DRAGON: A Framework for Service Provisioning in Heterogeneous Grid Networks

Thomas Lehman, University of Southern California Information Sciences Institute
Jerry Sobieski, University of Maryland/Mid-Atlantic Crossroads
Bijan Jabbari, George Mason University

This work is supported by the National Science Foundation (NSF) under grant nos. 0335300, 0335266, and 0335230.

ABSTRACT

Dynamic Resource Allocation in GMPLS Optical Networks (DRAGON) defines a research and experimental framework for high-performance networks required by Grid computing and e-science applications. The DRAGON project is developing technology and deploying network infrastructure which allows dynamic provisioning of network resources in order to establish deterministic paths in direct response to end-user requests. This includes multidomain provisioning of traffic-engineering paths using a distributed control plane across heterogeneous network technologies while including mechanisms for authentication, authorization, accounting (AAA), and scheduling. A reference implementation of this framework has been instantiated in the Washington, DC area and is being utilized to conduct research and development into the deployment of optical networks technologies toward the satisfaction of very-high-performance science application requirements.

INTRODUCTION

The Dynamic Resource Allocation in GMPLS Optical Networks (DRAGON) [1] project is developing technology and deploying network infrastructure that allows advanced e-science applications to dynamically acquire dedicated and deterministic network resources to link computational clusters, storage arrays, visualization facilities, remote sensors, and other instruments into globally distributed and application specific topologies. These advanced network services are motivated by the realization that science benefits when expensive resources such as radio telescopes, powerful computational clusters, and large data repositories can be easily shared among researchers regardless of location. The DRAGON project is addressing this by employing all-optical network technologies, Generalized Multi-Protocol Label Switching (GMPLS) [2, 3] routing and signaling protocols, advanced inter-domain-service routing techniques, and detailed application formalizations to deliver these advanced services. The objective is to enable the dynamic configuration of these large scientific resources and network infrastructures into application-specific network topologies in an automated response to domain scientists' requests. A reference implementation has been constructed in the Washington, DC area.

DRAGON is collaborative research project between University of Maryland (UMD) Mid-Atlantic Crossroads (MAX), University of Southern California Information Sciences Institute East (USC/ISI), and George Mason University (GMU). Other project partners include the Massachusetts Institute of Technology Haystack Observatory, NASA's Goddard Space Flight Center (GSFC), and MOVAZ Networks.

The remainder of this article is organized as follows. We describe our motivations for this work as well as our approach. We discuss our DRAGON architecture and present a description of a DRAGON network instantiation. We then present the status of our software implementation efforts. Finally, we present our conclusions.

MOTIVATIONS

HIGH-PERFORMANCE APPLICATIONS

The e-science applications noted above are examples of an increasing number of high-performance applications and users who require network services beyond what is typically available on best-effort infrastructures today. Other communities and applications whose interest in these types of services is anticipated include emergency response, mission/business-critical services, and building (or traffic engineering) of a best-effort IP network. These super users require what has been referred to as "deterministic" network services. In this context, "deterministic" implies a defined and guaranteed service level. These service-level parameters include bandwidth as a minimum. An ability to specify latency, jitter, packet loss, and other parameters is envisioned as well. These applica-

 0163-6804/06/$20.00

tions desire to obtain these network resources in such a manner that they do not have to engage in "sharing" of the resources allocated to them. This typically means they should be free to use any transport protocol that works best for their application, and not be restricted to a congestion-control scheme that is TCP friendly. In addition, these services need to be provisioned on an interdomain basis (across heterogeneous network technologies) and include features for authentication, authorization, accounting (AAA) and scheduling. It should also be noted that the nature of how these applications desire to utilize the network does not lend itself to the purchase or provisioning of dedicated network infrastructures, since these applications generally do not require these high-performance network services on a continuous basis. However, when these network services are required, there is often a large amount of cost and effort associated with the reservation of scarce time on expensive science resources, preparation for the observation of an infrequent physical event, or other considerations for which a well-defined and schedulable network service is critical.

As an example of such an application, the DRAGON project is working closely with global leaders in the field of electronic very-long baseline interferometry (e-VLBI). A more detailed description is provided below.

e-VLBI

VLBI has been used by radio astronomers for more than 30 years as one of the most powerful techniques for studying objects in the universe at ultra-high resolution and measuring earth motions with ultra-high accuracy. VLBI allows images of distant radio sources to be made with resolutions of tens of microarcseconds, far better than any optical telescope. VLBI also provides a stable inertial-reference frame formed by distant quasars to study the motions of the Earth in space with exquisite precision, revealing much information about both the surface and internal motions of the Earth system, including interactions with the dynamic motions of the atmosphere and oceans.

VLBI combines data simultaneously acquired from a global array of up to ~20 radio telescopes to create a single coherent instrument. Traditionally, VLBI data are collected at data rates close to ~1 Gb/s on magnetic tapes or disks that are shipped to a central site for correlation processing. This laborious and expensive data-collection and transport process now has the possibility of being replaced by modern global high-speed networks, potentially enabling important new capabilities, real-time data correlation and analysis, and scientific returns. By nature, VLBI data are digital representations of the analog signal arriving at a radio telescope and sampled at the Nyquist rate so that each sample is independent and the data are uncompressible.

The transmission of VLBI data via high-speed network is dubbed "e-VLBI." As part of earlier work in this area, we conducted several e-VLBI demonstrations across best-effort IP networks using antennas in Westford, MA and Greenbelt, MD with correlation at Haystack Observatory in Westford, MA. While these experiments were a successful first step for e-VLBI ,they also highlighted the many problems when conducting this type of experimentation via existing high-speed networks and protocol suites. Of primary concern was the nondeterministic performance response of the network. While nominally the network capacity was sufficient to support the ~1 Gb/s experiment, there was no way to guarantee a level of service for the duration of an experiment. Application developers understand that cost-effective capacity provisioning may dictate that a desired capacity is not always available. However, they would like to have a deterministic performance level. Mechanisms that would allow applications to determine (and reserve) end-to-end capability in advance would allow them to plan for a capacity they can count on for the duration of an experiment.

Utilizing the DRAGON control plane and network implementation described in this article, we are working with the VLBI domain experts to develop real-time e-VLBI correlation. We have demonstrated this with a data rate of 512 Mb/s from each of multiple radio telescopes. Real-time correlation provides a timeline for data processing (and the discovery of data recording errors) that is close to instantaneous, as compared to days with traditional methods.

Utilizing the DRAGON control plane and network implementation described in this article, we are working with the VLBI domain experts to develop real-time VLBI correlation. We have demonstrated this with a data rate of 512 Mb/s from each of multiple radio telescopes.

The DRAGON Approach

A key objective of the DRAGON project is to be able to provision these deterministic services or "lightpaths" on an interdomain basis and across heterogeneous network topologies. Interdomain is important because expensive scientific resources are geographically dispersed around the world and under the control of a variety of administrative organizations. Provisioning across heterogeneous network topologies is important because different organizations or administrative domains build their network infrastructures in a variety of ways. For instance, the DRAGON implementation in the Washington, DC area is based on an all-optical dense wave-division multiplex (DWDM) infrastructure with multiprotocol client interfaces supported. Other infrastructures may be based on SONET, Ethernet, or some other technology over which deterministic paths can be established. To maximize the ability to share resources in this environment, a control plane is needed which will allow end-to-end provisioning in this environment. We began with GMPLS as our basic building block in the development of this capability. GMPLS is addressing the traffic engineering (TE) constraints as applied to heterogeneous networks and includes issues associated with routing, path computation, and signaling in this environment. There were several unresolved problems that we needed to address, including issues associated with interdomain routing, interdomain signaling, and heterogeneous network technologies. In addition, integration of features for AAA and scheduling into the provisioning environment was also required to enable organization to effectively use these capabilities. The remainder of this article describes our architecture for an

The NARB is an entity that represents the local autonomous system or domain, and serves as a path computation engine from which end-systems or other devices can query to find out about availability of traffic-engineered paths between specified source and destination pairs.

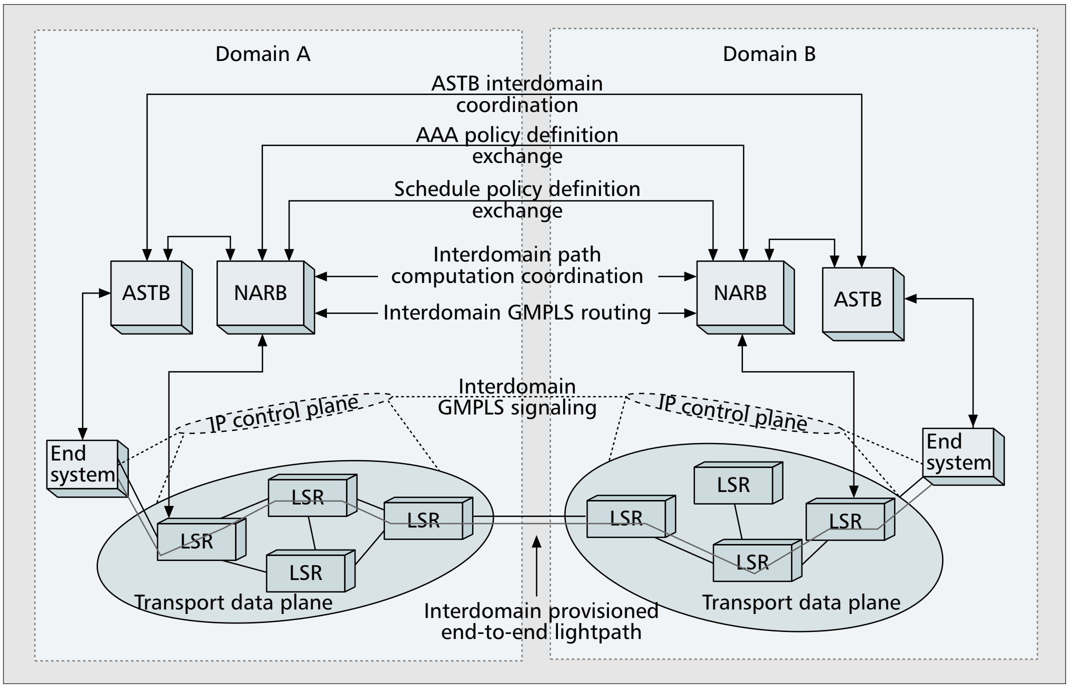

Figure 1. *DRAGON control plane architecture.*

optical control plane that permits dynamic provisioning of lightpaths across heterogeneous network technologies and includes features for AAA and scheduling. A description of our DRAGON network implementation in the Washington, DC area is also provided.

DRAGON ARCHITECTURE

A key factor in the development of the DRAGON control-plane architecture is the realization that network infrastructures are highly diverse, and this diversity will increase in the future. The diversity referenced here includes the type of network technologies, internal provisioning mechanisms, administrative ownership, use policies, and capabilities. As a result of this diversity, it follows that different networks will have different internal provisioning mechanisms. So while the DRAGON control-plane architecture utilizes GMPLS as a basic building block, it does not assume that GMPLS protocols will be used within a domain for actual provisioning. The issues that must be resolved for an interdomain control plane are the protocols and messages exchanged across a domain boundary. While the DRAGON network implementation also uses GMPLS for its internal provisioning, this is not a requirement. The interdomain architecture and provisioning mechanisms described here are GMPLS based in the context that they rely on a link-state protocol for interdomain topology exchanges, engage in multiconstraint path computation to determine appropriate network routes, and utilize RSVP-TE [4] for signaling. Should a domain decide to use a provisioning mechanism other than GMPLS internally, interoperability will require translation of those internal representations to external representations that are compatible with these control-plane definitions. The architecture we developed is depicted in Fig. 1. A description of this architecture is provided in the following subsections.

ARCHITECTURAL COMPONENTS

Network-Aware Resource Broker (NARB) — Routing, path computation, and signaling on an interdomain basis across topologies, which include a heterogeneous mix of network technologies and vendor equipment, is beyond what is defined in standards and also beyond the capability of current vendor equipment. To enable routing, path computation, and signaling in this environment, the network-aware resource broker (NARB) provides several important functions. The NARB is an entity that represents the local autonomous system (AS) or domain, and serves as a path computation engine from which end-systems or other devices can query to find out about availability of traffic-engineered paths between specified source and destination pairs. The NARB is also responsible for interdomain routing. NARBs peer across domains and exchange topology information so as to enable interdomain path computation and label switched path (LSP) provisioning. This interdomain topology exchange can be based on the actual topology, as discovered by listening to the local OSPF-TE [5] protocol or optionally based on an "abstracted" view of the domain topology (generated by the configuration file or automatic synthesis of the OSPF link state database). Domain abstraction provides mechanisms for an administrative domain to advertise to the outside world a highly simplified view of its topology. This allows domains to hide their real topologies as well as minimize the amount of external updates required. The trade-off is reduced accuracy for path computations. Each administrative domain can utilize configuration parameters to tailor its domain abstraction to the level desired. One of the goals of our project is to evaluate

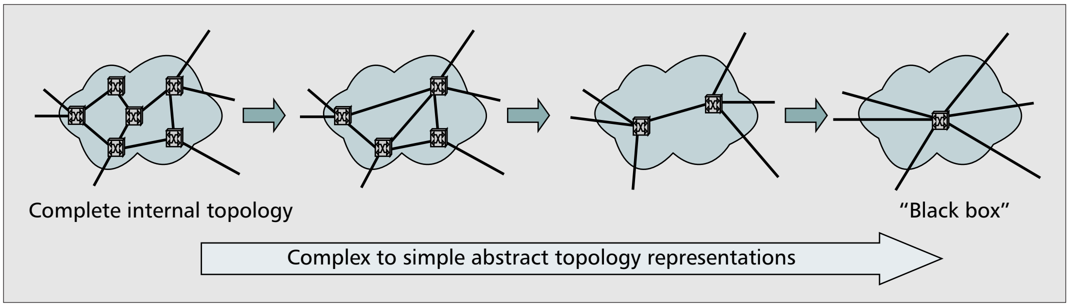

Figure 2. *Domain abstract topology representations.*

Domain abstraction provides mechanisms for an administrative domain to advertise to the outside world a highly simplified view of its topology. This allows domains to hide their real topologies as well as minimize the amount of external updates required.

various architectural issues. The domain-abstraction features of NARB are geared toward allowing experimentation with differing levels of topology hiding. Examples of various levels of domain abstraction are depicted in Fig. 2. The resulting interdomain architecture would most accurately be described as a hybrid between a peer-to-peer and overlay model.

The NARB also includes advanced algorithms, which allow path computation with multiple constraints. Some of these features are similar or comparable to those described in the IETF Path Computation Element Architecture [6, 7]. These path-computation algorithms allow network paths to be identified across multidomain environments. The NARB first calculates this path based on the interdomain abstract-topology representations as described above. This is generally sufficient to allow signaling and actual provisioning of the network path to occur. However, in some situations this initial path-computation result will be expanded into a higher fidelity path via coordination of NARBs in different domains. This provides for optimal traffic engineering and resource management.

The NARB design is such that the decision on the degree of path computation fidelity can be decided at provision time, based on policy and operational constraints. The set of GMPLS TE constraints can be fairly large and complex, especially for a multidomain topology, which includes different types of network technologies or regions. These regions are generally identified by the type of switching capabilities supported. GMPLS identifies several switching capabilities, including packet-switch capability (PSC), Layer-2 switch capability (L2SC), time-division multiplexing (TDM), lambda-switch capability (LSC), and fiber-switch capability (FSC). The NARB includes features to account for transitions across switching capability regions as well as to address specific vendor limitations or incompatibilities.

The standards-defined set of GMPLS traffic-engineering information and capabilities is advertised via a routing protocol such as OSPF-TE via protocol packets called link-state advertisements (LSAs). The accumulation of this data from all the nodes in the network is typically referred to as the traffic-engineering database (TEDB). In addition to the standard GMPLS TE constraints, constraints for AAA and scheduling are also included in the TEDB. This is accomplished via the use of some new OSPF-TE LSAs, which specifically associate AAA and scheduling information with network resources. We refer to this as a 3D TEDB. The three broad dimensions we refer to are standard GMPLS TE constraints, AAA constraints, and scheduling constraints. The combination of path computation with constraints for AAA and scheduling and the 3D TEDB is referred to as our "3D Resource Computation Element (3D RCE)."

The DRAGON control plane utilizes the 3D RCE to implement "policy-based provisioning." The objective of the policy-based provisioning is to incorporate AAA policy and schedule information into path computation, resource allocation, and signaling functions. This requires high-level associations of policy with users (or groups of users) as well as lower-level associations of policy with actual network elements at a fidelity sufficient to implement meaningful policy-based resource allocations. The NARB includes features to map high-level user-based AAA and schedule policy information to network resources per the 3D RCE as described above. Domain administrators can specify this higher-level policy in terms of user-specific bandwidth and schedule constraints which the NARB translates into specific GMPLS-TE AAA and scheduling LSAs. A similar process in other domains allows one NARB to learn about other domains' AAA and scheduling constraints via the GMPLS routing protocols. As with other routing information, domain-abstraction techniques allow for domain-specific minimization of AAA and scheduling information, which is advertised to the outside world. For the higher-level AAA functions, there is related work underway or completed in several communities. From the GRID community this includes the Open Services Grid Architecture (OSGA) developed by the Global Grid Forum (GGF), and the associated Global Security Infrastructure (GSI) developed by the Globus Consortium. From the Internet2 research community, Shiboleth is an architecture and product that provides similar AAA functions. While these have been generally geared toward grid applications, they are now being reevaluated for application to network-provisioning functions. Other related work from the IETF includes working drafts on the integration of AAA functions into the Session Initiation Protocol (SIP). These efforts are addressing many of the generic issues associated with AAA and include mechanisms to aggregate policies, develop virtual organizations, and include the use of certificate based and/or ticket based authentication mechanisms.

The DRAGON architecture includes the notion of establishing application-specific topologies. These are requested by an end user and are generally a set of LSPs, which an application domain desires to be set up as a group.

We are interested in experimenting with multiple architectures and techniques for AAA policy-based provisioning and intend to evaluate multiple higher-level AAA and scheduling methods. As a result, we have attempted to incorporate AAA and scheduling information into our 3D RCE in a manner which is flexible enough to adapt to changes in the higher-level AAA architecture. This includes a NARB API to allow the development of interfaces to other higher-level AAA mechanisms such as those mentioned above. We plan to evaluate several of these and report findings and recommendations once they are complete.

Application-Specific Topology Builder — The DRAGON architecture includes the notion of establishing application-specific topologies (ASTs). These are requested by an end user and are generally a set of LSPs, which an application domain desires to be set up as a group. The Application-Specific Topology Builder (ASTB) accepts requests from users or end systems for multiple network connections, and utilizes the services of the NARB to determine if the requested network paths are available with appropriate AAA and schedule constraints applied. The NARB views these requests as individual LSPs and the ASTB is responsible for the assembly of multiple LSPs in to a specific topology.

Virtual Label Switch Router — In order to provide end-to-end automated provisioning, it was necessary to provide the GMPLS protocols to cover switching components that did not have their own native GMPLS protocols. For this reason the virtual label switch router (VLSR) was developed. A non-GMPLS capable network device is converted to a VLSR by the addition of a small UNIX-based PC which runs a GMPLS control plane consisting of OSPF-TE [5] and RSVP-TE [4]. The VLSR PC acts as a GMPLS proxy agent for a device and translates protocol events into commands that the local switching element understands, such as SNMP, TL1, or even scripted CLI commands. This allows non-GMPLS devices to be included in end-to-end path instantiations. The primary use for VLSR on the DRAGON project is to control Ethernet switches via the GMPLS control plane. However, the VLSR has also been adapted to control TDM and optical switches. While a VLSR is not identified directly in the architectural diagram, any of the LSRs identified could in fact be a VLSR.

End-System Agent — The end-system agent (ESA) is software that runs on (or on behalf of) the end-system that terminates the data plane (traffic engineering) link of the provisioned service. This is the software that participates in the GMPLS protocols to allow for on-demand end-to-end provisioning from end system to end system. The ESA typically runs in peer-to-peer mode or overlay mode via a UNI protocol [8]. For a peer-to-peer mode interaction, the ESA includes an instance of OSPF-TE and RSVP-TE. For a UNI mode interaction, the ESA would include an instance of RSVP-TE only. In either case, the ESA allows the initiation of a provisioning action on behalf of the end system. The ESA may also interact with the ASTB if a more complicated topology is to be built, or if the end system does not desire to run any GMPLS protocols such as OSPF-TE or RSVP-TE. For a complicated topology, the ESA will send a request to the local ASTB describing a number of network paths to be set up at the same time. The ASTB will handle instantiating them or reserving them in a group as requested. This provides for a basic bulk setup and scheduling of network paths. Future plans include adding interactive and iterative processing to better enable the ASTB to identify acceptable alternatives should it not be possible to satisfy the full set from the initial request.

DRAGON NETWORK

The DRAGON testbed is situated in the Washington, DC metropolitan area. It is architected around a multiprotocol all-optical metro-area WDM infrastructure so as to allow for maximum flexibility in terms of the type of end systems and applications which can be supported. It consists of two interlocking dark fiber rings interconnecting five core points of presence. The total diameter is approximately 150 km, although it does support one channel across BossNet for a distance of 800 km. This network is depicted in Fig. 3 and consists of optical switching nodes and edge devices configured in a hybrid ring/mesh topology. This fiber infrastructure carries a DWDM system designed for 40 waves at 100 GHz spacing. At each of the five core nodes is a photonic-wavelength-switching capability. The photonic switches are MEMs-based reconfigurable OADMs (ROADMS) provided by Movaz Networks. These ROADMs are GMPLS capable, and include some early deployment of prototype equipment. Each WDM port on the ROADMS is connected either to another core switching node or to an edge OADM. These edge OADMs are where ITU wavelengths are created via conventional transponders, or where external waves are introduced into the core photonic-transport layer. The core photonic network is framing agnostic; that is, it does not care about framing formats. The core photonic engineering is designed to support modulation rates up to 13 Gb/s.

Edge OADMs are typically connected to campus switches that provide high-layer services such as SONET or Ethernet. The photonic layer is used to provision both transient service connections as well as more persistent connections that support upper-layer services. End users are encouraged to source ITU waves directly from their laboratory switches using ITU-compliant small-form-factor pluggable laser modules.

The increased flexibility offered by an all-optical switched core does present several challenges in terms of routing, path computation, and dynamic provisioning. Because there is no optical-to-electrical conversion across the core of the network, special attention must be applied to the optical physical impairments such as attenuation and dispersion. In addition, algorithms to calculate and minimize wavelength

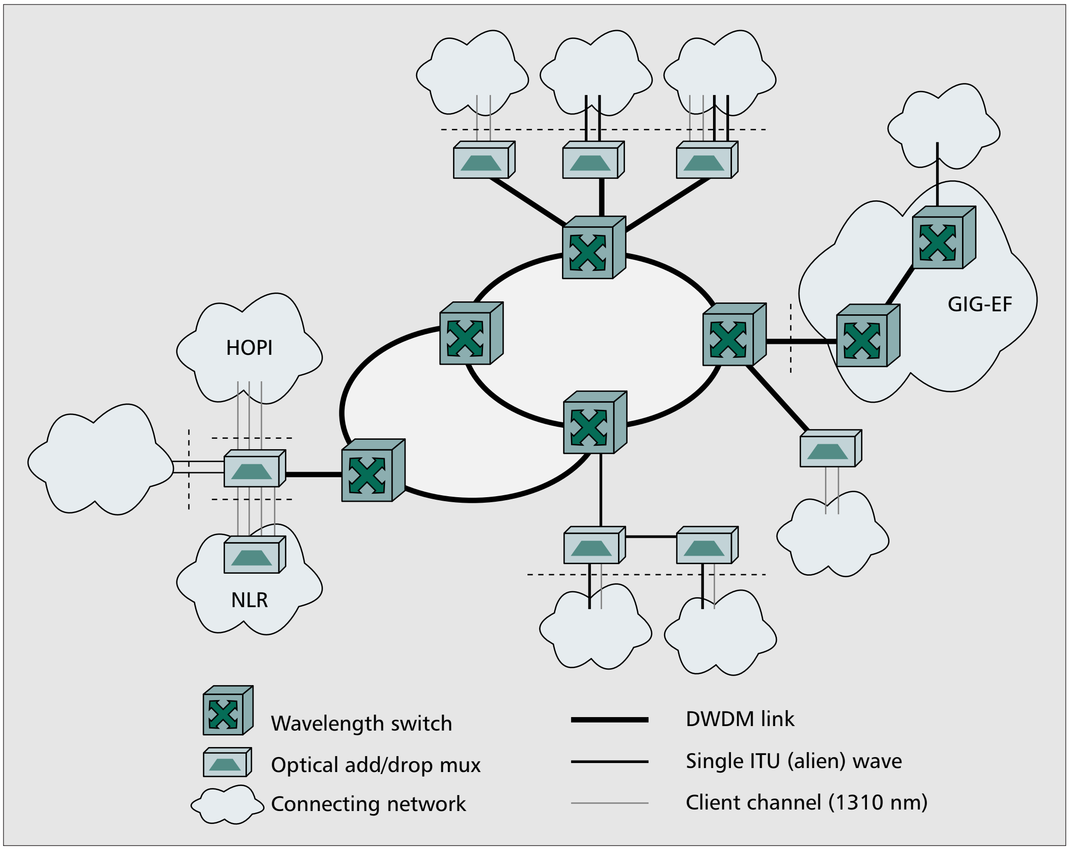

Figure 3. *DRAGON network.*

blocking are required. These parameters may change as different paths are provisioned. As part of our research agenda, we are developing techniques to dynamically include these parameters into path computation and provisioning algorithms. Additionally, some of the equipment which connects to the optical network core will be network equipment as well, which is also capable of provisioning based on GMPLS. This network equipment might be routers, ethernet switches, TDM switches, or a variety of other equipment options. This introduces the need to conduct traffic engineering, path computation, and provisioning across multiple network technologies or layers. The requires extra intelligence in the control plane to account for these situations. To facilitate provisioning in an all-optical environment, we are also planning to introduce other advanced capabilities. This includes the dynamic use of tunable lasers, tunable filters, wavelength translation, and dispersion compensation in direct response to individual provisioning request

Above the photonic layer, DRAGON also supports an Ethernet service layer. This has proven to be convenient in terms of development and testing of control-plane software, but we expect ultimately it will be subsumed by wavelength and TDM layers for transport with Ethernet framing used simply as a convenient access format. This is because dedicated waves and/or TDM circuits provide a more deterministic and predictable performance than shared Ethernet services. Indeed, while DRAGON does Ethernet switching, it is generally port-to-port Layer 2 switching of Ethernet-framed waves or fiber rather than switching of shared VLANs from one trunk to another. The latter is largely the same as MPLS LSPs and so seems less interesting from an advanced services and research standpoint.

The DRAGON network also has an IP layer. The IP layer provides for packet-based LSP capability. LSPs have been used quite extensively by commercial ISPs to provide both efficient utilizations of core IP resources as well as the more recent deployment of VPN services. DRAGON is exploring means to integrate such Layer 3 LSP capabilities with Layer 2 or lower transport capabilities.

Given these network capabilities, the DRAGON control plane reflects the principles described in this article and implemented alongside the DRAGON data plane.

The DRAGON network connects multiple academic, government, and other research institutions in this metropolitan area. In addition, there are peerings with other network infrastructures, which provide for national and global interconnections of lightpath services. This includes a peering with the Global Information Grid Evaluation Facilities and BossNet, which provide connectivity to the MIT Haystack Observatory, where the e-VLBI resources are located. In addition, there is a direct peering with the Hybrid Packet/Optical Infrastructure project from Internet2. This provides for national and global connections to a variety of advanced network infrastructures and organizations.

Above the photonic layer, DRAGON also supports an Ethernet service layer. This has proven to be convenient in terms of development and testing of control-plane software, but we expect ultimately it will be subsumed by wavelength and TDM layers for transport with Ethernet framing used simply as a convenient access format.

We are still working on some design and implementation issues associated with the integration of the higher-level AAA, scheduling, and policy information into the NARB. This will likely require additional features and modifications to the other parts of the control plane as well.

SOFTWARE IMPLEMENTATION

As part of the DRAGON project, we have developed software based on the architecture presented in this article. We have completed implementations of NARB, ASTB, ESA, and VLSR, as identified in this document. We are still working on some design and implementation issues associated with the integration of the higher-level AAA, scheduling, and policy information into the NARB. This will likely require additional features and modifications to the other parts of the control plane as well.

We have modified and extended open-source software for some of our implementation efforts. For GMPLS routing, we have extended the open-source ZEBRA software (www.zebra.org). For GMPLS signaling, we have extended the KOM RSVP implementation from Darmstadt University of Technology. This includes extending these software components to include the needed MPLS/GMPLS IETF standards.

CONCLUSION

In this article we have described an architecture and implementation for a network which can provide for dynamically provisioned deterministic-network paths. This includes mechanisms to allow provisioning across multidomains with AAA and scheduling features. These features have been implemented in the DRAGON network and are being utilized to provide advanced services to science applications.

REFERENCES

[1] DRAGON Web site, dragon.maxgigapop.net

[2] D. Awduche and B. Jabbari, "Internet Traffic Engineering using Multi-Protocol Label Switching (MPLS)," invited paper, *J. Comp. Net.*, vol. 40, no. 1, Sept. 2002, pp. 111–29.

[3] E. Mannie, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," RFC 3945, Oct. 2004.

[4] L. Berger, Ed., "Generalized MPLS Signaling — RSVP-TE Extensions," RFC 3473, Jan. 2003.

[5] D. Katz, D. Yeung, and K. Kompella, "Traffic Engineering Extensions to OSPF Version 3," Internet Draft (work in progress), Mar. 2005, available at draft-ietf-ospf-ospfv3-traffic-05.txt

[6] A. Farrel, J-P, Vasseur, and J. Ash, "Path Computation Element (PCE) Architecture," Mar. 2005, available at draft-ietf-pce-architecture-00.txt.

[7] J. Vasseur, A. Yyangar, and R. Zhang, "Inter-domain Traffic Engineering LSP Path Computation Methods," Jan. 2005, available at draft-vasseur-ccamp-inter-domain-path-comp-00.txt.

[8] G. Swallow *et al.*, "Generalize Multiprotocol Label Switching(GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model," Oct. 2004, available at draft-ietf-ccamp-gmpls-overlay-05.txt.

BIOGRAPHIES

THOMAS LEHMAN (tlehman@east.isi.edu) received his B.S. degree from Virginia Tech, Blacksburg, and an M.S. degree from The Johns Hopkins University, Baltimore, Maryland, in electrical engineering. He is a computer scientist in the Computer Networks Division at the University of Southern California's Information Sciences Institute (ISI). His research interests are in the areas of advanced network architectures, network control planes, research testbeds, end-to-end application performance, network monitoring, and network security.

JERRY SOBIESKI is the director of research initiatives for the Mid-Atlantic Crossroads (MAX), a consortium of 40+ research and higher education institutions in the Washington, DC region. He is responsible for developing strategic and multi-institutional network research programs that address the needs of the next generation of globally distributed "e-science" applications. Besides his work on DRAGON, he also heads up the Testbed Support Center for the Internet2 Hybrid Optical/Packet Infrastructure (HOPI) — a national research effort funded by the Internet2 organization exploring next-generation Internet architectures, and the Global Information Grid Experimental Facility (GIG-EF), a DoD-funded advanced technology testbed. Previous to his work in advanced networks, he was focused on high-performance parallel and distributed computing research at the University of Maryland Institute for Advanced Computer Studies.

BIJAN JABBARI received his Ph.D. degree from Stanford University in electrical engineering. He is a professor of electrical engineering at George Mason University, Fairfax, Virginia, and an affiliated faculty member with École Nationale Superieure de Télécommunications, Paris, France. He is the past chairman of the IEEE Communications Society Technical Committee on Communications Switching and Routing. He is a recipient of the IEEE Millennium Medal (2000) and the Washington Metropolitan Area Engineer of the Year Award (2003). He continues research on multi-access communications and high-performance networking.