Harmony - Advance Reservations in Heterogeneous Multi-domain Environments*

(Work in Progress)

Alexander Willner¹, Christoph Barz¹, Joan Antoni Garcia Espin²,
Jordi Ferrer Riera², Sergi Figuerola², and Peter Martini¹

¹ University of Bonn, Institute of Computer Science IV, Römerstr. 164,
53111 Bonn, Germany

`{willner,barz,martini}@cs.uni-bonn.de`

² i2CAT Foundation, c/ Gran Capità 2-4, Edifici Nexus I, 2^a planta, desp. 203,
08034 Barcelona, Spain

`{joan.antoni.garcia,jordi.ferrer,sergi.figueroles}@i2cat.net`

Abstract. Grid computing aims at offering standardized access to heterogeneous and distributed resources for scientific communities. In order to ensure certain quality of service requirements, the interconnecting networks have also to be considered as Grid resources and must be taken into account for the co-scheduling process. However, most current systems do not support co-allocation of heterogeneous network resource provisioning systems and malleable advance reservations in multi-domain, multi-technology, and multi-vendor environments. Our approach, called Harmony, provides a functional service plane to unify the underlying network management systems and supports advanced reservation capabilities to utilize the available capacity and network resources in an efficient manner. The developed prototype has been demonstrated on numerous conferences and a preliminary performance evaluation of the current implementation is given.

Keywords: Advance Reservations, BoD, GMPLS, MPLS, QoS.

1 Introduction

Since its first demonstration in the context of the I-WAY [2] project more than ten years ago, the concepts of Grid computing [3] have become increasingly popular in the field of high performance computing. The vision is to virtualize any kind of resource and to combine them to a single abstract entity that is as simple to use as the power grid. One important research area in this context is the co-allocation of these different resources. In particular, the underlying network has to be considered as a first-class Grid resource that must be managed to ensure certain quality of service (*QoS*) requirements.

* Research for this paper was partially financed by the EU in the IST-034115 project PHOSPHORUS [1]. We thank our project partners for their contributions and their collaboration to this research work.

As shown by Foster et al. [4] resource co-allocation in Grid computing environments should be done automatically and transparently. In order to integrate the network into this scheduling process, mainly two issues have to be faced. Meanwhile the infrastructure has to support different QoS features, such as a minimum bandwidth or maximum delay on a specified path, an agreement based resource management is needed to reserve resources immediately or in advance.

In 2006, Travostino et al. [5] gave an overview of the relevance of network research in Grid environments. In a nutshell, there have been two different approaches. Either low-level provisioning mechanisms for IP networks were used such as Integrated Services (*IntServ/RSVP*) [6] and Differentiated Services (*Diff-Serv*) [7]. Or, in order to meet the very high bandwidth demands of some Grid applications, recent projects have started to use circuit-switched optical networks. However, shortcomings of the existing approaches are mostly their limited applicability in real environments or missing support of advanced scheduling algorithms.

In this paper, we assume a set of independent administrative domains that run their local network resource provisioning systems (*NRPSs*). These domains might use different network technologies and are interconnected by static pre-defined links. In our approach, we use an abstract service plane to provide an end-to-end service and allow to co-allocate these networks in conjunction with other Grid resources. The primary goal of this work is to provide a system that reuses existing developments already running as NRPSs, and to allow independent administrative domains to be integrated into the Grid resource scheduling process, and therefore support advanced mechanisms to utilize the transport network efficiently.

The key contribution of this paper is to provide a proof-of-concept and to deliver measurements of a real implementation that builds an abstraction layer for various existing signalling protocols. Overall, we believe this paper helps to construct a foundation for further network research and developments in the field of resource co-allocation in multi-domain, multi-technology, and multi-vendor Grid environments.

The remainder of the paper is structured as follows. We give an overview of related work in the context of dynamic circuit provisioning in Section 2. In Section 3, we state our assumptions and the proposed architecture and we introduce the challenges that are given in co-allocated multi-domain advance reservation situations. Section 4 contains implementation details of our solution. In the subsequent Section 5 a preliminary performance evaluation of our implementation is given. Finally, we close with some conclusions and describe future work in Section 6.

2 Related Work

Grid applications typically need to allocate and reserve multiple types of resources, such as computation, data, instrumentation, and networks. The co-allocation problem for computational Grids has been defined by Czajkowski [8]

1999. In the same year, based on the techniques and concepts of the Globus Resource Management Architecture (*GRMA*) [9], a Distributed Resource Management Architecture (*GARA*) [10, 4] that used a Globus Resource Allocation Manager (*GRAM*) job scheduler to co-allocate the network with other resources in advance was proposed. By then, the mechanisms used for network provisioning were IntServ [6] and DiffServ [7].

While GARA is popular among the Grid community as a general purpose platform allowing reservations of numerous resources it is not specialized for networks. Its API and Resource Specification Language (*RSL*) do not take network specific attributes into account. Therefore, the Network Resource Scheduling Entity (*NRSE*) [11] was introduced in the Grid Resource Scheduling (*GRS*) project in 2002. Still IntServ and DiffServ mechanisms were used to deliver guaranteed throughput over packet-based networks.

Unfortunately, bulk data transfer oriented Grid computing often requires guaranteed minimum bandwidth and minimized packet loss, which are not easily achievable in packet switching networks. Moving Terabytes or Petabytes of data among multiple sites require dedicated optical networks, e.g., based on wavelength switching, that can provide guaranteed bandwidth and performance in terms of low bit error rates.

In 2005 the Exploitation of Switched Light paths for eScience Applications (*ESLEA*) project had started. It demonstrated the usefulness of circuit-switches networks for different application areas. The ESLEA Control Plane Software (*CPS*) was implemented as a modification of the afore mentioned NRSE and was integrated into the EGEE Bandwidth Allocation and Reservation (*BAR*) architecture. At the same time the DARPA DWDM-RAM project addressed similar issues and a Network Resource Scheduling (*NRS*) service was developed to enable the efficient use of optical networks as a primary Grid resource.

Associated projects. Harmony is based on existing solutions for intra-domain path provisioning and scheduling, developed within associated projects. In total, four different NRPSs are used to administer the underlying network: *ARGIA* [12] (former *UCLP* – User Controlled Light path Provisioning); Nortels proof-of-concept middleware called Dynamic Resource Allocation Controller (*DRAC*) [13]; the MPLS [14] based Allocation and Reservation of Grid-enabled Optical Networks (*ARGON*) [15] system that was developed within the German research project VIOLA¹ and also a thin adaption layer for the GMPLS control plane.

Related Projects. In fact, other related projects aim at similar challenges, although from a different perspective, while the Phosphorus view is from the user to the network: Originating from the GÉANT2² project, the Automated Bandwidth Allocation across Heterogeneous Networks (*AutoBAHN*) system; as an achievement of the DANTE-Internet2-CANARIE-ESnet collaboration (*DICE*), an inter-domain control plane based on OSCARS was developed where an

¹ <http://www.viola-testbed.de>

² <http://www.geant2.net/>

inter-domain controller (*IDC*) communicate in a decentralized way to provision end-to-end multi-domain network paths; the G-lambda project that develops an interface between Grid resource management systems and network resource management systems that also support advance reservations; the GMPLS based EnLIGHTened Computing project that focuses on dynamic optical light-paths between supercomputing sites which are created and torn down in advance or on demand based upon application needs; the Dynamic Resource Allocation in GMPLS Optical Networks (DRAGON) project that aims at both the dynamic intra- and inter-domain provisioning of packet and circuit switched network resources in response to user requests for high-performance e-Science applications; and finally, the Grid-enabled GMPLS (G²MPLS) Network Control plane, as an enhancement of the ASON/GMPLS Control Plane architecture that implements the concept of Grid Network Services (*GNS*).

3 Assumptions, Limitations and Architecture

Consider a network consisting of interconnected, independent domains that are used as transit systems or that contain different kinds of resources. A domain in this context is itself again a high performance (optical) network, with a control plane above and a network resource management system (*NRPS*) on top. Each level of this architecture performs different roles and tasks in order to allow creating on demand data plane paths for several users along a single domain. Interconnection with third party domains has to be proposed, agreed and set up by all involved parties. This leads to high delay in inter-domain path provisioning, bureaucratic overhead and hard scheduling of resource utilisation (considering both intra- and inter-domain resources) which most of the times discourage end users.

The problem we consider here in detail is how to find and reserve an-end-to end path between two or more resources on demand with certain QoS requirements, while supporting heterogeneous network technologies. As a consideration, each domain must not reveal its internal topology and the network resources should be utilized to capacity.

General Architecture. The proposed architecture extends the afore mentioned hierarchy by one additional layer. The network service plane (*NSP*) consists of at least one inter-domain broker (*IDB*) and one NRPS Adapter (*HNA*) for each subjacent domain. They both share the same Harmony Service Interface (*HSI*) and allow to exchange abstract topology information and to administer network paths. The architecture has started with a centralized design (Fig. 1a) and evolved to an hierarchical (Fig. 1b) and distributed (Fig. 1c) system.

Topology Knowledge. In particular, for reasons of confidentiality the internal topology information of a domain might not be opened to the service layer. This is the reason why the knowledge of the global topology is restricted to a set of basic information based on three main elements: the endpoint, the link, the domain itself, and the border points.

Fig. 1. The service plane architectures and the operating modes implemented. *White circles* are IDBs, *grey circles* are adapters, and *grey boxes* are administrative domains with their corresponding NRPS.

Based on the Transport Network assigned address (TNA) [16] endpoints in the NSP are identified by strings with IPv4 syntax. Each network port receives a unique TNA in the domain it is attached to and the domain itself serves one or more TNA address ranges. Two border endpoints identify uniquely an inter-domain link between two given domains.

Every domain exports its border endpoints connected to inter-domain links and the inter-domain links themselves. Hence, a transport network controlled by a single HSI capable system is seen as a cloud with a set of border endpoints. If two or more domains are controlled by a single IDB, this new super-domain will be seen as one cloud with a subset of the original border endpoints. Border endpoints connected to the other domains under control of the same IDB are kept as intra-domain endpoints and are not passed up.

Path Computation. End-to-end path provisioning requires dynamic intra- and inter-domain path computation. Since only an abstracted topology knowledge is distributed, the intra-domain path computation is performed by each NRPS. The inter-domain path computation is performed inside each IDB by using Dijkstra's shortest path algorithm. Only border endpoints and inter-domain links are considered.

A reservation itself contains one or more services and each service contains one or more connections. Whereby each connection is a requested path between one source TNA and one or more destination TNAs. Different fixed or malleable QoS demands can be defined on each level.

The actual reservation is following a non-blocking tree two-phase commit protocol scheme. After the reception of a reservation request, the path computer starts looping over all known network resources in order to find feasible paths. Upon the completion of this task, the availability of the resources within all involved domains in the calculated path is requested. In case the resources are not available, they are pruned within the path computer session and the loop will start again. Otherwise the resource will be reserved.

Fig. 2. The different types of reservations. Fading areas represent unspecified bandwidth or time constraints.

Malleable Multi-domain Network Resource Allocation. As stated in [17] and depicted in Fig. 2 reservations can be divided into two different types. First, reservations can start straightaway upon receipt of the reservation (Fig. 2a). Second (Fig. 2b), they can be scheduled with a specified starting time and a specified duration (fixed, FAR, STSD), a specified starting time and a unspecified duration (STUD), an unspecified starting time and a specified duration (deferrable, DAR, UTSD), and finally with an unspecified starting time and an unspecified duration (malleable, MAR, UTUD).

Of particular interest are the advance reservation requests with elastic starting times and an unspecified duration. These malleable reservations can be used by Grid middleware applications for file transfers. They allow the largest flexibility to schedule resources and to utilize the full network capacity.

When a malleable reservation request is received, the IDB starts looping over distinct inter-domain paths, with feasible start-times and bandwidths in order to find a set of available resources to fulfil the request. First of all, the possible paths are calculated. After that, the path computation algorithm starts looping over all the obtained paths. Inside the loop, the feasible bandwidths of the end-points involved in the connection are retrieved and the maximum bandwidth in the range of the bandwidth provided in the request is chosen. This bandwidth is adjusted in all endpoints to the given path according to the technology-induced granularity. In a final step, the algorithm checks for the availability of the network resources. In case any NRPS returns non-available result, the IDB begins to adapt the start time and bandwidth.

Formal Definitions. We define the data plane as digraph $G_d = (D, L)$. The elements of set D are called *domain*, and the elements of L are ordered pairs of vertices, called *inter-domain link* with $(u_d, v_d, id_d) \in L \subseteq D \times D \times ID$. An *inter domain path* in G_d between two domains d_{source} and d_{dest} is an adjacent sequence of inter domain links $p_d(d_{source}, d_{dest}) = (u_{d_0}, v_{d_0}, id_{d_0}), \dots, (u_{d_{m-1}}, v_{d_{m-1}}, id_{d_{m-1}})$, with $u_{d_0} = d_{source}$, $v_{d_{m-1}} = d_{dest}$, $m \in \mathbb{N}_0$ and $len(p_d) = m$. The number of parallel inter-domain links between two domains can be defined as $n(u_d, v_d) = |\{(u_{d_i}, v_{d_i}, id_{d_i}) \in L | u_d = u_{d_i} \wedge v_d = v_{d_i}\}|$. The successor domains of a domain referring to the data plane are defined as a function $s(u_d) = \{w_d | \exists (u_d, w_d, id_d) \in$

$L\}$. Let $p_d^i(d_{source}, d_{dest}) = (u_{d_i}, v_{d_i}, id_{d_i})$ be the i th link of a path. Then the adjacent domains of a path p_d are defined as $a(p_d) = \{d_{source}\} \cup_{0 \leq i < len(p_d)} \{v_{d_i} | p_d^i\}$.

The service plane is described as the digraph $G_s = (H, C)$ with H a set whose elements are called *Harmony Service Interface entity* (HSI entity), and C a set of ordered pairs of vertices, called *HSI interconnection* with $(u_s, v_s) \in C$. In the centralized or hierarchical model these HSI interconnections also reflect the hierarchy. A path in G_s is denoted as $p_s = (u_{s_0}, v_{s_0}), \dots, (u_{s_{m-1}}, v_{s_{m-1}})$, with $m \in \mathbb{N}_0$ and $len(p_s) = m$. The child of an HSI adapter is defined as $s_{u_s} = \{w_s | (u_s, w_s) \in C\}$. This means, that the HSI entity u_s delegates the reservation task to the set $s(u_s)$. Thus, we have a tree structure of HSI entities. In the distributed model we have a mesh of HSI entities which may be themselves the root of a HSI entity tree.

Let the set of Domains for which an HSI entity h is responsible be $r(h \in H) \in D$ and $R(H_i \subseteq H) = \{r(h) | h \in H_i\}$. Thus, the set of responsible HSI entities of a data path is $R(a(p_d))$.

4 Implementation

The communication between the entities is done by using SOAP messages and the Apache Muse framework. To support workflows and resource co-allocation, the API was derived from existing interfaces used to abstract and map network resources with most limited knowledge in intra-domain Grid environments and was modified to cover aspects of multi-domain conditions [18].

The underlying testbed (Fig. 3a and Fig. 3b) has a high similarity to real network environments in the National Research and Education Networks (*NRENs*). Due to its heterogeneity and variety of administration actors, testing of new developments in the service plane is more controllable when done in a virtual testbed (Fig. 4a and Fig. 4b).

Fig. 3. The data and service plane testbed architecture. Whereby *grey boxes* are administrative domains with their corresponding NRPS and *arrows* are preconfigured VLANs between them. *Grey circles* are domains controlled by a single NRPS and *white circles* represent IDBs.

Fig. 4. The service and the emulated data plane architecture within the virtual testbed. Whereby *grey boxes* are administrative domains with their corresponding NRPS and *arrows* represent emulated VLANs between them. *Grey circles* are domains controlled by a single NRPS and *grey circles* represent IDBs.

The data plane inter-connections within the testbed use dedicated light-paths from either GÉANT2 or GLIF³ infrastructure. The switching equipment in the local testbeds is composed among others by the following systems: Alcatel-Lucent 1678, 1850, and 7750; Calient DiamondWave FiberConnect; Cisco Catalyst 6509, 3750; Nortel Optera 5200, OME 6500, and HDXc; and Riverstone 15008. Furthermore, hardware used to run the NSP was compiled by standard personal computer components and virtual machines. The main IDB was running on an Intel Quad Core Xeon with 2.80 GHz, 1 GByte RAM, and Fedora Core release 6.

5 Evaluation

Formal Performance Analysis. First, we consider the centralised and hierarchical model for the formal performance analysis. We show, how the processing times of the different domains can be combined for the overall service time of a single path request. Based on this definition, we extend the approach to a distributed model. Subsequently, we show a worst case calculation of the necessary configuration steps if there are several alternative data links between the domains and reservation attempts fails.

The time a single HSI entity $h \in R(a(p_d))$ needs to process a request internally is denoted as t_h and delegates the reservation to its child HSI entities. If we assume that all children of h are requested consecutively, we can define t_h^{total} of an HSI entity recursively as:

³ <http://www.glif.is/>

Fig. 5. Processing of a reservation request within a single IDB (500 repetitions). (a) depicts the response time for each Adapter and (b) shows the delay that is added by every IDB hierarchy level.

$$t_h^{total} = \begin{cases} t_h & \text{if } |s(h)| = 0, \\ t_h + \sum_{h_i \in s(h) \cap R(a(p_d))} t_{h_i}^{total} & \text{if } |s(h)| \geq 1. \end{cases} \quad (1)$$

If the communication with the children is parallelized, we can define $t_{h_{total}}$ as:

$$t_h^{total} = \begin{cases} t_h & \text{if } |s(h)| = 0, \\ t_h + \max_{h_i \in s(h) \cap R(a(p_d))} t_{h_i}^{total} & \text{if } |s(h)| \geq 1. \end{cases} \quad (2)$$

In case of using the distributed model for the service plane is used there is a mesh of HSI entities on top level $H_{top} \in H$. Each top level HSI entity $h_{top} \in H_{top}$ can itself be the root of a hierarchical HSI entity structure. Thus, we get for path request p_d for a distributed service plane operating in a consecutive mode:

$$t^{total} = \sum_{h_i \in H_{top} \cap R(a(p_d))} t_{h_i}^{total} \quad (3)$$

For a path request for p_d for a distributed service plane operating in a parallel mode:

$$t^{total} = \max_{h_i \in H_{top} \cap R(a(p_d))} t_{h_i}^{total} \quad (4)$$

with $t_{h_i}^{total}$ to be evaluated with Eq. 3 or Eq. 2.

If there are several links between at least two adjacent domains of a path $p_d(d_{source}, d_{dest})$ and the reservation request fails in using one of these redundant links, there might still be a chance in reserving the data path via one of the alternative links. In a consecutive model, where a local reconfiguration is possible

Fig. 6. Processing of n reservation request per second within the Service Plane. (a) depicts the duration of each successful request and (b) shows the success rate.

as a crankback mechanism, the number of configuration steps c of a path p is limited by:

$$O(c) = \text{len}(p) + \sum_{0 \leq i \leq \text{len}(p)-1} n(u_i, v_i) - 1 \quad (5)$$

Here, the first summand represents the successful configurations and the second summand the reconfigurations because of a failed reservation request at one of the HSI entities.

Measurements. To emphasize and support the preceding analysis, measurements from the running prototype are given. As shown in Fig. 5a the mean response time of each NRPS adapter (t_h) varies from 200 ms to 410 ms and has a noticeably large number of outliers. The dummy adapter delay points out a 10 ms communication and processing overhead for each request. As seen in Fig. 5b each additional hierarchy level increases the total response time (t_h^{total}) by approximately 500 ms. Furthermore, Fig. 6a and Fig. 6b indicate that the current prototype's response time increases by arithmetic progression with reference to the amount of incoming requests. It also shows clearly a performance threshold at about 39 requests per second. After that the success rate drops dramatically and the request duration almost triples.

6 Conclusion and Future Work

We have shown that under the given assumptions and limitations the proposed architecture allows for multi-domain, multi-technology, multi-vendor network reservation and co-scheduling. We've tested it in a real scenario under the Phosphorus Europe wide testbed. Besides this proof-of-concept the crucial result is that even with a very limited knowledge of topology information and without an homogeneous control plane, complex malleable advance reservations are feasible and realizable.

Based on these results different network operators such as National Research and Education Networks (*NRENs*) would be able to establish bilateral agreements for end-to-end provisioning services without changing technologies or exposing confidential data.

Besides the implemented prototype that certainly is valuable to demonstrate the main functionality within a testbed with real hardware and users, simulations would provide the possibility to evaluate alternative algorithms in a more sophisticated manner. We expect to have first results using the discrete event simulation package SimJava [19] shortly. Moreover, we are currently implementing more progressive algorithms for malleable reservations in order to reduce the number of availability requests and increase the average network utilization.

Long-term objectives mainly focus on enhanced network resiliency capabilities. This includes as well the computation of disjoint paths as the delegation of responsibilities, path monitoring, and automatic re-routing and -connection.

References

1. Figuerola, S., Ciulli, N., de Leenheer, M., Demchenko, Y., Ziegler, W., Binczewski, A., et al.: PHOSPHORUS: single-step on-demand services across multi-domain networks for e-science. In: Wang, J., Chang, G.-K., Itaya, Y., Zech, H. (eds.) Network Architectures, Management, and Applications V. Proceedings of the SPIE, vol. 6784, (2007) 67842X
2. Foster, I., Geisler, J., Nickless, B., Smith, W., Tuecke, S.: Software Infrastructure for the I-WAY High-Performance Distributed Computing Experiment. In: HPDC, vol. 96, pp. 562–571 (1996)
3. Foster, I., Kesselman, C.: The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann, San Francisco (2004)
4. Foster, I., Jennings, N., Kesselman, C.: Brain Meets Brawn: Why Grid and Agents Need Each Other. In: International Conference on Autonomous Agents: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, vol. 1, pp. 8–15 (2004)
5. Travostino, F., Mambretti, J., Karmous-Edwards, G.: Grid Networks: Enabling Grids with Advanced Communication Technology. John Wiley & Sons, Chichester (2006)
6. Braden, R., Clark, D., Shenker, S., et al.: Integrated Services in the Internet Architecture: an Overview. IETF RFC 1633 (Proposed Standard) (1994)
7. Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W.: An Architecture for Differentiated Service. IETF RFC 2475 (Proposed Standard) (1998)
8. Czajkowski, K., Foster, I., Kesselman, C.: Resource Co-Allocation in Computational Grids. In: Proceedings of the Eighth IEEE International Symposium on High Performance Distributed Computing (HPDC-8), pp. 219–228 (1999)
9. Czajkowski, K., Foster, I., Karonis, N., Kesselman, C., Martin, S., Smith, W., Tuecke, S.: A Resource Management Architecture for Metacomputing Systems. In: Feitelson, D.G., Rudolph, L. (eds.) IPPS-WS 1998, SPDP-WS 1998, and JSSPP 1998. LNCS, vol. 1459, pp. 62–82. Springer, Heidelberg (1998)
10. Foster, I., Kesselman, C., Lee, C., Lindell, B., Nahrstedt, K., Roy, A.: A distributed resource management architecture that supports advance reservations and co-allocation. In: Seventh International Workshop on Quality of Service, 1999. IWQoS 1999, pp. 27–36 (1999)

11. Bhatti, S.N., Sorensen, S.A., Clark, P., Crowcroft, J.: Network QoS for Grid Systems. *International Journal of High Performance Computing Applications* 17, 219 (2003)
12. Grasa, E., Junyent, G., Figuerola, S., Lopez, A., Savoie, M.: UCLPv2: a network virtualization framework built on web services [web services in telecommunications, part II]. *Communications Magazine, IEEE* 46, 126–134 (2008)
13. Travostino, F., Keates, R., Lavian, T., Monga, I., Schofield, B.: Project DRAC: creating an applications-aware network (2005)
14. Rosen, E., Viswanathan, A., Callon, R., et al.: Multiprotocol Label Switching Architecture. IETF RFC 3031 (Proposed Standard) (2001)
15. Barz, C., Bornhauser, U., Martini, P., Pilz, M., de Waal, C., Willner, A.: ARGON: Reservation in Grid-enabled Networks. In: *Proceedings of the 1. DFN-Forum on Communication Technologies*. (2008)
16. Rajagopalan, B.: Documentation of IANA Assignments for Label Distribution Protocol (LDP), Resource ReSerVation Protocol (RSVP), and Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) Extensions for Optical UNI Signaling. IETF RFC 3476, Informational (2003)
17. Varvarigos, E., Surlas, V., Christodoulopoulos, K.: Routing and scheduling connections in networks that support advance reservations. *Computer Networks* (2008)
18. Andree, B., Hommes, F., Peeters, B., Recio, J., Snchez, A., Schon, J., de Waal, C., Willner, A.: PHOSPHORUS Deliverable D1.1.: Requirements and specifications of interfaces architecture for interoperability between NRPS, GMPLS, Middleware. European IST project PHOSPHORUS (2007)
19. Howell, F., McNab, R.: SimJava: A Discrete Event Simulation Package For Java With Applications In Computer Systems Modelling. In: *Proceedings of the First International Conference on Web-based Modelling and Simulation* (1998)