DOE Final Report

1.DOE award # and name of the recipient (Institution)
DOE award #: DE-FG02-09ER46622 (DE-SC0001277)
Name of Institution: University of Nebraska-Lincoln

2.Project Title and name of the PI
Project Title: Dynamic Optimized Advanced Scheduling of Bandwidth Demands for Large-Scale Science Applications
Name of PI: Byrav Ramamurthy

3.Period covered by the report
August 15, 2009 to August 14, 2013

4.A brief description of accomplishments. This can be in bullet form or whatever you think is useful and appropriate to indicate the sense of progress. Please limit this section to no more than 5 pages.

In this project, developed scheduling frameworks for dynamic bandwidth demands for large-scale science applications. In particular, we developed scheduling algorithms for dynamic bandwidth demands in this project. Apart from theoretical approaches such as Integer Linear Programming, Tabu Search and Genetic Algorithm heuristics, we have utilized practical data from ESnet OSCARS project (from our DOE lab partners) to conduct realistic simulations of our approaches. We have disseminated our work through conference paper presentations and journal papers and a book chapter.

In this project we addressed the problem of scheduling of lightpaths over optical wavelength division multiplexed (WDM) networks. We published several conference papers and journal papers on this topic. We also addressed the problems of joint allocation of computing, storage and networking resources in Grid/Cloud networks and proposed energy-efficient mechanisms for operatin optical WDM networks.

The following are the technical details of our approaches.
We addressed the problem of lightpath scheduling in high-speed optical networks [see A below]. It is an important capability in next-generation wavelength-division multiplexing (WDM) optical networks to reserve resources in advance for a specified time period while provisioning end-to-end lightpaths. We proposed a continuous and parallel optimization method to address the dynamic and deterministic bandwidth-scheduling problem in next generation wavelength-division multiplexing (WDM) networks. In this method a greedy algorithm and genetic algorithm are run in parallel in separate threads and both of them take the Dynamic Scheduled Bandwidth Demand (D-SBD) as their

input. The user gets his response only from the greedy algorithm and hence he will get a deterministic answer in a short amount of time. The genetic algorithm takes as one of its inputs the output of the greedy algorithm and does the optimization of the D-SBDs with minimizing blocking probability as its fitness function. The greedy algorithm copies the optimized reservation database of the genetic algorithm at regular intervals. The user submitting a D-SBD request is unaware of the optimization done by the genetic algorithm. This method is evaluated using both trace-driven simulation of real network traffic from the DOE ESnet network and stochastic traffic in ESnet network topology and a 24 node network topology. We also compared our approach with an earlier proposed method called re-optimization at blocking. Adding the genetic algorithm improves the performance of the network (in terms of blocking probability) compared to using only the greedy approach or the re-optimization at blocking method.

We also addressed the need [see B, D, E, G, H] for solving the resource scheduling problem in optical Grid networks. The need for large-scale computing, storage and network capabilities by the scientific community has resulted in the development of Grid networks. Recently, the emergence of Cloud computing is gradually shifting the trend of computing, storage and network resource usage by an organization from a local paradigm to a global paradigm. Cloud computing users are provided with platform, infrastructure or software by the Cloud service providers. Among these services, Infrastructure as a Service (IaaS) provides a hosting environment for running a user's application/software. An efficient resource scheduling mechanism for allocating computing and storage resources across the network has a significant impact in the Grid/Cloud environment. In this paper, we investigated the co-scheduling of computing, storage and network resources in a Lambda Grid/Cloud network. An Integer Linear Programming (ILP) formulation has been developed and a heuristic based on best fit first approach for node selection and minimum hop approach for route selection has been proposed. Our heuristic uses Shortest Job First (SJF) and First Come First Served (FCFS) for job selection. The objective of our work is to minimize the number of declined jobs, which is related to the revenue generated by the resource providers, and to minimize the total cost of the job execution, which is of importance to the users. The performance of our heuristic approach is found to be close to the optimum ILP solution. Our heuristic is also compared with heuristics adapted from previous work, and the results demonstrate the efficiency and feasibility of the proposed solution. In our most recent work [see G, H below], we proposed a detailed cost model to incorporate network resource costs into the problem formulation.

In other work [see C, F below], we investigated the important topic of energy efficient operation of optical networks for carrying large-scale scientific data. Current core networks are composed of high-end routers which are connected by high-speed fibers. These optical connections are commonly overprovisioned and in low utilization. Many of them are combined together to form bundle links or composite links and the component links are referred to as sublinks. These physical sublinks could be SONET connections, Ethernet circuits, wavelengths on a fiber, etc. And they could be shut down or brought up independently. Selectively shutting down sublinks during low traffic periods could save a large amount of energy while keeping the network topology unchanged. Based on this concept, we proposed a local heuristic threshold-based method to explore the potential energy savings in the backbone network by adjusting the number of active sublinks in bundle links. An experiment based on an Internet2 derived synthetic network was conducted to verify the performance of our method and the results showed that 86% of energy consumed on ports of core routers could be saved when setting 90.0% as link utilization threshold. The experiment also showed that setting 90.0% as threshold is safe enough to avoid data loss during extreme traffic increases in this case. Compared to previously proposed ILP (Integer linear programming) based global heuristic algorithms, our local heuristic algorithm can achieve energy savings close to the optimum and greatly reduce the response time and the risk of data loss.

We have had regular contact with ESnet researchers at Lawrence Berkeley Labs (our DOE-Lab partner on this project) who have kindly provided us with OSCARS schedule traces from the DOE ESnet networks. We regularly participate in DOE proposal review panels and DOE-sponsored workshops.

5.A list of papers (already published, in press, submitted) in which DOE support is acknowledged.

[A] Pragatheeswaran Angu and Byrav Ramamurthy, “Continuous and Parallel Optimization of Dynamic Bandwidth Scheduling for WDM Networks”, IEEE GLOBECOM 2010. (published)
[B] Anusha Ravula and Byrav Ramamurthy, A Tabu Search Approach for Joint Scheduling of Resources in a Lambda Grid Network, IEEE GLOBECOM 2010. (published)
[C] Lin Liu and Byrav Ramamurthy, “A Dynamic Local Method for Bandwidth Adaptation in Bundle Links to Conserve Energy in Core Networks,” Elsevier Optical Switching and Networking, vol. 10, no. 4, 2013 (published).
[D] Anusha Ravula and Byrav Ramamurthy, Grid Networks, Chapter contributed to the book “Next Generation Internet: Architectures and Protocols,” Cambridge University Press, 2011 (published).

[E] Pan Yi, Hui Ding and Byrav Ramamurthy, “Budget-Minimized Resource Allocation and Task Scheduling in Distributed Grid/Clouds (invited paper),” 22nd International Conference on Computer Communications and Networks (ICCCN), 2013 (published).
[F] Lin Liu and Byrav Ramamurthy, “Rightsizing bundle link capacities for energy savings in the core network,” IEEE Global Telecommunications Conference (GLOBECOM), 2011 (published).
[G] Hui Ding, Pan Yi and Byrav Ramamurthy, “Cost-optimized reservation and routing for scheduled traffic in optical networks,” IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol. 5, no. 11, 2013 (published).
[H] Pan Yi, Hui Ding and Byrav Ramamurthy, “Cost-optimized joint resource allocation in Grids/Clouds with multi-layer optical network architecture,” IEEE journal (in submission).

6. A list of people working on the project -graduate students, post docs, visitors, technicians, etc.
Indicate for each whether receiving full or partial support. In case of partial support indicate percentage of support.
Byrav Ramamurthy (PI, partial support, summer support)
Pragatheeswaran Angu (Graduate Student, Full RA, Graduated with MS 2011)
Lin Liu (Graduate Student, Full RA, Obtained MS degree, Currently PhD student)
Anusha Ravula (Graduate Student, No support, Graduated with MS 2010)
Mukesh Subedee (Graduate Student, No support, Graduated with MS 2010)
Anisha Kolasani (Graduate Student, Full RA, Graduated with MS 2012)
Pan Yi (Graduate Student, Full RA, Currently PhD student)

7.Planned activities for the future, which could be a short paragraph.

We are involved actively in deploying a regional optical network in the Great Plains region to transport scientific data to other campuses. We are also taking on a leadership role in deploying advanced network capabilities in the US Midwest (Great Plains region). We have received funding from the National Science Foundation (NSF) to upgrade our campus connectivity to Internet2 to 100 Gbps.

We are continuing to build on our partnerships with DOE Lab researchers (LBL) and domain scientists. We have been working closely with the DOE ANI Testbed staff (also at LBL) who have provided us access to this advanced network testbed following a competitive peer-reviewed selection process.