

ERCIM “Alain Bensoussan” Fellowship Scientific Report

Fellow: Roberto Podesta’

Visited Location : INRIA Saclay-Ile-de-France, Orsay (France)

Duration of Visit: 12 months

I - Scientific activity

The work started with the idea to use the proposal project I had submitted together my ERCIM application into the Private Virtual Cluster (PVC) project, an internal effort targeted to provide an abstraction of a virtual computational cluster on top of Peer-to-Peer collected (P2P) machines. My proposal was focused on a brokering system for meta- and grid- computing platforms. The initial idea was the design of a distributed brokering system for PVC replacing the centralized one which was included in its proof-of-concept implementation. The main target I was told to persecute was a structure not representing a single point of failure for PVC. Being PVC still an evolving project, the set-up time of the work very time consuming. At least five months have been necessary to identify the way to integrate efficiently the new component into the existing system. The distributed broker design after this set-up step was a lookup system with some not yet defined additional services. Therefore, the research was focused towards distributed data model system and the structured overlay network approach was selected because it is the best suited for this kind of requirements. Structured overlay network relies on Distributed Hash Tables (DHTs) which provide the same functional abstraction of a normal hash table storing data on set of cooperative nodes. The theoretical research on DHT started in 2001 and now several open-source implementations are available. The choice of which one is best suited to be adopted depends on several aspects but the main one is the verification of the real state of the art of a certain implementation and the level of completion of the provided functionalities. Therefore, I have spend three months in testing and trying DHT software. In the end of this phase I selected freePastry and Chimera as the most promising and complete DHT implementation. The two projects provide a good degree of maturation and important research group are active and open to cooperation to correct and improve the software updating and delivering periodically new releases. Moreover, they provide the same set of functionalities even if they are implemented in different way. I decided to design the new distributed broker based on DHT taking as the reference those two projects. However, I choose freePastry for the reference implementation since it is written in Java thus getting an immediate portability. The scientific direction was proposing to extend the design of the distributed broker in order to provide support to the HIPCAL project, a national Virtual Cluster project. However, a full integration with PVC had still to be completed. In fact, the network interface between PVC and the distributed broker was designed and implemented in the summer. The design took into account the last cues of the scientific direction by the adoption of a flexible, pluggable and modular architecture able to allow an easy compatibility with not-PVC platforms. In the middle of September I presented a paper at the Middleware for Web Services Workshop (part of EDOC 2008) about the adoption of DHT together the UDDI standard (“A Lightweight Inter-node Operation for UDDI Cloud”). Finally, the reference implementation of the distributed broker was integrated with a new version of PVC (a re-implementation from scratch performed by an other developer than the original one) and successfully tested just in the last week of the fellowship. A final publication with all the results is scheduled in the next months.

The developed prototype can be described briefly as follows: the broker is an autonomous, fault-tolerant, distributed and not-hierarchical system targeted to provide brokering service to resources which are, in general, parts of peer to peer application. I intend as brokering service a system providing a set of functionalities which can be used to discover and get information about distributed software entities registered into the system. In the most common usage scenario, this kind of service can be used in order to establish a direct connection between the entities. Conceptually it can be seen as the Service Broker entity in the SOA operation model . Actually, it has to be underlined the fact that with this distributed broker the notion of software service registered in the broker has a wide sense and it is extended to parts of cooperative applications even representing raw resources. It can be useful for a variety of systems requiring an entity providing extended lookup services. Good samples can be virtual cluster system based on P2P collected resources or even cloud computing system. Even if the design allows plugging in support for different technologies, current implementation provides support for PVC.

II- Publication(s) during your fellowship

A Lightweight Inter-node Operation for UDDI Cloud, Roberto Podesta',

Coordinated UDDI public registries disappeared almost two years ago. In 2006 the UDDI cloud, a registries federation managed by IBM, HP, Microsoft and SAP, was shut down with the end of the UDDI Business Registry (UBR) project. UDDI is still widely adopted for private registries. Isolated public registries are quite rare and an UDDI coordination mechanism is still missing. In URB a pure entries replication mechanism was allowing to access to all the entries published on each server. However, such pure replication does not scale and prevents the creation of dynamic federation of UDDI registries. This paper describes a work-in-progress solution based on a inter-node operation enabling the replication APIs to distribute the entries across a Distributed Hash Table (DHT). Thus, it reduces the number of replicated entries to a sub-set of nodes participating in the UDDI cloud and opens a new possible scenario for the adoption of the UDDI standard.

III -Attended Seminars, Workshops, and Conferences

*Enterprise Computing Conference 2008 (EDOC 2008), 15-19 September, 2008, Munich (GER)
Middleware for Web Services Workshop 2008 (part of EDOC 2008), 16 September, 2008,
Munich (GER)*

IV – Research Exchange Programme (12 month scheme)

*SZTAKI, 10 June 2008 – 19 June 2008, Budapest (Hungary)
FORTH, 27 June 2008 – 04 July 2008, Heraklion (Greece)*