are much larger at the edge than in the core. Also, the majority of applications today are designed to be support by a single communication service layer. The OptIPuter was designed to take advantage of multi-layer service integration.

To provide for a communications infrastructure that is significantly more flexible than traditional implementations, a new control pane is required to allow for edge processes to directly interact with the resources within that infrastructure. The primary existing tools used for manipulating L1 and L2 services are centrally administered control planes designed to establish static not dynamic resources. Consequently, a major objective of this research was to design and implement an innovative control plane, created to be invoked by distributed processes used for interacting continually with communication infrastructure resources. The OptIPuter environment was not developed for static paths established for very long periods, but instead for paths that are continually being configured and reconfigured as application requirements change over time. The majority of these paths are not implemented as a shared resource as they are with traditional services, but as dedicated resources, such that they become extended backplanes of large scale clusters that can be national and international in scale. Another problem that this research is exploring is the integration of edge processes and devices with large scale bandwidth by creating such backplanes with multiple parallel lightpaths. A related problem is optimization for resource allocation, utilization, and monitoring.

## 3. Control plane architecture

To address the requirements of next generation applications and communication services, a new architecture for control planes is required, one that will provide for functions that are abstracted from the specific characteristics of supporting physical devices. Such an architecture must be developed not only for existing devices but also for anticipated future devices and for hybrid environments that contain multiple devices that must be integrated into contiguous environments. In general, the architectural design is oriented to supporting large numbers of 1 Gbps and 10 Gbps streams. An expectation is that many core systems will exist that can provide support for hundreds of concurrent 10 Gbps streams as well as 40 Gbps and 100 Gbps streams. However, any current discussion of novel control plane architecture should be placed within the context of current macro trends in this area, which are forcing the reconsideration of traditional approaches.

For example, traditionally, each network layer has generally been provided with a separate control plane. However, current architectural developments are oriented toward designing unified control planes that can operate at all network layers. Traditionally, control planes have been designed as closed proprietary systems closely interlinked with the physical characteristics of specific devices. In contrast, open architecture is now being developed, based on services supported by IP. Also, control planes, especially those for optical networks, have been designed to implement fairly static resources, such as lightpaths. In addition, a basic approach has been using an overlay model, which is based on concepts of separate domains for each layer. More recently, signaled overlay models have been developed, which provide for separate domains, but also support a wide range of interactions between the IP layer and the optical layer.

This approach, which is used by the OptIPuter, allows for significantly more flexible networking, for example, supporting resources on-demand, including end-to-end lightpaths between any two or more edge points. This model can also use out-of-band signaling, for example, through a control plane channel that is physically separate from data planes, perhaps implemented on a separate lightpath. The majority of OptIPuter experiments used out-of-band signaling. Also, its architecture supports requirement signaling for types and levels of services, such as defining priority, protection class, direction (e.g., uni-directional vs bi-directional, etc.), resource discovery and availability, lightpath management (e.g., create, delete, change, swap, reserve), optimization and related performance parameters, survival, protection, and restoration processes, etc. The architecture also provides functions for implementation, management, and routing, and includes the use of link-state-protocols.

The OptIPuter control plane is an integral part of a distributed system designed for dynamically orchestrating reconfigurable resources. It provides for an API that allows the application to signal for resources. Experimentally, several such methods were used,