A View on Enabling-Consumer Oriented Grids through Optical Burst Switching

Marc De Leenheer, Pieter Thysebaert, Bruno Volckaert, Filip De Turck, Bart Dhoedt,
and Piet Demeester, University of Ghent
Dimitra Simeonidou, Reza Nejabati, Georgios Zervas, Dimitrios Klonidis,
and Mike J. O'Mahony, University of Essex

ABSTRACT

As Grid computing continues to gain popularity in the research community, it also attracts more attention from the enterprise and consumer levels. Applications in these domains generate large amounts of jobs, with individual jobs having only modest resource requirements. In this article, a novel architecture to realize a highly scalable and flexible platform for consumer-oriented Grids is proposed. The architecture is based on an optical burst switched network, complemented with an advanced control and signaling plane. The architecture, functionality, and interfaces of all the relevant entities are presented and issues, current initiatives, and future directions for the control and management of these Grid networks are discussed.

INTRODUCTION

GRID computing aims to provide a uniform interface to geographically distributed resources, all connected through a global network. Various resource types exist, including computational, storage, and network resources. Resources are heterogeneous in nature and belong to different organizations with various local management policies. The bundling of resources is required in order to process the large datasets frequently found in scientific problems. Transferring these datasets between the involved sites constrains the transport network considerably. Availability of large amounts of bandwidth coupled with low latency is required, making photonic networks the technology of choice.

As long as the datasets remain large, it is advantageous to use circuit switching where one or more complete wavelengths are used for the data transfer. However, if only small-to-medium amounts of data need to be transferred and processed (as is frequently the case in consumer applications), it is desirable to address bandwidth at a level below a full wavelength. Optical burst switching (OBS) [1] has been proposed as an enabling technology, allowing flexible and adaptable routing for unpredictable network traffic. A working group in the Global Grid Forum (GGF) is currently committed to the standardization of OBS in the context of Grid computing [2].

In this article, it is shown how OBS can enable a Grid platform to be accessible by a massive number of highly dynamic users, a concept denoted as consumer Grids [3]. General characteristics and requirements of the platform are discussed and the architecture, functionality, and interfaces of all relevant entities residing both in the core and at the edge of the network are presented. Furthermore, issues related to the control and management of these Grid networks are discussed, and current initiatives and overall directions for future investigation are presented.

FEATURES AND REQUIREMENTS

USE CASE

In what follows, some typical application requirements are presented and their impact on the underlying Grid system are discussed. Consider a multimedia editing application in which integrated audio and video manipulation programs are widely adopted and allow users to manipulate video clips, add effects, restore films, and so forth. Advances in recording, visualization, and effects technology will demand more computational and storage capacity, especially if the editing is to be performed within a reasonable time frame (allowing user feedback). More specifically, 1080 p high-definition television (HDTV) offers a resolution of up to 1920 x 1080 pixels, amounting to around 2 MPixel per frame. Suppose that applying an effect requires 10 floating-point operations (Flop) per pixel per frame, and the user would like to evaluate the effect for 10 different options; in this case, processing a 10 s clip (25 fps) requires more than 50 GFlop. This will take about 5 s to complete locally, assuming local processing power is 10 GFlop. However, if service providers offer resources having a hundred-fold capacity, execution time should only

take 50 ms. The transmission time of 10 s of compressed HDTV video (20 Mb/s bitrate or 25 MB file size) on a 10 Gb/s access link is 20 ms. A Grid user can thus create an optical burst containing the multimedia material to be processed, and hand over this burst to the network. The network is then responsible for delivering this burst to a resource with sufficient capacity. As such, important improvements in application response times can be achieved, thus making interactivity possible for applications that are otherwise too resource intensive. Also, we observe rather modest requirements of the Grid job for both the computational resource and the network resources, although a large number of such jobs will be generated at unpredictable times and locations in the network.

JOB CHARACTERISTICS

Several important properties of the job-generation process in consumer Grids can be identified:
* A large number of jobs are generated.
* Individual jobs have fairly modest resource requirements.
* Job sizes are quite small, usually on the order of a few megabytes. This means that network holding times are short in comparison to optical switching times. For instance, a 1 MB burst is transmitted in 0.2 ms over a 40 Gb/s link, while current photonic switches have setup times in the millisecond range.
* The time and location of job submissions are highly unpredictable. This can lead to frequent mismatches in available versus generated load, indicating that remote execution of jobs is a necessity.
* Several application types will be interactive in nature (e.g., immersive learning environments, haptic feedback systems, etc.), indicating a strict deadline must be met for successful completion. Other applications have different degrees of tolerance for various system parameters, which implies the existence of several quality of service (QoS) classes.

GENERAL REQUIREMENTS

The architecture as a whole and the control plane in particular should have the following properties:
* Because a large number of users and resources must be supported, *scalability* of the infrastructure is essential. This becomes even more relevant in light of the highly dynamic nature of the job submissions; spikes in the generated load should be handled gracefully by resources, which in turn have fluctuating load levels and limited availability themselves.
* It is economically unfeasible to build a dedicated network for each application. The basic infrastructure should consequently be able to support all application types, each with its own access and resource-usage patterns. The infrastructure should offer relatively *simple components* and have sufficient *flexibility* to deploy new applications quickly and efficiently. Support is also needed for

traditional OBS users requiring plain data-transfer functionality, whereby bursts contain, for instance, IP packets instead of a Grid job.
* Support for real-time applications is only possible through adequate levels of speed and flexibility in the control plane, implying that low-latency signaling is required. Also, the overhead imposed by the control plane on the involved entities should be minimized.

ASSUMPTIONS

This article contains several assumptions based on observations or technologies which are currently at an early stage:
* Actual implementations of OBS routers in hardware currently reside only in controlled laboratory environments. Notable prototype networks supporting Grid applications are currently being realized at the Research Triangle Institute (RTI) in the United States, the Photonic Networks Laboratory at Essex University in the United Kindgom, and the Beijing University of Posts and Telecommunications (BUP) in China. Nevertheless, it remains uncertain whether OBS will be adopted on a large scale to realize fine-grained access to bandwidth in optical transport networks. Finally, under certain circumstances, OBS introduces difficulties for supporting sessions between users and resources.
* As this article makes clear, the large number of functions an OBS router must support introduces additional complexity. This ultimately results in high infrastructure costs.
* Reliable network operation is crucial in telecommunications in general and in a consumer scenario specifically. The approach taken in this article lacks natively reliable services, although several measures to alleviate this problem are subsequently discussed.

ARCHITECTURE

OVERALL ARCHITECTURE

In this section, a network architecture utilizing OBS technology is proposed. The architecture (Fig. 1) is based on the novel concept of intelligent OBS routers (IOBSR) that perform on-the-fly resource discovery and routing of Grid jobs. Job requests are generated by client applications, and each request is transformed into a single optical burst. This burst is preceded by a burst control packet (BCP), which contains the job's characteristics. The user can transfer the (BCP, burst) pair to the optical network through the Grid user network interface (GUNI), while computational and storage resources can access network services through the Grid resource network interface (GRNI).

The proposed network solution has the potential to offer global access to computing and storage resources for a large number of users with different traffic profiles. One of the main advantages is that both traditional data traffic and

> The proposed network solution has the potential to offer global access to computing and storage resources for a large number of users with different traffic profiles. One of the main advantages is that both traditional data traffic and Grid traffic can be supported by one common infrastructure.
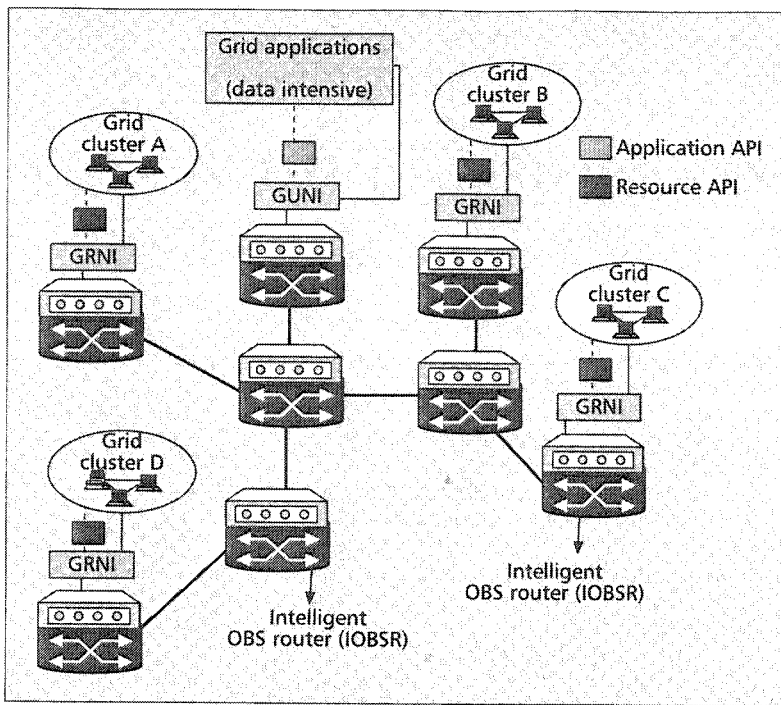
Figure 1. *OBS infrastructure for grid services.*

Grid traffic can be supported by one common infrastructure. All OBS routers perform conventional burst forwarding when normal traffic transits the network, while supporting intelligent transport of Grid traffic.

CORE NODES

Intelligent OBS Router — Central to the proposed Grid OBS network architecture is the possibility to deploy network processors (NPs) in OBS routers. The NPs can execute specific processing functions on Grid control parameters contained within the BCP at line rates (issues and solutions at 40 Gb/s have been discussed in [4]). IOBSRs are key enablers for the support of user-controlled networking functionalities:

* QoS provisioning: Grid applications are sensitive to QoS for both network and resource parameters [5]. The proposed network architecture can adopt any OBS control protocol (JIT, JET) and the IOBSR processing power can be used to deploy advanced burst scheduling algorithms. These algorithms are able to reduce delay while maintaining high bandwidth efficiency and low burst-loss rates.
* Reliable multicasting: Multicasting is important for the deployment of interactive and distributed applications. A reliable and scalable multicast protocol framework is deployed in order to minimize the traffic load and to reduce recovery latency.

Architecture — The advantage of the intelligent OBS network is the capability to provide on-the-fly resource discovery and therefore intelligent Grid traffic routing. In the IOBSR architecture (Fig. 2), the header (BCP) that is transmitted out-of-band is first converted to the electronic domain and informs the elec-

tronic interface about the type of incoming burst. When the BCP is marked as Grid data, the BCP is sent towards the network processor interface and processed for local resource discovery.

An NP performs a number of functions as summarized below:
* At the input stage, Pattern Matching extracts BCP data blocks for identification and classification.
* The Block Mapper provides an efficient means to pass the Grid control data carried by the BCP between several components of the NP. The mechanism is used to increase the speed of the burst handling on the node.
* The Data Manipulation module makes modifications to bursts, such as decrementing the Time to Live field, recalculation of checksums, and encryption of bursts.
* The Memory function stores lookup tables, packet data, and queue information.
* Queue Management schedules and stores bursts so as to provide traffic-shaping functionality and QoS priority queuing.
* The Switch Fabric Interface provides an interface to the switch fabric and manages the appropriate I/O ports for routing purposes.
* The Policy Scheduler is responsible for ensuring that bursts conform to relevant safety and security policies.

Quality of Service — The Grid job-classification mechanism facilitates on-demand resource discovery and allocation, based on Grid requirements specified by the BCP. The classification triggers the IOBSRs' intelligent mechanisms, including job scheduling, queuing, and resource discovery, for G-DiffServ provisioning. G-DiffServ provisioning will be based on job requirements such as processing power and type, bandwidth size, and time limitations. A more comprehensive description of typical QoS-related job attributes is given in the Job Submission Description Language (JSDL) [6]. G-DiffServ is vital for job processing efficiency, as well as network and computational utilization efficiency. Thus, Grid job classification will introduce QoS in the Grid infrastructure and users will send job parameters as separate fields within a BCP. These job specifications will be classified into the appropriate class of service (CoS) and in turn the CoS will drive the resource-discovery mechanism. Grid job classification must fulfill some important design requirements such as search speed, update rates, and scalability in the BCPs' fields and their sizes. Most importantly, job classification speed is important for intelligent services at the IOBSR level. Lack of classification at wire speed will result in queuing BCPs before they are processed, resulting in increased burst losses or unfair queuing [7].

EDGE NODES

Grid User Network Interface — To facilitate on-demand access to Grid services, interoperable procedures between Grid users and the optical network for agreement negotiation and Grid service activation have to be developed. These procedures constitute the GUNI, and its

functionalities and implementation will be influenced by:
- Service-invocation scenarios: the Grid user can request Grid services from the optical control plane either directly or through a Grid middleware platform.
- The optical-transport format determines the transmission format of the signaling and control messages as well as the data from the Grid user to the optical network.

Considering the above parameters in the proposed Grid-enabled OBS network, the GUNI needs to provide the following functionalities:
- Flexible bandwidth allocation: The GUNI must provide a mechanism for allocation of the required bandwidth for the Grid user or service. Flexibility indicates that various bandwidth granularities must be supported, for example, multiple wavelengths or a single wavelength, or on the subwavelength level.
- Support for claiming existing agreements: The GUNI must facilitate the incorporation of information that relates to an existing agreement. GUNI signaling would also be required to support the ownership policy of bandwidth and transport of authentication and authorization credentials.
- Automatic and timely lightpath setup: Grid users can automatically schedule, provision, and create lightpaths across the network through the GUNI.
- Traffic classification, grooming, shaping, and transmission entity construction: At the transport layer (physical layer), the GUNI must be able to map the data traffic to a transmission entity (optical burst).

Grid Resource Network Interface — Geographically distributed processing and storage resources constitute fundamental elements of the large-scale Grid network. These Grid resources can dynamically enter and leave the network, usually without a pre-established agreement. This imposes the necessity of a dedicated signaling and control interface between resources and the Grid network. Similar to the GUNI, the GRNI must perform interoperable procedures between external network elements and the optical network. But unlike the GUNI, the interface will be between the processing or storage resources and the optical network. The similarity between the GUNI and the GRNI makes it possible to extend the GUNI model to support the required functionalities for the resource network interface. The GRNI provides a signaling mechanism between resource managers and the optical network, and its main functionalities are:
- Support for existing agreements
- Job submission to local Grid resources
- Support for advance reservation schemes, which is challenging even for traditional networks
- Propagation of local resource states (available processing/storage capacity)
- Propagation of service-related events
- Flexible bandwidth allocation for returning bursts
- Returning results to the source or multiple alternative destinations
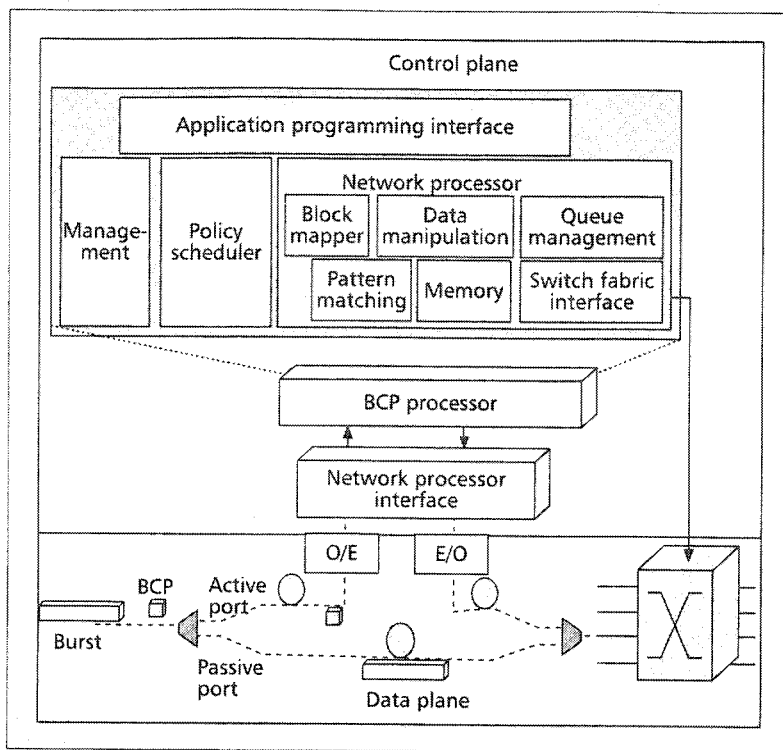
■ Figure 2. *Architecture of an intelligent OBS router.*

GRID CONTROL AND MANAGEMENT

SERVICE DISCOVERY AND RESERVATION

Traditionally, Grid service discovery has been enabled by a limited number of nodes responsible for collecting and sharing resource-state information. Resource discovery is then reduced to contacting this centralized system and performing a query for suitable resources. This approach can offer a uniform and consistent interface to state information, and can provide Grid scheduling algorithms with high-quality input parameters. However, it rapidly becomes unfeasible for an increasing number of highly dynamic users and resources. A more distributed approach for resource discovery is therefore needed. Additionally, because of the need for low-latency signaling, it is desirable to integrate service discovery with the reservation of resources. Figures 3 and 4 show two possible approaches: one-way (implicit) or two-way (explicit) discovery and reservation.

Implicit discovery and reservation (Fig. 3) uses a single (BCP, burst) pair to discover a suitable resource, reserve the relevant network resources along the way, and transfer the actual Grid task. The discovery BCP contains a description of the job's characteristics: most notably, computational and storage requirements, the data burst's size, and a time deadline. This information enables the core routers to successfully perform the job-classification mechanism and reserve bandwidth on the path from client to resource. The data burst contains the actual Grid job and is transmitted shortly after the BCP without waiting for confirmation. Integrated resource discovery and reservation offers low-latency signaling, making high job throughput
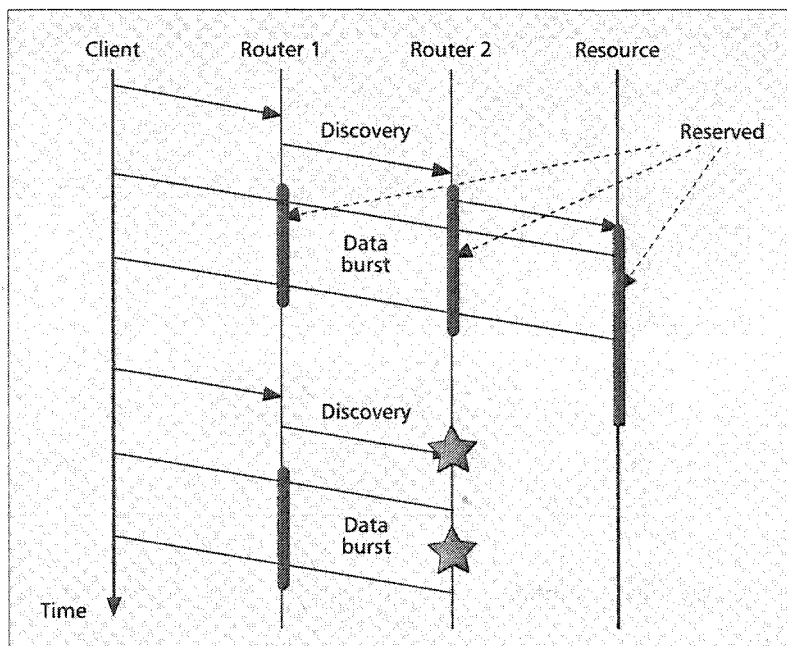
Figure 3. *One-way (implicit) service discovery and reservation offers low latency but possibly wastes bandwidth in case of failed reservations.*

possible. Unnecessary bandwidth reservations and increased job-blocking probabilities are possible, however, because data bursts are transmitted without explicit confirmation of the reservation. Also note that assigning an appropriate offset time between the BCP and the data burst is nontrivial and requires the use of application-specific hints combined with prediction techniques. Finally, instead of handing over the burst without an explicit destination, more intelligent destination assignment approaches are explored, as described in [8].

Explicit discovery and reservation (Fig. 4) is a natural extension to the implicit approach. A first BCP, containing the job's description, is sent out as discovery probe and triggers the discovery and reservation mechanisms on the visited intermediate routers. When a suitable resource is reached, an acknowledgment indicating that the transfer of the actual data burst can proceed is sent to the client. Failure to locate a resource results in the creation of a negative acknowledgment. The responses of multiple discovery probes can be collected for additional resource scheduling at the client side. Although explicit reservation can provide guaranteed reservations, implicit reservation has a much lower latency and thus reduces job turnaround times (i.e., the time between the creation of the job and the actual delivery of results). This can be of great importance to job types that must meet strict deadlines, most notably, interactive applications.

Note that when explicit reservation is used, the reservation of wavelengths can be made immediately after reception of the discovery burst, or can be delayed until the estimated time of arrival of the data burst. Delayed reservation has the advantage of improved network utilization, at the cost of increased complexity in the router's control logic. A similar time-delay

approach is possible for the computational and storage resources in the edge.

ROUTING

The routing protocol is responsible for the correct and timely delivery of a job to a resource. Its main function is forwarding bursts between core routers, based on both the job description in the BCP and the network state contained in the Memory of the router. This information, together with the Queue Management for QoS provisioning, drives the Switch Fabric Interface in the core nodes to route and forward bursts to a suitable destination (Fig. 2).

Explicit discovery and reservation indicates the client is fully aware which resource to reach. If wavelengths have been reserved during the discovery phase, the burst can travel over the constructed path. However, if no network reservations have been made in the discovery phase, traditional OBS routing protocols can be used. This usually means that shortest-path routing is employed to minimize the number of (and thus also the load on) intermediate core nodes. To reduce contention between bursts (i.e., when two different bursts need access to the same wavelength at some point in time), deflection is frequently used. The idea is to route around congested links on the path from source to destination, thus lowering the burst's blocking probability at the expense of an increase in network utilization.

When resource reservation is performed implicitly, the forwarding protocol drives the discovery and reservation mechanism by selecting the router that will receive the burst. Routing without an explicit destination at the source is frequently denoted as *anycast routing* [9]. It can provide end users with automatic resolving of services to suitable network destinations and, as such, allows services to be introduced in the network in a generic and flexible way. Service providers can deploy multiple resources for providing an identical service at various locations in the network. Clients simply contact the anycast address associated with the desired service and the routing protocol delivers the payload to a specific physical resource. The main challenge is the multitude of optimization criteria available. Indeed, traveling a minimal number of hops, reaching the most powerful resource, balancing resource utilization, and avoiding network congestion are all relevant goals for the forwarding protocol. Since jobs do not have a specific destination, deflection can be extended to increase the possibility of reaching a suitable resource (Fig. 5).

Finally, the response burst contains results that usually need to be returned to the client and differs from the request burst in two distinct ways:
* The destination of the burst is known in advance; thus, traditional routing protocols can be used.
* Because both resource capacity and time have already been consumed by a processed job, priority should be given to the response burst in case of contention with a request burst. Whenever contention occurs between two response bursts, the burst that received the largest capacity should be preferred.

SIGNALING

Signaling entails the transport of control messages between all involved network entities. Since our architecture is based on OBS, it inherits the out-of-band signaling approach in which the signaling channel is separate from the data channel. The following signaling messages are possible:

* As noted above, each data burst is preceded by a BCP. Since the loss of a BCP inevitably implies the loss of its associated data burst, the loss probability of these control messages must be minimized. Giving higher priority to BCPs in favor of other signaling messages makes this possible.
* Confirmation packets for (un)successful resource reservations, on which the client can act accordingly; the key here is to send as few notifications as possible.
* Resource-state packets carry state information to remote routers, where they are stored in the Memory block and used for resource discovery and routing. Resource states include free and total capacity of computational and storage resources, load levels of core routers, blocking probabilities of network links, and physical characteristics of the network links (e.g., BER, dispersion rate).

An important goal is to minimize the amount of traffic in the signaling layer, given the highly dynamic nature of resources. This can be achieved by minimizing the number of resource state advertisements, according to the following approaches:

* Realizing that, in most cases, the results need to be returned to the original client. As such, state information can be appended to the BCP, thus reducing the need for explicit resource state updates.
* Aggregation of resource states can significantly reduce the amount of status messages, although it inherently renders the contained information of lower quality.

RELIABILITY

Reliability in Grid computing is concerned with ensuring that failures of computational, storage, or network resources have no effect on the correct execution of jobs. The study presented in [10] investigates techniques to offer a reliable OBS network through the use of both loss-minimization and loss-recovery mechanisms.

The former involves techniques that can be deployed at the edge of the network or in core nodes. Notable edge-related techniques are intelligent-destination assignment for load balancing, and admission control or grooming policies to limit the amount of traffic in the network. To minimize Grid job losses, efficient anycast algorithms must be developed for implicit resource discovery. Finally, the Policy Scheduler and Queue Management blocks in the core routers allow prioritized handling of jobs with higher reliability requirements.

Loss-recovery mechanisms can be either proactive or reactive. Burst cloning at the client side can dramatically improve reliability, at the cost of increased resource utilization. This can
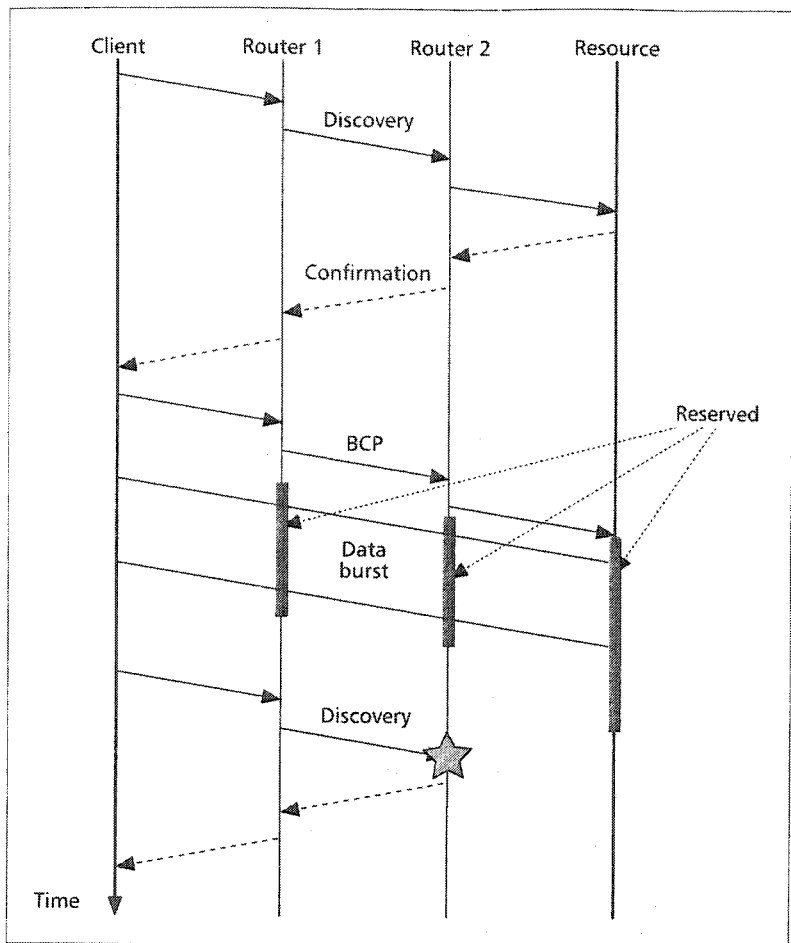
■Figure 4. *Two-way (explicit) service discovery and reservation provides guaranteed reservation but increases signaling latency.*

be partially alleviated by the use of the reliable multicast function offered by the IOBSR, whereby the cloning of bursts is delayed at a later stage in the resource discovery. If this resource discovery is performed explicitly, the traditional protection and restoration techniques of redundant resources can also be utilized. A primary path and resource are initially used for the transfer and execution of a job. If a failure occurs, the secondary resources take over; these backup resources can be assigned exclusively to a job, or shared by several job requests. Nevertheless, all techniques introduce additional complexity in the signaling layer because synchronization of identical jobs is necessary. Finally, job retransmission is a reactive approach that can be used as a final measure in case failures have occurred.

Note that support for reliable job completion can be complemented by application-level support. Currently available Grid middleware solutions can assist reliability strategies with monitoring, checkpointing, and automatic migration of failed jobs.

CONCLUSION

A novel architecture to realize a highly scalable and flexible platform for consumer Grid networks has been presented. Arguments why opti-

■ Figure 5. *Deflection strategies compared. Bursts have no explicit destination and check if local resources are available at each intermediate router. Bursts are deflected if local capacity is insufficient. None means the burst is dropped, while random uses a random output interface. Network selects the interface with the lowest blocking probability, and grid forwards the burst to the neighbor with the highest available resource capacity.*

cal burst switching (OBS) is an enabling technology for the architecture were given and the properties and requirements of the Grid architecture in general were discussed. The architecture's elements were described in detail and the functionality and interfaces of the nodes both in the core and at the edge of the Grid network were analyzed. Finally, issues concerning the control and management layer were presented, together with current initiatives and overall directions for future investigation.

REFERENCES

[1] C. Qiao and M. Yoo, "Optical Burst Switching: A New Paradigm for an Optical Internet," *J. High Speed Networks*, vol. 8, no. 1, Mar. 1999. pp. 69–84.
[2] D. Simeonidou and R. Nejabati (Eds.), "Grid Optical Burst Switched Networks (GOBS)," Global Grid Forum Draft, May 2005.
[3] M. De Leenheer et al., "An OBS-based Grid Architecture," *Proc. IEEE Globecom Wksp. High-Performance Global Grid Networks*, Dallas, TX, Nov. 2004.
[4] D. Simeonidou et al., "Dynamic Optical Network Architectures and Technologies for Existing and Emerging Grid Services," *IEEE J. Lightwave Tech.*, vol. 23, no. 10, Oct. 2005, pp. 3347–57.
[5] I. Foster et al., "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *Int'l. J. High Perf. Computing Applications*, vol. 15, no. 3, 2001, pp. 200–22.
[6] A. Anjomshoaa et al., "Job Submission Description Language (JSDL) Specification," Grid Forum Draft, Ver. 1.0, June 2005.
[7] V. P. Kumar et al., "Beyond Best Effort: Router Architectures for the Differentiated Services of Tomorrow's Internet," *IEEE Commun. Mag.*, vol. 36, May 1998, pp. 152–64.
[8] F. Farahmand et al., "A Multi-Layered Approach to Optical Burst- Switched Based Grids," *Proc. 5th IEEE/CreateNet Int'l. Wksp. Optical Burst/Packet Switching (WOBS)*, Boston, MA, Oct. 2005.
[9] C. Partridge et al., "Host Anycasting Service," Internet Engineering Task Force, RFC 1546, Nov. 1993.
[10] V. Vokkarane and Q. Zhang, "Reliable Optical Burst Switching for Next-Generation Grid Networks," *Proc. IEEE/CreateNet GridNets 2005*, Boston, MA, Oct. 2005, pp. 505–14.

BIOGRAPHIES

MARC DE LEENHEER (marc.deleenheer@intec.ugent.be) received his M.Sc. degree in computer science engineering from Ghent University, Belgium, in June 2003. He is now a research assistant and Ph.D. student affiliated with the Department of Information Technology at Ghent University, and has received a scholarship from the Institute for Innovation in Science and Technology (IWT), Flanders, Belgium. His main interests include modeling and optimization of grid management architectures, specifically in the context of photonic networks.

PIETER THYSEBAERT received his M.Sc. degree in computer science engineering from Ghent University in June 2001. He is now a research assistant and Ph.D. student affiliated with the Department of Information Technology at Ghent University, and has received a scholarship from the Fund for Scientific Research (F.W.O.-V.), Flanders, Belgium. His main interests include grid simulation and modeling of grid scheduling problems.

BRUNO VOLCKAERT received his M.Sc. degree in computer science from Ghent University in June 2001. He is now a research assistant and Ph.D. student affiliated with the Department of Information Technology at Ghent University, and has received a scholarship from the IWT. His main interests include grid management architectural designs and grid simulation.

FILIP DE TURCK received his M.Sc. degree in electrical engineering from Ghent University in June 1997. In May 2002 he obtained his Ph.D. degree in electrical engineering from the same university. From October 1997 to September 2001 he was a research assistant with the F.W.O.-V. At the moment, he is a part-time professor and post-doctoral fellow of the F.W.O.-V., affiliated with the Department of Information Technology at Ghent University. He is author or co-author of approximately 80 papers published in international journals or the proceedings of international conferences. His main research interests include scalable software architectures for telecommunication networks, service management, performance evaluation and optimization of routing, admission control, and traffic management in telecommunication systems.

BART DHOEDT received a degree in engineering from Ghent University in 1990. In September 1990 he joined the Department of Information Technology at the same university. His research, addressing the use of micro-optics to realize parallel free space optical interconnects, resulted in a Ph.D. degree in 1995. After two years of post-doctoral research in opto-electronics, he became a professor with the Faculty of Engineering, Department of Information Technology. Since then, he is responsible for several courses on algorithms, programming, and software development. His research interests are software engineering and mobile and wireless communications. He is author or co-author of approximately 100 papers published in international journals or the proceedings of international conferences. His current research addresses software technologies for communication networks, peer-to-peer networks, mobile networks, and active networks.

PIET DEMEESTER received his M.Sc. degree in electrical engineering and his Ph.D. degree from Ghent University in 1984 and 1988, respectively. In 1992 he started a new research activity on broadband communication networks resulting in the INTEC Broadband Communications Network (IBCN) research group. In 1993 he became a professor at Ghent University, where he is responsible for research and education on communication networks. His research activities cover various (IP, ATM, SDH, WDM, access, active, mobile) communication networks and include network planning, network and service management, telecom software, internetworking, network protocols for QoS support, and so forth. He is author of more than 400 publications in the area of network design, optimization, and management. He is a member of the editorial board of several international journals and has been a member of several technical program committees (ECOC, OFC, DRCN, ICCCN, and IZS).

DIMITRA SIMEONIDOU is currently a professor at the University of Essex, United Kingdom. She has more than 15 years of experience in the field of optical transmission and optical networks. She received B.Sc. and M.Sc. degrees from the
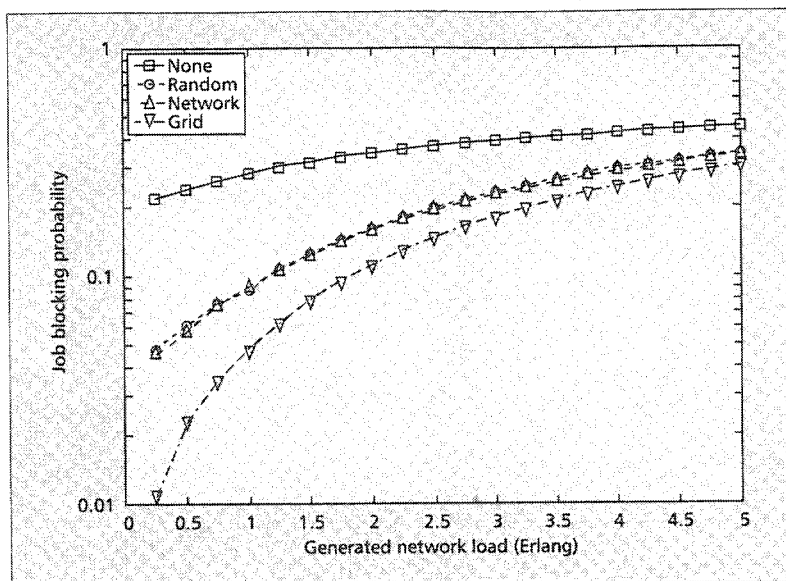
Physics Department of Aristotle University of Thessalonica, Greece, in 1987 and 1989, respectively, and a Ph.D. degree from the University of Essex in 1994. She joined Alcatel Submarine Networks in 1994 as principal engineer and contributed to the introduction of WDM technologies in submerged photonic networks. She is the author of more than 100 papers and holds 12 patents related to photonic technologies and networks. Her main research interests include optical wavelength and packet-switched networks, network control and management, and grid networking.

REZA NEJABATI is a senior research officer in the Photonic Network Research Group at the University of Essex. He received his B.Sc. degree in electrical engineering from Shahid Beheshti (Melli) University, Tehran, Iran, in 1997, and his M.Sc. degree with distinction in telecommunication and information systems in 2001 from the University of Essex. His main research interests are in the areas of optical subwavelength switching and grid networking.

GEORGIOS ZERVAS was awarded an M.Sc. degree in electronic and telecommunication systems engineering (with distinction) from the University of Essex in 2003. He is currently working toward a Ph.D. in optical networks and end systems for grid services. His research interests include high-speed digital system design, optical burst switched networks, and grid networks.

DIMITRIOS KLONIDIS was awarded a degree in electrical and computer engineering from Aristotle University of Thessa-

lonica in 1998 and an M.Sc. degree in telecommunication and information systems from the University of Essex in 2001. He has worked for the EPSRC project OPSnet since August 2004 and is currently involved with a number of research activities within the University of Essex, while working toward a Ph.D. in high-speed optical packet switching. His main research topic is the area of ultra-fast optical communication networks.

MIKE J. O'MAHONY [SM] graduated from the University of Essex in 1974 and received a Ph.D. degree for research in digital transmission systems in 1977. In 1979 he joined the Optical System Research Division of British Telecom, where he worked on research into fiber optic systems for undersea systems, in particular, experimental and theoretical studies of receiver and transmitter design. In 1984 he became group leader responsible for the study and application of optical amplifiers to transmission systems. In 1988 he became head of section, responsible for 50 graduates researching terrestrial long-haul optical systems and networks. In 1991 he joined the Department of Electronic Systems Engineering at the University of Essex as a professor of communication networks. He was head of department from 1996 until 1999. His current research is related to the study of future network infrastructures and technologies. He is principal investigator for grants supported by industry, national research councils, and the EU. He is the author of more than 200 papers related to optical communications and is a member of the IEE.

Support for reliable job completion can be complemented by application-level support. Currently available Grid middleware solutions can assist reliability strategies with monitoring, checkpointing, and automatic migration of failed jobs.

An Optical Control Plane
for the Grid Community

Radio Communications

Topics in Network and
Service Management