

Towards a Scalable Routing Architecture for Dynamic Substrates

Boris Drazic
University of Toronto, Canada
boris@comm.utoronto.ca

Jörg Liebeherr
University of Toronto, Canada
jorg@comm.utoronto.ca

ABSTRACT

We present a novel routing architecture for diverse collections of substrate networks with a mix of mobile and fixed nodes. A key element of the architecture is to utilize less volatile parts of the network as landmarks, and use path information from the landmarks as node locators. Our ongoing research seeks to improve the scalability of this routing approach.

1. INTRODUCTION

Data networks are typically built using an overlay approach, where new network services are build over one or more existing (underlying) substrate networks. As case in point, the Internet is an overlay which provides end-to-end paths over a vast collection of link layer substrate networks. A tenet of routing in the Internet is that substrates are static in the sense that, barring failures or explicit configurations, the number and positions of substrates in the larger collection does not change. Also, each substrates provides bidirectional paths between all nodes in the substrate. While many overlay concepts were developed since the Internet, the basic assumption of static substrate networks has not changed. In particular, recent application-layer overlays for network services assume the Internet a permanently available and universally accessible (thus, static) substrate network.

However, routing architectures for static substrate networks are not well suited for interconnected collections of substrates with a large fraction of mobile nodes. Our research seeks to find new addressing and routing approaches for interconnecting networks where a high degree of mobility of nodes causes entire substrate networks to become dynamic. Our research seeks to improve the ability of interconnecting infrastructure networks with dynamic substrates, such as,

Permission to make digital or hard copies of all or part 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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CoNEXT Student Workshop, December 6, 2011, Tokyo, Japan.
Copyright 2011 ACM 978-1-4503-1042-0/11/0012 ...\$10.00.

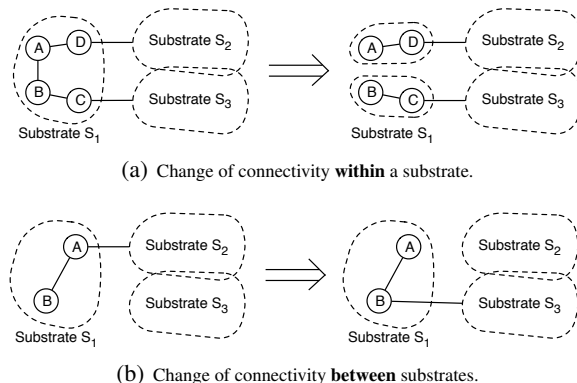


Figure 1: Dynamic substrates.

ad-hoc networks, community networks, Bluetooth piconets, or vehicular networks.

2. REACHABILITY DOMAINS

We view a substrate as a network that enables exchange of data between network nodes with compatible attachment points when a path of links between them exists in the network. Substrates can exist at the data link layer, e.g., Bluetooth or WiFi networks with compatible configurations, the network layer, e.g., Internet or network-layer ad-hoc networks, or the application layer. Substrate networks may include relay devices, e.g., access points in WiFi, routers or middleboxes for IP based substrates.

Dynamic substrates are characterized by exhibiting two types of frequent changes of the connectivity between nodes: (1) *Internal* – changes to the availability of end-to-end paths between nodes in the same substrate, and (2) *External* – changes to the connectivity between substrates. For illustration, refer to Fig. 1. In Fig. 1(a), four nodes in Substrate S_1 become partitioned into two groups without an internal path between the groups. In Fig. 1(b), nodes in substrate S_1 lose attachment to Substrate S_2 , and become connected to Substrate S_3 .

Devising addressing and routing schemes for dynamic substrate networks is challenging. First, nodes need to detect maximal sets of nodes reachable via a path of nodes within a substrate network. We refer to such sets as *reachability domains*. (In this sense, Fig. 1(a) depicts the creation of an




Figure 2: Routing with landmark domains.

additional reachability domain). A second, more difficult, problem is to locate and determine paths to remote reachability domains. A straightforward approach, which assigns a label to each reachability domain, and performs min-cost routing between the reachability domains is not viable, as, in dynamic substrates, the set of reachability domains may change faster than routing tables can stabilize.

To address these challenges, we propose a routing architecture built following the theoretical concepts of compact routing [?], which provide bounds on maximum path stretch and storage space per node. As compact routing was developed for static networks, a key insight when dealing with dynamic substrates is to abandon trying to create globally consistent routing tables. Rather, we take advantage of the fact that some substrates and reachability domains are more stable than others. In particular we designate the least volatile reachability domains as *landmark domains* and define them as reachability domains to which all nodes have routing paths, and only concern ourselves with routing to and from the landmarks. In our scheme, the routing information on the reverse path from a node to one or more landmarks serves as the locator address(es) of the node. With this, nodes merely need to deliver a message to a landmark using next-hop forwarding. We refer to this scheme as *landmark domain routing*. As nodes move and change their attachment points to substrates, they change their locators. A node can determine its locator by routing a message to a landmark domain and a suitable distributed hash table (DHT) realization, e.g., [?], can be used for locator lookup by asking the desired node (using its globally unique and persistent node identifier) for its locator. Note that locators can grow large, if the density of landmarks is too low. A hierarchical organization of landmarks can ameliorate the issue.

3. LANDMARK DOMAIN ROUTING

We present a simple example in Fig. 2, which depicts a network with six reachability domains, RD_1, \dots, RD_6 , where RD_4 is a landmark domain. Consider that node A wants to send a message to K . A has performed a lookup of K 's locator, consisting of the partial path $H \rightarrow J \rightarrow K$. Forwarding of the message is done in three steps: (1) Next-hop forwarding delivers the message from A to G , i.e., to a node in the landmark domain; (2) In the landmark domain, the message is forwarded on a (necessarily available) path from G to H ; Finally, (3) using the locator of K , the mes-




Figure 3: Path stretch.




Figure 4: Control messages.

sage is forwarded to its destination.

4. PRELIMINARY RESULTS

We perform an evaluation using a packet level simulator. Our first objective is to measure the viability of the routing paths and the incurred overhead. We present results for a sufficiently randomized network with up to 70 reachability domains and with 100 nodes in each substrate. A reachability domain is set as landmark domain with a probability of 0.2. In Fig. 3, we present the stretch of paths, i.e., the ratio to the lengths of min-hop paths. On average, routing paths are less than 50% longer than best possible paths. In Fig. 4, we show the total number of routing messages needed to establish consistent routing tables of paths to all landmark domains. While the overhead is limited, the linear increase of the overhead in the number of reachability domains indicates the need of a hierarchical organization of landmarks.

5. CONCLUSION AND FUTURE WORK

Introducing the concept of reachability domains, we have defined an addressing and routing scheme for interconnecting collections of dynamic and static substrates. The scheme abandons the goal of creating globally consistent routing tables, and, instead, is based on finding routes to and from more stable regions of the network. Future efforts seek to organize landmarks to achieve an overhead of control messages which grows at most logarithmically with the size of the network.

6. REFERENCES

- [1] B. Ford. Unmanaged Internet Protocol: Taming the Edge Network Management Crisis. *SIGCOMM Comput. Commun. Rev.*, 34:93–98, January 2004.