Topology Discovery for Large Ethernet Networks

Bruce Lowekamp

David R O'Hallaron

Thomas R Gross

SIGCOMM'01

Srikant Sharma

Ethernet Topology Discovery?

- Typical constituents of Link-layer network are Bridges, Hubs and NICs (hosts)
- Bridges form an acyclic network using spanning tree algorithm and perform source learning to build a forwarding database
- Topology discovery is the discovery of active interconnections between the bridges, hubs and the hosts connected to these bridges

Challenges

- No central knowledge of link level connections
- The bridging is transparent to the end hosts
- The only available information is the forwarding database, that too from smart (SNMP capable) switches only

Contribution of this Paper

- An algorithm to do topology discovery using the information from the forwarding databases of SNMP capable switches
- The work has been put to test to discovery the topology of networks comprising of 2000 nodes and 50 bridges

Forwarding Database

- Based on source learning
- Forwarding set of a port
  - For a bridge **B** the set of hosts reached through port **x**
  - Denoted by F_C^x
- Forwarding set F_C^x is complete if it includes all the hosts reachable from **x**

Connectivity for Bridges

- Direct Connectivity:

  Two bridges are directly connected if there are no bridges in between them and both have forwarding entries for each other on some port

- Simple Connectivity:

  Two bridges have simple connectivity with each other if both have forwarding entries for each other

- All directly connected bridges have simple connectivity as well

Discovery using Direct Connectivity

- Determine all entries in the FDB of all switches, ensure that they are complete
- Select a single bridge and determine its neighborhood. Next traverse through its neighboring bridges which are directly connected to it
- Uses *Direct Connection Theorem*

Direct Connection Theorem

- Assume that F_i^x and F_j^y are complete. Two bridges i and j are directly connected via the link connected to port x of i and port y of j if and only if $F_i^x \cap F_j^y = \phi$ and $F_i^x \cup F_j^y = N$
- N is the set of all nodes in the network

Problems with Shared Segments

- Shared segments at the periphery do not cause any problems
- Internal shared segments result in conflicting entries in the FDBs of multiple switches and can lead to the failure of topology discovery
- Internal hubs and non SNMP bridges can be seen as internal shared segments

Shared Segment Theorem

- A set of nodes *S* consists of a shared segment between some bridges S_B if and only of all the nodes in S can be reached from a single port for each bridge in S_B
- SST can be used in conjunction with DCT to carry out topology discovery

Limitations of DCT

- Completeness requirement
- All machines need to generate some traffic to have entries in FDB
- With large networks it may not be possible to have a complete forwarding set

Topology Discovery with Incomplete Knowledge

- Based on proof by contradiction
- Uses the concept of Through Set T_C^x, essentially complement of forwarding set F_C^x
- If for any two switches

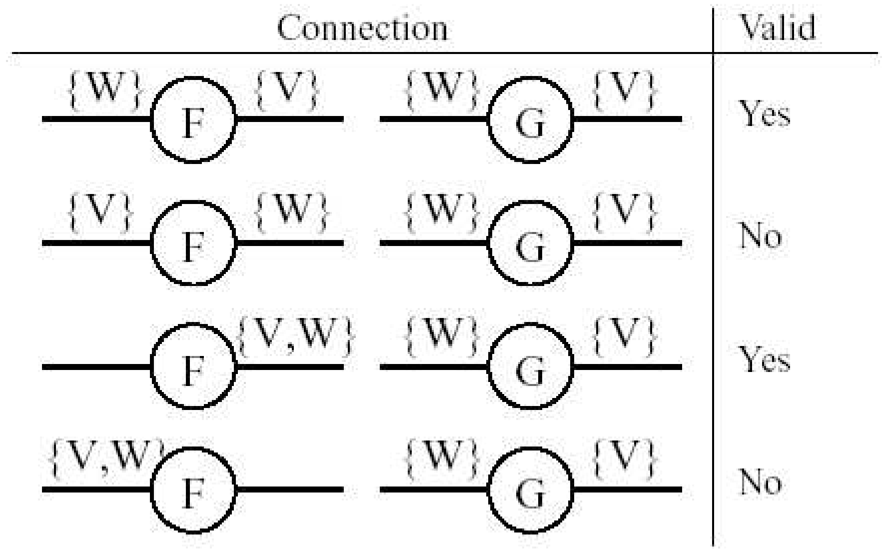

Figure 3: Examples of valid and invalid connections between two bridges with different forwarding entries on bridge F. In each case, the forwarding entries are examined to see whether those two edges can form a valid simple connection. Note that the only cases where there are conflicts occur when the bridges map the same address in opposite directions.

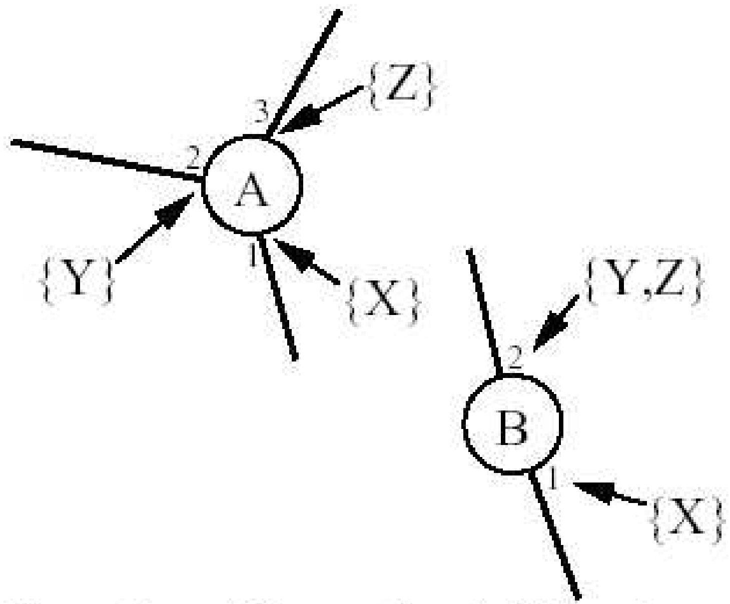

Bridge	Port	Forwarding Entries	Through Sets
A	1	X	Y, Z
	2	Y	X, Z
	3	Z	X, Y
B	1	X	Y, Z
	2	Y, Z	X

Figure 2: Example of two bridges for which contradictions can be used to determine the connections.

Ports A	Ports B	Mapping	Conflict
1	1	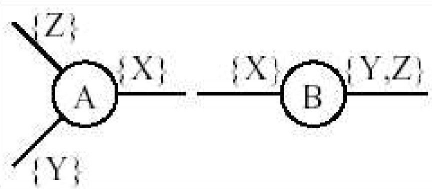	Y and Z
3	1	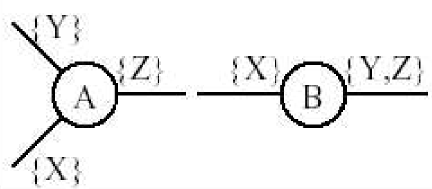	Y
2	2	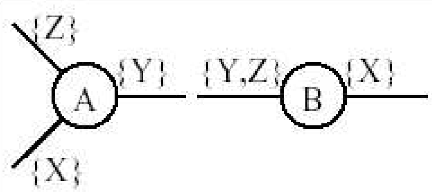	X
2	1	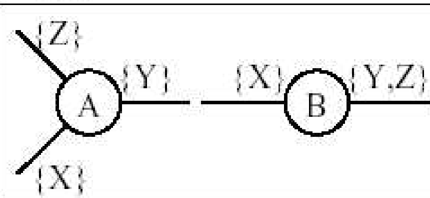	Z
1	2	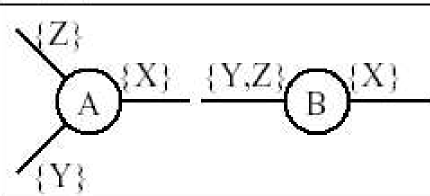	**NONE**
3	2	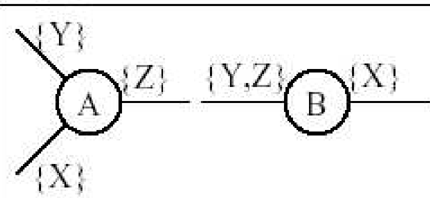	X

Figure 4: How contradictions can be used to eliminate impossible connections from the bridges in Figure 2.

Simple Connection Theorem

- If for two switches a & b suppose there exists exactly one pair of ports a_x and b_y such that $T_a^x \cap T_b^y = \phi$ then a_x and b_y are connected. Furthermore, if a_x and b_y are connected then $T_a^x \cap T_b^y = \phi$
- If there can some indeterminate solutions to this.

Indeterminate solutions

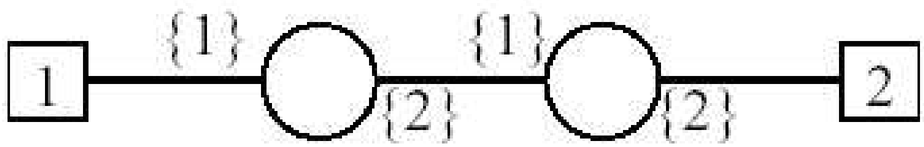

Figure 5: This network consists of two bridges and two hosts, labeled 1 and 2. The connection of the two bridges cannot be determined because there is insufficient information. Observe that the bridges can be placed in either position, as they are indistinguishable using their forwarding databases.

Minimum Knowledge Requirement

The ports x & y that connect bridges a & b are uniquely determined iff at least one of the following conditions is met

- Each bridge has an entry for the other in FDB; or
- Bridge a has an entry for b in F_a^x and $\exists k \neq x : F_a^k \cap F_b^y \neq \phi$; or
- Forwarding entries for three nodes are shared between a and b, divided among at least two ports of a or b and three ports of the other switch

Experiment Results

- Has been tested on networks scaling up to 2000 nodes and 50 bridges
- The mechanism works with wide variety of network hardware
- The topology discovery is run on P3 566MHz

Experiment Results…

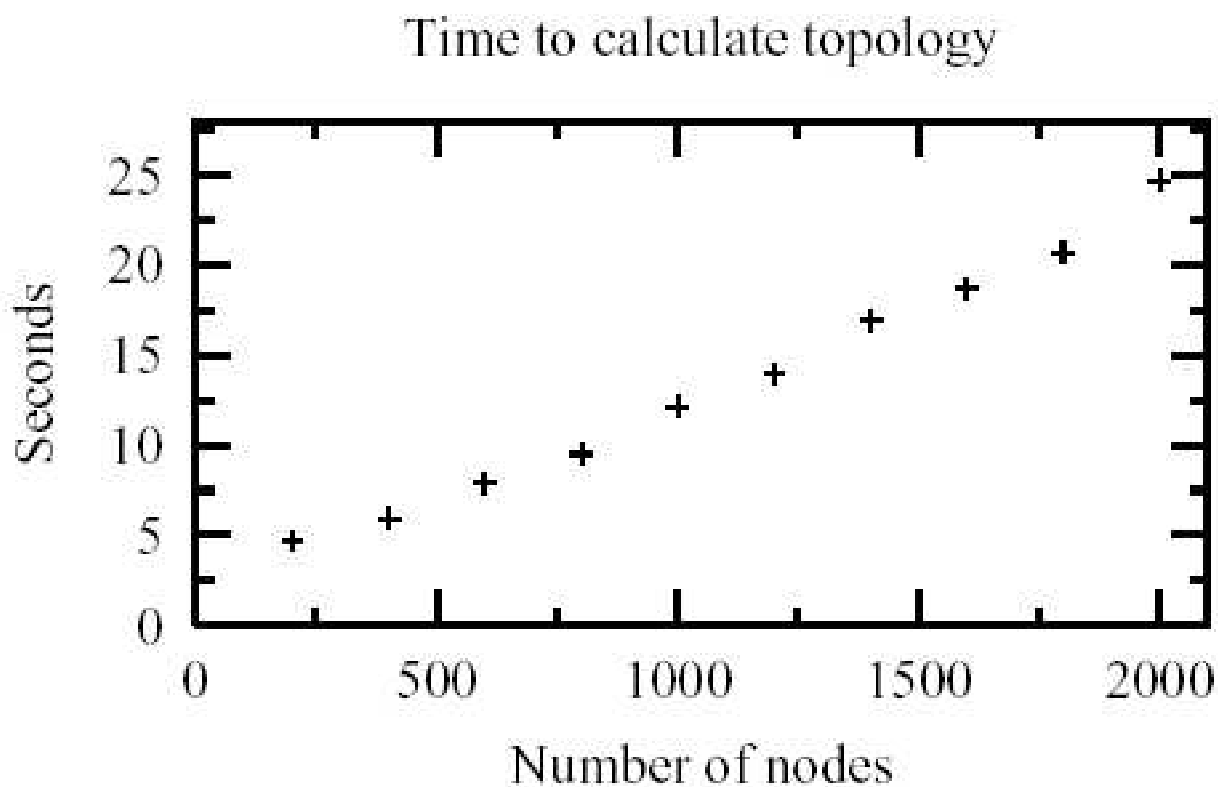

Experiment Results…

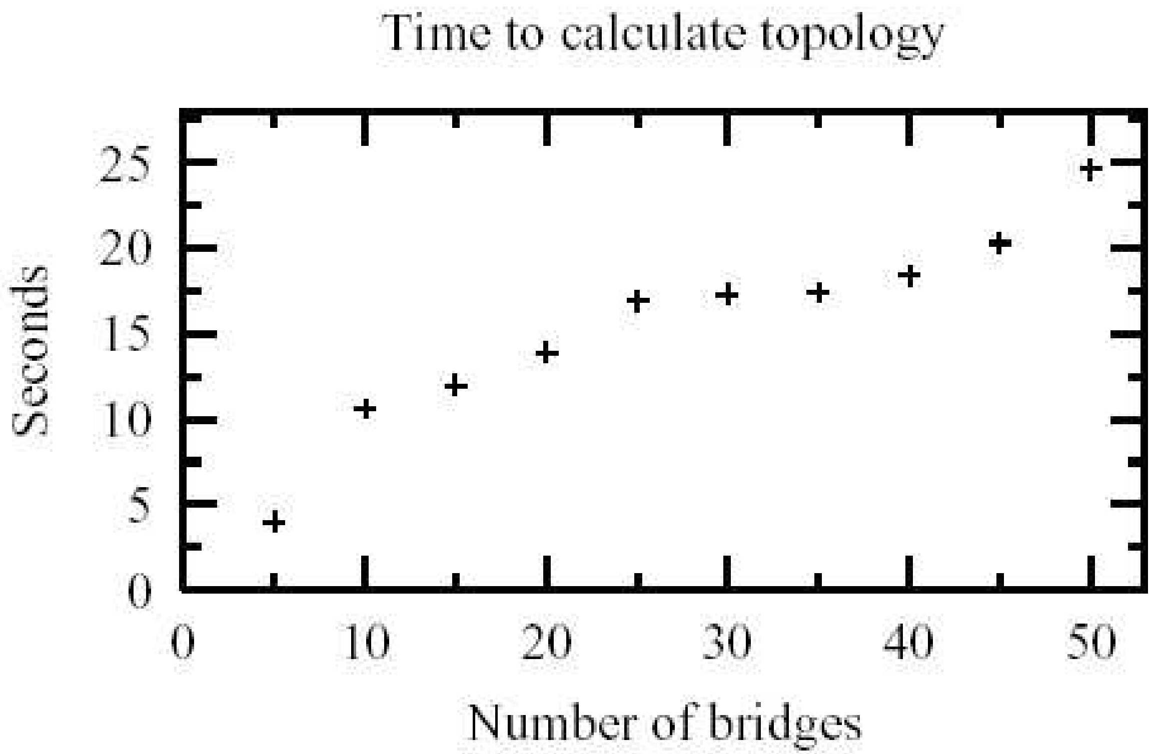

Opinions

- A novel approach to utilize scattered information in the network
- Does not talk about implementation complexities
- Carries out discovery of only active network topology