On the Advantages of Integrating Service Migration and GMPLS Path Restoration for Recovering Grid Service Connectivity [*]

Luca Valcarenghi and Piero Castoldi
Scuola Superiore Sant'Anna
Via Cisanello 145, 56124, Pisa (PI), Italy
Email: {valcarenghi, castoldi}@sssup.it

Abstract

Currently, grid computing fault tolerance, i.e. the capability of recovering from both hardware and software failures, is based on IP dynamic rerouting and on fault tolerant schemes implemented in the applications or in the middleware. However, despite their wide range of applicability, current grid fault tolerant schemes might not be as efficient as emerging resilient schemes, based on the Generalized Multi-Protocol Label Switching (GMPLS) protocol framework, in overcoming specific hardware failures, such as network infrastructure failures.

In this study a mathematical programming model for evaluating the performance of a grid fault tolerant scheme capable of overcoming single physical link failures is presented. The scheme is based on the integration of application/middleware fault tolerance, i.e., service migration, and GMPLS path restoration.

Numerical results show that the integrated scheme allows to increase the ratio of recovered inter-service connections with respect to the utilization of standard IP shortest path (OSPF) dynamic rerouting while limiting the number of service migrations. However, many other factors, such as the network throughput upon failure occurrence, the number of locations allowed for service migration, and the amount of bandwidth necessary for service migration, negatively affects the ratio of recovered connections.

1 Introduction

In grid computing a complex task is accomplished by leveraging services, such as the execution of a job (i.e., computational resources, virtual processors) or a replica of a frequently used data set (i.e., storage capacity, distributed virtual shared memories), distributed among different sites physically connected by the communications network infrastructure [3]. Emerging network architectures, such as Internet Protocol over Generalized MultiProtocol Label Switching over Wavelength Division Multiplexing (IP/GMPLS/WDM), are fostering the expansion of grid computing from Local Area Networks (LANs) (i.e., cluster grid) to Wide Area Networks (WANs) (i.e., global grid) by allowing different sites at which services are located to communicate through Quality of Service (QoS) guaranteed end-to-end connections [8].

Fault tolerance is emerging as a QoS requirement of paramount importance for grid computing. Indeed grid services collaborating in solving a complex task must be able to seamlessly interact in spite of both software and hardware failures, such as temporary unavailability of computational resources due to system crash, or communications network failures, such as optical fiber span failure. Current solutions for grid computing fault tolerance are based on three schemes that are directly implemented by the user in the applications: *checkpointing*, *migration*, and *replication* [10, 6, 7, 9]. To automatize response to failures, schemes directly implemented in the middleware are also emerging [5, 4, 1]. However multipurpose fault tolerance schemes, i.e. schemes able to address different type of failures such as the ones implemented in the application and in the middleware, might not efficiently overcome specific failures, such as communications network infrastructure failures. In addition IP based fault tolerance, i.e. dynamic shortest path rerouting, might not be able to maximize the number of recovered connections without degrading the connection quality, e.g. bandwidth.

The objective of this study is to develop a mathematical model to investigate the benefits of integrating application and middleware layer fault tolerant schemes with resilient schemes implemented below the Network layer of the TCP/IP reference model, e.g. at the GMPLS layer. In

---

[*] This work was supported in part by the Italian Ministry of Education and University (MIUR) under FIRB project "Enabling platforms for high-performance computational grids oriented to scalable virtual organizations (GRID.IT)".

particular the integration of *service migration* [1] with GMPLS *path restoration* for recovering inter-service connectivity upon single network link failure is considered. The integrated fault tolerant scheme consists in utilizing two alternative approaches for overcoming a single network link failure. On the one hand, upon failure, grid services whose connections are affected by the failure might migrate to some allowed locations. New connections among the locations where services migrated must then be set up to maintain inter-service connectivity. On the other hand, path restoration can be utilized to find a new (not necessarily the shortest) physical route for recovering the disrupted inter-service connectivity among the locations at which services where originally staged. By recovering, transparently to the services, the inter-service connectivity path restoration potentially avoids the service synchronization and restart required by service migration.

The mathematical model utilized to evaluate the integrated fault tolerant scheme is based on the Mixed Integer Linear Problem (MILP) formulation of maximizing the number of recovered inter-service connections, considering both the connections that must be set up between locations where services migrated and the connections recovered through path restoration. Constraints such as the number of allowed locations where services can migrate, and the amount of bandwidth necessary for service migration are also taken into account.

Numerical results show that service migration and path restoration integration allows to improve the number of recovered connections, i.e. the number of services whose communication seamlessly overcome the fault, with respect to the utilization of standard IP shortest path dynamic rerouting. On the one hand, path restoration allows to limit the number of utilized replica locations and, on the other hand, service migration helps the underlying network infrastructure in maintaining inter-service connectivity. Moreover the integrated fault tolerant scheme allows to limit the additional computational and storage resources required by service migration and the burden of service synchronization and restart.

2 Service Migration and Path Restoration Integration

The generic scenario in which the proposed scheme based on the integration of service migration and path restoration can be applied is depicted in Fig. 1(a). A service S hosted at node 0 originates a task and it utilizes services, i.e., job running services and data storage services, hosted at different network nodes. Connections among the network nodes involved in the computation are set up to generate the desired inter-service connectivity pattern, i.e., the inter-service logical topology. Connections are assumed to be long-lived and permanent, i.e., flow based, to guarantee the required QoS to the inter-service communication. Upon failure of a physical network link, e.g. link $(0, 1)$, some inter-service connections are disrupted.

The integrated fault tolerance scheme is utilized to successfully recover full inter-service connectivity even in case of limited number of allowed replica locations (i.e., locations where services are allowed to migrate) and maximum number of allowed connections (because of QoS constraints) along the network links. For instance assume that, in the scenario depicted in Fig. 1(a), node 5 is equipped to run only data storage services, no more than one job running service can run on each network node, and all the network links can carry at most two connections (except link $(1, 4)$ whose maximum capacity is 3 connections) [2]. In the integrated scheme (Fig. 1(b)) full inter-service connectivity is recovered if service A migrates from node 1 to node 2 and the connection between node pair $(0, 1)$ is rerouted along the route spanning nodes 0, 5, 3, and 4. If instead either service migration only or path restoration only had been utilized full inter-service connectivity could not have been recovered.

To evaluate the performance of the integrated fault tolerant scheme a simplified model (Fig. 2) of the scenario depicted in Fig. 1(a) and in Fig. 1(b) is considered. A *client-server* scenario is considered where only service pairs communicate. Bidirectional connections, $f^{s,d}$ and $f^{d,s}$, are established between each task emitter service, i.e. the client s, and each utilized service, i.e. the server d. Thus the number of connections requested between node pairs equals the number of service pairs that communicate. Only utilized services, i.e. services hosted at the server node d, can migrate. The locations to which services can migrate, e.g. l_i and l_j, are not predefined but found upon failure occurrence. To migrate, services set up unidirectional connections (*replication connections*), e.g. fr^{s,d,l_i} and fr^{s,d,l_j}, between the node where they are currently hosted to the node chosen for migration. Bidirectional connections (*recovered connections*), e.g. the unidirectional connection pairs f^{s,d,l_i} and fb^{s,d,l_i}, are then established between the client node s and the location chosen for service migration, e.g. l_i. Both replication and recovered connections must be routed along sub-network links supporting a limited number of QoS guaranteed connections.

The mathematical programming model is based on the MILP formulation of the problem of maximizing the number of recovered inter-service connections after each single network link failure. The recovered connections consist of both disrupted connections recovered through path restora-

---

[1] In this study the term service migration is utilized to refer to both service checkpointing and migration and service dynamic replication schemes.

[2] In this example all links and connections are assumed bidirectional.

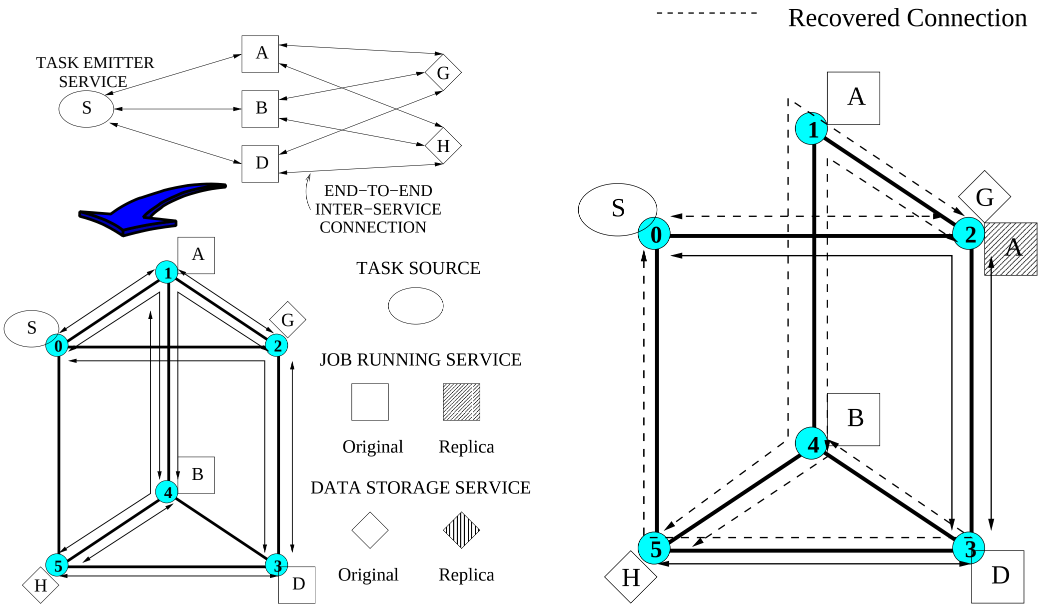

(a) Logical (top) and physical (bottom) network topology

(b) Integrated fault tolerant scheme

Figure 1. Integrated fault tolerance example.

tion and new connections to locations where services migrate. Thus solving the MILP problem formulation corresponds to maximize the number of services whose communication seamlessly overcome the fault by jointly solving the problem of where services migrate, where to route connections to migrated services, where to route connections utilized by services for migration, and where to re-route, through path restoration, disrupted connections so that the number of recovered connections is maximized. The communications network infrastructure (i.e., the physical network topology) is modeled by a graph $G(\mathcal{N},\mathcal{L})$ consisting of $\mathcal{N}$ nodes and $\mathcal{L}$ unidirectional links. The considered failure event e represents the disruption of both the unidirectional links connecting a node pair. For each failure event e the MILP problem formulation is the following.

A number of constants are defined:

- $N = |\mathcal{N}|$, number of network nodes;
- $L = |\mathcal{L}|$, number of network unidirectional links;
- $c_{i,j}$, capacity of link (i,j);
- R_l, total number of (s,d) pairs for which location l can be utilized for service migration;
- $R^{s,d}$, total number of locations allowed for the migration of services hosted in d and communicating with services hosted in s;
- $f^{s,d}$, number of failed connections between pair (s,d);
- Φ_e, set of (s,d) pairs whose connections are disrupted by failure event e;
- $\mathcal{L}_e$, set of network unidirectional links surviving failure event e.

The variables utilized in the MILP formulation are:

- $f^{s,d,l}$, number of connections, originally between the pair (s,d), recovered by migrating service to location l[3];
- $fr^{s,d,l}$, number of connections for service migration from location d to location l;
- $f^{s,d,l}_{i,j}$, number of connections, originally between the pair (s,d), recovered by migrating service to location l passing through link (i,j);
- $fb^{s,d,l}_{i,j}$, number of backward connections corresponding to connections $f^{s,d,l}_{j,i}$ passing through link (i,j);
- $fr^{s,d,l}_{i,j}$, number of connections for service migration from location d to location l passing through link (i,j);

---

[3] The value of $f^{s,d,l}$ equals the number of services that migrated from location d to location l.

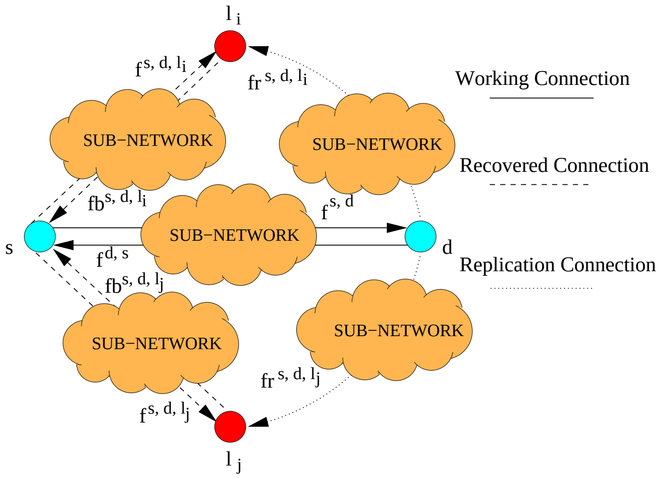

Figure 2. Integrated fault tolerance model with migration flow.

- $r_l^{s,d}$, binary variable indicating whether at least one service hosted in d and communicating with services hosted in s migrates to location l.

2.1 Problem formulation

Objective:

$$max \sum_{l=0,l\neq s}^{N-1} \sum_{(s,d)\in\Phi_e:s>d} f^{s,d,l} \tag{1}$$

Subject to:

$$\sum_{l=0,l\neq s}^{N-1} f^{s,d,l} \leq f^{s,d} \quad \forall s,d : s>d \tag{2}$$

$$f^{s,d,l} \leq f^{s,d} \cdot r_l^{s,d} \quad \forall s,d,l : l\neq s, s>d \tag{3}$$

$$r_l^{s,d} \leq f^{s,d,l} \quad \forall s,d,l : l\neq s, s>d \tag{4}$$

$$fr^{s,d,l} \geq \alpha \cdot f^{s,d,l}$$
$$\forall s,d,l : l\neq s, l\neq d, s>d, 0\leq\alpha\leq 1 \tag{5}$$

$$fr^{s,d,l} \leq f^{s,d,l}$$
$$\forall s,d,l : l\neq s, l\neq d, s>d \tag{6}$$

$$fr^{s,d,l} = 0 \quad \forall s,d,l : l=d, s>d \tag{7}$$

$$\sum_{l=0,l\neq s}^{N-1} r_l^{s,d} \leq R^{s,d} \quad \forall s,d : s>d \tag{8}$$

$$\sum_{(s,d)\in\Phi_e:s>d} r_l^{s,d} \leq R_l \quad \forall l : l\neq s \tag{9}$$

$$\sum_{j:(i,j)\in\mathcal{L}_e} f_{i,j}^{s,d,l} - \sum_{h:(h,i)\in\mathcal{L}_e} f_{h,i}^{s,d,l} = $$
$$= \begin{cases} f^{s,d,l} & \text{if } i=s \\ 0 & \text{if } i\neq s,l \\ -f^{s,d,l} & \text{if } i=l \end{cases}$$
$$\forall s,d,i,l : l\neq s, s>d \tag{10}$$

$$\sum_{j:(i,j)\in\mathcal{L}_e} fr_{i,j}^{s,d,l} - \sum_{h:(h,i)\in\mathcal{L}_e} fr_{h,i}^{s,d,l} = $$
$$= \begin{cases} fr^{s,d,l} & \text{if } i=d \\ 0 & \text{if } i\neq d,l \\ -fr^{s,d,l} & \text{if } i=l \end{cases}$$
$$\forall s,d,i,l : l\neq s, l\neq d, s>d \tag{11}$$

$$\sum_{(s,d)\in\Phi_e:s>d} \sum_{l=0,l\neq s}^{N-1} \left(f_{i,j}^{s,d,l} + fb_{i,j}^{s,d,l} + fr_{i,j}^{s,d,l} \right) \leq c_{i,j}$$
$$\forall (i,j)\in\mathcal{L}_e \tag{12}$$

$$fb_{j,i}^{s,d,l} = f_{i,j}^{s,d,l}$$
$$\forall i,j : (i,j),(j,i)\in\mathcal{L}_e, s,d,l : s>d, l\neq s \tag{13}$$

$$r_l^{s,d} = \{0,1\} \quad \forall s,d,l : s>d, l\neq s \tag{14}$$

$$r_l^{s,d} = 0 \quad \forall s,d,l : s>d, l=s \tag{15}$$

$$\text{int } f^{s,d,l}, f_{i,j}^{s,d,l}, fb_{i,j}^{s,d,l}$$
$$\forall s,d,l,(i,j)\in\mathcal{L}_e : s>d, s\neq l \tag{16}$$

$$\begin{aligned} & f^{s,d,l}, f^{s,d,l}_{i,j}, fb^{s,d,l}_{i,j} \geq 0 \\ & \forall s,d,l,(i,j) \in \mathcal{L}_e : s > d, s \neq l \end{aligned} \quad (17)$$

$$\begin{aligned} & fr^{s,d,l} \geq 0 \\ & \forall s,d,l : l \neq s, l \neq d, s > d \end{aligned} \quad (18)$$

The problem formulation assumes that bidirectional connections consists of unidirectional connections routed in opposite direction along links connecting the same node pairs. The objective of the MILP problem formulation, see Eq. 1, is to maximize the number of recovered inter-service connections. Eq. 2 limits the number of recovered connections for each (s,d) pair to the number of disrupted connections. Eq. 3 and Eq. 4 state that for each (s,d) pair a connection $f^{s,d,l}$ to a replica location l exists *iff* location l is used by at least one service hosted at d. Eq. 5 states that, if a service migrates to a location l, a service migration connection of bandwidth at least $\alpha f^{s,d,l}$ is set up between the current service location d and l. Eq. 6 states that the amount of flow required by the migration connection is limited by the flow carried by the connection $f^{s,d,l}$ to the replica location l. However if path restoration is utilized, i.e. $l = d$, no service migration connection is necessary (see Eq. 7). Eq. 8 limits to $R^{s,d}$ the total number of locations to which services hosted at location d and communicating with services at location s can migrate. Eq. 9 limits to R_l the total number of (s,d) pairs allowed to used location l for service migration. Eq. 10, Eq. 11, and Eq. 12 represents the flow conservation constraint for the recovered connections, the service migration connections, and the link capacity constraint, respectively. Eq. 13 imposes the bidirectionality constraint on the recovered connections. Eq. 14 states that, for each (s,d) pair, the location l can either be utilized or not for service migration. Eq. 15 states that utilized services cannot migrate to nodes hosting task emitter services, i.e. the client node. Eq. 16 and Eq. 17 constrain the related variables to assume positive integer values. Eq. 18 allows $fr^{s,d,l}$ to assume positive real values. The MILP problem formulation complexity is function of the number of problem variables. In the worst case scenario, i.e. connections are set up between any node pair and the link failure causes the disruption of all the connections set up, the number of problem variables can be approximated as $O(N^3 \times L)$.

3 Performance Evaluation Criteria

Four parameters are utilized to evaluate the performance of the integrated fault tolerant scheme. The *expected network blocking probability* is defined as the ratio between the number of unrecovered connections after failure event e and the total number of connections disrupted by the failure event e averaged among all the possible failure events e:

$$E\{Pr_b\} = \sum_{e \in \mathcal{E}} P_e^f \left(1 - \frac{\sum_{l=0, l \neq s}^{N-1} \sum_{(s,d) \in \Phi_e} f^{s,d,l}}{\sum_{(s,d) \in \Phi_e} f^{s,d}}\right), \quad (19)$$

where $\mathcal{E}$ is the set of failure events and P_e^f is the probability of failure event e.

The *expected replica utilization ratio* is defined as the ratio between the number of locations utilized for service migration and the number of locations allowed for service migration averaged among all the possible failure events:

$$E\{\eta_r\} = \sum_{e \in \mathcal{E}} P_e^f \frac{1}{|\mathcal{R}_{\Phi_e}|} \sum_{(s,d) \in \mathcal{R}_{\Phi_e}} \sum_{l, l \neq s} \frac{r_l^{s,d}}{R^{s,d}}, \quad (20)$$

where $\mathcal{R}_{\Phi_e} = \{(s,d) \in \Phi_e : \sum_{l=0, l \neq s}^{N-1} f^{s,d,l} \geq 0\}$ is the set of (s,d) pairs for which at least one service migrates from location d to location l and the connection between node pair (s,l) is successfully set up.

The *expected path restoration utilization* is calculated as the average number of times the original server location node is utilized as replica location normalized to the number of replica locations utilized:

$$E\{\rho_r\} = \sum_{e \in \mathcal{E}} P_e^f \frac{1}{|\mathcal{R}_{\Phi_e}|} \sum_{(s,d) \in \mathcal{R}_{\Phi_e}} \frac{r_d^{s,d}}{\sum_{l, l \neq s} r_l^{s,d}}. \quad (21)$$

The *expected connection route length* is the number of hops spanned by recovered connections averaged among all the possible failure events:

$$E\{H\} = \sum_{e \in \mathcal{E}} P_e^f \frac{1}{|\mathcal{R}_{\Phi_e}|} \sum_{(s,d) \in \mathcal{R}_{\Phi_e}} \sum_{l, l \neq s} \frac{\sum_{(i,j) \in \mathcal{L}} f_{i,j}^{s,d,l}}{f^{s,d,l}} \quad (22)$$

4 Evaluation Setup and Results

The physical network considered is depicted in Fig. 3 and it is assumed to consist of optical fiber links. Each of the 54 unidirectional links connecting the 16 network nodes has a capacity c_l of 32 wavelengths.

Inter-service connectivity patterns are obtained by generating bidirectional connection between uniformly distributed node pairs until a target network achievable throughput θ_{ac} is reached. Each connection occupies one wavelength. The value of θ_{ac} represents the ratio between the amount of capacity (i.e., wavelengths) utilized for routing the connections between node pair (s,d) along the shortest path and the total network capacity (i.e., wavelengths) available:

$$\theta_{ac} = \frac{\sum_{s,d} f^{s,d} \cdot h_{s,d}^{SP}}{\sum_{(i,j) \in \mathcal{L}} c_l}, \quad (23)$$

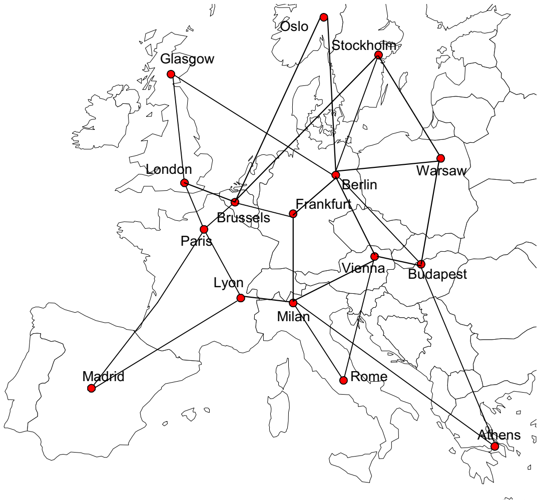

Figure 3. Pan-European network.

where $h_{s,d}^{SP}$ is the number of links spanned by the shortest path between the node pair (s,d). Numerical results are obtained by averaging the solutions, found by CPLEX [2], of the MILP problem for 100 different inter-service connectivity patterns achieving the same network throughput θ_{ac}.

Fig. 4(a) compares the expected blocking probability obtained by different fault tolerant schemes with respect to the proposed integrated scheme (i.e., integrated service migration and GMPLS path restoration scheme). The IP shortest path (OSPF) dynamic rerouting scheme reroutes disrupted connections through the shortest path computed on the physical network topology without the disrupted links. The GMPLS path restoration scheme reroute disrupted connections by solving the MILP problem formulated in section 2.1 without the possibility of migrating services to locations different from the original service location. The service migration scheme solves the MILP problem formulated in section 2.1 without the possibility of recovering inter-service connectivity by finding a new route to the original service location (i.e., by utilizing path restoration) but just by migrating it to a different node. As shown in Fig. 4(a) the utilization of service migration allows to considerably decrease the number of recovered connections with respect to IP shortest path (OSPF) dynamic rerouting.

Fig. 4(b) shows that, in the integrated scheme, for average values of the network achievable throughput, the expected path restoration utilization is about 10%. Thus the utilization of the integrated scheme allows to achieve the same expected network blocking probability obtained by utilizing only service migration (the plots for the service migration scheme and the integrated service migration and GMPLS path restoration scheme overlap in Fig. 4(a)) but it reduces the need for service synchronization and restart required by service migration.

In Fig. 5 and Fig. 6 the comparison of the integrated scheme performance against the path restoration scheme (i.e., $l = d$) performance is shown in function of several integrated scheme parameters.

Fig. 5(a) shows that, if the bandwidth required by service migration increases, i.e. if α increases, the performance of the integrated scheme degrade. Moreover, while $R^{(s,d)}$ does not affect, in the considered scenario, the expected restoration blocking probability, a low value of R_l causes the expected restoration blocking probability to increase. In addition Fig. 5(a) and Fig. 5(b) shows that, as long as the services are allowed to migrate to any network node, only one allowed migration location is sufficient for achieving both the optimal expected network blocking probability and the optimal replica utilization ratio. Indeed allowing all the nodes but the source node s to be utilized for service migration for any service pair (s,d), i.e. $R^{s,d} = 15$, does not improve $E\{Pr_b\}$ (see Fig. 5(a)) and decreases the efficiency in utilizing allowed replica locations, $E\{\eta_r\}$ (see the five overlapping plots for $R^{s,d} = 15$ in Fig. 5(b).

Fig. 6(a) shows that the expected path restoration utilization is negligibly affected by $R^{s,d}$ and R_l while it heavily depends on the value of α. In particular Fig. 6(a) shows the path restoration is more utilized when a high fraction of bandwidth is required by service migration ($E\{\rho_r\} > 60\%$ for $\theta_{ac} \leq 0.6$ if either $\alpha = 1.0$ or $\alpha = 0.5$).

Fig. 6(b) shows that highest expected connection route length is reached when only path restoration is utilized as fault tolerant scheme. This is also in accordance with the higher expected network blocking probability achieved by path restoration alone with respect to the integrated scheme (see Fig. 5(a)): the longer are the restoration paths the higher is their likelihood to be blocked. In addition if the fraction of bandwidth required by service migration increases also the expected connection route length increases due to the fact that more links must be utilized to successfully route a connection.

5 Summary

In this work the efficiency of a new approach to design resilience schemes for grid computing applications operating in a Wide Area Network has been investigated. Specifically the proposed scheme consists of the integration of application/middleware layer resilient schemes, i.e. service migration, and of GMPLS layer (i.e., layer 2/3) resilient schemes, i.e. path restoration. A mathematical programming model has been developed for evaluating the performance of the integrated fault tolerant scheme under the assumption of single link failure in a grid computing network

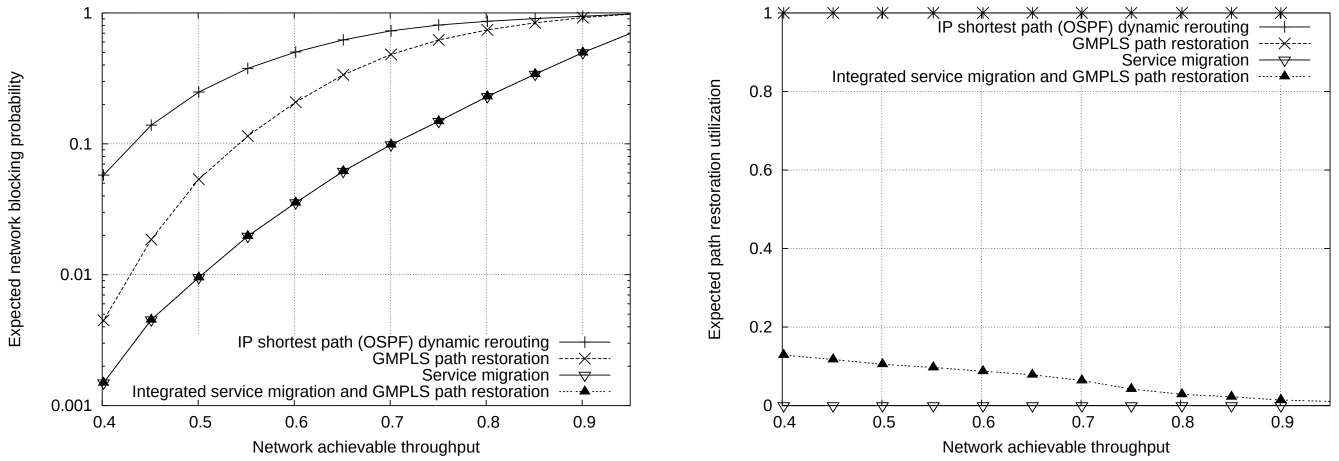

(a) Expected network blocking probability comparison

(b) Expected GMPLS path restoration utilization comparison

Figure 4. Fault tolerant scheme comparison.

infrastructure.

Numerical results showed that the integration of service migration and path restoration allows to engineer fault tolerant schemes that for medium-high throughput increase the number of recovered inter-service connections after a single network link failure and decrease the need for service synchronization and restart. However the integrated scheme performance depend on the network throughput and on service migration parameters, such as the number of locations allowed for service migration and the amount of bandwidth necessary for service migration.

References

[1] L. Alvisi and et. al. Wrapping Server-Side TCP to Mask Connection Failures. In *INFOCOM 2001*, volume 1, pages 329–337, 22-26 April 2001.

[2] CPLEX. http://www.ilog.com/products/cplex/.

[3] I. Foster and et. al. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. Open Grid Service Infrastructure WG, Global Grid Forum, June 22 2002.

[4] B. Lee and J. B. Weissman. An adaptive Service Grid Architecture Using Dynamic Replication. In *IEEE 2nd International Workshop on Grid Computing*, Nov. 2001.

[5] H. M. Lee and et. al. Grid Fault Tolerance Service for Quality of Service. In *The 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003)*, 2003. Posters & Research Demos.

[6] D. S. Milojicic and et. al. Process Migration Survey. *ACM Computing Surveys*, Sep. 2000.

[7] K. Ranganathan and et. al. Improving Data Availability through Dynamic Model-Driven Replication in Large Peer-to-Peer Communities. In *Global and Peer-to-Peer Computing on Large Scale Distributed Systems Workshop*, page 376, Berlin, May 2002.

[8] D. Simeonidou and (Ed.). Optical Network Infrastructure for Grid. GRID WORKING DRAFT, work in progress, Informational Track, Sep. 2003. draft-ggf-ghpn-opticalnets-0.

[9] A. C. Snoeren, D. G. Andersen, and H. Balakrishnan. Fine-Grained Failover Using Connection Migration. In *Proc. of the Third Annual USENIX Symposium on Internet Technologies and Systems (USITS)*, March 2001.

[10] J. B. Weissman. Fault Tolerant Wide-Area Parallel Computing. In *IEEE Workshop on Fault-Tolerant Parallel and Distributed Systems, International Parallel and Distributed Processing Symposium IPDPS*, May, 2000.

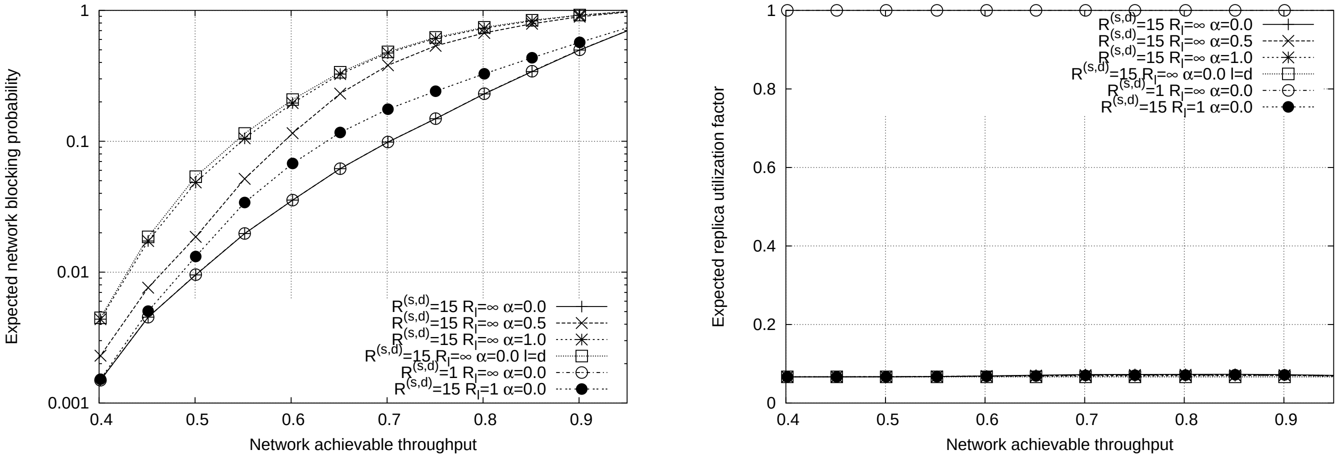

(a) (b)

Figure 5. Integrated scheme expected restoration blocking probability and replica utilization ratio in function of network achievable throughput.

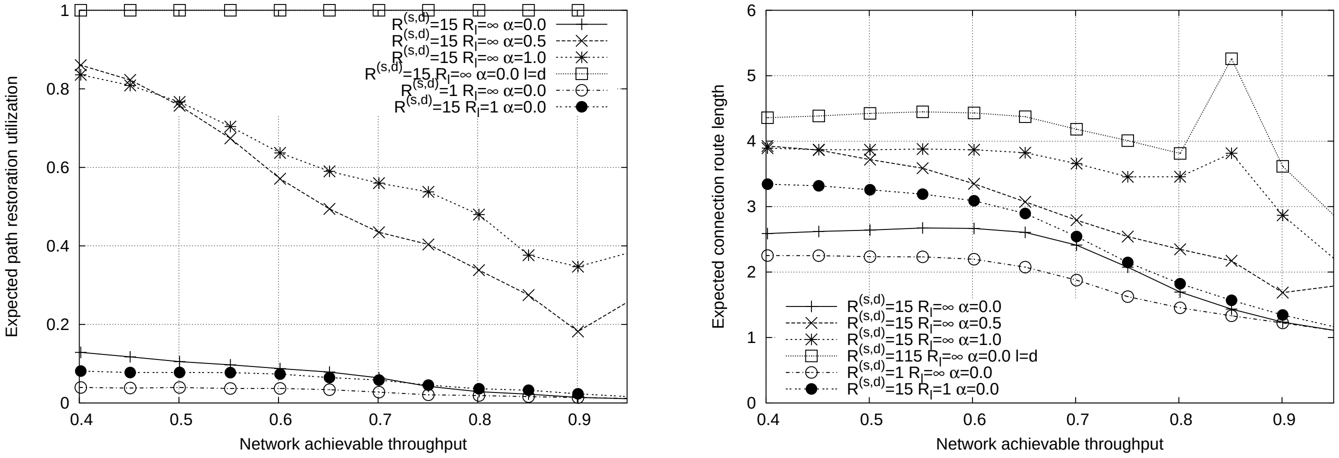

(a) (b)

Figure 6. Integrated scheme expected path restoration utilization and expected connection route length in function of network achievable throughput.