Computer Communications 46 (2014) 76–86

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Intra-domain routing with pathlets ☆

Marco Chiesa, Gabriele Lospoto, Massimo Rimondini *, Giuseppe Di Battista

Department of Engineering, Roma Tre University, Via della Vasca Navale 79, 00146 Rome, Italy

ARTICLE INFO

Article history:
Available online 20 March 2014

Keywords:
Computer networks
Intra-domain routing protocols
Pathlet routing

ABSTRACT

Internal routing inside the network of an Internet Service Provider (ISP) affects the performance of lots of services that the ISP offers to its customers and is therefore critical to adhere to Service Level Agreements (SLAs), achieve a top-quality offer, and earn revenue. Existing technologies (most notably, MPLS) offer limited (e.g., with RSVP-TE), tricky (e.g., with OSPF metrics), or no control on internal routing paths. Recent research results address these shortcomings, but miss a few elements that would enable their application in an ISP's network.

We introduce a new control plane, based on pathlet routing (Godfrey et al., 2009) [2], designed to operate in the network of an ISP and offering several nice features: it enables steering of network paths at different levels of granularity; it is scalable and robust; it supports independent configuration of specific network regions and differentiation of Quality of Service (QoS) levels; it can nicely coexist with other control planes and is independent of the data plane used in the ISP's network. Besides formally introducing the messages and algorithms of our control plane, we propose an experimental scalability assessment and comparison with OSPF, conducted in the simulation framework OMNeT++.

1. Introduction

In a never-ending effort to offer top-quality services, Internet Service Providers (ISPs) strive to distribute traffic loads in their networks with clever strategies that not only ensure satisfaction of Service Level Agreements (SLAs), but also realize competitive performance levels that earn them market shares and, therefore, revenue. Fine-grained control of internal routing paths is essential to achieve these goals, and several technologies (e.g., OSPF, RSVP) have been introduced and widely deployed to gain this control. However, their complexity of setup, scarce predictability of dynamic behavior, and limited degree of control of routing paths pushed the research community to seek for alternative solutions along approaches like source routing, multipath routing, and hierarchical routing. To the extent of our knowledge, none of these solutions yet succeeded in combining these approaches to obtain routing control while supporting other features that ISPs yearn for, like configuration simplicity, robustness, compatibility with deployed routing mechanisms, and Quality of Service (QoS) differentiation, to mention a few.

In this paper we propose the design of a new control plane for internal routing in an ISP's network which aims at achieving these goals. Our control plane supports control of routing paths at different levels of granularity, envisions several kinds of routing policies, and allows computation of multiple paths for resilience and, possibly, QoS differentiation. It reacts efficiently to topological changes and administrative reconfigurations, enables administrators to independently configure different network portions, and it can be incrementally deployed. We build our control plane on top of pathlet routing [2], one of the most convenient approaches introduced so far to tackle an ISP's requirements. By integrating this contribution and combining it with other suitably adapted approaches from the literature, we define a complete pathlet-based routing solution that is applicable to intra-domain routing, filling a gap that, as far as we know, is still open.

In the control plane we propose, routers exchange path fragments called *pathlets* and are grouped into *areas*: within a single area routers exchange all information about the available links, in a much similar way to what a link-state routing protocol does; when announced outside the area, such information is summarized in a single pathlet that goes from an entry router for the area directly to an exit router, without revealing routing choices performed by routers that are internal to the area. This special pathlet, which we call *crossing pathlet*, is considered outside the area as if it were a single link. An area can enclose other areas, thus forming a hierarchical structure with an arbitrary number of levels.

☆ A preliminary version of this paper appeared in [1].

* Corresponding author. Tel.: +39 06 57333215; fax: +39 06 57333612.
E-mail addresses: chiesa@dia.uniroma3.it (M. Chiesa), lospoto@dia.uniroma3.it (G. Lospoto), rimondin@dia.uniroma3.it (M. Rimondini), gdb@dia.uniroma3.it (G. Di Battista).

http://dx.doi.org/10.1016/j.comcom.2014.03.011
0140-3664/

The rest of the paper is organized as follows. In Section 2 we review the state of the art on routing mechanisms that could match the requirements of ISPs. In Section 3 we introduce a model for a network where nodes are grouped in a hierarchy of areas. Based on this model, in Section 4 we define the mechanisms for the creation and dissemination of pathlets in the network. We then describe in Section 5 how network dynamics are handled, including the specification of the messages of our control plane and of the algorithms executed by network nodes to update routing information. In Section 6 we elaborate on the practical applicability of our control plane in an ISP's network in terms of possible deployment technologies and illustrate an incremental deployment scenario. In Section 7 we present an experimental assessment of the scalability of our approach and compare its performance with those of OSPF, using the OMNeT++ simulation framework. Conclusions and plan for future work are presented in Section 8.

2. Related work

Many contributions in the literature propose methodologies, algorithms, and protocols that address the scalability, robustness, and controllability requirements faced by an ISP in managing its network. Commonly adopted approaches to satisfy these requirements include source routing, hierarchical routing, and multipath routing. For example, hierarchical routing has for long been known to be provably effective in reducing the size of routing tables [3]. On the other hand, multipath routing is widely used in sensor networks [4], where reachability of the various nodes must be guaranteed even under frequent connectivity variations.

However, none of the contributions we are aware of succeeds in proposing a complete routing solution that fits the requirements of an ISP in an intra-domain scenario: either they apply to inter-domain routing, where the degree of control offered by the available technologies, as well as the goals that ISPs are interested in pursuing, are different than those we focus on, or they fail to address some basic requirements, most notably simplicity of setup or compatibility with already deployed configurations and technologies. We now review the state of the art on the most relevant control plane mechanisms, using Table 1 as a reading key to classify the contributions we mention.

In terms of technologies, OSPF [5] is the state of the art for interior routing and has a wide deployment base. However, it offers limited control of routing paths, because they can only be affected by assigning costs and it is very hard to influence a single path without affecting others; it imposes restrictions on the configuration of areas, because they must adhere to a precise structure with a single backbone and multiple stubs/transits; it is not designed to support source routing; and it has limited options to handle multiple alternative paths, typically consisting in a set of possible load balancing policies. Although not a true routing protocol, RSVP [6] has been conceived with traffic engineering in mind, and yet it shares many of the shortcomings mentioned for OSPF.

MIRO [7] is a routing solution that supports the negotiation of multiple routing paths to satisfy the diverse requirements of end users, but no complete control can be enforced on these paths. A similar drawback is shared by path splicing [8], a mechanism designed to realize fault tolerance (see also [18]): it exploits multipath routing to ensure connectivity between network nodes as long as the network is not partitioned, but actual routing paths are not exposed and cannot therefore be controlled. The route discovery mechanism envisioned in NIRA [9] makes routing paths more controllable, but this solution is designed only for an inter-domain routing architecture, like MIRO, and it relies on a constrained address space allocation, a hardly feasible choice for an ISP that is taken also by Landmark [10]. Slick packets [11] achieves a combination of fault tolerance and source routing, obtained by encoding in the forwarded packets a directed acyclic graph of different alternative paths to reach the destination. Besides the intrinsic difficulty of this encoding, this solution inherits the limits of the dissemination mechanisms it relies on: NIRA or pathlet routing (discussed below). BGP Add-Paths [12] and YAMR [13] also address resiliency by announcing multiple routing paths selected according to different criteria, but they only adopt multipath routing, they offer very limited or no support for hierarchical routing, and they have some dependencies on the BGP technology. A completely different approach is taken by HLP [14], which proposes a hybrid routing mechanism based on a combination of link-state and path-vector protocols. In this paper the authors present an in-depth discussion of the routing policies that can be implemented in such a scenario. Although HLP matches more closely our approach, it is not conceived for internal routing in an ISP's network, it constrains the hierarchical network structure to reflect inter-ISP agreements, and it has limits on the configurable routing policies. A similar hybrid routing mechanism called ALVA [15] offers more flexibility but, like Macro-routing [16], it does not explicitly envision source routing and multipath routing. HDP [17] is a variant of this approach that, although natively supporting Quality of Service and traffic engineering objectives, is closely bound to MPLS and accommodates source routing and multipath routing only in the limited extent allowed by this technology. Some contributions, like LIPSIN [19], adopt a completely different routing approach based on Bloom filters to gain efficiency. However, these solutions are more oriented to multicast forwarding and do not offer a complete control on routing paths because they are based on a probabilistic model. Pathlet routing [2] is definitely the contribution that is closest to our control plane approach, because it introduces a data plane that supports a very flexible handling of routing paths. Its most evident drawback is the lack of a completely defined mechanism for the dissemination of pathlets, which the authors only hint at. Our control plane approach, which is based on pathlet routing, shares some routing principles with those adopted in wireless sensor and mobile networks: among the others, the existence of clusters (which we call areas) and the selection of routing paths based on some quality metrics. However, there are some differences. For example, while energy constraints, handover mechanisms, and evolution of network clusters are not a concern in the scenario we consider, we provide any vertex in the network with enough information to perform source routing, removing the need for central nodes that hold forwarding information (e.g., Cluster Heads).

As a general remark, previous contributions highlight how path-vector protocols typically support complex information hiding and path manipulation policies, whereas link-state protocols typically offer fast convergence with a low overhead. Therefore, a suitable combination of the two mechanisms, which is considered in our approach, should be pursued to inherit the advantages of both.

Table 1
A classification of the state of the art according to the adoption of some relevant routing techniques.

	Source routing	Hierarchical routing	Multipath routing
MIRO [7]	Limited	No	Yes
Path splicing [8]	Limited	No	Yes
NIRA [9]	Yes	No	No
Landmark [10]	No	Yes	Yes
Slick packets [11]	Yes	Limited	Yes
BGP Add-Paths [12]	No	No	Yes
YAMR [13]	No	Limited	Yes
HLP [14]	No	Limited	No
ALVA [15]	No	Yes	Limited
MACRO [16]	No	Yes	No
HDP [17]	Limited	Yes	No