ERCIM "Alain Bensoussan" Fellowship Scientific Report

Fellow:	Etienne Riviere
Visited Location :	NTNU Trondheim, Norway
Duration of Visit:	9 months

I - Scientific activity
(1 page at maximum)

At the Norwegian university of sciences of Trondheim (NTNU), I was part of the databases group, led by Prof. Nørvag. Prior to this term, I was with Prof. Felber at the distributed system and dependability group at the university of Neuchatel. During my stay at NTNU, I worked with the researchers and students of NTNU on the following topics: (1) efficient distributed computation of skyline queries on non-relocalizable data; (2) automatic discovery of conceptual relations in linked data by the means of activation mechanisms, and their representation.

During this period, I also worked on several research projects with colleagues abroad, including those I met during my first term in Switzerland: (1) gossip-based data dissemination; (2) network awareness for epidemic protocols; (3) distributed protocols evaluation frameworks (Splay); (4) distributed system support for information retrieval and recommendation engines (information retrieval).

I will briefly sketch the work that has been done at and with people from NTNU below.

1. Efficient distributed computation of skyline queries on non-relocalizable data: SkyLine queries are best-compromise queries amongst data elements that are positioned in discrete naming or attribute spaces. For instance, in a database, if all elements have two attributes x and y, an element is part of the result of the SL query iff this element has a better (e.g. smaller) value than all elements that have a strictly better value than itself on all other attributes. Efficiently supporting these kind of queries in a centralized setting is no longer a research problem, but their support in large scale distributed systems remain a challenge. A sub-class of the problem is when the data elements are present at different locations/servers but where they cannot be moved from one such node to the other. It is therefore necessary, to avoid unnecessary computations and to parallelize partial skyline computation at the storage sites, to plan the query execution and reduction. We worked with João da Rocha and Kjetil Nørvag on adaptive histogram, based on adaptive B-trees, for representing both SkyLine queries results and query execution plan. This work has not led to publications.
2. Automatic discovery of conceptual relations in linked data by the means of activation mechanisms, and their representation: the Web is composed of millions of pages in which one navigates by the means of link, defining naturally a directed graph. Particularly, sites such as Wikipedia directly map concepts (definitions) to pages, and links between such pages usually represent semantic relations. We worked with Dr. Marek Ciglan on the use of activation techniques onto the Wikipedia semantic graph for automatically discovering relations between concepts given by a user (i.e. chains of concepts linking one concept to the other, or which has common precedence with both input concepts). Moreover, we worked on the proper representation (graphical mapping onto 2 dimensions graphs) of such set of relations.

II- Publication(s) during your fellowship

LSDS-IR 2008: Regular paper

Managing Collaborative Feedback Information for Distributed Retrieval
Pascal Felber, Toan Luu, Martin Rajman, Étienne Rivière

Despite the many research efforts invested recently in peer-to-peer search engines, none of the proposed system has reached the level of quality and efficiency of their centralized counterpart. One of the main reasons for this inferior performance is the difficulty to attract a critical mass of users that would make the peer-to-peer system truly competitive. We argue that decentralized search mechanisms should not aim at replacing existing engines, but instead complement them by adding novel functionalities that would be difficult to provide in a centralized manner. This paper introduces an example of such a complementary search mechanism and presents the design of a distributed collaborative system for leveraging user feedback and document/user profiling information.

IEEE P2P 2008: Demo paper

P2P experimentations with SPLAY: from ideas to deployment in 30 minutes
Lorenzo Leonini, Etienne Rivière and Pascal Felber.

SPLAY is an integrated system that facilitates the complete chain of distributed systems evaluation, from design and implementation to deployment and experiments control. Algorithms are expressed in a concise yet very efficient language based on Lua. Implementations in SPLAY are highly similar to the pseudo-code usually found in research papers. SPLAY eases experiment deployment of any kind of testbeds, e.g., PlanetLab ModelNet clusters, or non-dedicated platforms such as networks of workstations. Using SPLAY and PlanetLab, this demonstration highlights the complete evaluation chain of an epidemic protocol and a churn-driven experiment using the Pastry DHT.

USENIX NSDI 2009: Regular paper

SPLAY: Distributed Systems Evaluation Made Simple (or how to turn ideas into live systems in a breeze)
Lorenzo Leonini, Etienne Rivière and Pascal Felber.

This paper presents SPLAY, an integrated system that facilitates the design, deployment and testing of large scale distributed applications. Unlike existing systems, SPLAY covers all aspects of the development and evaluation chain. It allows developers to express algorithms in a concise, simple language that highly resembles pseudocode found in research papers. The execution environment has low overheads and footprint, and provides a comprehensive set of libraries for common distributed systems operations. SPLAY applications are run by a set of daemons distributed on one or several testbeds. They execute in a sandboxed environment that shields the host system and enables SPLAY to also be used on non-dedicated platforms, in addition to classical testbeds like PlanetLab or ModelNet. A controller manages applications, offering multi-criteria resources selection, deployment control, and churn management by reproducing the system's dynamics from traces or synthetic descriptions. SPLAY's features, usefulness, performance and scalability are evaluated using deployment of representative experiments on PlanetLab and ModelNet clusters.

Computer Networks journal Volume 53, Issue 13, 28 August 2009, Pages 2304-2320 (Elsevier)

Rappel: Exploiting Interest and Network Locality to Improve Fairness in Publish-Subscribe Systems
Jay A. Patel, Etienne Rivière, Indranil Gupta, Anne-Marie

In this paper, we present the design, implementation and evaluation of Rappel, a peer-to-peer feed-based publish-subscribe service. By using a combination of probabilistic and gossip-like techniques and mechanisms, Rappel provides noiselessness, i.e., updates from any feed are received and relayed only by nodes that are subscribers of that feed. This leads to a fair system: the overhead at each subscriber node scales with the number and nature of its subscriptions. Moreover, Rappel incurs small publisher and client overhead, and its clients receive updates quickly and with low IP stretch. To achieve these goals, Rappel exploits "interest locality" characteristics observed amongst real multi-user multi-feed populations. This is combined with systems design decisions that enable nodes to find other subscribers, and maintain efficient network locality-aware dissemination trees. We evaluate Rappel via both trace-driven simulations and a PlanetLab deployment. The experimental results from the PlanetLab deployment show that Rappel subscribers receive updates within hundreds of milliseconds after posting. Further, results from the trace-driven simulator match our PlanetLab deployment, thus allowing us to extrapolate Rappel's performance at larger scales.

LSDS-IR 2009: Regular paper

Peer-to-Peer clustering of Web-browsing users.

Patrizio Dazzi, Pascal Felber, Le Bao Anh, Lorenzo Leonini, Matteo Mordacchini, Raffaele Perego, Martin Rajman, Étienne Rivière

For most users, Web-based centralized search engines are the access point to distributed resources such as Web pages, items shared in file sharing-systems, etc. Unfortunately, existing search engines compute their results on the basis of structural information only, e.g., the Web graph structure or query-document similarity estimations. Users expectations are rarely considered to enhance the subjective relevance of returned results. However, exploiting such information can help search engines satisfy users by tailoring search results. Interestingly, user interests typically follow the clustering property: users who were interested in the same topics in the past are likely to be interested in these same topics also in the future. It follows that search results considered relevant by a user belonging to a group of homogeneous users will likely also be of interest to other users from the same group. In this paper, we propose the architecture of a novel peer-to-peer system exploiting collaboratively built search mechanisms. The paper discusses the challenges associated with a system based on the interest clustering principle. The objective is to provide a self-organized network of users, grouped according to the interests they share, that can be leveraged to enhance the quality of the experience perceived by users searching the Web.

ALGOTEL 2009: Regular paper (national peer-reviewed conference)

PULP: Un protocole épidémique hybride
Pascal Felber, Anne-Marie Kermarrec, Lorenzo Leonini, Etienne Rivière et Spyros Voulgaris

Les protocoles épidémiques offrent une solution simple et robuste pour la diffusion d'informations dans les systèmes répartis à grande échelle. Ils utilisent des communications

périodiques entre couples de pairs, choisis de manière aléatoire. Leur efficacité et leur robustesse résultent de la multiplicité des échanges et de communications redondantes. Néanmoins, les algorithmes proposés jusqu'à présent tiennent peu compte de la charge sur le réseau, induite par les communications redondantes et par les opérations périodiques, et des compromis implicites entre cette charge et l'efficacité de la diffusion. Cet article décrit PULP, un protocole qui combine l'efficacité des opérations de type fournir (push) et la légèreté de communications de type demander (pull) dont la fréquence s'adapte au rythme de publication.

SSS 2009: Regular Paper

Network-Friendly Gossiping
Sabina Serbu, Étienne Rivière, and Pascal Felber

The emergence of large-scale distributed applications based on many-to-many communication models, e.g., broadcast and decentralized group communication, has an important impact on the underlying layers, notably the Internet routing infrastructure. To make an effective use of network resources, protocols should both limit the stress (amount of messages) on each infrastructure entity (routers, links), and balance as much as possible the load in the network. Most protocols use application-level metrics such as delays to improve efficiency of content dissemination or routing, but the extend to which such application-centric optimizations help reduce and balance the load imposed to the infrastructure is unclear. In this paper, we elaborate on the design of such network-friendly protocols and associated metrics. More specifically, we investigate randomness-based gossip dissemination. We propose and evaluate different ways of making this representative protocol network-friendly, while keeping its desirable properties (robustness and low delays). Simulations of the proposed methods using synthetic and real network topologies convey and compare their abilities to reduce and balance the load while keeping good performance.

Submitted to P2P Networking and Applications journal (Springer), under reviewing.

Joining forces in gossip-based dissemination: PULP, A Pull-Push protocol
Pascal Felber, Anne-Marie Kermarrec, Lorenzo Leonini, Etienne Rivière, Spyros Voulgaris

Gossip-based protocols are well known to provide a simple, scalable and extremely robust way to disseminate messages in large-scale systems. In such protocols, messages are spread in an epidemic manner so that, with high probability, they are received by all nodes in the system. Gossip may take place using push, pull, or any combination of push/pull techniques between the nodes. Yet, so far practical systems have focused either on push-only or pull-only approaches, with the former ones dominating. Push-based systems achieve reasonable latency and high resilience to failures but impose a large, often unnecessary, redundancy and overhead in the system. At the other extreme, pull-based protocols impose a lower overhead to the network at the price of increased latencies. A few hybrid approaches have been proposed—typically pushing control messages and pulling data—to avoid the redundancy of high-volume content such as video streaming. Yet, no pure hybrid scheme combining push and pull applied to the same content has been proposed so far. In this paper, we propose a generic efficient push-pull dissemination protocol, PULP, which combines the best of both worlds. PULP exploits the reliability and efficiency of push approaches, while limiting redundant messages and therefore imposing a low overhead, as pull protocols do. PULP relies on the dissemination of a stream of messages to achieve this: by pushing some messages, PULP enables an efficient pulling of other messages, which in turn help with the dissemination of the stream. We have deployed PULP both on a cluster and on PlanetLab. Our results demonstrate that PULP achieves an appealing trade-off between robustness, redundancy, and delays.

Submitted to WWW 2010 (World Wide Web conference)

Collaborative Ranking and Profiling: Exploiting the Wisdom of Crowds in Tailored Web Search
Pascal Felber, Lorenzo Leonini, Toan Luu, Martin Rajman, Etienne Rivière

Popular search engines essentially rely on information about the structure of the graph of linked elements to find the most relevant results for a given query. While this approach is satisfactory for popular interest domains or when the user expectations follow the main trend, it is very sensitive to ambiguous queries, that is, queries that have answers over several different domains. Elements pertaining to an implicitly targeted interest domain of overall lower popularity are usually ranked lower than expected by the user. This is a consequence of the poor use of user-centric information in search engines. Leveraging semantic information can help avoid such situations, by proposing complementary results that are carefully tailored to match user interests. This paper proposes a collaborative search companion system, CoFeed, that collects user searches and access feedbacks to build user- and document-centric profiling information. Over time, the system constructs ranked directories of elements that maintain the required information diversity and enhance the user search experience, by presenting additional results tailored to the user interest space. This collaborative search companion requires a supporting architecture adapted to large user populations generating high request loads. To that end, it integrates mechanisms for ensuring scalability and load balancing of the service under varying loads and user interest distributions. Experiments with a deployed prototype highlight the efficiency of the system by analyzing improvement in search relevance, computational cost, scalability and load balance.

Submitted to ECIR 2010 (32nd European Conference on Information Retrieval)

Visualizing the connections in Wikipedia
Marek Ciglan, Étienne Rivière, and Kjetil Nørvåg

Visualization is often helpful to communicate complex information to users. In our recent work, we exploit Wikipedia link graph as a knowledge base for nd interesting connections between set of input concepts. The results have a form of a set of connections, which is often quite hard to communicate using plain text. We have therefore proposed a method for visualization of those connections; it depicts our results in a compact and user friendly way and visualize the strength of the presented relations at the same time. The paper makes two contributions; we describe modi cation of standard Spreading Activation algorithm, that allows us to discover connections between concepts and we describe how to use of spring-mass relaxation technique to visualize semantic connectedness between concepts.

We are currently writing a second paper on the activation mechanism itself, probably for submission to SIGIR, the major IR conference (with Marek Ciglan and Kjetil Nørvag).

III -Attended Seminars, Workshops, and Conferences

Please identify the name(s), date(s) and place(s) of the events in which you participated during your fellowship period.

Only the conferences that took place during the stay in Norway are presented. See the administrative report for a comprehensive list of activities, conferences, invited talks etc.

- CicLing 2009, Mexico city, March 1 to 7, 2009

- IPTPS 2009, Boston, US, April 21, 2009
- NSDI 2009, Boston, US, April 22 to 24, 2009
- Algotel 2009, Carry-le-rouet, France, June 16 to 19, 2009

IV – Research Exchange Programme (12 month scheme)

Please identify the name(s), date(s) and place(s) of your Research Exchanges during your fellowship period and detail them .

Not applicable.