DOE COMMON Project 2010-11 Annual Report

DOE Award number: DE-SC0004909
Name of recipient: University of Massachusetts, Dartmouth
Project Title: Coordinated Multi-layer Multi-domain Optical Network (COMMON) for Large-Scale Science Applications
Principal investigator: Vinod Vokkarane
Date of Report: June 1, 2011
Period covered by the report: September 1, 2010 – May 31, 2011.

Introduction

We intend to implement a Coordinated Multi-layer Multi-domain Optical Network (COMMON) Framework for Large-scale Science Applications. In the COMMON project, specific problems to be addressed include 1) anycast/multicast/manycast request provisioning, 2) multi-layer, multi-domain quality of service (QoS), and 3) multi-layer, multi-domain path survivability. In what follows, we outline the progress in categories 1) and 2) (Year 1 deliverables).

The work conducted in this report was performed at the University of Massachusetts, Dartmouth in the Computer and Information Science department. Dr. Vinod Vokkarane (principal investigator) is the director of the Advanced Computer Networks Lab (ACNL), which comprises of the following research team:
Post-Doctoral Research Associate: Dr. Arush Gadkar
Visiting Fulbright Scholar: Joan Triay (Universitat Politècnica de Catalunya (UPC), Spain)
Graduate Students:
Jeremy Plante (BS/MS)
Bharath Ramaprasad (MS)
Derek Rousseau (MS)

Progress/Accomplishments

In this section we describe the progress and accomplishments in each of the tasks (labeled T1 and T2) as outlined in the COMMON project proposal:

- **T1: Anycast/Multicast/Manycast Request Provisioning**

  We introduce algorithms for provisioning anycast, multicast, and manycast calls for both, the immediate and advance (IR/AR) reservation systems [1]. Anycast can be supported directly at the optical-layer without need for special hardware or at the higher layers as a virtual circuit service. Multicast and manycast would require multicast-capable optical cross-connects (MC-OXCs) to implement at the optical-layer. Because ESnet does not currently have these devices, multicast and manycast will be implemented as overlays. We use anycast and unicast at the optical-layer to support multicast and manycast at higher layers. Our algorithms are aimed to minimize Layer 1 resources when the point-to-multipoint paradigms are implemented as network overlays. Note that a naïve method of supporting the multicast functionality in a multicast incapable environment (such as the ESnet) is by creating a virtual topology consisting of wavelength routes from the multicast source to each destination of the multicast session. For large multicast groups however, the network bandwidth consumed by such a scheme may become unacceptable because of the unicasting nature of the lightpaths. In the context of IR/AR systems, we also investigated a novel technique called lightpath switching (LPS) to reduce the blocking probability of a request and increase network utilization.

T1.1: Multicast Overlay in WDM Unicast Networks

Work Performed:

In our initial work [2], we addressed the problem of routing and wavelength assignment for multicast advance reservation in all-optical wavelength-routed WDM Networks. Here it was assumed that the optical crossconnects were capable of splitting an incoming signal on one input port to multiple output ports (i.e., they are multicast capable).

To support users' multicast requests (from the higher electronic layers) in a multicast incapable network, we proposed two overlay solutions: Drop at Member Node (DAMN) and Drop at Any Node (DAAN). In these solutions, we achieved multicasting by creating a set of lightpath routes (possibly multiple hops) in the overlay layer from the source node of a request to each destination member. In the DAMN case, we allow a lightpath route to originate/terminate only at source and destination members of a request, whereas in the DAAN model we pose no such restrictions. We considered a static traffic model, wherein the set of multicast requests is known ahead of time. We presented integer linear programs (ILPs) to solve these problems with a goal of minimizing the total number of wavelengths required to service the set. We also proposed an efficient heuristic and compared its performance to the ILP for a small network, and ran simulations over real-world, large-scale networks. We also formulated an ILP for the naïve solution of achieving multicasting and compared its performance to that of DAMN and DAAN. We also considered a dynamic traffic model, and proposed efficient heuristics to solve the DAMN and DAAN problems with a goal of minimizing the total number of wavelengths required to satisfy the request. Moreover, we presented a simple Multicast via WDM Unicast (MVWU) heuristic to implement the baseline unicast approach (naïve method).

We have performed extensive simulations on three real-world large scale networks (24-node network, NSFnet, and ESnet) for the above mentioned heuristics by also considering a dynamic traffic scenario, wherein multicast requests arrive to the network following some stochastic process.

Findings:

We first compared the DAMN/DAAN ILPs to the MVWU ILP (for a static traffic matrix), and observed that the DAMN/DAAN uses $2-5$ wavelengths fewer than MVWU at higher loads. Also, we observed that our proposed heuristic performed well as compared to the ILPs. Note that for this comparison we used a simple six-node network, as the run-times of ILPs was unacceptable for larger networks. We then compared the heuristics for the MVWU, DAMN, and DAAN on real world large-scale networks. It was observed that DAMN clearly outperformed the MVWU and achieved $20-50\%$ improvement in average wavelength usage. We also noted that the DAAN achieved an additional $4-5\%$ improvement as compared to DAMN, but at the expense of a slightly higher average number of logical hops. This work resulted in a conference paper [3] in the *IEEE International Conference on Computer Communication Networks (ICCCN) - 2011.*

For the dynamic traffic case, a similar trend in results was observed. It was observed that DAMN clearly out-performed the MVWU, and achieved a 44 to 60% improvement. The DAAN achieved a further 3 to 5% improvement over the DAMN at the cost of higher number of logical hops. This work has been submitted to the IEEE GLOBECOM-2011 conference [4], and is currently under review. A complete set of the multicast overlay results can be found in [5].

T1.2: Manycast Overlay in WDM Unicast Networks
We have also extended the above work and solved the respective problems for the many-casting communication paradigm which finds applications in real-world bandwidth intensive environments.

Work Performed:

With the increasing number of high-bandwidth applications, energy consumption of networks has become an important issue that needs to be addressed. Manycasting is a communication paradigm that finds applications in such high-bandwidth environments. To support manycasting functionality in multicast incapable optical networks, we proposed two overlay approaches: Manycasting with Drop at Member Node (MA-DAMN) and Manycasting with Drop at Any Node (MA-DAAN). In these solutions, we support manycasting by creating a set of lightpath routes (possibly multiple hops) in the overlay layer.

We considered a static traffic model and present integer linear programs (ILPs) to solve these problems with a goal of minimizing the total number of wavelengths required to service the request set. Note that the reduction in the number of wavelengths will lower the network's energy consumption. We also developed an efficient heuristic and compared its performance to the ILP for a small network, and ran simulations over real world, large-scale networks. We compared our proposed approach to a baseline unicast approach, i.e., by creating unicast (single hop) lightpaths to each destination.

Moreover, we have also performed extensive simulations on three real-world large scale networks for the above mentioned heuristics by also considering a dynamic traffic scenario, wherein multicast request arrive to the network following some stochastic process. The work on manycasting has been submitted to two conference papers [6, 7], which are currently under peer-review. Currently, we are working on a journal submission for the above mentioned work.

Findings:

Our results showed that MA-DAMN and MA-DAAN achieved a $25-45\%$ improvement (reduction in the number of wavelengths) compared to the baseline unicast approach. We further weighed our manycast approaches against comparable multicast approaches, and observed that manycasting achieves a $27 - 44\%$ improvement in terms of reduction in the number of required wavelengths.

T1.3: Dynamic Lightpath Switching in WDM Networks
Finally, we describe the work conducted in this category by using a technique called lightpath-switching (LPS).

Work Performed:

We examine provisioning holding-time-aware dynamic circuits using a technique called lightpath switching (LPS). Instead of using the same lightpath for the duration of the data transmission, in LPS we allow a request to switch lightpaths over time. Data transmission may begin on one lightpath from the source to destination, then at a later time a different lightpath from the source to the destination may be selected to continue data transmission. The lightpath switches are transparent to the user and are managed by the network. Allowing LPS creates a number of segments that can use independent lightpaths. We proposed a heuristic to solve the LPS problem and compared its performance to a simple All-Segments (AS) heuristic which does not allow LPS.

Findings:

We first observed that the blocking probability is significantly reduced when lightpath switching is allowed. When the average request duration is about 12 timeslots, the number of lightpath switches ranged from two to four, meaning each segment is on average four timeslots at low loads and over two timeslots at higher loads. These results show there is a trade-off between reduced blocking and increased network signaling (number of lightpath switches). We also ran simulations for different k (Yen's k shortest path algorithm parameter) values and for different networks. With $k = 1$, the performance improvement between AS and LPS is not as significant because here, only wavelength switching would be possible. For all other values of k, the relative improvement of LPS to AS is about the same. As the value of k increases, the number of lightpath switches also increases while the blocking probability decreases with LPS. We also note that depending on the network's average nodal degree, there is a maximum value of k for which no further improvement occurs. This work resulted in a conference paper in the *Fourth IEEE International Symposium on Advanced Networks and Telecommunication Systems (IEEE ANTS)* conference, [8]. This work [8] received the Honorable Mention in the Best Paper Category (first runner-up) and was invited for a potential publication in the *Elsevier Journal on Optical Switching and Networking*, which is currently in preparation.

Note that the concept of Lightpath Switching can be applied to switch between multiple virtual circuits, if a single virtual circuit cannot satisfy the bandwidth requirement of a single user request.

T1.4: Dynamic Anycast Lightpath Switching in WDM Networks

We also investigated the problem of provisioning dynamic holding-time-aware (HTA) lightpaths in all-optical wavelength division multiplexed (WDM) networks for the anycast communication paradigm.

Work Performed:

We proposed two heuristics to solve the anycast routing and wavelength assignment (RWA) problem: anycast lightpath switching (ALPS) and anycast continuous segment (ACS). In ALPS, we allow a request to switch lightpaths to the closest anycast destination. If we cannot accommodate the data transmission request to the selected destination with lightpath switching we try other candidate destinations, thereby increasing the probability of the user's request being provisioned. In ACS, we do not allow a connection request to switch lightpaths. We find a lightpath using traditional RWA to each anycast destination and select the one that best utilizes network resources.

Findings:

We compared the performance of ALPS to that of ACS. Our results suggest that ALPS achieves better blocking performance compared to ACS. Furthermore, we also compared the performance of these two anycast RWA algorithms to the traditional unicast RWA algorithms. We observed that the anycast RWA algorithms significantly outperform the traditional unicast RWA algorithms. This work has resulted in a conference paper in the upcoming conference *IEEE International Conference on High Performance Switching and Routing (HPSR) 2011* [9].

- **T2: Multi-Layer Multi-Domain Quality of Service (QoS)**

In this project we introduce different priority levels to ensure QoS for high priority traffic (circuit requests). Our algorithms will ensure that high priority traffic will get the required bandwidth, possibly at the expense of lower priority traffic. For example, we may incorporate preemption of lower priority traffic by higher priority traffic in order to satisfy the QoS requirements of the high priority traffic.

We will investigate the use of non-continuous transmission rates to improve the quality of service (QoS) of other low priority traffic in the network. If we find that a request cannot be scheduled, we may be able to adjust the transmission rate of certain requests in order to accommodate the other requests, as long as we do not violate deadline or other QoS constraints. This may require coordination between the OSCARS system scheduler and the user.

T2.1: Hybrid Immediate and Advance Reservation QoS Framework
In what follows, we examine the multi-layer QoS in optical WDM networks. The purpose is to deliver a QoS framework to map input connection requests to a certain number of classes, wherein each of these classes gets a different treatment in the network.

Work Performed:

We proposed a QoS framework for multi-layer networks and for hybrid immediate and advance reservation systems. The framework maps different types of connection requests to some predefined classes. By making use of diverse immediate and advance reservation scheduling algorithms, we are able to differentiate these classes within the network and provide to them, based on their requirements, the required quality level. In particular, in this task we performed an evaluation with 4 different classes, two for immediate reservation and the other two for advance reservation. The two immediate classes, which are devised for delay-sensitive applications use immediate reservation: the lowest priority class uses a simple IR holding-time-aware algorithm, while the other class uses lightpath switching. For the delay-tolerant applications, we make use of advance reservation. Again, we separate these into two classes: the lowest uses a fixed advance reservation algorithm, while the highest priority class runs a flexible advance reservation algorithm.

Further, we proposed to extend the framework by providing a partition-based QoS. The idea is to set different network resource partitions to different classes. In order to provide some extra flexibility, we have initiated some work on defining some policies over these partitions. Instead of setting a fixed partition, our aim is to make this partition flexible in the way that connection requests mapped to one class may also use resources from another class, as long as they are re-mapped at a lower priority. To this end, two different policies are expected to be evaluated: one based on switching connection requests to other class resource partitions before preempting resources within its original class, and another that first performs the preemption of resources, and then (should these resources be unavailable) try to use resources from another partition. We also proposed a mathematical model for hybrid immediate and advance reservation. The purpose is to compute, in advance, the blocking probability into the system, depending on whether the input call is for immediate or advance reservation. We tested our model on a simple link scenario.

Findings:

With the aforementioned QoS framework we performed several simulations to check its performance. We found that the proposed framework is able to achieve a very well-defined relative quality differentiation among the four classes. Due to the book-ahead of AR, delay-tolerant applications obtain a better blocking probability from the network. Also, the highest priority class, which makes use of the flexible AR yields better blocking than the other AR class. The delay-sensitive applications however, result in higher blocking than their AR counterparts. However, by using the LPS algorithm, we are also able to provide differentiation within the IR classes. This work will be presented as an invited paper in IEEE PMECT 2011 workshop, collocated with IEEE ICCCN 2011 [10].

As introduced, we proposed a simple link model for hybrid IR/AR. The results show its good fitting in comparison with the simulations for an scenario with 3 different traffic classes, one IR and two AR. We also tested the model under different traffic percentage distributions among the three classes, and in almost all of them, the results fitted quite well those obtained from simulation. The results were submitted as a conference paper to GLOBECOM 2011 [11] and it is currently under review. We plan to extend the model to a network-wide model in the Q4.

NOTE: All the above work which has resulted in journal and conference papers have explicitly acknowledged the support of the DOE COMMON project.

Cost Status & Unexpected Funds

See attached document.

Next Year Deliverables

- Develop and integrate our anycast algorithms into an OSCARS module.
- Develop and integrate our new multi-layer QoS framework into an OSCARS module.
- Investigate hybrid IR and AR request systems in a multi-domain scenario.
- Investigate survivability techniques for ESnet using existing OSCARS capability.

References

[1] N. Charbonneau, C. Guok, I. Monga, and V. M. Vokkarane, "Advance Reservation Frameworks in Hybrid IP-WDM Networks," *IEEE Commun. Mag, Special Issue on Hybrid Networking: Evolution Towards Combined IP Services and Dynamic Circuit Network Capabilities*, May 2011.

[2] N. Charbonneau and V. Vokkarane, "Static Routing and Wavelength Assignment for Multicast Advance Reservation in All-Optical Wavelength-Routed WDM Networks," *IEEE/ACM Trans. Networking*, April 2011.

[3] A. Gadkar, J. Plante, and V. Vokkarane, "Static multicast overlay in WDM unicast networks for large-scale scientific applications," in *Proc. ICCCN*, July 2011.

[4] A. Gadkar and J. Plante, "Dynamic Multicasting in WDM Optical Unicast Networks for Bandwidth-Intensive Applications," in *Proc. of IEEE GLOBECOM 2011*, Houston, TX, USA, Dec. 2011, under review.

[5] A. Gadkar, J. Plante, and V. Vokkarane, "Static multicast overlay in WDM unicast networks for large-scale scientific applications," 2011. [Online]. Available: http://www1.umassd.edu/engineering/cis/reports/report.cfm?r=39

[6] ——, "Manycasting: Energy-Efficient Multicasting in WDM Optical Unicast Networks," in *Proc. of IEEE GLOBECOM 2011*, Houston, TX, USA, Dec. 2011, under review.

[7] J. Plante, A. Gadkar, and V. Vokkarane, "Dynamic Manycasting in Optical WDM Networks for High-Bandwidth Applications," in *Proc. European Conference on Optical Communication (ECOC) 2011*, Geneve, Switzerland, Sep. 2011, under review.

[8] N. Charbonneau and V. Vokkarane, "Dynamic circuits with lightpath switching over wavelength routed networks," in *Proc. IEEE ANTS*, Mumbai, India, Dec. 2010.

[9] B. Ramaprasad, A. Gadkar, and V. Vokkarane, "Dynamic anycasting over wavelength routed networks with lightpath switching," in *Proc. IEEE Conf. on High Performance Switching and Routing*, Cartagena, Spain, July 2011.

[10] J. Triay, D. Rousseau, C. Cervelló-Pastor, and V. Vokkarane, "Dynamic Service-Aware Reservation Framework for Multi-Layer High-Speed Networks," in *Proc. 5th Workshop on Performance Moldeling and Evaluation in Computer and Telecommunication Networks (PMECT) collocated with IEEE ICCCN 2011*, Maui, HI, USA, Jul. 2011, to appear.

[11] J. Triay, C. Cervelló-Pastor, and V. M. Vokkarane, "Analytical Model for Hybrid Immediate and Advance Reservation in Optical WDM Networks," in *Proc. of IEEE GLOBECOM 2011*, Houston, TX, USA, Dec. 2011, under review.