Optical Burst Switching for Consumer Grids

M. De Leenheer, E. Van Breusegem, J. Cheyns, P. Thysebaert, B. Volckaert, F. De Turck

Abstract—**Advances in both computer hardware and software has given end users access to increasingly complex applications. However, this evolution imposes a continuous cycle of replacement and renewal of infrastructure, which is both in terms of required resources as from an economical viewpoint a wasteful process. We aim to eliminate this inefficiency through the use of consumer grids, in which users are given access to powerful computing and storage resources, all connected through a high-performance optical burst switched network. Since these consumer grids should be accessible by the public at large, the main requirements are massive scalability, a robust design, and the delivery of non-trivial amounts of computing resources. We show how the basic functional requirements of our architecture can be realised, and demonstrate several extensions which permit the full spectrum of end user applications on our platform.**

Keywords—**grid computing, optical burst switching, anycast routing**

I. Introduction

THE term grid computing [1] suggests a paradigm similar to the operation of the electric power grid; a variety of resources contribute power into a shared pool for consumers to access on an as-needed basis. A grid is composed of different types of server resources, such as computational resources which offer processing power, and storage resources which allow data to be stored and retrieved. Likewise, network resources form the interconnections between these server resources. As shown in Fig. 1, server resources are typically geographically distributed, and consequently have different access and control policies. Also, the heterogeneous nature of the resources (e.g. differing CPU types and memory sizes) form an additional problem for the straightforward deployment of grid platforms.

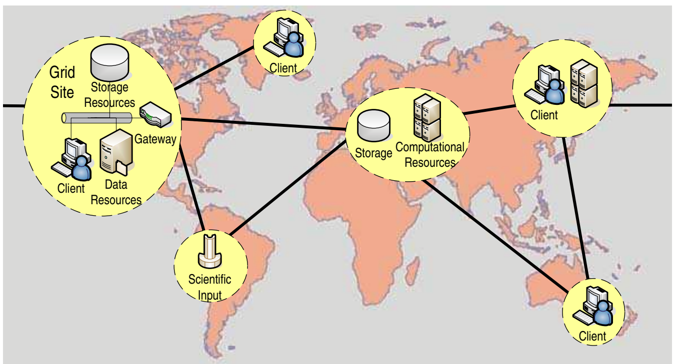

Fig. 1. General Grid concept

The high performance network characteristics, as required by grid applications, can only be fulfilled by photonic network technology [2]. Basically there are three methods for communication through optical networks. The first one is called Optical Circuit Switching (OCS) and is similar to the current telephone system. It works by setting up a dedicated circuit (lightpath), so the endpoints are assured of their connection, but at the same time bandwidth is wasted when neither endpoint is transmitting data. Also, the time to construct a circuit is not negligible, thus making this method unsuited for a large user base. The second technique is called Optical Packet Switching (OPS) and is similar in operation to the Internet. All data is encapsulated in packets of limited size and each packet is routed independently through the network. Unfortunately, the current state of optical technology is unable to create a fully functional implementation of this method. The third technique, called Optical Burst Switching (OBS) [3], combines the advantages of both preceding methods. Like OCS a circuit is constructed, but data is sent in variable sized packets called bursts, before the completion of the setup phase. Also, as soon as the data burst has passed, the circuit is torn down again, improving the efficiency of bandwidth usage. Consequently, this method is able to support a large number of concurrent users while providing each individual user with a sufficient amount of bandwidth.

II. Basic Architecture

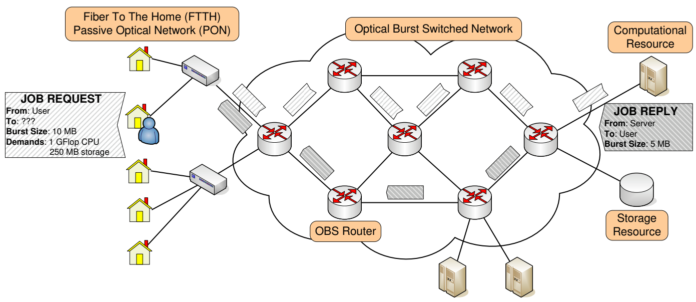

Fig. 2. OBS-based grid

Fig. 2 shows an overview of our consumer grid concept [4]. Suppose a user is running an application, and he would like to do the processing in the network, so he can obtain results faster than he would when executing the application locally. For that purpose he constructs a control packet, which contains a description about the job, and is responsible for reserving bandwidth for the burst at each intermediate router. The actual code and/or data which is to be processed, is placed inside an optical burst; clearly we perform the mapping of one (application-level) job on one (physical-level) optical burst. The user then sends the control packet and the actual burst in the network, both separated by a specific time interval. It is very important to note that the control packet does *not* contain a destination, a concept also known as *anycast* routing [5]. Once the control packet arrives at an optical router, the router will decide through which interface the control packet and the burst will leave. This decision is based on the job description present in the control packet, and the current state of the grid, which includes both network and server resources.

Because of the scalability requirement, it is unfeasible to provide every router with a detailed view of the complete grid.

All authors are with the Department of Information Technology, Ghent University (UGent), Belgium. Email: marc.deleenheer@intec.UGent.be.

Therefore, we introduce algorithms to perform *aggregation of information*; a router located close to a resource will have much more detailed state information at its disposal than one located farther away. In effect, a nearby router will know precisely where the resource is located and what its characteristics are. A distant router however, only knows the direction in which to send a job.

Once the job reaches the resource, it is submitted to the local scheduler which will handle the actual processing of the job. When execution finishes, the results need to be transmitted back to the original sender; obviously the destination is now known! Since the path followed to the resource is not necessarily the optimal return path back to the user, it is desirable to develop a routing algorithm for the transmission of results. State information on the used server resource can be inserted in the control packet going back to the original sender; this data can in turn be used to update the state information contained in the routers.

III. Routing Algorithms

Although the basic architecture satisfies the essential functional requirements, we now discuss the routing algorithms in further detail, as to allow a broader range of application scenarios.

A. Robustness

Robustness deals with providing a reliable service to users, even in case of failing network or server resources. Typically, a distinction is made between *which* resource types can fail (network or server), and *how many* resources are allowed to fail concurrently. One way to achieve the goal of robustness is to provide more capacity than strictly necessary; in effect this means we overdimension the grid network (see Section IV). For instance, if every network link has a backup link, the operation of the grid is guaranteed for at least one link failure. Another method to provide a reliable service is to send each job multiple times in the grid network. In this case we use more capacity than strictly necessary, but the probability that at least one job is processed succesfully increases significantly with the number of identical job requests.

B. Quality of Service

The fact that grid technology will be offered to the general public, incurs a shift in what computing represents. Instead of viewing it as an *asset* (e.g. owning a machine park), it should be viewed as a *service* (e.g. process this dataset). This naturally leads to a differentiation in the services offered to the end users, a concept also known as *Quality of Service* (QoS). Just like in a market-based setting, most users are satisfied with the standard service, which in our case is a *best effort* job execution. However, by paying a premium, an end user can receive an extended service. For instance, some users require guarantees that their jobs are executed, while others will simply resubmit a failed job request. Also, several applications (e.g. online collaboration) demand that results are received in real-time, while other applications have far less stringent time requirements.

C. Burst Correlation

As stated before, there exists an explicit mapping of one job on one optical burst. However, because of limitations imposed by optical technology, an optical burst is limited in size, which implies our architecture would be unsuited for applications generating large jobs. This problem can be solved by dividing a large burst into smaller sub-bursts and making sure each subburst arrives at the *same* resource. The basic anycast routing algorithm does not support this operation, since it only provides burst routing to *any* available resource. We can however make use of the first burst to indicate that a known number of bursts will follow and that they should leave through the same interface. Finally, this functionality is desirable, though not required, by applications which need the presence of preceding input and output data in each processing step. Sending consecutive bursts to the same resource ensures all data will be available at the processing site.

IV. Grid Dimensioning

Suppose we are given an OBS-based network, and are asked to build a consumer grid capable of handling a given amount of users. By measurement or predicition of the access patterns of different users, we are able to optimally place the server resources in the network. Moreover, we can assign capacities to both the server and network resources such that the grid architecture can handle the load generated by the users. Optimality is usually guided by the grid's operator main concern, i.e. cost, but could also focus on efficiency of resource usage or users' acceptance rate.

V. Conclusions

We presented a novel architecture to enable the adoption of grid technology on the consumer level, offering massive scalability and robustness in service. We showed how the basic functional requirements of our architecture can be realised, and demonstrated several extensions which permit the full spectrum of end user applications on our platform.

Acknowledgments

E. Van Breusegem and B. Volckaert would like to thank the IWT for financial support through their Ph.D. grant. P. Thysebaert and J. Cheyns are Research Assistant of F.W.O.-V., and F. De Turck is Postdoctoral Fellow of the same institution.

References

[1] I. Foster and C. Kesselman, editors. *The Grid: Blueprint for a Future Computing Infrastructure*. Morgan Kaufmann, 1999.

[2] D. Simeonidou, R. Nejabati, et al. Optical Network Infrastructure for Grid. Working Draft, Global Grid Forum (GGF), 2004.

[3] C. Qiao and M. Yoo. Optical Burst Switching - A New Paradigm for an Optical Internet. *Journal of High Speed Networks*, 8(1):36–44, Jan 2000.

[4] M. De Leenheer, E. Van Breusegem, et al., An OBS-based Grid Architecture, To appear in the *Proc. of the Workshop on High Performance Global Grid Networks (Globecom 2004)*, Dallas, Texas, USA, Nov 2004.

[5] C. Partridge, T. Mendez, W. Milliken, Host Anycasting Service, IETF RFC 1546, Nov 1993.