

Partial Adaptive Name Information in ICN: PANINI Routing Limits FIB Table Sizes

Thomas C. Schmidt
HAW Hamburg
t.schmidt@acm.org

Sebastian Wölke, Nora Berg
HAW Hamburg
{first.last}@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlich@fu-berlin.de

ABSTRACT

Name-based routing as proposed in Information Centric Networking encounters the problems of (a) exploding routing tables, as the number of names largely exceeds common routing resources, and (b) limited aggregation potentials, as names are commonly independent of content locations. In this poster, we introduce PANINI, an approach to scale routing on names by adapting FIB tables simultaneously to available resources and actual traffic patterns. PANINI introduces routing hierarchies with respect to aggregation points, bimodal FIBs, and confined flooding. First evaluations show promising results in theory and experiments.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—*Routing Protocols*

Keywords

Scalable adaptive forwarding; NDN; confined flooding

1. INTRODUCTION

Information Centric Networking has introduced a new, promising communication paradigm, but continues to struggle with severe challenges [1]. NDN [2] (among others) unifies routing with names at a high level of maturity. However, the multitude and complexity of distributed content names has not been treated convincingly [3]. Even though several original approaches have been presented, the sheer number of (delocalized) names prevents a striking step forward.

Routing on identifiers can generally be achieved by aggregation or mapping to topology, the latter may be dynamically obtained from broadcasts. In the following, we will present a hybrid combination of (artificially enhanced) name aggregation at a rendezvous point, an adaptive (static) mapping by FIBs, and a dynamic on-demand flooding of Interests towards content suppliers. We sketch the PANINI routing scheme in Section 2 and give a brief evaluation in Section 3.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author(s). Copyright is held by the owner/author(s).

ICN'15, Sept.30–Oct. 2, 2015, San Francisco, CA, USA.

ACM 978-1-4503-3855-4/15/09

DOI: <http://dx.doi.org/10.1145/2810156.2812603>.

2. PANINI ROUTING

The PANINI approach starts from fixing an aggregation point for names in the network¹, which typically would be a larger cache repository in the fixed Internet, or a gateway in the IoT. We assume topology building mechanisms in place that generate a shortest path tree rooted at the aggregation point. This is in full analogy to the current Internet, where standard routing protocols can construct shortest paths on the inter- and intra-domain level. Given this basic topology, every node can identify up- and downward paths with respect to the aggregating root—with upward paths serving as default.

The objective of name-based routing is to link content requesters with content suppliers. In PANINI, this is achieved via the (name-specific) aggregation points that can be reached via (prefix-specific) default routes. Every node that offers content under a routable name advertises this name in a Name Advertisement Message (NAM). Per default, NAMs travel hop-by-hop towards the aggregation point, and every intermediate router can harvest the content advertisement for including in its own routing table. Filling all FIBs will generate a complete routing path from the aggregation point to the content source.

A consumer requests content by transmitting the Interest up to the aggregating root (default), and down along the previously installed paths. Data forwarding will follow the regular pending Interests of NDN on the reverse path. Routing and forwarding are thus aligned to a network hierarchy that resembles the current Internet with aggregation points located at the transit tier.²

Up to this point, we have required names in all FIBs, which is known to be infeasible in ICN. We now weaken this requirement as follows. Complete routing tables shall only be required at the aggregation points. This is a significant relaxation since aggregation points are designed to facilitate name aggregation and largely reduce routing table space. In addition, providers may select strong devices to serve as aggregation points. From complete, aggregated FIB tables, the (transit) root can thus always tell which branch (or lower tier ISP) holds the requested content. Without further FIB entries, flooding may lead the Interest down this branch.

Intermediate nodes are not required to carry a full FIB, but rather aim at adapting selected entries to minimize Interest flooding. In analogy to caching content, each node autonomously

¹Notably, multiple aggregation points for different lexicographic ranges are possible.

²Extensions to multiple transits per prefix, as well as peering shortcuts are subject of future work.

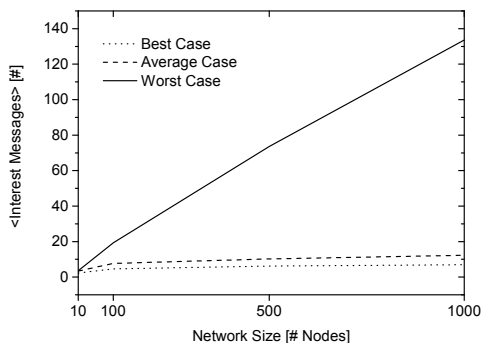


Figure 1: Average number of interest messages for different scenarios and network sizes

ously decides about (a) its memory resources available for the FIB, and (b) the forwarding logic it applies within its vicinity. Traffic flows can be continuously used to adapt the FIB to relevant traffic patterns. For example, a node can hold more specific information for frequently requested names, while it may erase entries for traffic rarely seen.

To optimize Interest guidance at partial forwarding information even further, we introduce a *bimodal FIB*. This extends the FIB structure to hold two modes—`include` and `exclude`. In `include` mode, all Interests that match a FIB prefix will be forwarded on the associate Face, while all Interests that match a FIB `exclude`-prefix will be blocked on that Face. The initial state of an empty FIB reads `include *` which leads to a transparent forwarding (flooding) of all incoming Interests. A node that has seen no routable names from NAMs in a subtree of his may as well switch to `exclude *`.

3. EVALUATION AND OUTLOOK

In our brief evaluation, we concentrate on the penalty PANINI inherits from only partially filling FIBs, i.e., flooding. In the best case, PANINI is optimal, while in the worst case, an Interest is flooded down an entire branch of the routing tree. The actual efficiency depends on (i) the available FIB sizes, (ii) the adaptation strategy of names, and (iii) the actual distribution of content within the topology.

Shortest path trees are theoretically well described by Uniform Recursive Trees (URTs) [4]. In a first evaluation step, we can use this analogy to theoretically estimate the flooding overhead inherited from the topology prior to adaptive FIB optimisations. In particular, we can derive the worst, best, and average case scenarios for flooding branches of subtrees.

A Uniform Recursive Tree of N nodes has an average depth of $\log(N)$, which is the optimal number of dntree Interest messages. In the worst case, an entire branch is flooded—on average N/H_{N-1} nodes [5], where H_N is the N -th Harmonic number. For the average scenario, we consider a random FIB on path empty and its corresponding subtree flooded. We omit the corresponding (complex) expression here and visualize the outcome in Figure 1. As can be seen from the graph, the average number of Interest messages needed for locating content follows closely the logarithmic behavior of the best case. In contrast, the worst case scenarios grow only slightly sublinearly ($\approx N/\log N$).

In a second evaluation step, we analyze the effect of popularity-based FIB adaptation. We build a (virtual) test network of 100 ICN nodes and harvest a frequency distri-

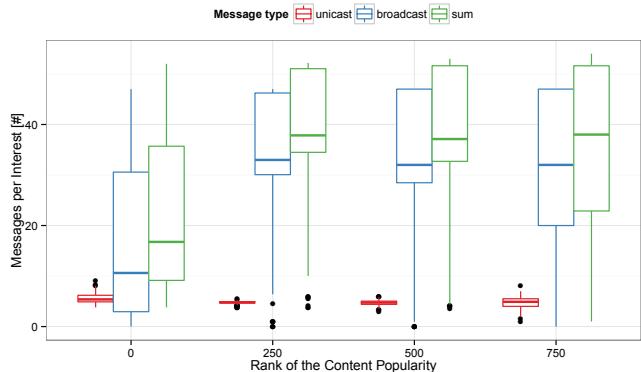


Figure 2: Measurements of (typed) message frequencies for four selected popularity ranks

bution of name usages from Quantcast³. During name advertisement from randomly selected content sources, we add names to FIBs according to these popularity distributions. Content is then requested accordingly from random nodes.

Figure 2 depicts measurement results for message frequencies per type of four selected ranks of content popularity. Within this limited experimental setup, our measurements clearly reveal for routing (i) a constant unicast message effort for all content types, and (ii) a much enhanced routing routing determinism (i.e., reduced flooding) for popular content. Given the highly skewed popularity distribution, average efforts for flooding remain close to popular content behaviour and reproduce the theoretical average scenario fairly well. In particular, we cannot encounter large deviations in single events or other indications of unexpected effects. This clearly indicates that the topological structure of routing trees (rather wide than deep) conjointly with a reasonable adaptation strategy will confine Interest flooding to rather limited subtrees.

In summary, we could show that PANINI routing is a promising hybrid approach to mitigate between FIB sizes and interest flooding for locating content. Our future work will concentrate on to elaborate and evaluate the missing details of the PANINI routing scheme. It is our intend to show its feasibility even for large-scale inter-provider set-ups.

4. REFERENCES

- [1] D. Kutscher, et al., “ICN Research Challenges,” IETF, Internet-Draft – work in progress 01, February 2015.
- [2] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, “Networking Named Content,” in *Proc. of the 5th Int. Conf. on emerging Networking EXperiments and Technologies (ACM CoNEXT’09)*. New York, NY, USA: ACM, Dec. 2009, pp. 1–12.
- [3] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, “Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure,” *Computer Networks*, vol. 57, no. 16, pp. 3192–3206, Nov. 2013.
- [4] P. Van Mieghem, *Performance Analysis of Communications Networks and Systems*. Cambridge, New York, 2006.
- [5] C. Su, Q. Feng, and Z. Hu, “Uniform Recursive Trees: Branching Structure and Simple Random Downward Walk,” *Journal of mathematical analysis and applications*, vol. 315, no. 1, pp. 225–243, 2006.

³<http://www.quantcast.com/top-sites> — this distribution is even more highly skewed than the common Zipf law.