Resource Provision and Allocation in Shared Network Testbed Infrastructures

Panayotis Antoniadis, Timur Friedman, and Xavier Cuvellier
Université Pierre et Marie Curie - LIP6/CNRS, France

1. INTRODUCTION

In this position paper we motivate the use of a rule-based incentive mechanism relating resource provision (contribution) with resource allocation in shared network testbeds architectures, like PlanetLab, as an alternative to market-based approaches [4][5]. We describe its main attributes and provide insights for its configuration in realistic conditions based on data collected through the CoMon monitoring infrastructure of PlanetLab[1]. Our goal is to ensure that organizations acquiring more value from the system will contribute more (and be able to consume more). If designed appropriately such a mechanism would improve the system's operating point without introducing significant complexity.

Note that in addition to their complexity, market mechanisms could harm the inherent "community spirit" and also suffer from the standard problems of virtual currencies (inflation, deflation, etc.). Moreover, the fact that consumers and producers are now the same entities significantly complicates the economic modelling of the system and it is not straighforward to compute the corresponding equilibria and ensure that the nice theoretical properties of market mechanisms will still hold. This is especially so since in general such mechanisms fail to capture the existing externalities related to the *public good* aspect of the system[2].

A large part of the research on incentives mechanisms for p2p systems, considers the design of simple rules such as reciprocity or fixed contribution [2]. Along these lines, PlanetLab currently uses a fixed contribution rule for enforcing resource provision. That is, each participant, a *site*, is required to contribute two nodes. Concerning resource allocation, sites are constrained to create a maximum of 10 *slices* (tasks). But since each slice has unrestricted access to the system resources, which are allocated equally among all competing slices, the current constraints are fairly weak[3].

2. A RULE-BASED APPROACH

PlanetLab's simple fixed contribution rule has enabled the system to grow. However, it is highly loaded, especially before dealines of major conferences, even if the majority of the sites do not consume significant amounts of resources (see Figure 1)[4]. Moreover, sites that contribute more cannot materialize their increased contribution to a better quality of service, which reduces their incentive to do so.

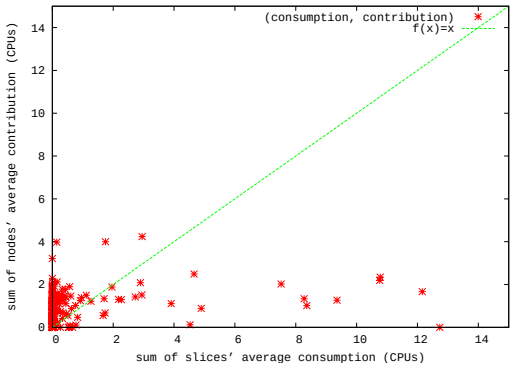

Figure 1: Sites' (consumption, contribution)

We propose to define meaningful levels of service and design a simple mechanism to ensure that sites receive the appropriate level of service according to their overall contribution. Clearly, the most challenging problems towards this direction is, first, to define the notion of contribution and the different levels of service, and, second, to decide the mapping between them. Notice that from an economic point of view, the selected contribution level is used to express utility (instead of the willingness to pay assumed in market-based approaches) and it is thus the measure for implementing an efficient resource allocation (i.e., resources are allocated to users that value them the most). So, organizations that wish to have a better quality of service will have to contribute more and thus increase the whole system's capacity.

This is particularly important in the case of OneLab and PlanetLab Europe [1], in the context of which this work is being carried out. The reason is that OneLab aims to introduce more specialized network resources to be shared amongst participants, such as wireless testbeds (e.g. WiMAX, UMTS, etc.) for which an efficient way to allocate must be devised since they will probably be highly congested. Our proposed mechanism will ensure their efficient allocation but also exploit the increased value that they will introduce so as to further encourage contribution.

2.1 Effective contribution

Currently, contribution in PlanetLab is defined as the number of contributed nodes, independently of their actual performance. However, many of these nodes are often down resulting in less than 50% of them being always available.

[1]for the period 1/1/2006 – 31/6/2006.
[2] Although the underlying resources are rivalrous the contribution of each participant increases the value of the system as a whole.
[3]PlanetLab's service Sirius allows for reservations granting a percentage of a node's CPU but it treats all sites equally.
[4]The largest consumers (i.e., ucb, nyu, and princeton) are not shown for readability purposes.

Moreover, they differ in terms of the value they add to the system, which is depicted by the large difference in terms of the utilization of different PlanetLab nodes (see Figure 2)[5]. Of course this difference depends on several reasons: the reliability of the node, the number of other nodes in the same area, since the geographical heterogeneity of nodes is desirable for many networking experiments [3], the limited use of resource allocation mechanisms, and many more. In general, the real value of a node depends on the actual use of the system (for example unstable nodes could be useful for experiments that wish to take into account such phenomena). The same holds for the system resources (CPU, memory, bandwidth, etc.), since those that will be congested, and thus more valuable, will vary.

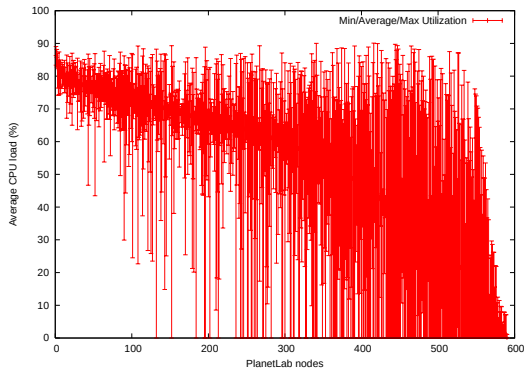

Figure 2: Average CPU load (%) of all PlanetLab nodes in decreasing order

So, it would be reasonable to extend the notion of a site's contribution to include such system-specific characteristics and try to evaluate its *effective contribution*. A simple way to do this is to assume that an objective measure for a node's value could be its popularity (its actual utilization) over time. And define contribution based only on this metric. However, such an approach would lead to high uncertainty and possible gaming of the incentive mechanism. We believe thus that the first step should be to identify the most important aspects as derived from the actual use of the system and assign the appropriate thresholds and weights, which could be further refined based on the system dynamics.

We distinguish between two main parts of a site's contribution: 1) a *static* one, which will include the number of nodes, the amount of resources (with possibly different weights), and maybe their "uniqueness" (which could be evaluated by some predefined objective measures such as the density of a specific geographic area) and 2) a *dynamic* one, which will include the total up-time and/or the rate of failures, the actual popularity, and possibly the utilization of services made available for everybody's use (like CoMon).

2.2 Different levels of service

On the resource allocation end there are two relatively simple ways to differentiate the quality of service received. First, at a *database level*. More specifically, sites could have different treatment in terms of the maximum allowed number of slices, the maximum number of nodes can be used, the maximum number of nodes from which the selection is made, or even exclusive access rights to highly congested or unique resources, such as wireless testbeds.

However, this kind of service differentiation does not address the issue of efficient resource allocation upon contention. So, ideally, one should also provide differentiated treatment *during service provision*. More specifically, additional system rules could be encoded in the operating sytem and assign different priorities to sites belonging to different "contribution classes". This could be performed on the fly or in the context of a resource reservation service like Sirius.

2.3 Mapping

Then our final step will be to define appropriately the mapping between the defined levels of effective contribution with the different levels of service. Ideally, one could compute the parameters that would optimize the efficiency of the system based on appropriate economic models. However, such models are not yet available.

Hence, we wish to follow a bottom-up incremental approach considering first only static contributions and service differentiation at a database level, and gradially introduce mechanisms to reward good dynamic behaviour[6]. The configuration of the mechanism's parameters will be based on on-going measurements of the past and current activity of PlanetLab, while we will also provide the means to the system administrator to tune them as the system's use and value evolves over time. Hopefully we will use this experience to formulate a suitable economic model, which will provide further insights.

3. CONCLUSION

We have argued in favor of a rule-based approach for encouraging resource provision and managing resource allocation in systems like PlanetLab and OneLab. Our short-term next step is to complete our measurements of the past and current activity of PlanetLab and describe and implement a specific rule-based incentive mechanism as presented above in the context of OneLab.

4. REFERENCES

[1] OneLab: . `http://www.one-lab.org/`.

[2] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing Economic Incentives in Peer-to-Peer Networks. *Special Issue on Network Economics, Computer Networks, Elsevier*, 45(1):133–146, 2004.

[3] B. Chun and A. Vahdat. Workload and Failure Characterization on a Large-Scale Federated Testbed. Inter Research Berkeley Technical Report IRB-TR-03-40, November 2003.

[4] M. Feldman, K. Lai, and L. Zhang. A Price-Anticipating Resource Allocation Mechanism for Distributed Shared Clusters. In *ACM Conference on Electronic Commerce*, June 2005.

[5] A.A. Young, B.N. Chun, A.C. Snoeren, and A. Vahdat. Resource Allocation in Federated Distributed Computing Infrastructures. In *Proceedings of the 1st Workshop on Operating System and Architectural Support for the Ondemand IT InfraStructure*, October 2004.

[5] (see plotted also the minimum and maximum monthly average loads).

[6] A rewarding instead of a punishing approach could help sustain the community spirit and encourage participation, and also facilitate the federation of PlanetLab Europe with systems without such mechanisms (e.g., PlanetLab USA).