Inferential network monitoring, also known as Network Tomography, aims at reconstructing various properties of large communication networks from indirect measurements to help oversee the performance of these networks. Network inference can be achieved by two general approaches. In the internal approach, one takes measurements directly at the edges and nodes of the network. However, this approach suffers from several drawbacks: direct measurements may create an extra computational burden as well as congestion in the network; the network operator may not allow access to internal devices of the network or may not make public any measurements on them; the routers may not have the technological capabilities of performing the required measurements. This has led some in the networking community to consider instead the external approach. In this case, one uses so-called end-to-end measurements, e.g. measurements of delays or rate of packet drops between nodes in the network, and seeks to infer the desired network properties from them. This gives rise to an inverse problem similar to tomographic image reconstruction, whereby the name Network Tomography.

We focus on multicast-based inference. Multicast routing consists in sending a packet from a source to a set of receivers through a routing tree. The packet is duplicated at each branch point and sent further down the tree. The routing tree is generally unknown to the user. The idea is to use inherent correlation of measurements between different receivers to reconstruct the topology of the routing tree as well as to estimate link properties of this tree.

We design a novel algorithm for multicast-based delay inference, i.e. the problem of reconstructing the topology and delay characteristics of a network from end-to-end delay measurements on network paths. Our inference algorithm is based on additive metric techniques widely used in phylogenetics. It runs in polynomial time and requires a sample of size only poly(log n).