Techniques for Designing Survivable Optical Grid Networks

Ying Chen, Ataul Bari, Arunita Jaekel
University of Windsor, ON N9B 3P4, Canada
Email: {chen13r, bari1, arunita}@uwindsor.ca

***Abstract*—Grid computing involves high performance computing with resource sharing to support data-intensive applications, and requires high speed communications. Wavelength division multiplexing (WDM) optical networks become a natural choice for interconnecting the distributed computational and/or storage resources due to their high throughput, high reliability and low cost. This has led to increased research attention on techniques for developing fault tolerant optical grids. Different solutions have been proposed in the literature to provide protection and fault-tolerance in both grid computing and optical networks. However, the design of a resilient optical grid network should consider the inter-relation between computing and networking resource usage and how they are affected by potential faults. Hence, it is necessary to develop integrated, sophisticated algorithms that jointly allocate both computing and networking resources to improve the resource utilization while guaranteeing service availability. In this paper, we investigate the survivability of optical grids and review the state-of-the-art techniques and approaches in this area. We also identify some open problems and point out a number of promising approaches for future research direction to achieve fault tolerance in optical grid networks.**

***Index Terms*—survivable optical grids, fault-tolerant optical grids**

I. Introduction

Grid computing is motivated due to the growing demand for the computational and the storage resources, and to the fact that a large amount of such resources, geographically distributed around the globe, remain underused [1]. The grid paradigm offers computation and storage using these resources, such that the exact geographic location of the physical resource remains transparent to the user. As opposed to cluster computing, grid computing does not employ a central administrative entity and the topology changes frequently due to the dynamic nature of the resources [1]. Grid computing can potentially be used in a number of areas [2], such as:

- Large scale, data intensive applications: Grid can facilitate secure access to massive global data, allow high-speed data movement and replication within geographically distributed locations.
- eHealth remote applications: Grid can facilitate on-demand, secure and high-speed data transfer of large images used for screening.
- Bio-informatics applications: Grid can facilitate computationally highly-intensive phytogenetic analysis, comparative analysis, synthesis and pattern matching application widely used in the field of bio-informatics.

As mentioned above, grid computing targets large scale, data-intensive applications using geographically distributed resources. While integrating these geographically distributed resources into a grid, bandwidth and delay-guarantee become major issues, especially for real-time applications [3]. Optical networks have been effectively used to meet the capacity and the global connectivity requirements of the wide-area backbone networks, due to their ability to handle large amounts of data, at high speeds and with low latency. Therefore, optical networks become a natural choice for the infrastructure to enable grid computing as well. The optical network architecture used for grid computing is generally known as *optical* grid, or *photonic* grid [3], and a typical example of such a network is shown in Fig. 1.

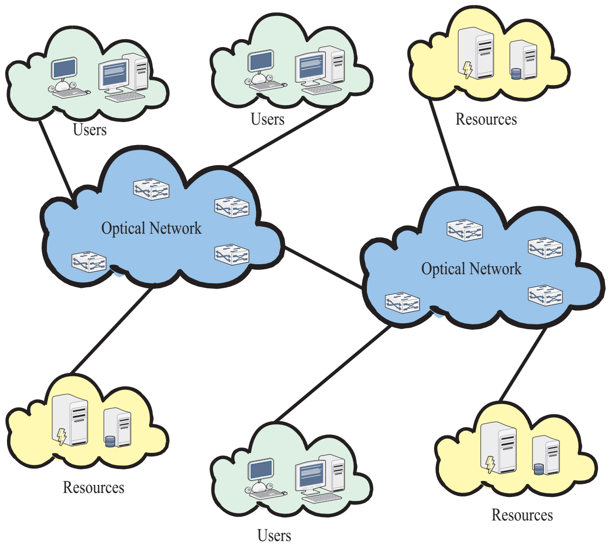

Fig. 1. An example of optical grid network.

Optical grid networks adopt a distributed model, where an user can set up peer-to-peer network that includes deciding routing and establishing/tearing-down lightpaths [4]. The optical switching techniques used in data transmission paradigm of a optical grid can greatly influence its architecture. A number switching schemes has been considered for the optical grid (e.g., Wavelength switching, Hybrid router-wavelength switching, Optical burst switching, Optical flow switching, and Optical packet

Manuscript received February 28, 2012; accepted March 25, 2012.

doi:10.4304/jcm.7.5.391-399

switching [4]). Among these, Optical burst switching (OBS) and Optical packet switching (OPS) have drawn significant attention due to their ability to meet the demand for fast and dynamic network connections [5].

Optical grid networks require a scalable and reconfigurable architecture [4], along with cheap, variable bandwidth and support for high capacity data transfer. An implementation for such networks requires tools to provision ene-to-end and on-demand network services that may span multiple administrative and network technology domains [5]. For a grid to deliver meaningful services, the following topics need to be addressed appropriately [5]:

- scheduling, i.e., spatial and temporal assignment of the tasks to the required resources. The scheduling scheme in optical grid networks attempt to minimize the *scheduling span* [3], i.e., the time the last task finishes.
- design/redesign of a flexible optical layer architecture.
- development of design techniques, e.g. for dimensioning.
- development of reliable and survivable networks, including routing and control schemes that offer both QoS and resilience.

In this paper, we focus primarily on the last topic, i.e. the design of reliable and survivable optical grid networks. An earlier version of this article appears in [6]. This is an emerging area that has started receiving research attention in the last few years. We first classify and discuss the existing techniques and then identify a number of promising approaches and open problems for survivable optical grid network design.

II. Survivability in Optical Grid Networks

Fault tolerant techniques for both grid computing and optical networks have been widely investigated in the literature. For example, job checkpointing and replication are two widely used recovery methods in general grid systems [7]. Checkpointing is used to store the state of a process running on a computational resource periodically. If the current system fails, the process can be rolled back and restarted from the saving point on the other resources. On the other hand, replication provides one or more identical copies on different resources. It guarantees that at least one replica would survive from the failure. A hybrid scheduling algorithm combining the aforementioned techniques is proposed in [8] for designing a fault-tolerant grids. Similarly, path protection schemes can be used at the design phase to ensure survivability in optical networks [9]. In path protection, two link-disjoint lightpaths, primary path and backup path are reserved for each connection request. If a link on a primary path fails, the traffic will be re-directed to the pre-assigned backup path.

It is possible to make use of the demonstrated advantages of existing techniques, such as checkpointing and replication or path protection, by extending and adapting them for fault management in optical grids. However, applying such techniques without considering the inter-relation between computing and networking resource usage and how they are affected by potential faults will lead to inefficient utilization of available resources. An effective and comprehensive strategy for handling faults should use an integrated approach that jointly considers faults in both computing clusters and network links. It has been shown in [10] that such a combined approach leads to improved performance in terms of network survivability.

The data transmission paradigm, i.e. whether optical burst switching (OBS) or optical circuit switching (OCS) is used, can have a significant effect on how fault tolerance is implemented in the network. OCS typically incurs higher signalling overhead, particularly for small jobs, but can provide a reliable service. On the other hand, OBS may be ideally suited for certain grid applications, due to its lower latency and better resource utilization, particularly for smaller jobs [5]. The focus for achieving fault tolerance in optical grids can be diverse, ranging from dimensioning and provisioning the network with adequate resources, to designing a survivable logical topology, improving resource availability, or minimizing loss. In this section, we outline the various approaches that have been proposed in recent years for designing survivable optical grid networks, with various objectives, for both OBS and OCS networks.

Based on the selected data transmission paradigm, survivability techniques in optical grid networks can be classified into two main streams, OBS and OCS, as shown in Fig. 2. In addition to the two main streams, hybrid schemes, using OBS and OCS in different layers, have also been proposed.

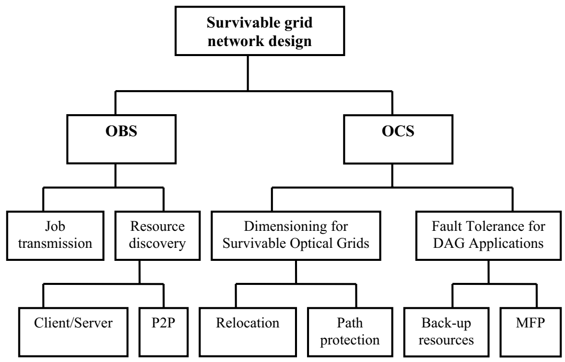

Fig. 2. A taxonomy of survivability in optical grid networks.

A. Fault Tolerance in OBS Based Optical Grids

A number of paradigms to support OBS-based grid networks is presented in [11]. In [12] the authors show that due to the inherently bursty nature of most inter-task communication, OBS tends to result in lower job completion times, compared to OCS for such short-lived dynamic communication.

In [2], several schemes for loss minimization and recovery, for reliable OBS over grid networks are discussed.

In particular, two mechanisms, based on forward error correction (FEC) and composite burst assembly (CBA) are shown to reduce loss and increase throughput.

In addition to ensuring survivability of job transmission, in case of link failures, it is also important to consider survivability for the resource discovery mechanism. Most resource discovery schemes employ the client/server model [13], [14], where grid resources publish their information to a central server and the server is then responsible for assigning resources to a job. However, such a scheme is vulnerable to "single point of failure" problem, which would cause the entire grid to lose functionality in case the central server fails. In order to address this, a Peer-to-Peer (P2P) based resource discovery scheme can be used, as proposed in [15]. The idea is to introduce a P2P protocol layer that eliminates the need for a centralized resource information server. Instead, each user is provided sufficient intelligence to manage resource discovery requests, based on its own information about the network.

1) Hybrid Schemes: Rather than focusing solely on the reliability in the OBS layer, a hybrid OBS/GMPLS architecture is proposed in [16]. Contention resolution schemes such as burst cloning [17] and burst retransmission [18] are used to implement survivability in the OBS layer; standard protection, or restoration schemes are used in the GMPLS layer. Depending on the QoS requirements of the grid job, different levels of service, e.g. 1+1 protection, preplanned restoration or full rerouting can be implemented for the job.

B. Dimensioning for Survivable Optical Grids

The goal of dimensioning a network is to determine the minimum amount of network resources required for a given set of traffic demands. In traditional optical networks, the source and the destination of a connection are known in advance. So, a traffic matrix can be formed and given as an input to the dimensioning algorithm. In contrast, in optical grids, the *anycast* [19] principle is often applied, where a job generated at the source site i can be processed at any site j that has sufficient computational resources, of the appropriate type. This means that the destination for a job to be executed is not predetermined, but can be freely chosen by the scheduling algorithms. This adds another degree of freedom and complexity on the network dimensioning problem, since the traffic matrix for the network is not known in advance.

The complete dimensioning problem in optical grids has the following subproblems that have to be solved strategically for optimizing the network cost or the resource utilization: For a given set of demands, i) Network Design: figure out how much transport capacity is needed for data transfer, ii) Server Site Design: decide the minimum number of grid server sites and their locations that needs to be deployed in order to support computational tasks, iii) Job Scheduling: distribute jobs among the selected server sites, and iv) Route and Wavelength Assignment (RWA): find a suitable route over the physical network and assign a specific channel on each link in the route, for each connection request.

In [20], the authors analyze the dimensioning cost on an optical-switched transport network for grid applications. They take into account the possible resource failures, including computational resource failures, optical cross-connect failures and network link failure, and use *Divisible Load Theory* to solve the combined Grid network dimensioning and workload scheduling problem. In [21], [22], the authors integrate the anycast principle with the classical shared path protection scheme to design a fault-tolerant optical grid network. The authors first derive a static source-destination based traffic matrix using an iterative dimensioning approach adopted from [19]. They show that the grid site location problem is equivalent to a K-means clustering problem where cluster centers represent the server sites. Then, they propose a relocation strategy to minimize the network cost in terms of the total number of wavelengths used in the topology. Unlike the existing path protection schemes, a job is allowed to be relocated to another available resource (possibly, close to the original resource) in case of a network failure, instead of reserving a backup path between the source and the original destination. This leads to reduced load on network resources, at the cost of extra load for the server which receives the relocated jobs [21]. ILP formulations for both classical shared path protection and relocation schemes are presented for small-scale case studies (up to 20 connections).

In [23], the authors present another ILP formulation to address the dimensioning problem, using relocation with path protection. This approach differs from [21] in that the primary destination site for each job is also determined by the ILP, and not given as input to the ILP. A heuristic algorithm is presented to handle large network instances. Initially, two link-disjoint shortest paths are computed as primary and backup paths for each job. In the optimization phase, both primary path and backup paths are possibly rerouted in order to reduce the number of wavelengths for primary path and maximize the backup wavelength sharing. It has been shown that the shared path protection with relocation approaches consistently outperform the traditional shared path protection scheme in terms of the network utilization in optical grids [21] - [25].

To address the scalability problem, Jaumard *et al* [24], [25] solve the same dimensioning problem in wavelength convertible optical grids, using column generation approach for the relocation protection scheme. By decomposing the original ILP in [21] into two sub-problems, a Restricted Master Problem and a Pricing Problem, the column generation method reduces the complexity by limiting the number of variables, and hence is able to obtain near optimal solutions for large networks with hundreds of connections within a reasonable time. Heuristic approaches are also proposed in [25] to find faster solutions with an acceptable optimality gap. In [26], the authors further extend their earlier work in [24] with an integrated ILP-based solution to optimize the offline dimensioning problem while providing resilience against both single link failure and node failure.

Generally, two different objective functions have been considered in the area of resource provisioning for classical WDM optical networks, either minimizing the total network cost to accommodate a given set of traffic requests or maximizing the network throughput under certain resource limitations. In optical grid networks, the existing approaches for network dimensioning typically focus on minimizing both network resources and server capacity. In [27], the authors expand on the relocation idea to deal with the anycast routing and wavelength assignment problem with the objective of maximizing network throughput, measured by the ratio of successfully processed jobs.

The aforementioned works all assume that full wavelength conversion is able to be performed at each network node. However, full range all-optical wavelength conversion is generally not feasible due to both cost and technological restrictions. The routing constraints on lightpaths, such as wavelength continuity constraint, need to be considered in the network design phase.

C. Fault Tolerance for Directed Acyclic Graph (DAG) Applications

A job J to be scheduled over a distributed system can be viewed in terms of a *task graph* [28], which is modeled by a directed acyclic graph (DAG) $J = (V, E, d)$. Each node $v_i \in V$ of the DAG represents a *computation task* to be scheduled on an available grid computing resource, each edge $e_i \in E$ represents a *communication task* that can be scheduled to a lightpath for data transfer, and d is the job deadline [29].

All computation and communication tasks in the DAG must satisfy the following constraints:

- A computation task can be processed only after all preceding tasks and required data transfers have been completed and
- A communication task can start only after the preceding computation task is completely finished.

An example of a DAG application (adapted from [29]) is shown in Fig. 3, where each node is assigned an average execution cost based on its processing needs and each edge is assigned a communication cost based on the amount of data to be transferred. For instance, it takes 2 units of time to complete the first computation task N0, and an additional 2 units of time to transfer its data to N3 (after N0 has finished execution). The minimum amount of time needed to finish the entire job (assuming resources are available as soon as they are needed) would be 12 units.

The problem of task scheduling and lightpath establishment (TSLE) for DAG applications in optical grids has been addressed in [28], without considering fault tolerance. In [10], the authors first propose an *overlay policy* that implements survivability for DAGs by directly applying existing techniques, i.e., replication for computing nodes and path protection [9] for link failures. However, such an approach introduces excessive and unnecessary resource redundancy. So, the authors propose a second *joint policy* that considers backups for tasks and communications at the same time. The joint policy clearly achieves better results, both in terms of reduced job completion time and lower resource usage.

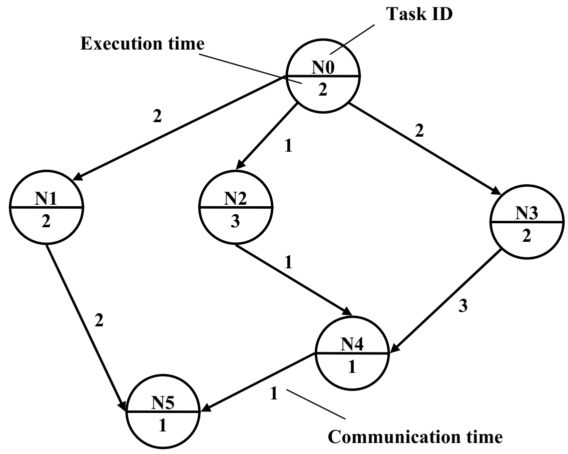

Fig. 3. An example of Directed Acyclic Graph.

A distributed DAG application can be envisioned as implementing a virtual infrastructure, consisting of nodes (e.g. processing or storage nodes) connected by dedicated circuits (e.g. lightpaths). In this scenario, task scheduling in optical grid can be viewed as analogous to a logical topology design problem. In [30], the authors address the problem of establishing lightpaths between clusters, so that a job can tolerate the failure of any single cluster and/or physical link. This is formulated as a survivable logical topology design problem using Mixed Integer Programming (MIP), and a heuristic algorithm is also provided.

In [12] only link failures are considered, and traditional protection and restoration schemes in optical networks are modified for optical grids, with the goal of minimizing the job completion time. Instead of protecting each lightpath independently, multiple lightpaths for a single job are protected, according to an optimal schedule. Path protection tends to increase the job schedule length, so backup paths are selected to minimize the job completion time. In case of restoration, in addition to rerouting, it may be necessary to reschedule some tasks, if communication is disrupted for a sufficiently long time.

Faults in optical grids can seriously delay the job completion time, even if recovery is ultimately achieved. So, rather than simply providing backup resources for handling faults, an alternative approach is to try to reduce the occurrence of faults for a job, in the first place. This approach is used in [31], where goal is to assign tasks, based on the fault rates for each processor, in a way that assures minimum fault probability (MFP) for the entire job. The proposed MFP algorithm not only achieves lower fault probability, but also reduces the average job completion time compared to an algorithm that only tries to minimize the finish time, without considering fault rates. The availability-driven scheduling (ADS) algorithm presented in [29] also tries to assign a DAG job with the goal of maximizing link availability, while meeting the time constraints. The availability a_j of a link j is calculated as:

$a_j = \frac{MTTF}{MTTF+MTTR}$, where MTTF and MTTR are the mean-time-to-failure and mean-time-to-repair respectively, for link j.

III. Challenges and Open Problems

The development of new strategies for survivable optical grid network design, gives rise to a number of different options and possible alternatives for handling faults. One of the most important factors is to ensure efficient utilization of available resources. This in turn reduces the cost for provisioning the network (including processing nodes, switching nodes, links etc.) with sufficient capacity to meet user requirements, even in the presence of faults. In this section, we suggest two specific approaches that can be used to enhance survivability in future optical grids. The focus in both cases is to ensure efficient implementation of the survivable optical grid networks. In addition to the strategies mentioned here, there is also a great potential for improvement in other areas, such as fault localization, energy-efficient resource utilization and protection against security breaches in optical grids. The two approaches we consider offer advantages and potential benefits for efficient implementation of survivable optical grids. These are:

- The use of restoration techniques (as opposed to path protection) for addressing link failures, and
- The effective use of advance reservation (AR) to guarantee resource availability, in case of faults.

A simple taxonomy, showing our proposed approaches, is given in Fig. 4, with detailed explanations provided next.

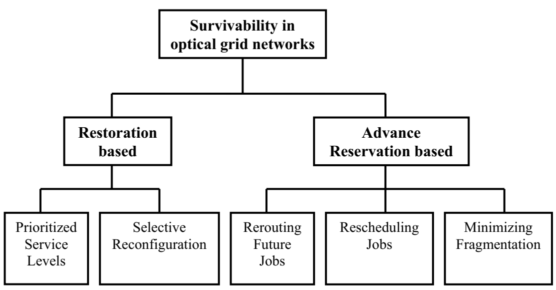

Fig. 4. Some approaches for improving survivability in optical grid networks.

A. Restoration Techniques for Optical Grids

As reviewed in the previous section, much of the work on fault management in OCS-based optical grids use preassigned relocation for node failures in conjunction with path protection for link failures. Protection schemes can guarantee survivability, but are inefficient in terms of resource utilization. The use of restoration based techniques, on the other hand, typically results in more efficient implementations. However, if restoration is used, it is quite possible that due to resource constraints some disrupted lightpaths cannot be rerouted and are blocked. Since it is unlikely that all affected lightpaths have the same priority and require the same level of service guarantees, it is desirable to have some control over which lightpaths survive and which ones are dropped. It is important to explore and develop different strategies and algorithms for achieving this, and some possible approaches are discussed here.

- Prioritized Service Levels: One straightforward option is to assign a priority level to each communication task. The priority can be pre-assigned, or calculated based on the bandwidth requirements, distance from source to destination, importance of the communication and a combination of other parameters. For example, a communication that must be guaranteed would automatically be assigned the highest priority level, regardless of other factors. Once a fault occurs, the disrupted connections are sorted based on their priority levels and processed. Intelligent algorithms to properly assign priorities and perform routing can greatly reduce the impact of faults.

  We consider a small optical network with five nodes as given in Fig. 5 where solid lines represent fiber links, directed lines stand for lightpaths established over the physical fiber links, and dashed directed lines indicate rerouted lightpaths. For simplicity, we assume that each fiber link can at most accommodate two WDM channels. Fig. 5 (a) shows an initial configuration plan with lightpath L_1 from node 1 to 4, L_2 from node 1 to 2, and L_3 from node 1 to 3. If the link between node 1 and 2 fails, without considering priority service levels, a possible restored configuration is to reroute lightpath L_1 through node 3. In this case, lightpath L_2 will have to be dropped due to lack of capacity (See Fig. 5 (b)). On the other hand, if the three lightpaths are prioritized as shown in Fig. 5 (a), lightpath L_3 would have to be preempted with higher priority lightpaths L_1 and L_2 rerouted through node 3 as shown in Fig. 5 (c). In both cases, one connection is dropped, the the second option allows greater control in selecting which connections must be routed, and which ones can be preempted.
- Selective Reconfiguration: Restoration schemes typically assume that when a fault occurs, only disrupted lightpaths will be rerouted; lightpaths that are unaffected by the fault can continue along their original routes. In this scenario, a disrupted lightpath that cannot be rerouted successfully is typically blocked. But rather than simply blocking such a connection, a limited amount of reconfiguration can be allowed.

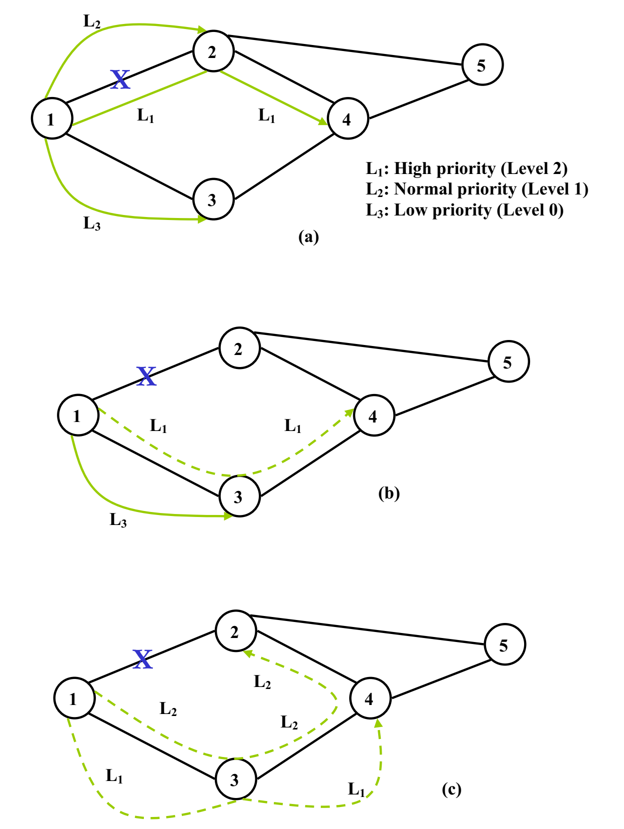

Fig. 5. Prioritized service levels (a) initial configuration (b) restored configuration without priority (c) restored configuration with priority.

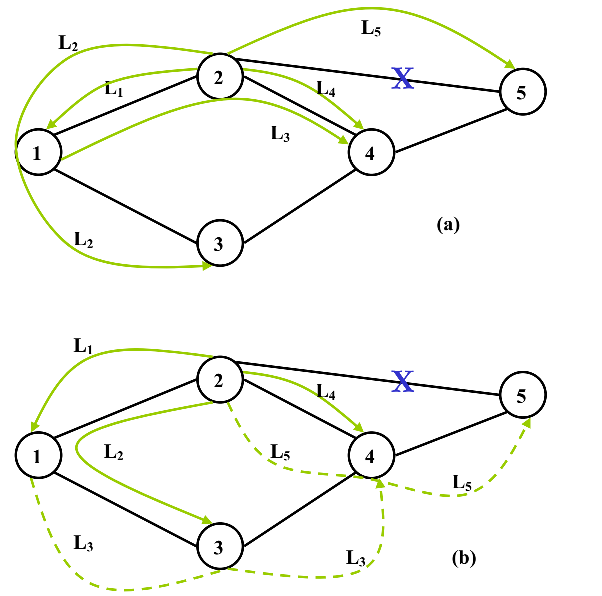

Fig. 6. Selective reconfiguration (a) Initial network with five lightpaths, (b) Limited reconfigured network.

This could include recently rerouted connections, as well as those that were originally unaffected by the fault. Reconfiguration could also involve *relocating* a job, based on the anycast principle. Reconfiguration of *all* (or a large number of) existing lightpaths would not only incur high overhead costs, it would likely be too disruptive for the network as a whole. So, it is critical to develop well-designed and effective techniques that target selected lightpaths to maximize the benefits of reconfiguration, but minimize the associated costs and penalties. In addition, appropriate control-plane strategies must be implemented to carry out the reconfigurations in an efficient and seamless manner.

Fig. 6 shows the advantage of using selective reconfiguration technique. If each link can accommodate at most two lightpaths and the link between node 2 and node 5 fails, based on the traditional restoration scheme, lightpath L_5 would have to be blocked, because the other two outgoing links from node 2 ($2 \rightarrow 1$ and $2 \rightarrow 4$) do not have any available wavelengths. However, lightpath L_5 can be restored through node 4 if the unaffected lightpath L_3 is first rerouted through node 3 under the selective reconfiguration approach. This illustrates how selective reconfiguration can improve performance in case of faults.

B. Enhanced Advance Reservation Schemes

In an optical grid environment, job requests may arrive for *immediate* scheduling or *advance reservation* (AR). Much of the currently available work on survivable optical grids focus on immediate scheduling of jobs. However, it can be argued that in a shared grid infrastructure, advance reservations will become increasingly necessary, in order to guarantee resource availability with the required QoS [32]. In the AR model, also called scheduled traffic model (STM), in addition to the resource and bandwidth requirements, the duration of a job and its start time (or allowable time window) are also specified [33], [34]. There has been considerable research interest in resource allocation strategies for STM in conventional wavelength-routed optical networks [33] - [36]. A few papers have recently extended the concept to address resource reservation in optical grids [32], and used standard path protection for scheduled lightpaths [37]. Novel restoration schemes used in conjunction with the AR model can be used to improve survivability in optical grids, in a variety of ways.

- Rerouting Future Jobs: When a fault occurs, the set of ongoing connections that are affected by the fault must be rerouted. Depending on the estimated recovery time, it may be useful to also consider rerouting scheduled jobs that have not yet started, but are due to commence in the near future. These jobs have a higher probability of becoming affected by the current failure, if it is not resolved in time. One immediate question that arises is how far into the future should rerouting take place? If the exact recovery time is known, a decision can be made relatively easily. But, in practice it is not always possible to accurately predict the fault duration. If the estimated duration is too high, then it will result in a lot of additional rerouting and increased overhead costs. On the other hand, if it is too low, some connections may be unnecessarily delayed. In [32], the authors have used a moving average of the historical failure durations to estimate the actual fault duration. Lightpaths that are scheduled to start within this time are automatically selected for rerouting. It would be useful to develop other appropriate metrics and estimation methods, for calculating the fault duration, possibly taking into consideration the nature of the fault or its location. Furthermore, additional parameters such as network load, traffic distributions or the priority level of a lightpath can be used in order to identify and target selected lightpaths for rerouting.

Fig. 7 shows two scheduled jobs L_1 and L_2 with different starting and ending times (S_1,E_1) and (S_2,E_2) respectively. If the fiber link from node 1 to 2 fails at time T_f with an expected recovery time at T_r, the currently transmitted job L_1 will have to be rerouted through node 3. In addition, as the future job L_2 will start before the failed link is repaired at T_r, we should reroute L_2 through node 3 as well.

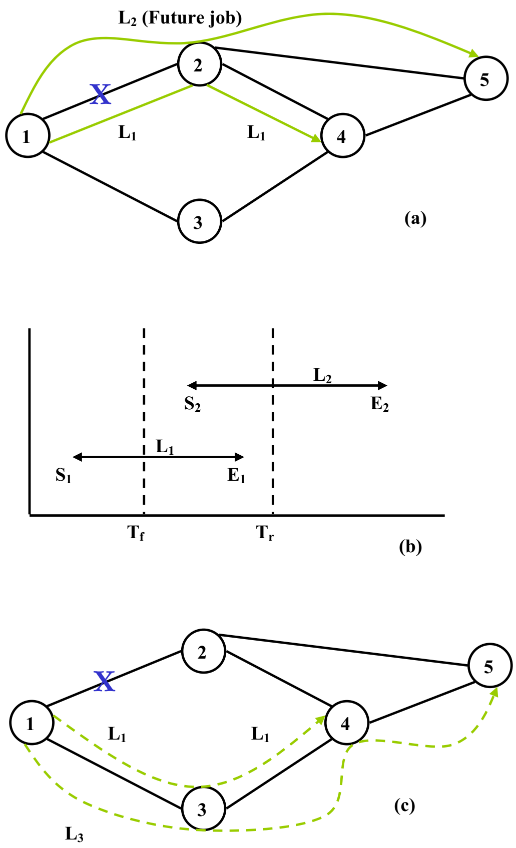

Fig. 7. Rerouting future jobs (a) Initial network with two scheduled lightpaths, (b) Demand schedule, (c) Rerouted lightpaths.

- Rescheduling Jobs: Existing techniques for fault management in AR models typically only consider *rerouting* of affected lightpaths. Network performance can be improved by allowing rescheduling of jobs (in time), in addition to rerouting. This added degree of flexibility can reduce blocking, but also increases the complexity of the process. There is a need for efficient algorithms to i) select appropriate jobs for rescheduling and ii) carry out the rescheduling (possibly in conjunction with rerouting).

  For example, let us consider the same network and set of lightpaths as given in Fig. 7 (a) and (b). Fig. 8 shows that only the interrupted job L_1 is rerouted through node 3, while the future job L_2 can continue to use the originally assigned route, if it is rescheduled to start at a later time (S_2_new), after the expected repair time T_r.

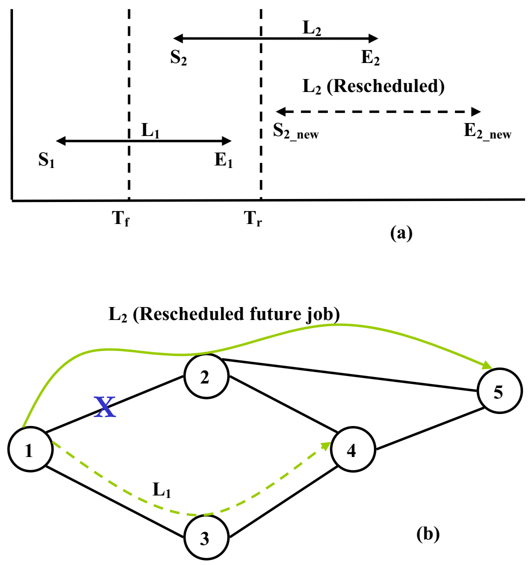

Fig. 8. Rescheduling jobs (a) Originally scheduled and rescheduled jobs, (b) Rerouted and rescheduled jobs.

- Minimizing Fragmentation: Rerouting and rescheduling of disrupted lightpaths can cause increased fragmentation of resources (in both time and space), leading to a degradation of performance. In order to address this, the AR model can be augmented to the more generalized *segmented model* [38], [39], where a single communication task may be decomposed into two or more components and each component can be sent separately. In addition to the usual RWA issues involved in scheduling lightpath demands, design strategies under the segmented model also need to take into consideration a number of other important factors such as: i) which demands (if any) should be divided into segments, ii) the number and sizes of the segments for each demand, and iii) RWA for individual segments. This technique is suitable for some applications where continuous data transmission is not strictly required, like large file transfers for grid computing. For these applications, the segmented model can allow more demands to be accommodated, compared to traditional AR model, which can be viewed as a special case where the maximum number of segments is limited to 1. Relatively little work has been done on the segmented model and there is great potential for enhancing the network performance, in the presence of faults, using this model.

IV. Conclusion

Optical grid networks can provide the necessary infrastructure for a wide range of scientific, business, and other applications, which have a large amount of compute/storage and data transfer needs. There has been significant interest in the design of high-performance optical grids over the past decade, and recently the issue of fault tolerance and survivability in such networks has been receiving increased attention. We have reviewed some important progress in this area, ranging from survivable network dimensioning to fault management for OBS and OCS based transmission. We have also outlined a number of promising approaches that merits further research attention. This includes, for example, exploiting the inherent efficiencies of restoration based techniques, or using advance reservation for guaranteed resource availability and QoS.

In order to respond to a fault scenario, it is necessary first to detect the fault. For optical grid networks, the fault may occur in the computing resources or the networking resources. The development of efficient fault discovery and localization schemes for both components is essential for achieving a survivable optical grid infrastructure. Another area which has received relatively little attention in survivable optical grids is the issue of physical layer impairments and their effects on the quality of transmission (QoT). The tremendous growth in high-bandwidth applications and devices used in backbone networks has led to a corresponding increase in power consumption. In this context, it is extremely important to utilize the available power efficiently. The development of energy-aware algorithms for routing and resource allocation can be of great potential benefit. Much work still remains to be done and there is tremendous opportunity for the development of new models and strategies for handling faults in optical grid networks.

Acknowledgment

The work of A. Jaekel has been supported by research grants from the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

[1] B. Volckaert, P. Thysebaert, M. De Leenheer, F. De Turck, B. Dhoedt and P. Demeester, "Grid Computing: The Next Network Challenge," *The Journal of The Communications Network*, in: *Proceedings of FITCE 2004*, pp. 159-165.

[2] V.M. Vokkarane and Q. Zhang, "Reliable optical burst switching for next-generation grid networks," in: *BroadNets 2005*, pp. 1428 - 1437.

[3] Y. Wang, Y. Jin, W. Guo, W. Sun and W. Hu, "Joint scheduling for optical grid applications," *J. of Optical Networking*, vol. 6 (3) (2007) pp. 304 - 318.

[4] D. Simeonidou (Ed), Optical Network Infrastructure for Grid, GRID WORKING DRAFT, Grid High Performance Networking Research Group, 2006. http://forge.gridforum.org/projects/ghpn-rg/

[5] M. De Leenheer *et al.*, "Design and control of optical Grid networks," (Invited), in: *Proc. Broadnets 2007*, pp. 107 - 115.

[6] Y. Chen, A. Bari, and A. Jaekel, "Strategies for Fault Tolerance in Optical Grid Networks," in: *Proceedings of ICNC 2012*.

[7] C. Dabrowski, "Reliability in grid computing systems, Concurrency and Computation," *Practice and Experience*, vol. 21 (8) (2009) pp. 927 - 959.

[8] M. Chtepen, F. H. A. Claeys, B. Dhoedt, F. De Turck, P. Demeester, and P. A. Vanrolleghem, "Adaptive task checkpointing and replication: Toward efficient fault-tolerant grids," IEEE Transactions on Parallel and Distributed Systems, 20 (2) (2009) 180 - 190.

[9] S. Ramamurthy, L. Sahasrabuddhe and B. Mukherjee, "Survivable WDM mesh networks," *J. of Lightwave Technology*, vol. 21 (4) (2003) pp. 870 - 883.

[10] Z. Sun, W. Guo, Y. Jin, W. Sun and W. Hu, "Fault-Tolerant Policy for Optical Network Based Distributed Computing System," in: *Cluster Computing and the Grid, 2008, CCGRID 2008*, pp. 704 - 709.

[11] R. Nejabati (Ed.), "Grid Optical Burst Switched Networks (GOBS)," Global Grid Forum, 2007.

[12] X. Liu, C. Qiao and T. Wang, "Survivable Optical Grids," in: *Optical Fiber communication/National Fiber Optic Engineers Conference, 2008, OFC/NFOEC 2008*, pp. 1 - 3.

[13] F. Travostino, J. Mambretti and G. Karmous-Edwards, (Eds.), *Grid networks: enabling grids with advanced communication technology*, John Wiely & Sons Ltd, 2006.

[14] G. Zervas, R. Nejabati, Z. Wang, D. Simeonidou, S. Yu and M. O'Mahony, "A fully functional application-aware optical burst switched network test-bed," in: *Proc. OFC, OWC2*, 2007.

[15] L. Liu, X. B. Hong, J. Wu and J. T. Lin, "Experimental demonstration of P2P-based optical Grid on LOBS Testbed," in: *Proc. OFC, JWA88*, 2008.

[16] X. Hong *et al.*, "A resilient OBS/GMPLS network for survival optical grids," in: *Communications 2009, APCC 2009*, pp. 498 - 501.

[17] X. Huang, V. M. Vokkarane and J. P. Jue, "Burst cloning: a proactive scheme to reduce data loss in optical burst-switched networks," in: *Proc. IEEE ICC 2005*, vol. 3, pp. 1673 - 1677.

[18] Q. Zhang, V. M. Vokkarane, Y. Wang and J. P. Jue, "Analysis of TCP over optical burst-switched networks with burst retransmission," in: *Proc. IEEE GLOBECOM 2005*, vol. 4, pp. 1978 - 1983.

[19] C. Develder, B. Dhoedt, B. Mukherjee and P. Demeester, "On dimensioning optical grids and the impact of scheduling," *Photonic Network Communications*, vol. 17 (3) (2009) pp. 255 - 265.

[20] P. Thysebaert, M. De Leenheer, B. Volckaert, F. De Turck, B. Dhoedt and P. Demeester, "Scalable dimensioning of resilient Lambda Grids," *Future Generation Computer Systems*, vol. 24 (2008) pp. 549 - 560.

[21] J. Buysse, M. D. Leenheer, C. Develder, and B. Dhoedt, "Exploiting relocation to reduce network dimensions of resilient optical grids," in: *Design of Reliable Communication Networks (DRCN)*, 2009.

[22] J. Buysse, M. De Leenheer, B. Dhoedt, C. Develder, "Providing resiliency for optical grids by exploiting relocation: A dimensioning study based on ILP," *Computer Communications*, vol. 34 (12) (2011) pp. 1389 - 1398.

[23] J. Buysse, M. D. Leenheer, B. Dhoedt, and C. Develder, "On the impact of relocation on network dimensions in resilient optical grids," in: *Optical Network Design and Modeling - ONDM*, 2010.

[24] B. Jaumard, J. Buysse, A. Shaikh, M. De Leenheer and C. Develder, "Column Generation for Dimensioning Resilient Optical Grid Networks with Relocation," in: *GLOBECOM 2010*, pp. 1 - 6.

[25] A. Shaikh, J. Buysse, B. Jaumard and C. Develder, "Anycast Routing for Survivable Optical Grids: Scalable Solution Methods and the Impact of Relocation," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 3 (9) (2011) pp. 767 - 779.

[26] C. Develder, J. Buysse, A. Shaikh, B. Jaumard, M. De Leenheer and B. Dhoedt, "Survivable Optical Grid Dimensioning: Anycast Routing with Server and Network Failure Protection," in: *Communications (ICC) 2011*, pp. 1 - 5.

[27] B. Jaumard and A. Shaikh, "Maximizing access to IT services on resilient optical grids," in: *ICUMT 2011*, pp. 1 - 6.

[28] X. Liu *et al.*, "Task scheduling and lightpath establishment in optical grids," *Journal of Lightwave Technology*, vol. 27 (12) (2009) pp. 1796-1805.

[29] M. Zhu *et al.*, "Availability-Driven Scheduling for Real-Time Directed Acyclic Graph Applications in Optical Grids," *Optical Communications and Networking, IEEE/OSA Journal of* vol. 2 (7) (2010) pp. 469 - 480.

[30] X. Yu, C. Qiao, T. Wang and X. Liu, "Survivable logical topology design for distributed computing in wdm networks," in: *OFC 2009*, pp. 1 - 3.

[31] W. Guo *et al.*, "Task scheduling considering fault probability for distributed computing applications over an optical network," *Journal of Optical Networking*, vol. 7 (11) (2008) pp. 947 - 957.

[32] S. Tanwir, L. Battestilli, H. Perros, and G. Karmous-Edwards, "Dynamic scheduling of network resources with advance reservations in optical grids," *Int. J. Netw. Management*, vol. 18 (2) (2008) pp. 79 - 105.

[33] J. Kuri, N. Puech, M. Gagnaire, E. Dotaro and R. Douville, "Routing and wavelength assignment of scheduled lightpath demands," *IEEE JSAC*, vol. 21 (2003) pp. 1231 - 1240.

[34] B. Wang, T. Li, X. Luo, Y. Fan and C. Xin, "On service provisioning under a scheduled traffic model in reconfigurable WDM optical networks," in: *BroadNets 2005*, pp. 15 - 24.

[35] T. Li and B. Wang, "Path protection based routing and wavelength assignment in wavelength division multiplexing optical networks under a scheduled traffic model," *J. Opt. Networking*, vol. 5 (2006) pp. 575 - 588.

[36] A. Jaekel and Y. Chen, "Resource provisioning for survivable WDM networks under a sliding scheduled traffic model," *Optical Switching and Networking*, vol. 6 (1) (2009) pp. 44 - 54.

[37] H. Nguyen, M. Gurusamy and L. Zhou, "Scheduling network and computing resource for sliding demands in optical grids," *Journal of Lightwave Technology*, vol. 27 (12) (2009) pp. 1827 - 1836.

[38] Y. Chen, A. Jaekel and A. Bari, "A New Model for Allocating Resources to Scheduled Lightpath Demands," *Computer Networks*, vol. 55 (13) (2011) pp. 2821-2837.

[39] N. Charbonneau and V. M. Vokkarane, "Dynamic Circuits with Lightpath Switching over Wavelength Routed Networks," in: *Proceedings, Fourth IEEE ANTS 2010.*