Anycast Routing Algorithms for Effective Job Scheduling in Optical Grids

Tim Stevens, Marc De Leenheer, Chris Develder, Filip De Turck, Bart Dhoedt, Piet Demeester
Dept. of Information Technology, Ghent University – IBBT – IMEC, Gaston Crommenlaan 8, 9050 Gent, Belgium

Abstract *Effective job scheduling in optical grids requires concurrent optimisation of resource and network path selection. To enable this functionality, practical anycast routing algorithms are proposed, and simulation analysis is used to compare their performance to a pseudo-optimal bound and to several heuristics.*

Introduction

Grid computing aims to provide a uniform interface to heterogeneous and geographically distributed resources, all connected over a global network. To realize its full potential, a high performance network is required, making photonic networking the technology of choice. This interest in optical grids is, among others, being confirmed by the Global Grid Forum (GGF) [1]. Efficient scheduling of all user requests necessitates a combined optimisation of both network and computational resource consumption. For this purpose, practical multiple-constraints anycast routing algorithms and a maximum flow upper boundary are introduced. Simulation results reflect the effectiveness of the algorithms.

Anycast Routing

Currently deployed networks employ shortest-path routing for transferring data from source to destination. In a grid scenario however, a user's interest typically lies in successful job execution subject to certain predetermined requirements. Since multiple processing locations exist in the network, the exact location and network route used is of less importance to the end user. Anycast routing specifically enables users to transmit data for processing and service delivery, without assigning an explicit destination. This approach is especially useful for delivering consumer-oriented services over an optical network to a large number of users, as centralized job scheduling and grid status monitoring on client devices can be avoided [2].

Anycast routing naturally leads to a multi-constrained routing problem, with constraints emerging from both network parameters (e.g., bandwidth, delay, BER) and resource parameters (e.g., available and total capacity, processing speed). We present several algorithms to perform the combined optimisation of network and resource parameters. These algorithms can be deployed on grid networks based on both optical circuit-switching (OCS) and burst-switching (OBS) technologies. Both photonic control plane solutions (e.g., GMPLS) and network service provisioning systems (e.g., UCLP) can incorporate the proposed algorithms.

Algorithms

To enable efficient job delivery in optical grids, several routing algorithms are proposed:

- SAMCRA*, an update of the SAMCRA algorithm;
- Maximum flow pseudo-optimal bound;
- Best Server and Best Delay heuristics.

SAMCRA or Self-Adaptive Multiple Constraint Routing Algorithm is an online algorithm to determine the shortest path subject to multiple constraints [3]. Unfortunately, its traditional method of ordering subpaths (based on a non-linear length function) can cause sub-optimal results, eventually leading to routing loops [4]. A novel path ordering, which guarantees optimality, is therefore introduced and the resulting algorithm is named SAMCRA*. If **L** represents the m-dimensional path constraints vector and **P** represents an m-dimensional path weight vector, let $\rho(\mathbf{P})$ be a rescaled vector with components $\rho(P)_i = P_i / L_i$. Let $o(\rho(\mathbf{P}))$ be the vector $\rho(\mathbf{P})$ with its indices reordered so that the components are in non-increasing order. The vector ordering $\mathbf{P} < \mathbf{Q}$ holds if the first nonzero component of $o(\rho(\mathbf{P})) - o(\rho(\mathbf{Q})) < 0$. In contrast, the original SAMCRA algorithm only considers the outcome of the first component of $o(\rho(\mathbf{P})) - o(\rho(\mathbf{Q}))$. Application of SAMCRA(*) is only possible for a unicast routing problem. Anycast routing requires the introduction of a virtual topology, consisting of a virtual resource linked to all physical resources. Each client will then route towards that virtual destination. SAMCRA(*) is available as a source-based, centralized algorithm, making routing decisions for the whole network on the edge routers, or as a sub-optimal, distributed hop-by-hop version, executed on each participating network router.

Maximum flow, due to Ford and Fulkerson, is an optimal, offline technique to determine the maximum amount of flows between a given source and destination. It essentially locates paths between source and destination with free capacity (referred to as augmenting paths), and routes as many flows as possible over these paths. Similar to SAMCRA(*), supporting the anycast scenario also requires the incorporation of a virtual resource, whereby the capacity of the virtual links is proportional to the processing rate of the attached resource. In case job characteristics of individual clients (e.g., required processing capacity and average runtime) remain identical, a virtual source can be introduced in the network, together with links connecting the virtual node to the physical clients. Virtual link capacities are proportional to the job arrival rate of the attached client, and the classical, single-commodity maximum flow algorithm can be employed. However, in case job characteristics differ between clients, a virtual client cannot be introduced and a multi-commodity, maximum flow algorithm needs to be used between all clients and the single, virtual destination. The remainder of this paper only considers the single-commodity, maximum flow algorithm. Finally, the incorporation of a deadline as job constraint causes the pseudo-optimal behaviour of the maximum flow technique. Indeed, paths violating the deadline constraint are not considered as a possible augmenting flow path, and thus the true maximum flow is not attainable.

Heuristic techniques, implementing straightforward strategies for resource and path selection, are introduced for comparison purposes. First, in Best Server, the client selects the server with the highest available capacity, and uses fixed shortest path routing to reach that server. In contrast, the client selects the server that can be reached within the smallest network delay in the Best Delay approach.

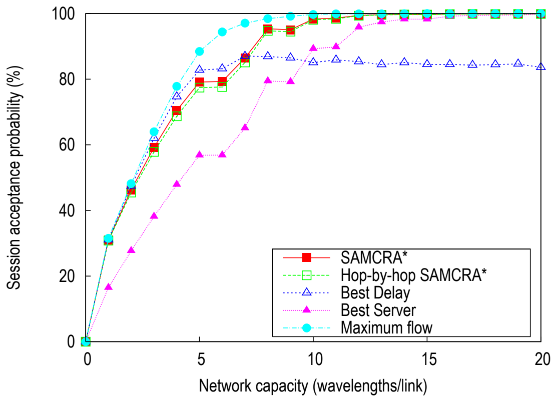

Figure 1: Session acceptance probability vs network capacity for different anycast algorithms

Simulation Results

A basic European network, composed of 28 routers and 41 bidirectional links, was considered as simulation topology [5]. The group of anycast servers consists of 4 randomly chosen resources, each attached to a different router node. All simulation results shown are averages for 1000 different resource instantiations.

Each resource is capable of handling 20 job aggregates in parallel. Jobs running on the same resource unit are assumed not to interfere with each other (e.g., they do not influence the execution time). For each job aggregate, the necessary bandwidth is reserved along the path from source to destination during the total execution time and each aggregate is assumed to consume a single wavelength on all links that belong to the end-to-end path. Additionally, intermediate routers support full lambda conversion. The end-to-end path delay constraint is set to 30ms. Furthermore, the average lightpath reservation time is assumed to be 40 minutes, and new job aggregates are generated according to a Poisson process (λ = 1.5 requests per minute).

As shown on Fig. 1, the acceptance rate of the intuitive heuristics Best Server and Best Delay is much lower than both SAMCRA* variants. When wavelengths are sparse, Best Delay can approach SAMCRA*'s acceptance probability. Unfortunately, as network capacity increases, job requests are frequently scheduled on overloaded resources. The Best Server heuristic consumes too much network resources (lambdas), and therefore converges only slowly to a maximum acceptance rate for an over-dimensioned network. The close match between the SAMCRA* scheduling results and the maximum flow pseudo-optimal bound emphasizes the effectiveness of this algorithm. Fig. 2 illustrates that SAMCRA* steers a middle course from the path delay perspective, while still satisfying the end-to-end delay requirements.

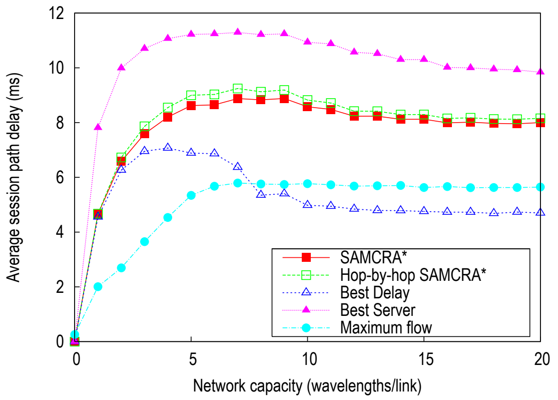

Figure 2: Average session delay vs network capacity for different anycast algorithms

Conclusions

A CPU offloading network service, based on an abstraction of the routing topology for network layer anycast and multi-constrained routing, is proposed in this paper. Several practical algorithms were presented, and simulations were performed, taking into account the path delay and bandwidth from source to destination and the servers' current loads. Results show that both the centralized and hop-by-hop variant of the updated SAMCRA algorithm approach the acceptance probability of the maximum flow upper boundary, which proves the practical optimality of the proposed approach.

References

1 D. Simeonidou (ed.), "Optical Network Infrastructure for Grid", GGF Draft, 2004.

2 M. De Leenheer et al., IEEE Commun. Mag., 44:3 (2006), pp.124-131.

3 P. Van Mieghem et al., IEEE/ACM Trans. on Networking, 12:5 (2004), pp. 851-864.

4 T. Stevens et al., Submitted to Broadnets, 2006.

5 S. De Maesschalck et al., Phot. Net. Comm., 5:3 (2003), pp. 203-225.