

# ERCIM “Alain Bensoussan” Fellowship Scientific Report

Fellow: Sarunas Girdzijauskas  
Visited Location : SICS, Sweden  
Duration of Visit: 15 Aug 2009 – 14 Aug 2010

## I - Scientific activity

During the stay at SICS, Sweden, the fellow Sarunas Girdzijauskas was involved in research on large scale publish/subscribe systems, efficient data dissemination within data clouds, partitioning graph databases etc. Together with the colleagues from SICS and IBM Haifa labs the fellow wrote two papers on two novel publish/subscribe systems: Magnet and Artemis. One of the papers has been accepted to DEBS 2010 conference. The details on the papers are given in the abstracts below.

The fellow was also supervising one PhD student and two master students.

## II- Publication(s) during your fellowship

**Magnet: Practical Subscription Clustering for Internet-Scale Publish/Subscribe,**  
Sarunas Girdzijauskas, Gregory Chockler, Ymir Vigfusson, Yoav Tock, Roie Melamed  
DEBS 2010, The 4th International Conference on Distributed Event-Based Systems,  
July 12-15, 2010, Cambridge, UK.

Abstract:

An effective means for building Internet-scale distributed applications, and in particular those involving group-based information sharing, is to deploy peer-to-peer overlay networks. The key pre-requisite for supporting these types of applications on top of the overlays is efficient distribution of messages to multiple subscribers dispersed across numerous multicast groups. In this paper, we introduce Magnet: a peer-to-peer publish/subscribe system which achieves efficient message distribution by dynamically organizing peers with similar subscriptions into dissemination structures which preserve locality in the subscription space. Magnet is able to significantly reduce the message propagation costs by taking advantage of subscription correlations present in many large-scale groupbased applications. We evaluate Magnet by comparing its performance against a strawman pub/sub system which does not cluster similar subscriptions by simulation. We find that Magnet outperforms the strawman by a substantial margin on clustered subscription workloads produced using both generative models and real application traces.

## **Artemis: A Topic-based Publish/Subscribe System**

Fatemeh Rahimian, Sarunas Girdzijauskas, Amir Hossein Payberah, Seif Haridi  
(Under preparation)

Abstract:

P2P overlay networks are an attractive solution for building Internet-scale publish/subscribe systems. They provide scalable means to disseminate events to all the subscribers in the system while relying on limited node neighbor sets. This allows the nodes to participate in such systems with moderate bandwidth budgets while incurring only minimal overlay maintenance overhead. However, scalability comes with a cost: a message published on a certain topic often needs to traverse a large number of uninterested (unsubscribed) nodes before reaching all of its subscribers. This might sharply increase resource consumption for such relay nodes (in terms of bandwidth transmission cost, CPU, etc) and could ultimately lead to a rapid deterioration of the system's performance once the relay nodes start dropping the messages or choosing to permanently abandon the system. In this paper, we introduce Artemis, a gossip based publish/subscribe system, which significantly decreases the number of relay messages, while preserving the scalability of the system. This is achieved by constructing an overlay which resembles a navigable Small-World network that spans along the clusters of nodes exhibiting similar subscriptions. The properties of such overlay make it an ideal platform for efficient data dissemination. We perform extensive simulations and evaluate Artemis by comparing its performance against a base-line publish/subscribe system, which is oblivious to the actual node subscriptions. Our measurements show that Artemis significantly outperforms the base-line solution on various subscription and churn scenarios, from both synthetic models and real-world traces.

### **III -Attended Seminars, Workshops, and Conferences**

- "Cloud Futures 2010", Microsoft Research, Redmond, WA, USA, April 8-9
  - Invited talk on "Cognitive Publish/Subscribe for Heterogeneous Clouds"
- DEBS 2010, The 4th International Conference on Distributed Event-Based Systems, July 12-15, 2010, Cambridge, UK
  - Accepted paper
- Guest lecturer for "Distributed Computing" course at KTH (<http://www.ict.kth.se/courses/ID2210/>)
  - Lecture on "Small world and Non-uniform DHT" (14-15 April 2010)
  - Lecture on "Publish Subscribe Systems" (05 May 2010)

### **IV – Research Exchange Programme (12 month scheme)**

- Visit at INRIA, Rennes, ASAP group lead by Anne-Marie Kermarrec (16-20 Nov 2010)
  - Seminar on "Interest Driven Large Scale Pub/Sub Systems" (17 Nov 2010)