Physical Communication 6 (2013) 142–163

Contents lists available at SciVerse ScienceDirect

Physical Communication

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

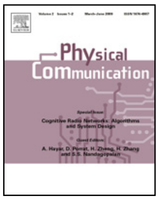

Full length article

Active topology inference using network coding

Pegah Sattari [a,*], Christina Fragouli [b], Athina Markopoulou [a,1]

[a] Calit2 Building, Suite 4100, UC Irvine, Irvine, CA 92697-2800, United States
[b] EPFL IC ARNI, Building BC, Station 14, CH - 1015 Lausanne, Switzerland

ARTICLE INFO

Article history:
Received 15 April 2011
Received in revised form 19 February 2012
Accepted 21 February 2012
Available online 3 March 2012

Keywords:
Tomography
Network coding
Topology inference

ABSTRACT

Our goal, is to infer the topology of a network when (i) we can send probes between sources and receivers at the edge of the network and (ii) intermediate nodes can perform simple network coding operations, i.e., additions. Our key intuition is that network coding introduces topology-dependent correlation in the observations at the receivers, which can be exploited to infer the topology. For undirected tree topologies, we design hierarchical clustering algorithms, building on our prior work in Fragouli et al. (2006). For directed acyclic graphs (DAGs), first we decompose the topology into a number of two-source, two-receiver (2-by-2) subnetwork components and then we merge these components to reconstruct the topology. Our approach for DAGs builds on prior work on tomography in Rabbat et al. (2006), and improves upon it by employing network coding to accurately distinguish among all different 2-by-2 components. We evaluate our algorithms through simulation of a number of realistic topologies and compare them to active tomographic techniques without network coding. We also make connections between our approach and alternatives, including passive inference, `traceroute`, and packet marking.

1. Introduction

Knowledge of network topology is important for network management, diagnosis, operation, security and performance optimization. Depending on the context, one may be interested in the topology at different layers, such as the Internet's router-level topology, an overlay network topology, the topology of an ad-hoc wireless network, etc..

There is a large body of prior work on measurements and inference of network topology. One family of techniques assumes the cooperation of nodes in the middle of the network, and uses `traceroute` measurements to collect the ids of nodes along paths and use them to reconstruct the topology. Another family of techniques, referred to as *network tomography*, assumes no cooperation from internal nodes and relies on end-to-end probes to infer internal network characteristics, including topology. More specifically, multicast or unicast probes are sent/received between sets of sources/receivers at the edge of the network, and the topology is inferred based on the number and order of received probes.

In this paper, we revisit the problem of topology inference using end-to-end probes, in networks where intermediate nodes are equipped with simple network coding capabilities. We show how to exploit these capabilities in order to perform active topology inference in a more accurate and efficient way than existing tomographic techniques.

Our key intuition is that network coding introduces topology-dependent correlation in the content of packets observed at the receivers, which can then be exploited to reverse-engineer the topology. For example, a coding point (that combines multiple incoming packets into one or more outgoing packets) introduces correlation between packets coming from different sources, in a similar way

---

[*] Corresponding author.
E-mail addresses: psattari@uci.edu (P. Sattari), christina.fragouli@epfl.ch (C. Fragouli), athina@uci.edu (A. Markopoulou).

[1] This work has been supported by AFOSR FA9550-10-1-0310, NSF CAREER 0747110, AFOSR MURI Award FA9550-09-0643, NSF CDI 1028394, Swiss National Science Foundation Award No PP00P2-128639, ERC Starting Grant Project NOWIRE ERC-2009-StG-240317.

1874-4907/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.phycom.2012.02.006

that multicast introduces correlation in the packets sent by the same source and observed by several receivers. In fact, the correlation introduced by multicast has been the starting point and the main idea underlying tomographic topology inference. Subsequent schemes made this idea more practical, by emulating multicast with back-to-back unicast probes [1,2]. In contrast, relating probes from different sources to reveal intermediate nodes, also referred to as multiple-source tomography, has been a challenge [2–4]. Using simple network coding operations at coding points solves this problem and allows accurate and fast topology inference.

Our approach is general and can be applied to infer the topology in a range of scenarios, including but not limited to wireless multi-hop networks. Wireless multi-hop networks are able to support simple network coding operations (additions are sufficient for our schemes), as demonstrated in [5], and can therefore benefit from our techniques. Furthermore, there is a good match between some properties and constraints of such networks and our schemes. First, there is natural variability in the delay of wireless links, which (if appropriately used—as explained in later sections) can expedite inference. Second, our schemes keep internal nodes simple (moving processing for inference to dedicated nodes at the edge) and anonymous (revealing the logical topology but not the identities of nodes). Finally, improving the speed of inference may prove important to keep up with changes, *e.g.*, due to mobility.

Our contributions are as follows. First, we consider undirected trees, where leaves can act as sources or receivers of probes, and we design hierarchical clustering algorithms that infer the topology, building on our prior work in [6]. Then, we consider directed acyclic graphs (DAGs) with a fixed set of M sources and N receivers and a pre-determined routing scheme. We first decompose the topology into a number of two-source, two-receiver subnetwork components and then we merge these components to reconstruct the topology. Our approach for DAGs builds on prior work on tomography [4], and improves upon it by employing simple network coding operations at intermediate nodes to deterministically distinguish among all possible 2-by-2 subnetwork components, which was impossible without network coding [2,4]. We evaluate our algorithms through simulation over a number of topologies and we show that they can infer the topology accurately and faster than tomographic approaches without network coding. We present our schemes as active probing: special probes are sent by the sources, specifically for the purpose of inference, and are treated in special ways by intermediate nodes and eventually received by the receivers and processed at a fusion center. We believe that our active probing approach with network coding provides one more building block, in the already large space of topology inference techniques, with core strength and ability to identify joining points. We also compare and make connections between our active probing approach and alternatives, such as passive inference, `traceroute`, and packet marking.

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 presents our assumptions, notation, and problem statement. Section 4 summarizes the main results of the paper. Section 5 presents algorithms for inferring tree topologies. Binary trees are discussed in Section 5.1, in the absence (Section 5.1.1) or presence (Section 5.1.2) of packet loss. General trees are discussed in Sections 5.2.1 and 5.2.2. Section 6 presents algorithms for inferring directed acyclic graphs (DAGs). Section 6.2 presents algorithms for inferring 2-by-2 subnetwork components, in the absence (6.2.1) or presence (6.2.2) of packet loss. Section 6.3 explains how to merge these components to reconstruct the topology. Section 7 provides simulation results for some realistic topologies. Section 8 discusses two possible deployment scenarios (one as an active probing scheme and another one using packet marking), and makes connections between our approach and alternative topology inference approaches. Section 9 concludes the paper. Appendices A and B analyze the probability of error of our inference algorithms in trees and DAGs, respectively.

2. Related work

One body of related work is network tomography in general, and topology inference in particular. A good survey of network tomography can be found in [7]. An early work on topology inference using end-to-end measurements is [8], where the correlation between end-to-end multicast packet loss patterns was used to infer the topology of binary trees. The correctness of this idea was rigorously established in [9], and was extended to general trees and to measurements other than loss, such as delay variance [10], or more generally any metric that grows monotonically with the number of traversed links. The idea has also been extended to unicast probes [1,2]. In summary, tomographic schemes for topology inference use end-to-end active probing and feed the number, order, or a monotonic property of received probes as input to statistical signal-processing techniques. Inference of link characteristics [11] can also be combined with topology inference [2]. In a different context, similar problems have been studied in the context of *phylogenetic* trees [12]. The work in [13] uses such algorithms [12], for topology inference in sparse random graphs.

In addition, inference of congested links has been studied from the angles of compressive sensing [14–16] and group testing [17–19]. The work in [17] formulates the problem as a *graph-constrained* group testing, where the items correspond to edges, some of them being defective, and the goal is to identify the defective edges given that the test matrix conforms to constraints imposed by the graph, *e.g.*, the path connectivities. The work in [14] recovers sparse vectors, representing certain parameters of the links over the graph, through ℓ_1 minimization. It improves the number of required measurements over [17], as compressive sensing allows real numbers for the link characteristics and measurements instead of true/false binary values in group testing problems.

Most tomographic approaches rely on probes sent from a single source in a tree topology [1,8–10,20–26]. Coates et al. [27], Rabbat et al. [2,4] introduced the multiple-source multiple-destination (M-by-N) tomography problem, by sending probes between M sources and N