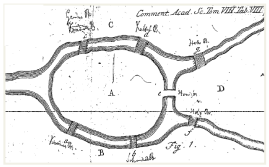

THE EULER NEWSLETTER

In this issue

Are we hitting the memory wall?

The foundational principles of the Internet routing system are i) distribution (local computation of the routing table entries), ii) adaptivity to topology and policy dynamics, and iii) policing (including decision process, routing updates filtering, etc.). Whereas static routing aims at optimizing the memory space required to store the routing table entries and the stretch of the routing paths, dynamic routing introduces a third dimension, referred to as the adaptation cost. Altogether the most fundamental challenges faced nowadays by the architecture of the Internet routing systems are i) *scalability*: the memory space consumption by the routing table entries stored locally at each node part of the routing system, ii) *convergence time*: the time required for each local routing table entryto reach a new stable state upon occurrence of a external and/or internal perturbation event, and iii) *adaptation cost*: which combines the number and rate of routing information exchanges between individual nodes (referred to as communication cost) together with their processing (referred to as processing cost) for the local routing function to properly operate.

These challenges result from i) the increasing number of routing table entries (and thus routing states) amplified by the design and usage of the addressing system (including prefix de-aggregation practices for traffic engineering purposes, and site multi-homing), ii) the short-term topology and policy dynamics and iii) the longer-term topology evolution (in particular, the increasing meshedness of the topology initially structured hierarchically). Their combination together with the intrinsic limits of the Border Gateway Protocol architecture and its underlying properties lead to a critical question: as the current growth rate of the DRAM memory capacity (55% per year) is lower than the routing table size increase (20% per year) can we eliminate the memory space dimension out of the problem at hand?

We argue that the effects of memory access time (which improves between 5 to 10% per year) should be taken into account in dynamic routing as the latter influences the convergence time and that memory space consumption should be considered as a system design constraint (instead of a design-independent performance objective). We also argue that minimizing the memory space consumption without accounting for the dependency of the dynamic routing scheme on the spatial and temporal properties of the information/input and its running conditions may lead to detrimental effects. Indeed, minimizing the memory space consumption under stationary conditions may lead up to the point that the adaption cost and convergence time objectives become unachievable but also limit the benefits of CPU cache memory.

The following quantitative estimation provides the baseline of our reasoning and subsequent arguments in favor of these fundamental observations:

- The combined effects of increasing node degree (thus increasing routing adjacencies), address prefix de-aggregation, site and domain multi-homing but also growth of the number of address prefixes leads to a growth of the routing table of about 20% per year (at about 450k entries end of 2012). Even if the routing table size would grow by 50% every year, the DRAM memory capacity growth (55% per year) would still be able to accommodate such increase. Moreover, the ratio routing table growth (1.2)/memory capacity growth (1.55) indicates that improvement is theoretically achievable over time.
- Decreasing the computation time by increasing the CPU speed further increases the memory wall effect as the rate of improvement of CPU speed exceeds the rate of improvement in DRAM speed; once there are more than 8 processor cores conventional memory architecture slows down application performance (for instance, with 64 cores time performance decreases by a factor 4 compared to 16 cores processors) because of the lack of memory bandwidth as well as contention between cores over the memory bus available to each core.
- The memory access time improves only by 5 to 10% per year, leading to a theoretical degradation of the convergence time up to 10% per year (routing table growth/memory access time improvement). Increasing the size of the cache memory (which is up to 100x faster than access to the main memory) so that it can act as main memory and in turn mitigate the memory wall effects by using for instance, the Intel Itanium processor has 54MB of cache (but 10x more expensive than conventional CPUs). However, this solution is of limited use for data intensive applications (as the entire dataset can't fit entirely in the cache).

Compared to the actual memory capacity consumption, it turns out that the memory access time becomes a major concern to avoid degradation of the routing scheme performance that is exacerbated when the CPU speed reaches its maximum with respect to time performance. Resolving the problem with conventional computer architecture would in turn require storing the entire routing information bases in fast memory caches (assuming 10x more expensive CPUs could be afforded). Consequently, memory space consumption should be considered in dynamic routing as a system-dependent constraint and convergence time as the main performance objective together with the adaptation cost.

How the EULER project is progressing?

The EULER project, started on October 1st, 2010, completed the fifth quarter of the 3-year project duration.

A plenary project meeting took place on 10-12 December 2012 at UCL, Louvain-la-neuve, Belgium. It consisted of two-days project meeting and the second Techical Advisory Board (TAB).

The second TAB meeting has been organized on December 11, 2012 at UCL, Louvain-la-neuve, Belgium with invited members Jerome Galtier (Orange Labs), Nicolai Leymann and Thomas Beckhaus (both from Deutsche Telecom). The goal of this second meeting was to discuss the operators' interest in the project proposals in terms of Internet routing. On the one hand, the discussion focuses on the improved BGP routes selection process where routes stability is incorporated as a new attribute. On the other hand, we debate the needs of a multicast routing in inter-domain environment.

The first full day has been dedicated to reviewing the second audit meeting, discussing the project status and planning the last year project. The discussions mainly focus on the methodology for the experimental activities for both the simulation and emulation prototyping. Amomng the activities planned for the last year, it is worth to mentioning two activities. First, the preparation of the **Measurement Tool Repository** where researchers will be able to download/upload tools and data sets for experimental-driven activities. Second, we discussed the current status of the **EULER Summer School 2013**. A brief summary of this event is presented in the second column and full details will be included in Newsletter No.9.

The second day of the plenary meeting focused on preparing the experimental demonstrations and the forthcoming events. In particular, EULER will propose a **demo at the Hands On FIRE** to be held during the FIA Dublin 2013 (May 8-10th, 2013). This demo will be remotely conducted in the Virtual Wall experimental facilities (see Newsletter No.7), and will show the operation of the CGMR multicat routing scheme when performing local/global discovery of an active multicast distribution tree (e.g., a multicast video streaming session), and dynamic join/leave procedures.

Forthcoming EC and FIRE events

FIRE at Celtic-Plus event **06-07/03/2013**
http://www.celticplus.eu/Events/Event-Kayseri-2013/default.asp
Kayseri, Turkey

Future Internet Assembly (FIA) **08-10/05/2013**
http://www.fi-dublin.eu/
Dublin, Ireland

Future Network & Mobile Summit 2013 **03-05/07/2013**
http://www.futurenetworksummit.eu/2013/
Lisbon, Portugal

EULER related publications

Publications

S. Sahhaf, W. Tavernier, D. Colle, M. Pickavet, P. Demeester, "Link failure recovery technique for greedy routing in the hyperbolic plane", Computer Communications, Sept. 2012.

D. Papadimitriou, L. Fàbrega, P. Vilà, D. Careglio, P. Demeester, "Measurement-based research: methodology, experiments and tools", ACM SIGCOMM Computer Commun. Review, 42(5):62-68, Oct. 2012.

D. Papadimitriou, A. Cabellos-Aparicio, F. Coras, "Stability metrics and criteria for path-vector routing", in Proc. of ICNC 2013, San Diego (CA), USA, Jan. 2013.

Deliverables

D2.2, "Routing scheme design and specifications", September 2012

D3.3, "Graph analysis/mining", September 2012

D3.4, "Measurement data analysis/mining", September 2012

D5.3, "Dissemination Intermediate report", September 2012

Summer School 2013
Graph and routing dynamics: models and algorithms

A summer school will be sponsored by EULER project and will be held on July 1-5, 2013 in Barcelona, Catalunya, Spain. This event focuses on current research and related challenges on Internet routing paradigms for distributed and dynamic routing schemes applicable to the current Internet and its evolution. The goal of this summer school is i) to stimulate research in the interdisciplinary area that lies at the intersection of graph theory, distributed routing algorithmic and network dynamics modeling, and ii) to provide a forum for active discussions among teachers/researchers and students.

This intensive summer school is addressed to MSc, PhD and Post-Docs but also researchers with a background in algorithmic graph theory, graph dynamics modeling, routing models and algorithms as well as proficiency in English. The Summer School course is expected to provide an excellent preparation for entering a PhD in Computer Science.

A full-week (5 days, July 1-5) duration is forecast for the summer school with a target of registered participants of around 30/35 students.

This summer school will consist of 9 lectures of 3 hours each. It will be organised in three sessions:

- Session 1: Introduction to the Internet routing
- Session 2: Algorithmis graph theory and graph dynamics modelling
- Session 3: Routing models and algorithms

The daily schedule will be:

- 09h30 - 13h00 Morning talk (30' break in the middle)
- 13h00 - 18h00 Lunch break
- 14h30 - 18h00 Afternoon talk (30' break in the middle)

One afternoon will be dedicated to students works with oral presentations and a session poster. A social event will be held on July 3 with a city sightseeing and a visit to a museum.

For additional information see EULER wiki at the following web: **http://bit.ly/TMZyTD**.

Call for papers

12nd IFIP TC6 Int. Conf. Networking **14/01/2013**
http://networking2013.poly.edu/
May 22-24, 2013, Brooklyn, USA

14th IEEE Conf. High Perf. Switch. Rout. (HPSR) **31/01/2012**
http://www.ieee-hpsr.org/
July 7-11, 2013, Taipei, Taiwan

19th Int. Eur. Conf Parallel Distr. Comp. (EuroPar) **31/01/2013**
http://www.europar2013.org/
August 26-28, 2013, Aachen, Germany

ACM Sigcomm **01/02/2013**
http://conferences.sigcomm.org/sigcomm/2013/index.php
August 12-16, 2013, Hong Kong

15th AlgoTel **04/02/2013**
http://algotel2013.sciencesconf.org/
May 28-31, 2013, Pornic, Loire-Atlantique, France

3rd ETSI Future Networks workshop **07/02/2013**
http://www.etsi.org/news-events/events/617-2013-future-networks
April 9-11, 2013, Sophia Antipolis, France

22nd Int. Conf. Comp. Commun. Netw. (ICCCN) **08/02/2013**
http://www.icccn.org/icccn13/
July 30-August 2 , 2013, Nassau, Bahamas

Future Network & Mobile Summit **08/02/2013**
http://www.futurenetworksummit.eu/2013/
July 3-7, 2013, Lisbon, Portugal

32nd ACM Symp. Princ. Distr. Comp. (PODC) **10/02/2013**
http://www.podc.org/
July 22-24, 2013, Montréal, Canada