Combining Grid Computing and Internet Measurements

Artur Ziviani , Bruno Schulze

[1]Laboratório Nacional de Computação Científica (LNCC)
Av. Getúlio Vargas, 333 – 25651-075 – Petrópolis, RJ

{ziviani,bruno}@lncc.br

Abstract. *Grid computing and Internet measurements are two areas that have taken off in recent years, both receiving a lot of attention from the research community. In this position paper, we argue that these two promising research areas have a strong synergy that bring mutual benefits. Based on such considerations, we propose a measurement middleware service for grid computing. By defining the architecture and the methods of this service, we show that a promising symbiosis may be envisaged by the use of the proposed measurement middleware service for grid computing.*

1. Introduction

Grid computing appears as a new trend in supercomputing and distributed computing [Foster, 2003]. Distributed platforms, such as PlanetLab [Planetlab, 2002], for developing and deploying global-scale network services are also emerging. Applications using ditributed grid computing demand the simultaneous allocation of resources. The quality of this service depends directly on the network conditions between computational clusters [Ferrari and Giacomini, 2004]. In a distributed grid computing infrastructure, we have to deal with different computational capacity among the clusters, but also with an heterogeneous network support linking them [Feng et al., 2003, Schulze and Madeira, 2004]. Therefore, a monitoring infrastructure may contribute in supporting a highly dynamic environment where operational conditions are constantly changing [Quéma et al., 2004]. Meanwhile, measurement-based approaches are emerging in the Internet, leading to a better knowledge of the network [Barford, 2001, Ziviani et al., 2004]. Moreover, measurement infrastructures may benefit of the high processing capacity of grid computing to analyze large amounts of data.

We roughly summarize the situation as follows. In one hand, efforts in the grid computing community consider taking into account measurements to improve resource management. At the other hand, efforts in the measurement community require both a distributed infrastructure and a large processing capacity to deal with large amounts of data collected at different network points.

Based on these considerations, we advocate in this position paper that the areas of grid computing and Internet measurements have a strong *synergy* that has not yet been widely explored. Therefore, we propose a measurement middleware service that can bring mutual benefits to both areas of grid computing and Internet measurements, which constitutes a promising symbiosis [Ziviani and Schulze, 2004].

This paper is organized as follows. Section 2 briefly reviews grid computing concepts while pointing out where Internet measurements can contribute. Section 3 discusses the current techniques for Internet measurement and how an existing distributed grid infrastructure may bring benefits. In Section 4, we propose a measurement middleware service for grid computing. Finally, we present our concluding remarks in Section 5.

2. Distributed grid computing

Large-scale distributed grids provide a virtual platform for computation and data management by integrating networking and computation. Foster *et al.* [Foster et al., 2003] specify the grid architecture using a set of various layers as illustrated by Figure 1. We build upon the original figure by indicating to which grid layers Internet measurements can contribute as we will discuss in more details along this section.

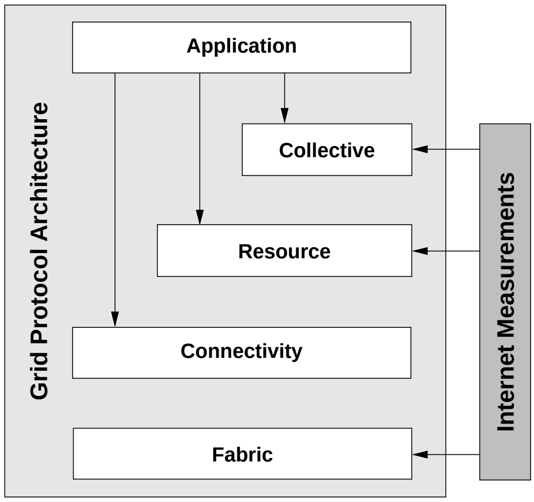

Figure 1: Contribution of Internet measurements to grid layers.

2.1. Fabric

Typically, the Fabric layer implements management mechanisms that allow control over computational, storage, and network resources. In a distributed grid computing scenario, a key point is the adequate allocation of resources to code or data transfers over the underlying network. Internet measurements allow the query of network characteristics, such as available bandwidth or current packet loss rate. This can be used to evaluate the expected quality of service provided in the access to a given resource.

2.2. Connectivity

The Connectivity layer deals with communication and authentication protocols to support grid-specific network transactions. The grid communication protocols include transport, routing, and naming. The TCP/IP protocol stack contributes directly in this context offering direct solutions to distributed grid computing implementations. Since the standard TCP/IP protocols are directly applied here, Internet measurements do not have an influence on this grid layer, although they certainly help on upper and lower layers that use and offer services to this layer, respectively.

2.3. Resource

Management protocols at the Resource layer deal with the provision of an individual resource. This includes the specification of the resource requirements and the operations to be carried out. The resource requirements may include advanced reservation of network resources and the quality of service to be currently expected from the network. Internet measurements play a key role in order to provide these services.

2.4. Collective

As opposed to the Resource layer that focuses on interactions with a single resource, the Collective layer rather controls the coordinated use of multiple resources available at a distributed grid. The scheduling of tasks throughout the distributed grid as well as monitoring services are typical services provided by this grid layer. Feedback provided by Internet measurements taken over the network provide elements to better evaluate the costs of allocating resources on a remote cluster for instance. Likewise, monitoring services take benefit of similar feedbacks to detect failures or attacks.

2.5. Application

Finally, the Application layer encompasses the user applications and provides them access to the services offered by the remaining layers. Although Internet measurements do not directly influence this layer, they enhance services provided by lower grid layer to the user applications.

3. Internet measurements

In the last decade, measurement-based investigation of Internet behavior have allowed a better understanding of fundamental issues in networking [Barford, 2001]. In general, measurement-based approaches to network problems use a distributed infrastructure with either active or passive measurement techniques. Figure 2 illustrates the use of both active and passive techniques, which will be explained along this section. In the following, we describe in further detail the differences between these measurement techniques and we point out how they can benefit from (or enhance) an existing distributed grid computing infrastructure.

3.1. Passive measurements

Passive measurements refer basically to the process of measuring the network traffic, without creating or modifying any traffic on the network. This is done at one or some points of the network. Passive measurement can provide a detailed set of information about the points in the network that are being measured (see the passive measurement points in Figure 2). This information can be for instance the current distribution of the observed traffic in terms of used protocols, the bit or packet rates, inter-arrival packet timing, and so on. Since passive measurements demand the administrative intervention, either by the need of privileged access or the need of dedicated monitoring hardware, it is applied only by either network operators or research institutions. Even if sampling techniques are applied, analyzing large amounts of data still remains a challenging problem. Grid computing easies this high demanding task. Moreover, if the grid computing infrastructure is coupled with the measurement infrastructure, collected raw data may then be sampled, filtered, aggregated, or submitted to any pre-processing action in a distributed manner before a centralized analysis.

3.2. Active measurements

In contrast to passive measurements, active measurements introduce dedicated packets into the network and these packets are timed as they travel through the network to infer network characteristics. Active measurement systems provide little information about single points in the network, but they provide instead a representation of the network path linking two hosts. Network performance metrics that are usually inferred using active measurements are round trip time (RTT), packet loss rate, and connection bandwidth. The active measurement probes sent from host A to host B in Figure 2 give insight about the routing path that links these two end hosts. From a passive perspective, only one of the passive measurement points sees the passage of the active probes.

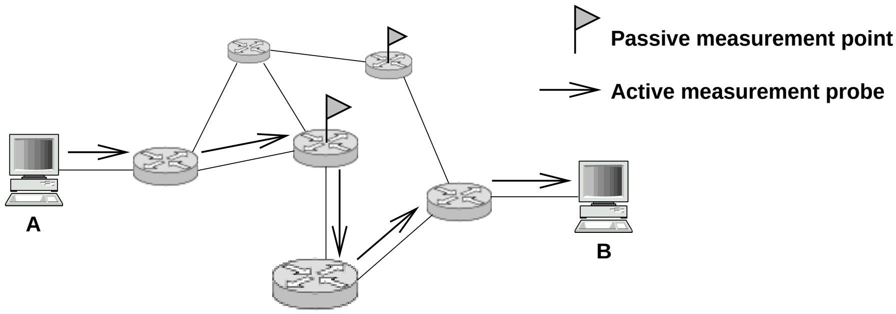

Figure 2: Example of active and passive measurements.

Active measurements are more useful to infer the network conditions along the path linking remote sites. Data from these measurements may then be used to evaluate the service to be expected from the network in a data transfer for instance. Grid scheduling mechanisms may then use such information to influence the decision on how to allocate resources among the remote computational clusters.

4. Measurement service for grid computing middleware

4.1. Architecture

The synergy between the areas of grid computing and Internet measurements points out the importance of a specific measurement middleware service. In Figure 3, we present an architecture for the measurement service for grid computing middleware. This architecture features an arbitrary number of different input channels, each representing the results provided by either passive measurement points or active measurement probes. Different administrative domains might be involved in the communication between each measurement system that is represented by the described architecture. If these systems know about each other and can query information from each other, one may envisage the formation of a peer-to-peer network to exchange measurement information. Thus, each system can operate as a loosely coupled cooperating system. Nevertheless, since different administrative domains may be concerned, a filtering firewall may then tailor the data as a function of the authorization level of the requesting party.

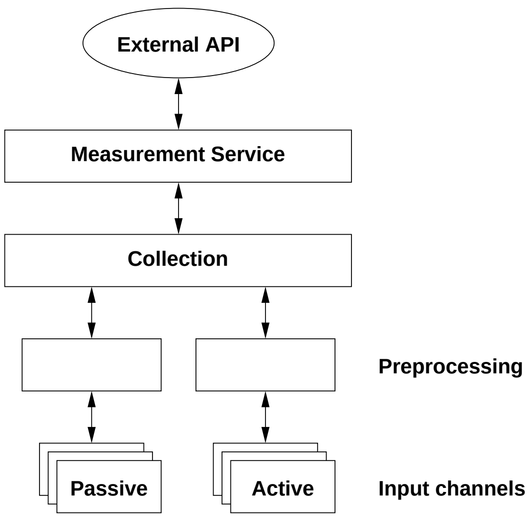

Figure 3: Measurement service architecture.

The heterogeneity nature of the different passive or active measurement techniques requires a preprocessing of the raw data tailored to the particular adopted technique. For example, the minimum observed delay in a sequence of active probes may be used to infer the network delay in the path between two points without taken into consideration network congestion. Similarly, at this point, large amounts of passive data may be filtered by certain criteria or sampled.

The collection module receives the preprocessed data. The task of this module is to integrate the collected information. Measurement scans are time stamped here, if they have not been stamped in their originating system. Time stamps become crucial to compare data from different sources and to determine if there is any lag among the measurement-based estimates of network conditions.

The function of the measurement service module is two-fold. In one hand, it can offer results inferred from measurements in the network connecting computational clusters of a distributed grid. In this case, knowledge of the current network conditions influences the scheduling of grid resources at different levels. In the other hand, measurements may be consulted or requested by authorized entities. The associated grid infrastructure may be then used to process the measurement data. Both functionalities are requested using specific services provided by the external API of the measurement service. These services are described in the following.

4.2. External API

4.2.1. get_measurement

In the case of distributed computational clusters that have links connecting them under periodic monitoring, the *get_measurement* operation allows the current state of the path between any pair of cluster to be known. This information can then be used by resource management to decide on the allocation of resources among the participating computational clusters. The quality of service to be expected on the data transfers between clusters may be also evaluated.

4.2.2. perform_measurement

In contrast to *get_measurement*, the *perform_measurement* operation requests a explicit active probe to be sent. The user has to specify the active measurement point and the target host. In this way, arbitrary measurement experiments can be performed and the associated grid resources can be explored to process the gathered data.

4.2.3. get_raw_data

Given a deployed measurement infrastructure, one important service is to allow users to directly gather raw data for further processing. This is useful for registering data traces that can be used to evaluate the performance of new technologies using real traffic. These traces can also be kept save for accounting and verification of accordance with a given profile.

4.2.4. process_data

As seen in Section 3.1, Internet measurements may generate large amounts of data to be processed. One important service is to allow the processing of this data to evaluate a specific network condition. For example, one may evaluate what is the volume of control data on a given portion of the network or determine the typical delay to be expected between two points on a given weekday from previously gathered data.

5. Conclusion

In this position paper, we point out the existing synergy between the areas of distributed grid computing and Internet measurements. Based on the potential mutual benefits of combining these two research areas, we propose a measurement middleware service for grid computing.

The proposed measurement middleware service brings benefits for the areas of grid computing and Internet measurements. From the grid computing viewpoint, Internet measurements provide estimation of the current network conditions to assure optimized resource management. The quality of service to be expected by the grid applications may also be evaluated. Meanwhile, from the measurement perspective, the grid infrastructure comprising distributed computational clusters allow the performance of a variety of Internet measurement experiments. Since some of these experiments may generate large amounts of data to be processed, grid computing helps to perform such a task in a distributed manner. Therefore, jointly adopting the recent advances in both areas can bring mutual benefits to them. A promising symbiosis may then be envisaged by the use of the proposed measurement middleware service for grid computing.

References

Barford, P. (2001). Global Internet measurement infrastructure. In *Position paper for Large Scale Networking Workshop 2001*, Vienna, VA, USA.

Feng, W.-C., Gardner, M. K., Fisk, M. E., and Weigle, E. H. (2003). Automatic flow-control adaptation for enhancing network performance in computational grids. *Journal of Grid Computing*, 1(1):63–74.

Ferrari, T. and Giacomini, F. (2004). Network monitoring for grid performance optimization. *Computer Communications*, 27(14):1357–1363.

Foster, I. (2003). *The Grid: A new infrastructure for the 21st century science*, in Grid Computing: Making the Global Infrastructure a Reality, pages 51–63. John Wiley & Sons.

Foster, I., Kesselman, C., and Tuecke, S. (2003). *The anatomy of the Grid*, in Grid Computing: Making the Global Infrastructure a Reality, pages 171–197. John Wiley & Sons.

Planetlab (2002). *PlanetLab: An open platform for developing, deploying, and accessing planetary-scale services*. `http://www.planet-lab.org`.

Quéma, V., Lachaize, R., and Cecchet, E. (2004). An asynchronous middleware for grid resource monitoring. *Concurrency and Computation: Practice & Experience*, 16(5):523–534.

Schulze, B. and Madeira, E. R. M. (2004). Grid computing and active services. *Concurrency and Computation: Practice & Experience*, 16(5):535–542.

Ziviani, A., Fdida, S., de Rezende, J. F., and Duarte, O. C. M. B. (2004). Toward a measurement-based geographic location service. In *Proc. of the Passive and Active Measurement Workshop - PAM'2004*, Lecture Notes in Computer Science (LNCS) 3015, Antibes Juan-les-Pins, France.

Ziviani, A. and Schulze, B. (2004). Measurement middleware service for grid computing. In *Proc. of the 2nd International Workshop on Middleware for Grid Computing – MGC 2004*, Toronto, Canada. Poster.