Evaluating the Impacts of Network Information Models on Applications and Network Service Providers

Nut Taesombut
Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Dr., La Jolla, CA 92093-0404
1-858-534-5486
nut@cs.ucsd.edu

Andrew A. Chien
Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Dr., La Jolla, CA 92093-0404
1-858-822-2458
achien@cs.ucsd.edu

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations– *network management.*

General Terms

Experimentation

Keywords

Configurable optical network, information model, Lambda-Grids

1. INTRODUCTION

Configurable optical networks and wide-area resource sharing in the form of Grid computing provide intriguing opportunities for new application capabilities and resource efficiencies. The resulting Lambda-Grid is a collection of geographically dispersed compute, storage and visualization resources that can be interconnected on-demand with private, high speed optical circuits (i.e., lambdas). Such as a capability enables novel, large-scale distributed applications such as scientific data sharing and distribution, and collaborative data visualization, which require large resource aggregations and high-quality network service. Further, dynamic provisioning enables distributed and network resources to be efficiently utilized and shared by different applications over distinct time periods.

A key challenge for configurable optical networks is definition and widespread acceptance of Network Information Models (NIMs). NIM provides information about network capabilities and resources (e.g., topology, connectivity and link state information), allowing applications to dynamically select and configure network and distributed resources for their private use. When a network is composed of multiple domains, a critical tension is between Internet Service Providers (ISPs) who are business competitors and not willing to share sensitive information of their internal network. By publishing internal network information, it could make an ISP vulnerable to a range of security threats and point a way to other ISPs to gain competitive advantages by offering better network services. At the same time, detailed network information is crucial for efficient path computation and traffic engineering. Therefore, there is the need for intelligent information sharing that not only enables effective path selection for Grid applications, but also maintains competitive advantages of individual ISPs.

This paper explores a spectrum of network information models (NIMs) and evaluates their impacts on applications' and service providers' ability to utilize network resources using a trace-driven simulation.

2. NETWORK INFORMATION MODELS

We assume that a configurable optical network is made up a collection of optical switches interconnected by DWDM links. Each end compute and storage resource has one or more optical interface to an optical network. A private optical circuit can be configured on-demand between a pair of end resources through a set of optical switches which forward data along the established path. The network is partitioned into domains; distinct groups of optical switches are managed independently by different ISPs.

Network information can be broadly categorized into domain-level and interdomain-level information. First, *domain network information* includes domain topology and link state information. *Domain topology information* specifies interconnectivity between optical switches and links within a domain and the latency of each link. *Domain link state information* specifies the capacity and usage of domain links. Second, *interdomain network information* includes interdomain topology and connectivity information. *Interdomain topology information* specifies interconnection between domains, including the latency, capacity and usage of interdomain links. *Interdomain connectivity information* specifies network reachability among ISPs using BGP-like "distance-vector" information.

We define six network information models below.

1. Open Interdomain, Open Domain (*Open*): provides full interdomain and domain network information. This assumes complete trusts amongst ISPs in an open infrastructure, and allows an external agent to control the selection and configuration of an entire network path across domains.

2. Open Interdomain, Topology Domain (*TopoDom*): provides domain topology and full interdomain network information. The rationale is that ISPs are not willing to publish information about their internal network usage/capacity because it can be used by other ISPs to gain competitive advantages.

3. Open Interdomain, Connectivity Domain (*ConnDom*): provides interdomain topology and summarized domain connectivity. The key idea is that when internal network information cannot be shared, some abstraction of domain connectivity can be useful to enable more efficient path computation [1]. Our specific approach provides approximate latency of connectivity between border switches and between pairs of each border and internal switch

HPDC'07, June 25–29, 2007, Monterey, California, USA.
ACM 978-1-59593-673-8/07/0006.

within a domain. This approach allows us to estimate the latency of an end-to-end circuit path across domains.

4. Topology Interdomain (*TopoInter*): provides interdomain topology information. While this model hides all domain information, it offers diverse interdomain paths.

5. Connectivity Interdomain (*ConnInter*): provides interdomain connectivity information. Information hiding provides a means for an ISP to enforce interdomain routing policies. This model reflects the philosophy of BGP interdomain routing to disseminate interdomain route instead of topology information.

6. No Information (*None*): provides no network information.

3. EVALUATION

3.1 Methodology

To evaluate the proposed NIMs, we used synthetic workloads, trace-driven simulations and metrics across application communication latency, system throughput and lambda utilization. *Application latency* is the average latency of the allocated circuit path for each request, *system throughput* is the average number of active requests and *system lambda utilization* is the fraction of available lambdas allocated for use.

The simulated infrastructure consists of an optical circuit-switched network and Grid resources. Our resource model is based on recent work on realistic resource modeling for today's Grids [2]. Our optical network model is derived from a realistic map of the Internet backbone network composed of ten leading ISPs in the world – AT&T, BT, Cogent, Global Crossing, Level3, NTT, Qwest, Sprint, Time Warner and Verizon. We derived the current network map of individual ISPs from their websites and inferred their peering points from the Rocketfuel's traceroute data [3]. The resulting network contains 1,219 switches and 2,351 links. For each link we assigned 20 lambdas, each with 1 Gbps.

Our workload is a synthetic trace of movie content delivery requests. Each request is a client at a known location requiring the streaming of a certain movie object via a 1 Gbps private link. We assume that the simulated resource infrastructure contains a set of replica servers, each maintaining a collection of movie contents. 4,876 servers were randomly chosen from the Grid resource pool. We generated 5,000 distinct movie objects with average size of 200 GB which runs for approximately 1 hour and 50 minutes. Our decisions to replicate movie objects to replica servers are based on the popularity replication heuristic algorithm [4]. The popularity of movie objects following the Zipf distribution. In our simulation, requests were removed from a queue one-by-one and each is allocated with a "shortest" network path to the replica server that stores a particular data object.

3.2 Preliminary Results

Table 1 compares the average application communication latency, system throughput and lambda utilization achieved by the six network information models (Open, TopoDom, ConnDom, TopoInter, ConnInter and None). These results were reported using the request rate of 30 requests/min which is high enough for all metrics to be measured at their saturation regions.

We find that Open and TopoDom achieve comparable application latency and significantly outperform other models. This is attributed to domain topology information which enables efficient network path computation and selection, a key driver for low application latency. Further, TopoInter's and ConnInter's application latency is comparable and slightly higher than that of ConnDom. This shows that when internal network information cannot be shared, approximate domain connectivity information in ConnDom can be useful. Additionally, interdomain topology information provides little benefit over interdomain connectivity information. This is because the studied network contains a small number of large ISPs and most of these ISPs are directly peered.

Table 1. Comparison of application performance, system throughput and lambda utilization of different network information models

Network Information Models	Application communication latency (msec)	System throughput (number of applications)	System lambda utilization (percent)
Open	12.966	908.553	10.665
TopoDom	12.513	743.839	9.025
ConnDom	25.585	446.682	9.415
TopoInter	31.328	399.514	9.435
ConnInter	30.864	394.462	9.220
None	43.727	189.581	6.833

We see that different information models achieve varying degrees of system throughput. While all network information has positive impact on system throughput, we find that the two key factors are domain topology and link state information. Domain topology information allows us to make efficient use of lambdas, thus increasing the probability of satisfying subsequent requests. Link state information is useful when a network becomes congested, necessitating the ability to identify and avoid highly-utilized links.

Lastly, we see that all models produce comparable and low system lambda utilization. This is a limit due to the nature of our workload that contains many requests requiring long lambda paths across the network that cannot be simultaneously satisfied. However, Open achieves the highest lambda utilization due to its domain link state information which allows us to exploit diverse domain paths in the face of resource contention.

4. ACKNOWLEDGMENTS

Supported in part by the National Scientific Foundation under awards NSF Cooperative Agreement ANI-0225642 (OptIPuter), NSF CCR-0331645 (VGrADS), NSF ACI-0305390 and NSF Research Infrastructure Grant EIA-0303622.

5. REFERENCES

[1] Bernstein, G. M., Sharma, V. and Ong, L. Interdomain Optical Routing. *Journal of Optical Networking*, Vol. 2(1), February 2002.

[2] Kee, Y., Casanova, H. and Chien, A. A. Realistic Modeling and Synthesis for Computational Grids. In *Proceedings of SC'04*, November 2004.

[3] Sprint, N., Mahajan, R. and Wetherall, D. Measuring ISP Topologies with Rocketfuel. In *Proceedings of SIGCOMM'02*, August 2002.

[4] Kangasharju, J., Roberts, J. and Ross, K. W. Object Replicaiton Strategies in Content Distribution Networks, *Computer Communications*, 25(4), April 2002.