December 2013

No. 11

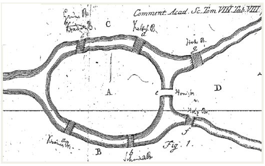

The EULER Newsletter

In this issue

Geometrizing Information-Centric Networks

In response to the increasing traffic volume in the Internet for applications such as (mobile) video and cloud computing, various proprietary technologies enabling content distribution have been developed that rely on caching and replication. Being deployed in silos, it is not possible to uniquely and securely identify named information independently of the distribution channel; moreover, these different content distribution technologies are typically implemented as an overlay, leading to unnecessary inefficiency. By introducing uniquely named data and name-based data access, Information-Centric Networking (ICN) [1] enables data to become independent from their network location, application, storage support but also means of content exchanges enabling in turn in-network caching and replication. However, content name spaces have not been designed to sustain forwarding performance and forwarders scaling contrary to IP addresses which can be efficiently aggregated, summarized and translated. Consequently, alternatives such as name-based routing have been proposed where routing function locates a content object based on its name which is initially provided by a requestor. Such process aims at better accommodating information/content routing in the Internet but would also become the scaling and performance bottleneck. Indeed, all known name-based routing approaches emphasize the major and well-known tradeoff experienced when designing routing systems: the first alternative (which omits name resolution) exacerbates the main drawback of the push model, i.e., storage capacity, and the second (which relies on the name resolution function to translate the name of the requested content object into its locator) the main drawback of the pull model, i.e., latency. The survey [2] produced under the auspices of the Information-Centric Networking Research Group (ICNRG) and the analysis of its associated challenges [3] demonstrate that all name-based routing approaches share common scaling problem.

To address this fundamental problem, the EULER project proposes a third alternative referred to as geometric information routing on universal content locators. This technique operates by associating to content identifiers (names) a content locator taken out of a geometric coordinate space from which a routing path (more precisely, a geodesic path) can be derived without requiring non-local information. Upon querying specific content multiple locators can be received enabling the receiver to select the (geometrically) closest locator. The principle underlying this alternative is thus relatively simple, perform information routing decision on locators avoiding name-to-locator resolution by intermediate nodes. Moreover, since it is based on local information only, routing on such locator space is less memory consuming than non-local information dependent routing.

The salient feature of this routing model comes from the property of coordinate-based content locators: these coordinates can be used by the distributed routing function to perform geometric routing decisions. Conventional geometric routing operates by assigning to each node virtual coordinates in a metric space (X,d) that are used as locators to perform point-to-point routing decisions in this space. Instead of assigning (virtual) coordinates and compute distances from these coordinates, geodesic geometric routing as proposed by the EULER project [4] operates by computing the distances between vertices from the length of the corresponding geodesic drawn out of negatively curved geometric space (the hyperbolic plane). It then derives the vertex coordinates from the selected geodesics. Thus, content locators substitute to network locators (stricto sensu, the routing function still performs on locators) but they can also be used in combination with other network locator spaces, e.g., IP addresses. The situation where content locators would require resolution to network locators ensures interoperability when messages are transmitted across IP-only forwarding networks. Moreover, as there is no distinction between "server" and "cache" locators, i.e., if an intermediate node determines it has to keep a local copy of a content object, it can decide to apply the registration step described here above. Subsequently, a requestor could receive the content locator associated to an intermediate node. Up to certain extent, in geodesic routing the addressing space follows the topology.

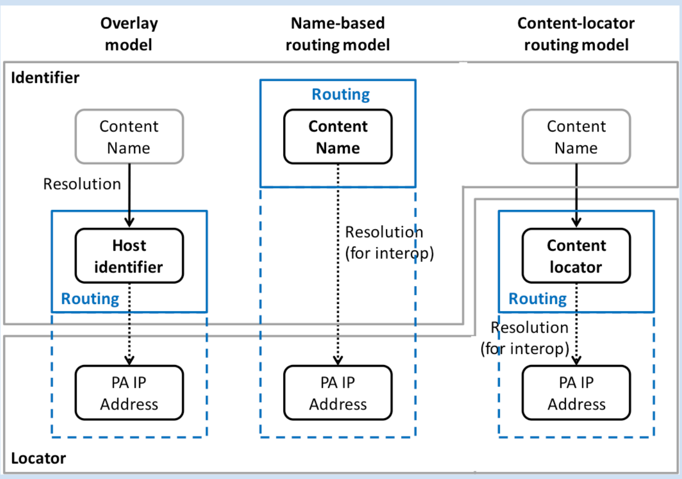

Content naming and addressing models

Comparison with the Border Gateway Protocol (BGP) shows that geodesic geometric routing provides remains competitive in terms of memory-stretch tradeoff.

- Routing state: BGP stores O(f(n)) routing states per node where the function f(n) = (n-1)! for a complete graph. Yet the topology underlying the Internet does not form a complete graph, we can thus relax this upper bound by assuming that each node accumulates from its neighbors $\nu(n-1)$ routing states per destination, where ν is the size of the neighbor set N(u) of node u, $|N(u)| = \nu$. Hence, each node stores $O(\nu n(n-1)^2)$ states. On the other hand, assuming that each geodesic routing process stores at most $O(\nu)$ states per destination, the total number of states stored per node is $O(\nu n)$.
- Memory space: assuming that the average BGP path length λ determines the size of each routing entry, the memory-space

Subscribe to electronic EULER newsletter: https://sympa.inria.fr/sympa/info/euler-news

(measured in memory-bits unit) required at each node reaches O(ν n(n-1)²λ). Assuming that each geodesic is represented by a succinct coordinate pair, geodesic routing requires O(ν n log(n)) memory bits to store locally the routing table entries.

- Stretch: for simplicity, we consider here the additive stretch which measures the difference between routing path length and the topological path length. BGP belonging to the class of shortest-path routing the additive stretch of this routing algorithm is 0. On the other hand, the additive stretch of geodesic routing upper bound is determined by δ log(n) where δ characterizes the hyperbolicity of the graph underlying the topology.

The above analysis shows (as expected) that decreasing the memory space consumption comes at the detriment of the routing path stretch. However, assuming in-network caching is enabled along routing paths, its effect would be further limited. As previously stated, it also explains the importance of proper characterization of the value δ for the network environments under consideration. Observe this upper bound fulfills the expectation of routing path stretch being polylogarithmic in the number of nodes n.

This property leads to deep implications in terms of routing path stretch, succinctness but also robustness. In particular, it elevates the fundamental drawback of conventional geometric routing as topology changes result in locator changes (a.k.a. renumbering).

Call for papers

IEEE Infocom Demo Session 2014 **15/12/2013**
http://www.ieee-infocom.org
April 27-May 2, 2014, Toronro, Canada

6th IEEE Int. W. NetSciCom **22/12/2013**
http://www.ctr.kcl.ac.uk/netscicom14/
April 27-May 2, 2014, Toronto, Canada

5th Int. W. TRAC **15/01/2014**
http://trac2014.ftw.at/
August 4-8, 2014, Nicosia, Cyprus

26th ACM Symp. SPAA **25/01/2014**
http://www.cs.jhu.edu/~spaa/
June 23-25, 2014, Prague, Czech Republic

ACM Sigcomm **31/01/2014**
http://conferences.sigcomm.org/sigcomm/2014/index.php
August 17-22, 2014, Chicago, IL, USA

IEEE Comm. Mag. **01/02/2014**
Special Issue on Disaster Resilience in Communication Networks

15th Int. Conf. HPSR **02/02/2014**
http://www.ieee-hpsr.org/
July 1-4, 2014, Vancouver, Canada

20th Int. EuroPar **06/02/2014**
http://europar2014.dcc.fc.up.pt/
August 25-29, 2014, Porto, Portugal

Forthcoming EC and FIRE events

CREW training days 2nd edition **14-15/01/2014**
http://www.crew-project.eu/trainingdays2014
Ghent, Belgium

Future Internet **18-20/03/2014**
http://www.fi-athens.eu/
Athens, Greece

FI-PPP – 1st Eur. Conf. on Future Internet **02-03/04/2014**
http://www.fi-ppp.eu/ai1ec_event/1st-european-conference-on-the-future-internet-ecfi/
Brussels, Belgium

Celtic-Plus event **23-24/04/2013**
http://www.celticplus.eu/
Montecarlo, Monaco

EULER progress report

The routing algorithmic work conducted during the third year of the project follows the routing system architecture designed during the first year of the project. After having identified the root causes for the absence of suitable alternatives to BGP, the feedback obtained from the second Technical Advisory Board (TAB) has been considered in order to understand the necessary and sufficient conditions for the development of new germs/seeds of routing processes/models taking into account operational considerations including policy. Following this analysis, our research work has been dedicated to i) possible improvements/enhancements of path-vector based routing (most of which are also applicable to BGP) and new classes of path-vector based routing, and ii) new routing paradigms (for the design of a genuine alternative to BGP) in particular for what concern multicast routing.

Regarding Improvements/enhancements to BGP, following the work initiated in Task 3.2, we have defined several stability metrics to characterize the local effects of BGP policy- and protocol-induced instabilities on the routing tables. We have also defined a differential stability-based decision criterion that can be taken into account as part of the BGP route selection process.

Regarding new routing paradigms, as described above, we have proposed geometric information routing on universal content locators. We have also developed a distributed name-independent compact routing scheme: under random coloring and synchronous hypothesis, and for weighted n-node graphs of poly-logarithmic hop-diameter D, the proposed scheme produces a stretch-7 routing algorithm but with high reduction in entries in the routing tables and less communication costs than path vector schemes. Finally, two new features have been designed for the Greedy Compact Multicast Routing (GCMR) scheme (see Newsletters No. 3 and No. 10) and implemented. One feature deals with the adaptability of the scheme, meaning that, after any change in the network (link/node failure, topology change, etc.), the scheme reacts to readjust the Multicast Distribution Tree (MDT). The second one is an Any Source Multicast (ASM) solution for GCMR. In particular, Rendezvous Points (RP) are strategically placed in the network to simplify the construction and maintenance of the MDT.

EULER selected publications

D. Papadimitriou, D. Careglio, P. Demeester, “Performance analysis of multicast routing algorithms”, in Proc. ICNC, Honolulu, USA, Feb. 2014.

D. Papadimitriou, D. Colle, P. Audenaert, P. Demeester, “Geometric Information Routing”, in Proc. 7th ANTS, Chennai, India, Dec. 2013.

D. Careglio, D. Papadimitriou, F. Agraz, S. Sahhaf, J. Perelló, W. Tavernier, S. Spadaro, D. Colle, "Development and experimentation towards a multicast-enabled Internet", Research report UPC-DAC-RR-CBA-2013-5, Dec. 2013.

C. Gavoille, C. Glacet, N. Hanusse, D. Ilcinkas, “On the communication complexity of distributed name-independent routing schemes”, in Proc. 27th DISC, Jerusalem, Israel, Oct. 2013.

C. Magnien, A. Medem, S. Kirgizov, F. Tarissan, “Towards realistic modeling of IP-level routing topology dynamics”, Networking Science, Oct. 2013.

F. Giroire, R. Modrzejewski, N. Nisse, S. Perennes, “Maintaining balanced trees for structured distributed streaming systems”, in Proc. 20th SIROCCO, Ischia, Italy, Jul. 2013.

D. Coudert, L. Hogie, A. Lancin, D. Papadimitriou, S. Pérennes, I. Tahiri, “Feasibility study on distributed simulations of BGP”, Research Report RR-8283, INRIA, Apr. 2013.

S. Sahel, W. Tavernier, D. Colle, M. Pickavet, P. Demeester, “Link failure recovery technique for greedy routing in the hyperbolic plane”, Computer Communications, vol. 36, no. 6, Mar. 2013.