Direct Queries for Discovering Network Resource Properties in a Distributed Environment

Bruce Lowekamp
David O'Hallaron
Thomas Gross

Carnegie Mellon University

Question:

Can network-based queries:

- provide accurate performance predictions?
- scale well?
- provide topology information?

Existing Techniques

Application-based

- Run the application itself on the network
- Most direct measurement

Benchmark-based

- Use set of “standard” applications
- Convert from benchmark to application

Limitations

Application-based

- Only one application
- Many parameter/resource combinations

Benchmark-based

- Hard to get quantitative predictions
- Still $O(P^2)$ pairs

Neither determines topology

- link sharing
- resource selection

Outline

- Remos
- SNMP
- Prediction Techniques
- Experimental Verification
- Topology Discovery
- Practical Considerations

Remos Architecture

Design goals.

- **Scalability: deal with large networks and queries.**
- **Diversity: wide range of technologies and use.**
- **Privacy: operators limit access to internal info.**

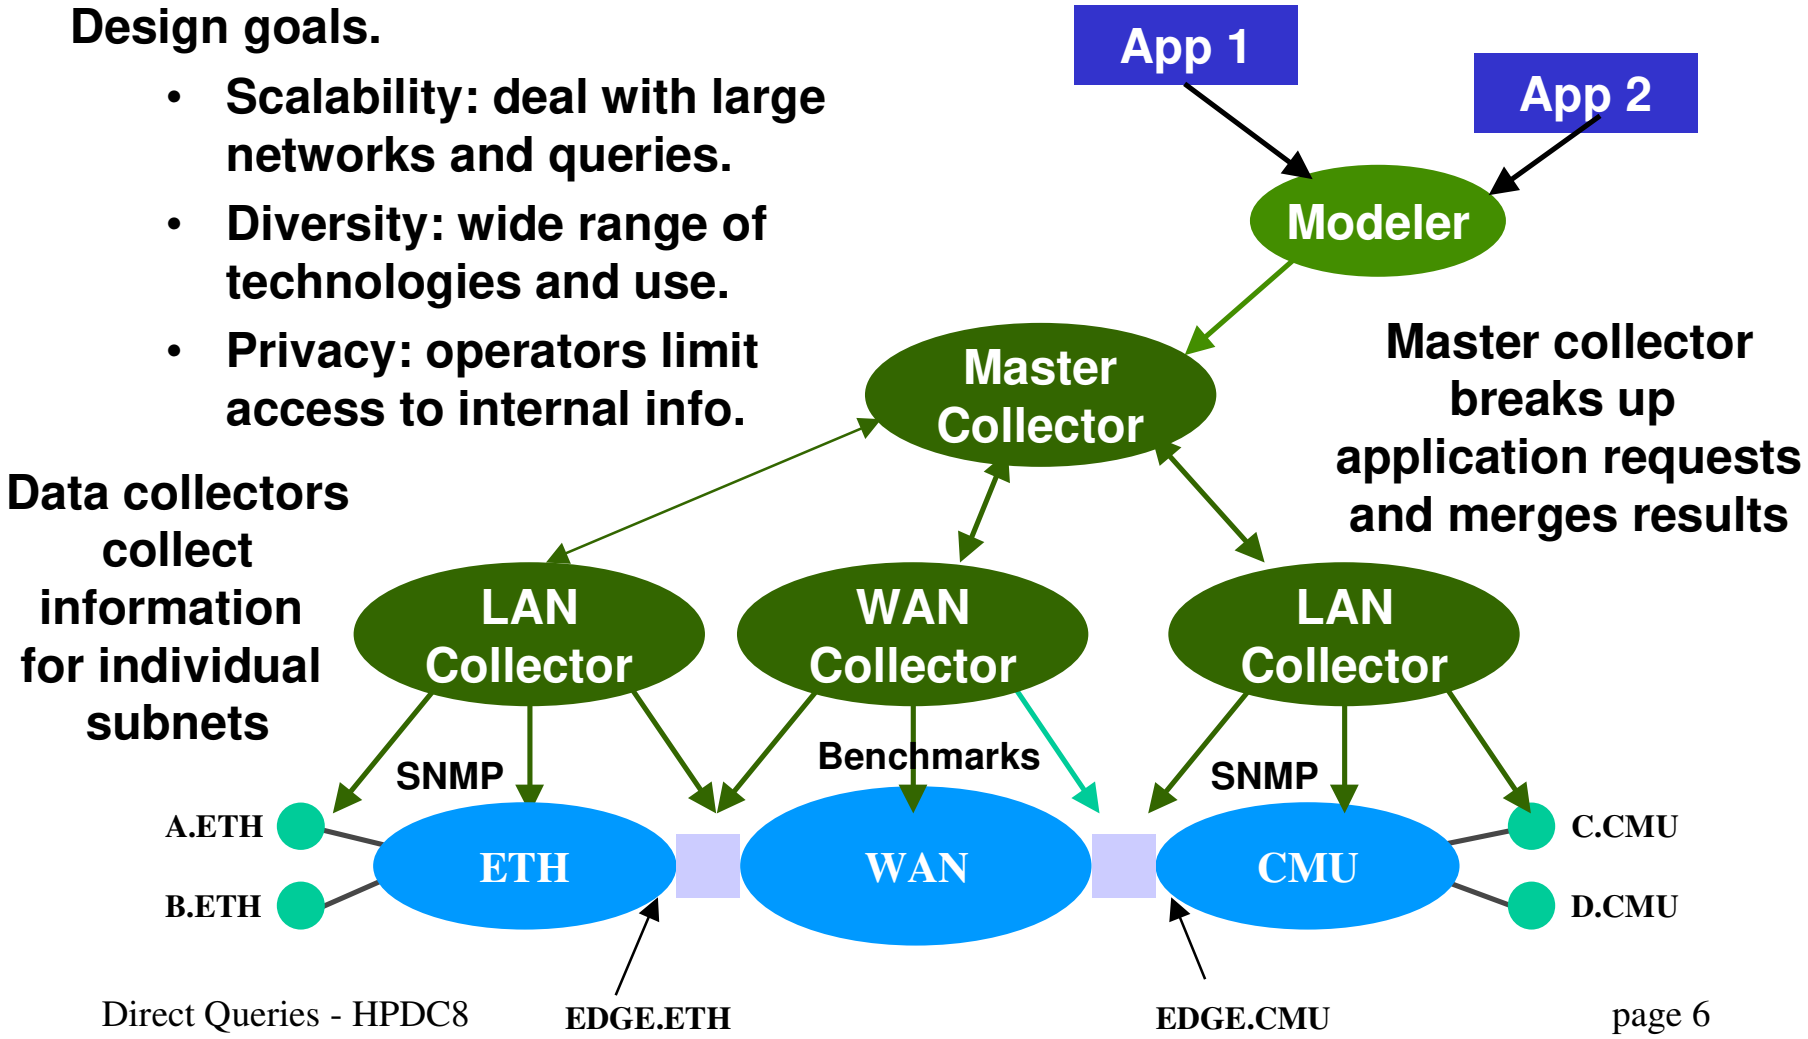

SNMP Overview

Database protocol for network managers

MIB-II

- interface table - octet counters
- routing table

Bridge-MIB

- forwarding database

Prediction Techniques

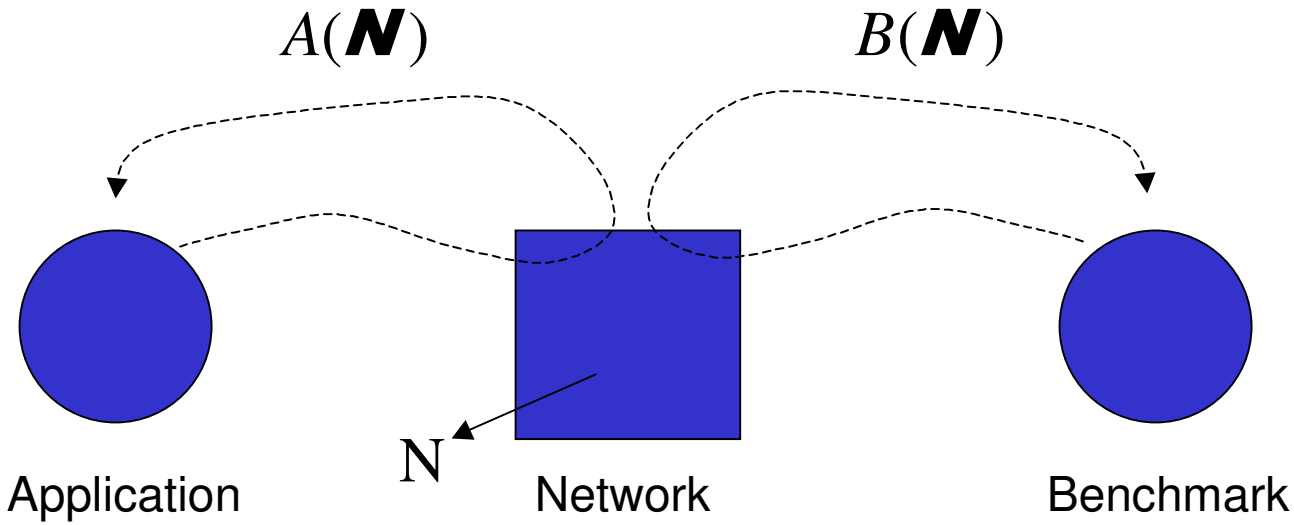

Prediction Techniques

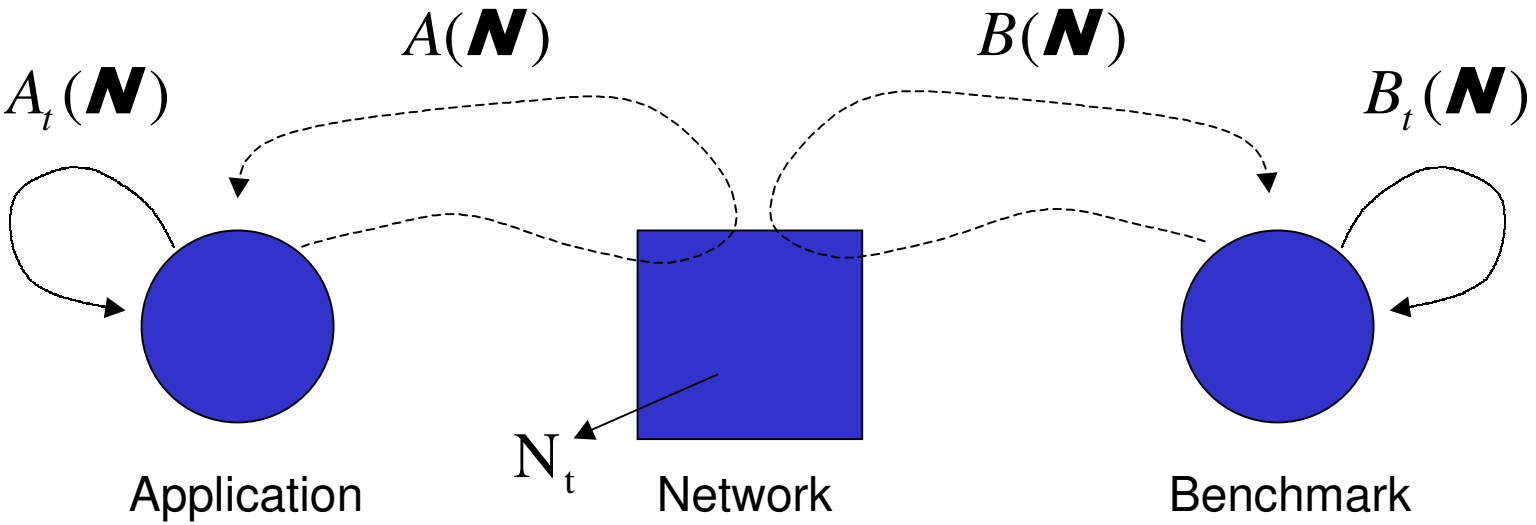

Prediction Techniques

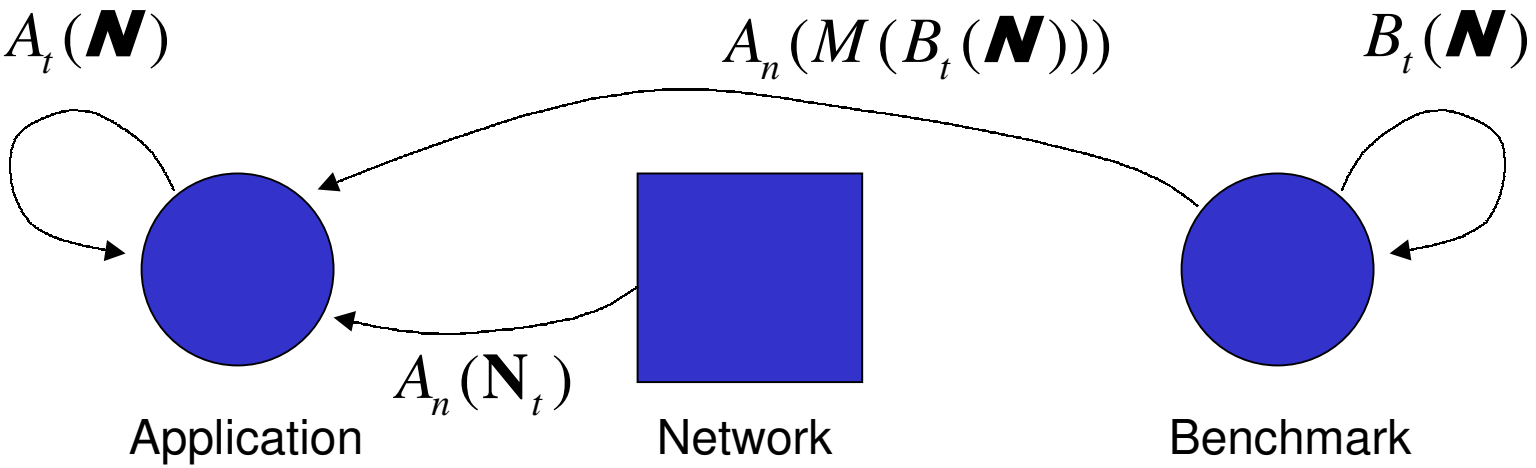

Prediction Techniques

$$\hat{A}(\boldsymbol{N}) = \begin{cases} A_t(\boldsymbol{N}) & \text{Application-based} \\ A_n(\mathbf{N}_t) & \text{Network-based} \\ A_n(M(B_t(\boldsymbol{N}))) & \text{Benchmark-based} \end{cases}$$

Motivations for network-based

Scales better

- O(N) vs O(P^2)

More protocols

- Easier to use network than convert

Quantitative predictions

- Harder to get quantitative predictions out of benchmarks

Experimental Verification

How accurate is network-based technique?

- TCP/UDP protocols
- Different traffic
- End-to-end measurement

Are these offset by benefits?

Is accuracy similar to application- or benchmark-based techniques?

Experimental Procedure

Vary rate of FGN background traffic

Each observation:

- Get available bandwidth with SNMP
- Transfer 1MB data

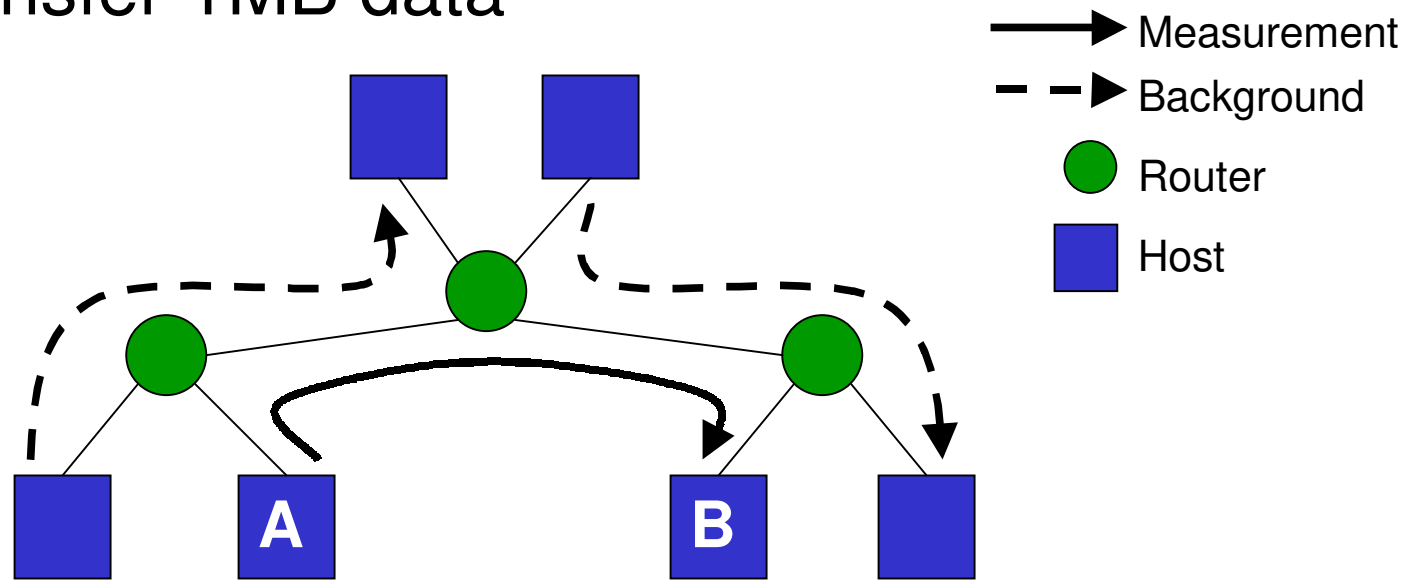

Experimental Procedure

Build up mapping from network model to communication time, $A_n()$

Use time-series prediction on $A(\boldsymbol{M})$ and $\mathbf{N}$, $A_t(\boldsymbol{M})$ and $\mathbf{N}_t$

Compare relative error of $A_t(\boldsymbol{M})$ and $A_n(\mathbf{N}_t)$ with next observation of $A(\boldsymbol{M})$

Experimental Results

Accuracy of two techniques is similar

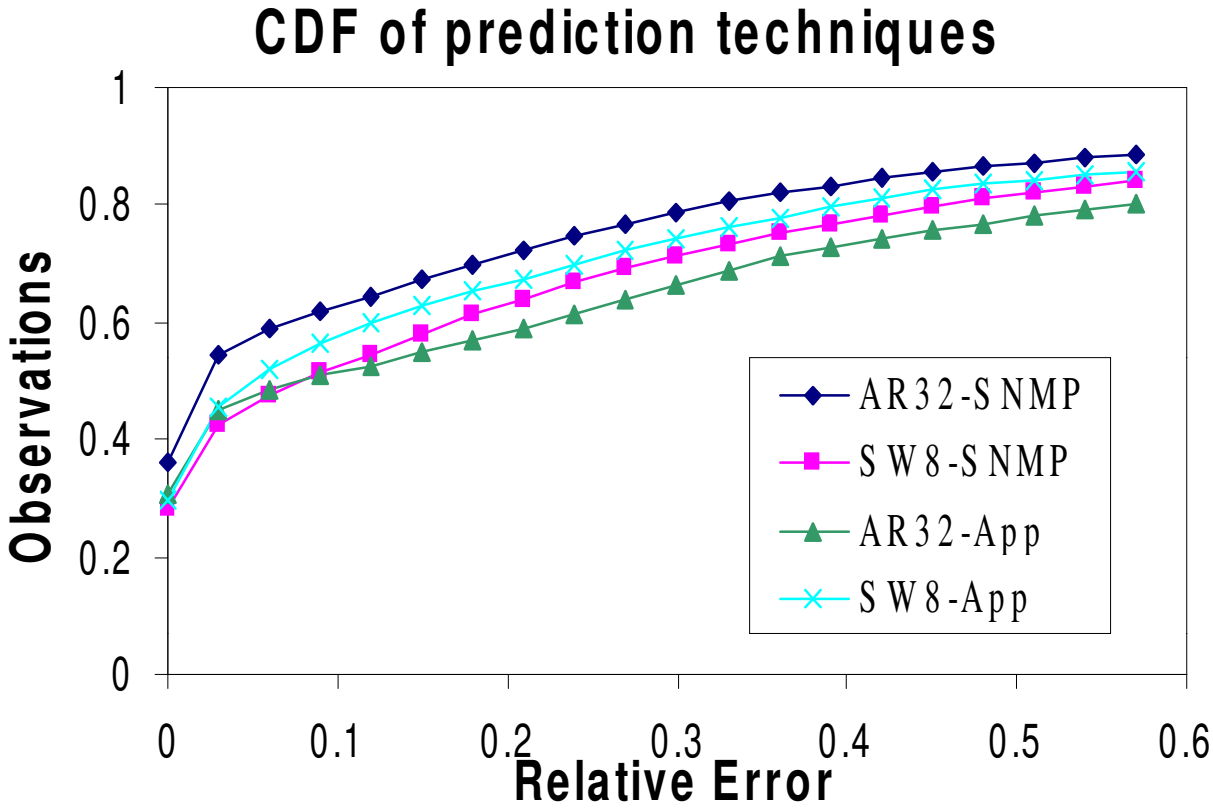

Topology Discovery

Several layers of topology:

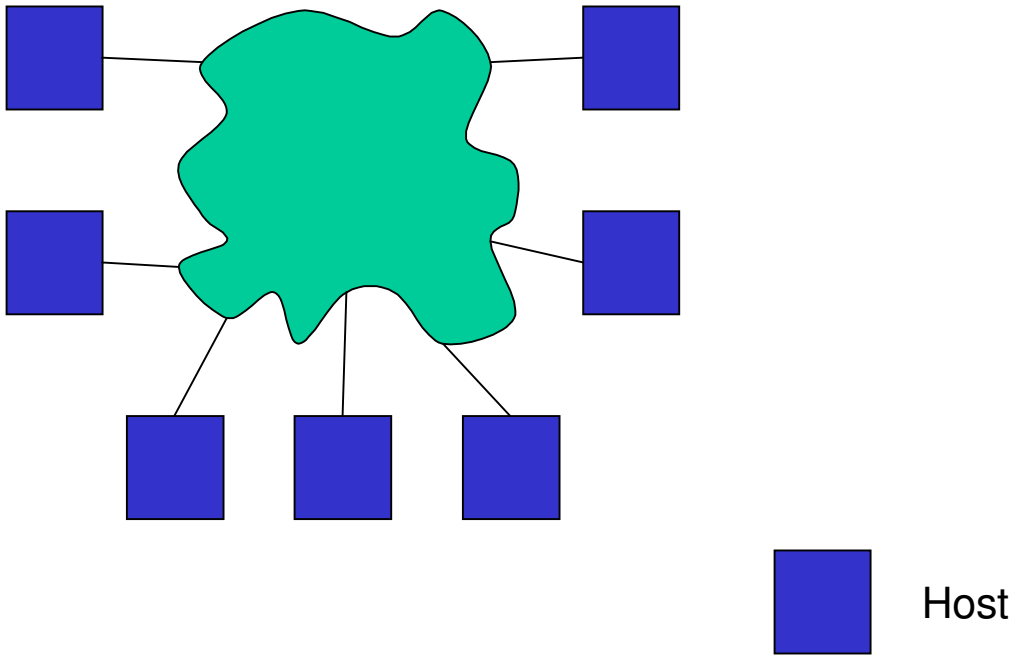

Topology Discovery

Routers

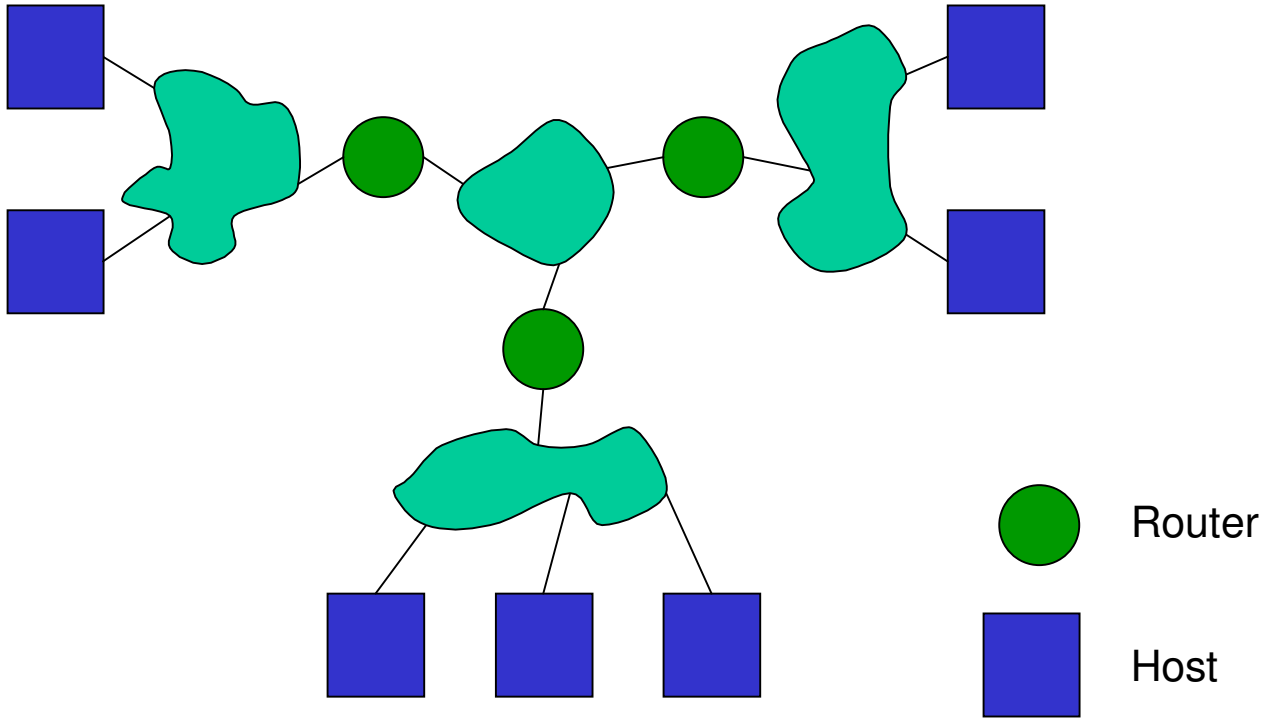

Topology Discovery

Bridges

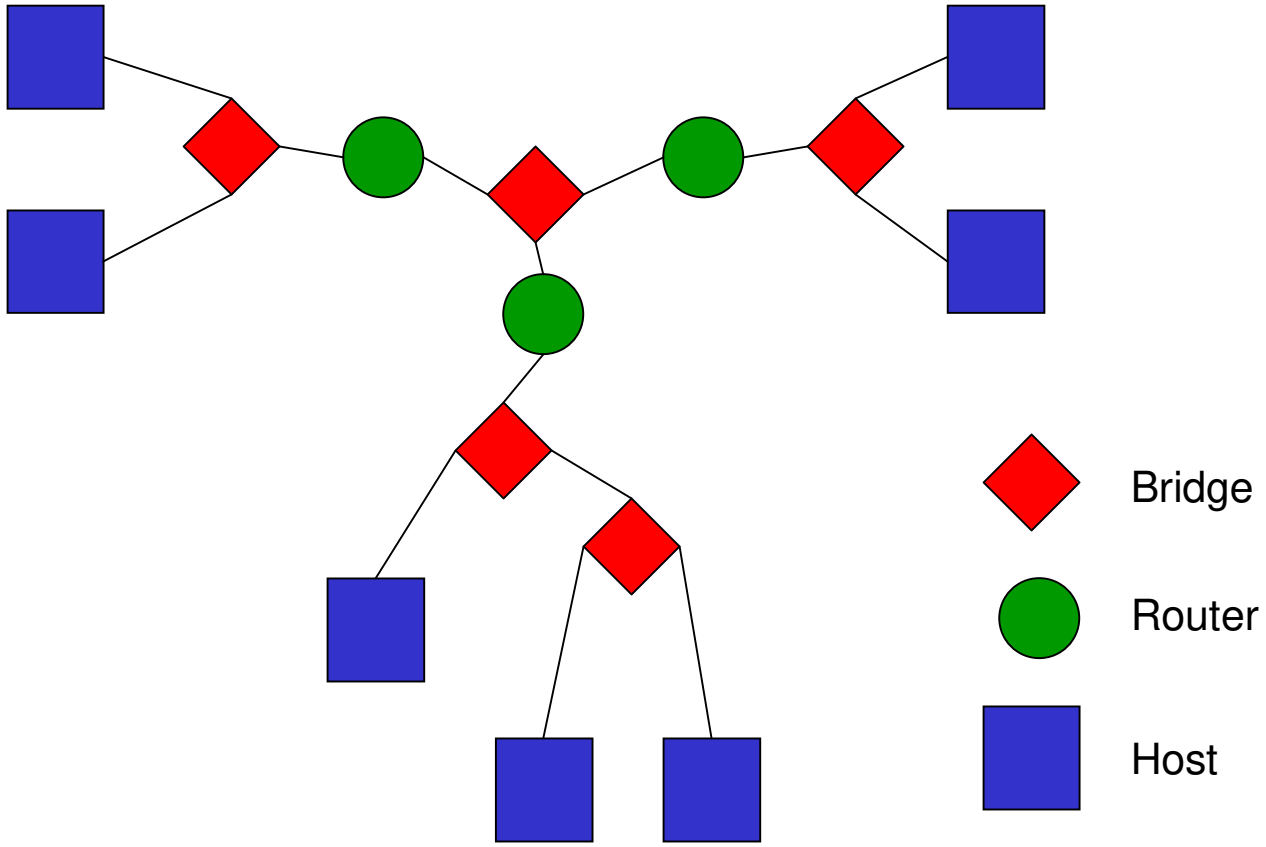

Topology Discovery - Routing

Routing is easy:

- Each machine has next-hop field in routing table
- Follow route hop-to-hop from source to destination

Transparent Bridging

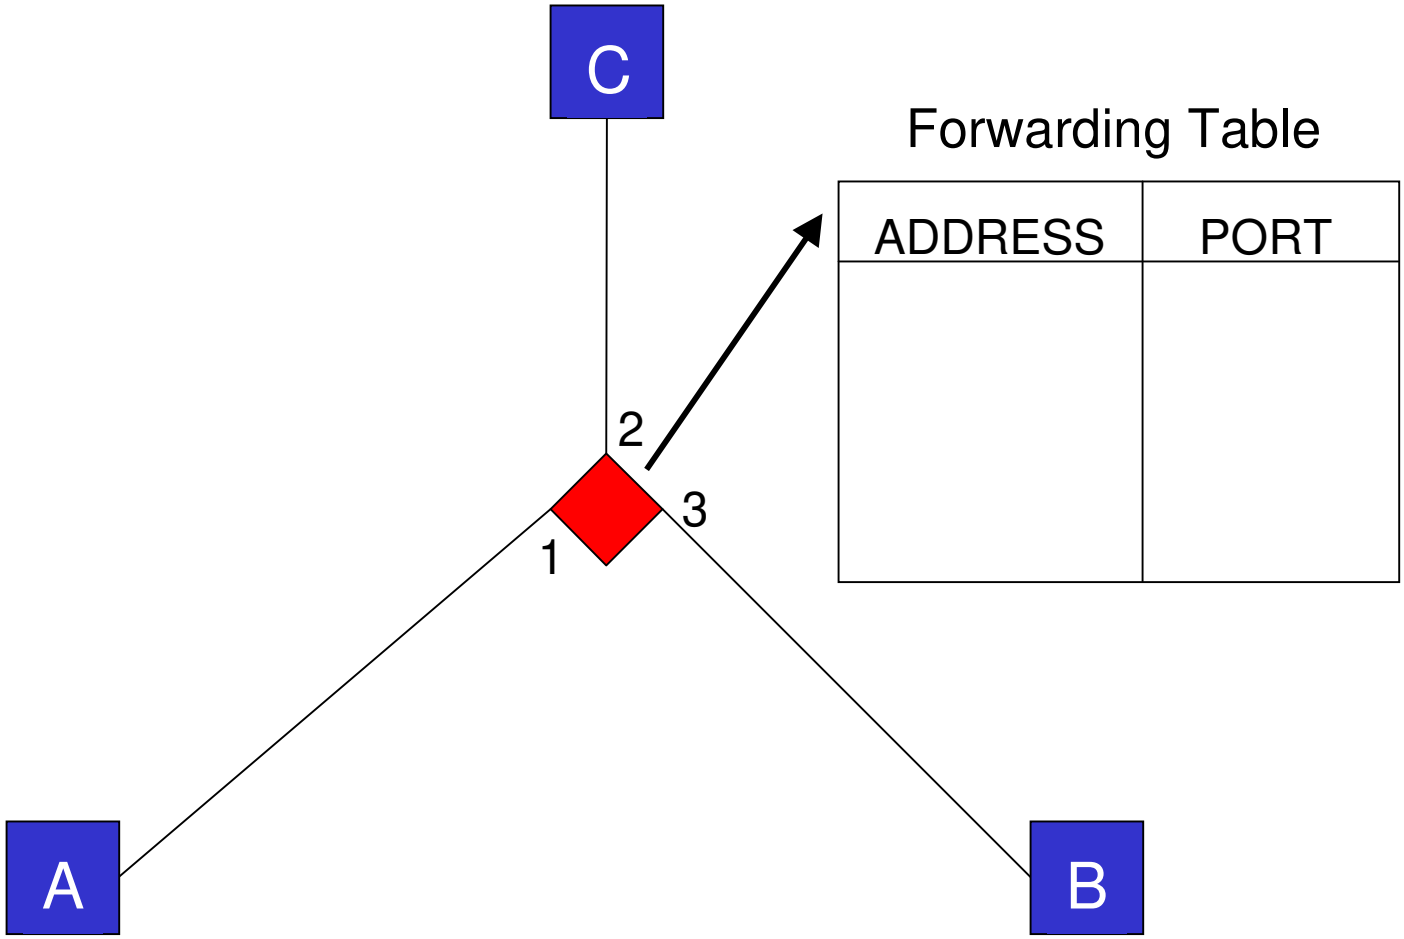

Transparent Bridging

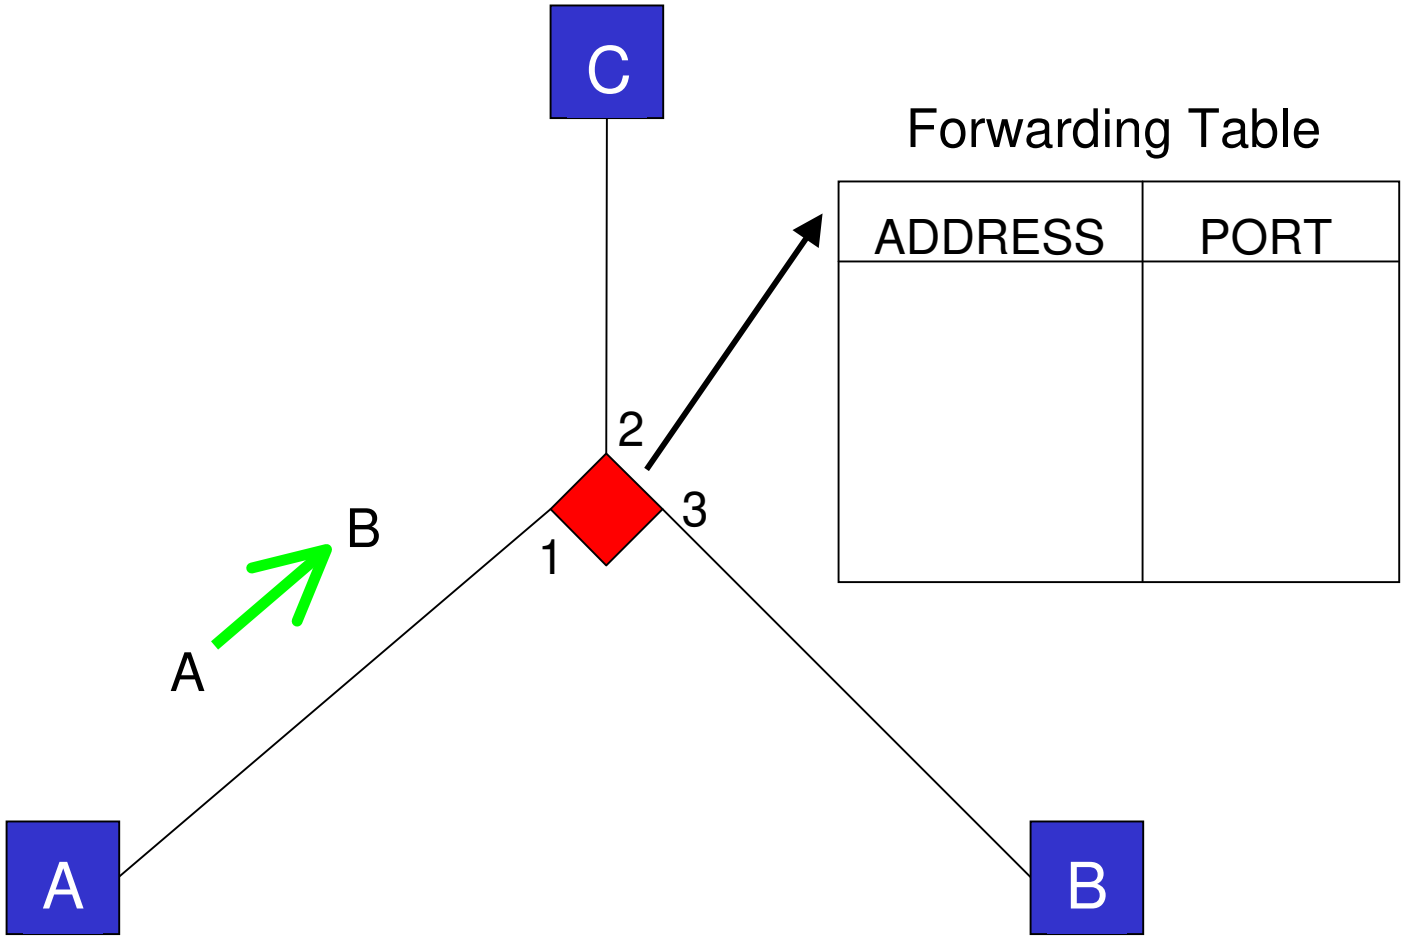

Transparent Bridging

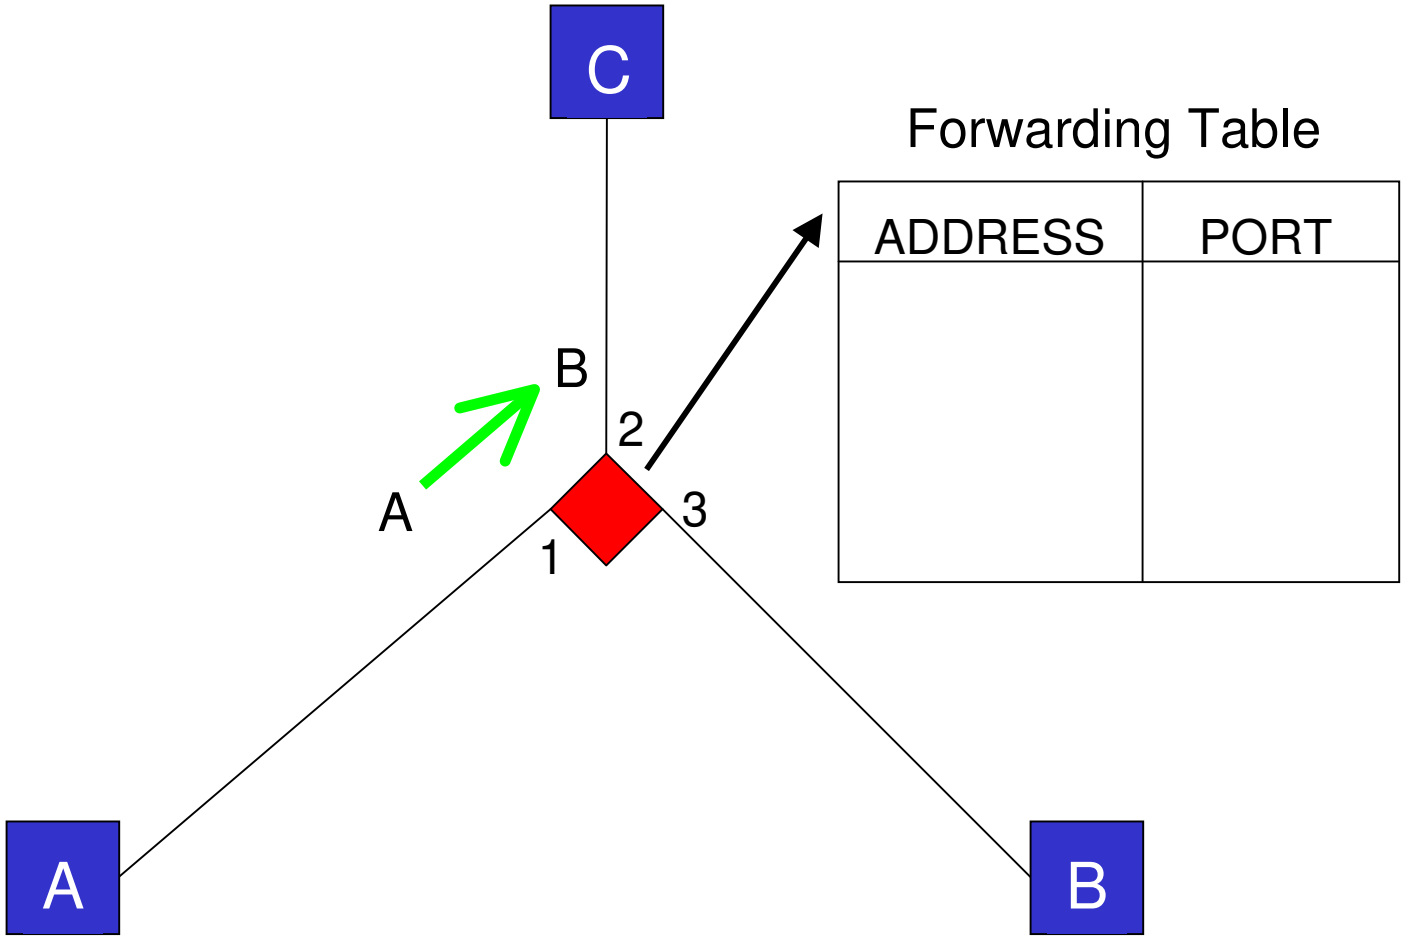

Transparent Bridging

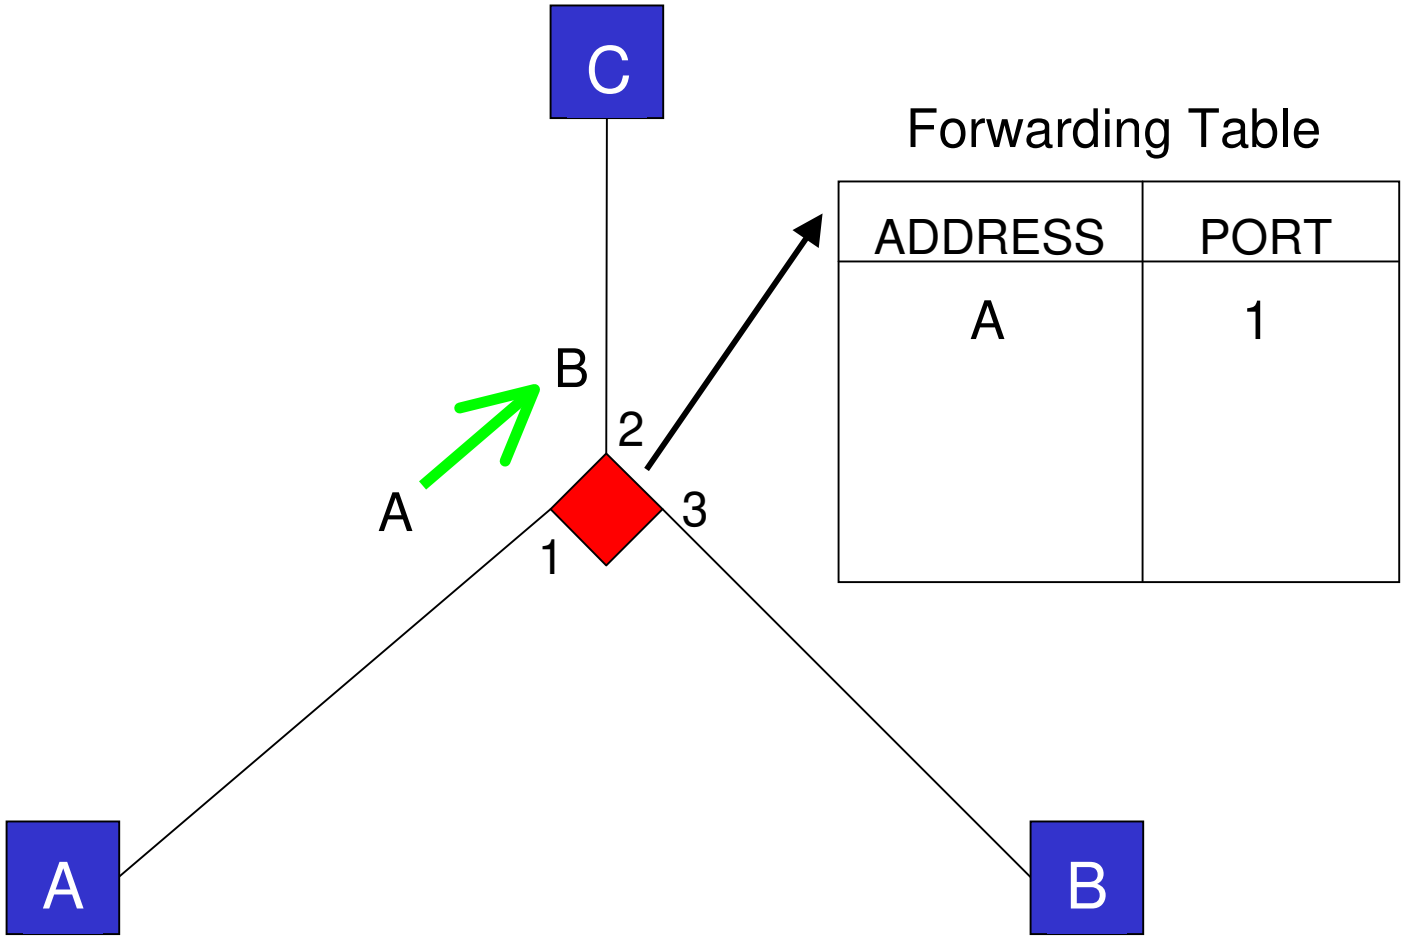

Transparent Bridging

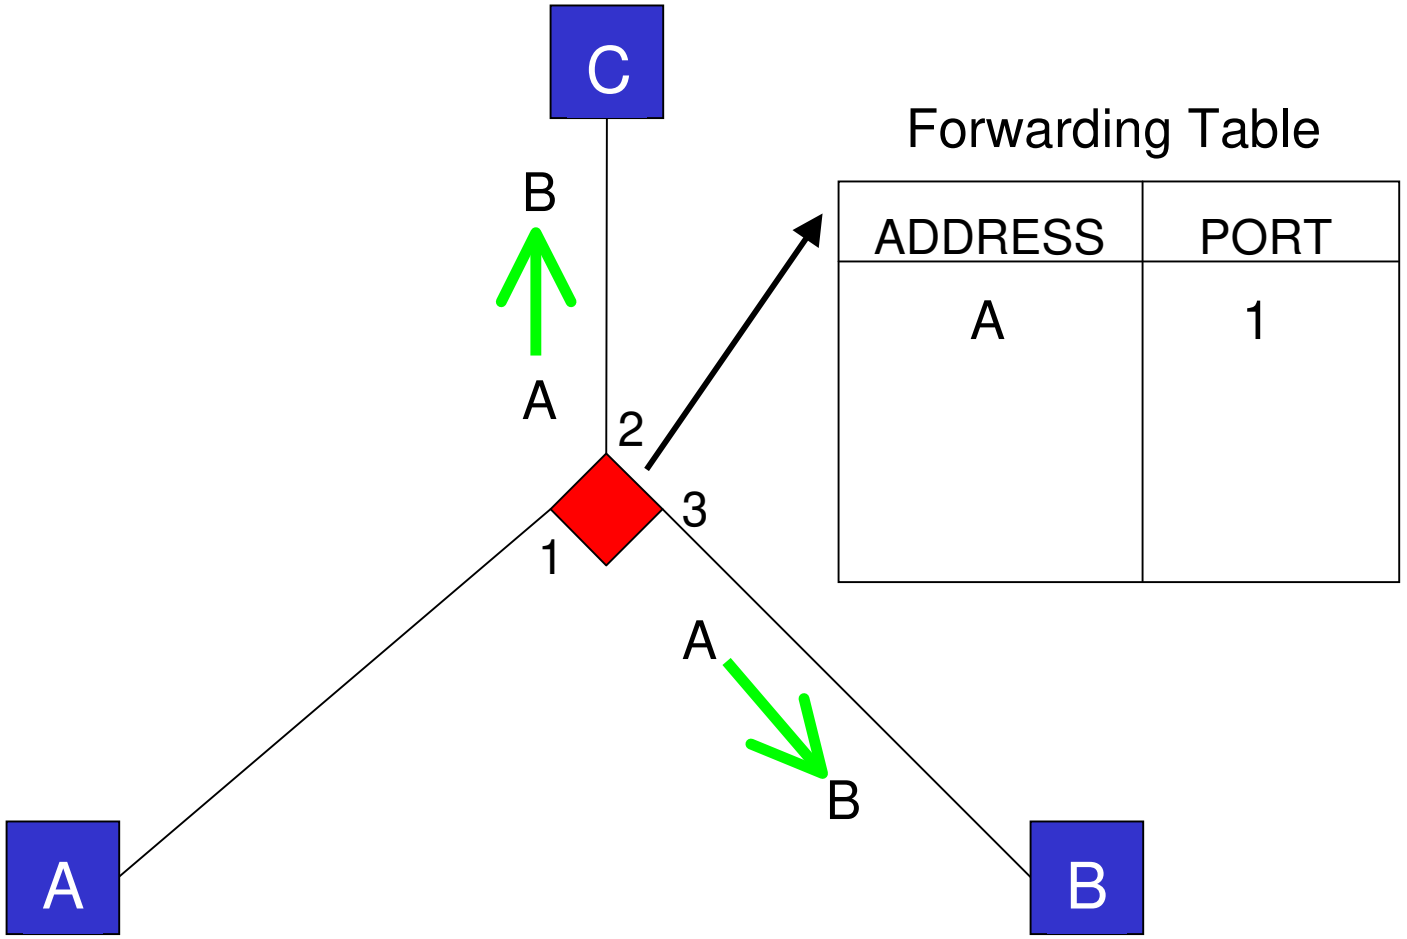

Transparent Bridging

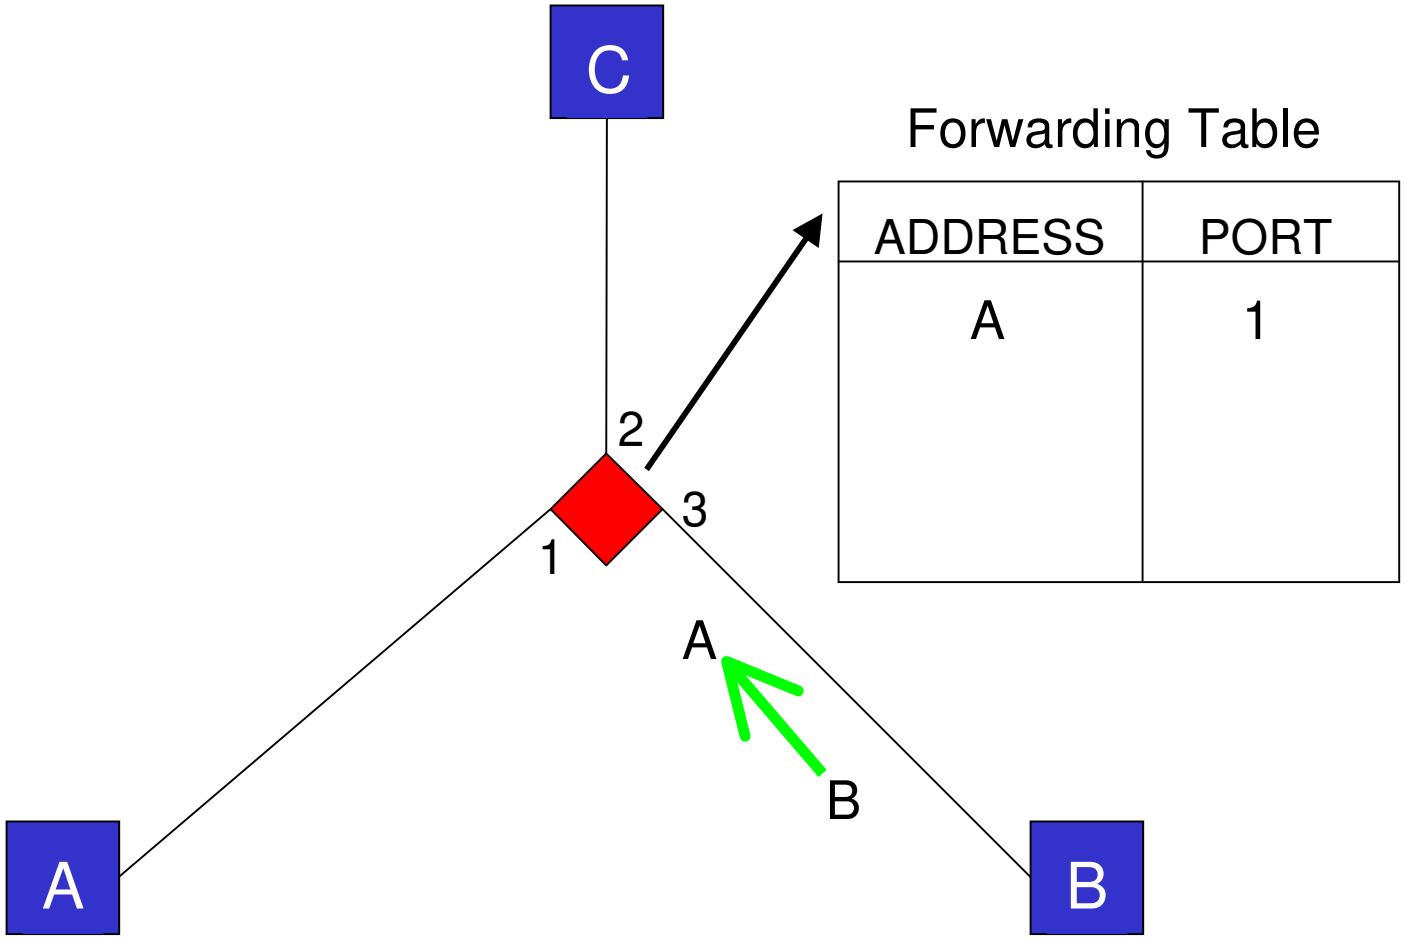

Transparent Bridging

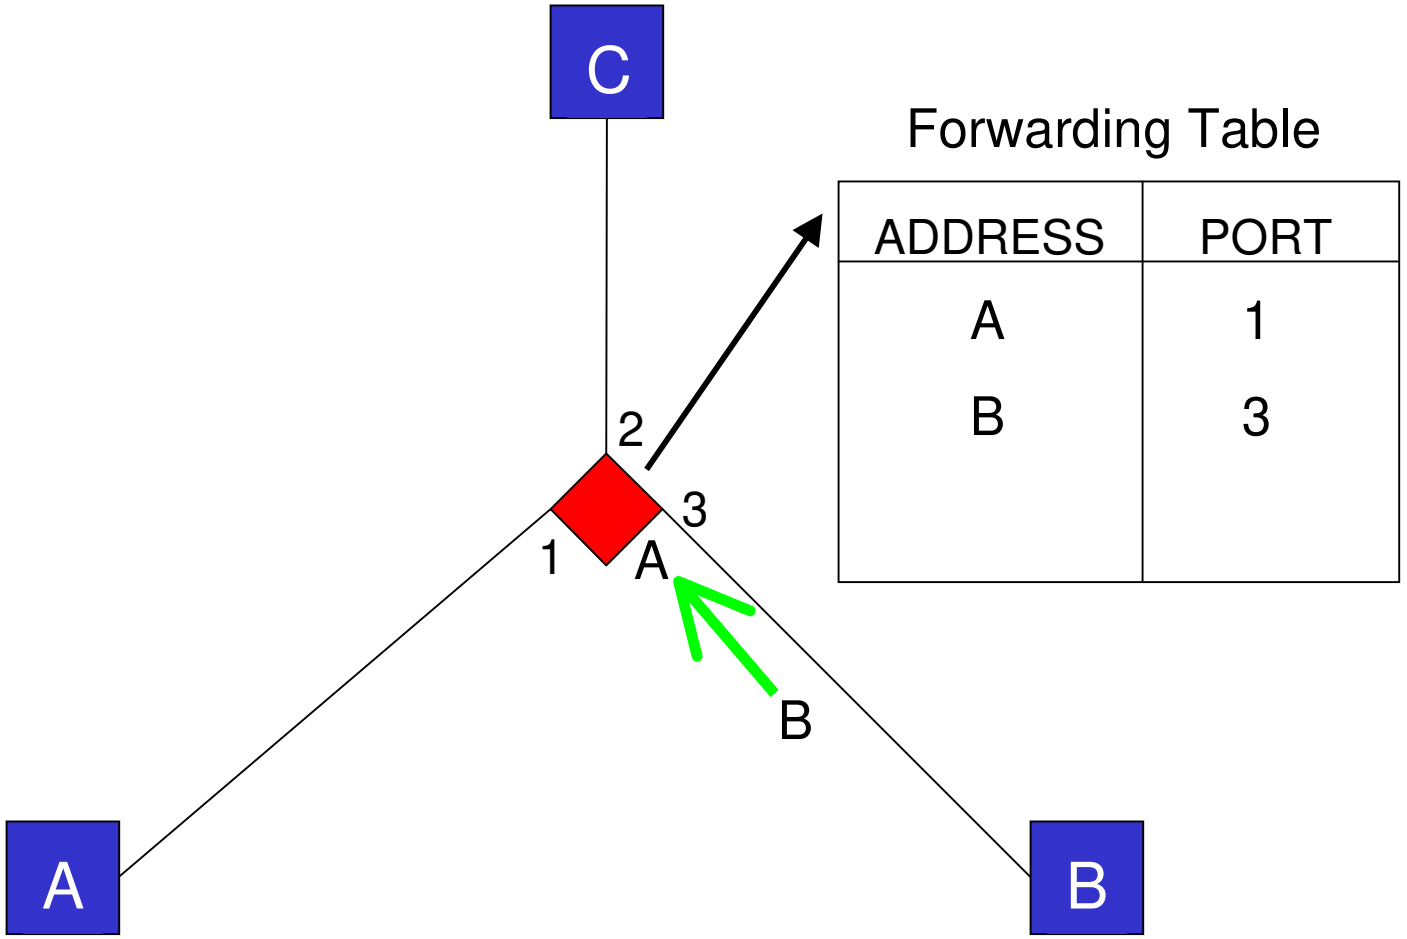

Transparent Bridging

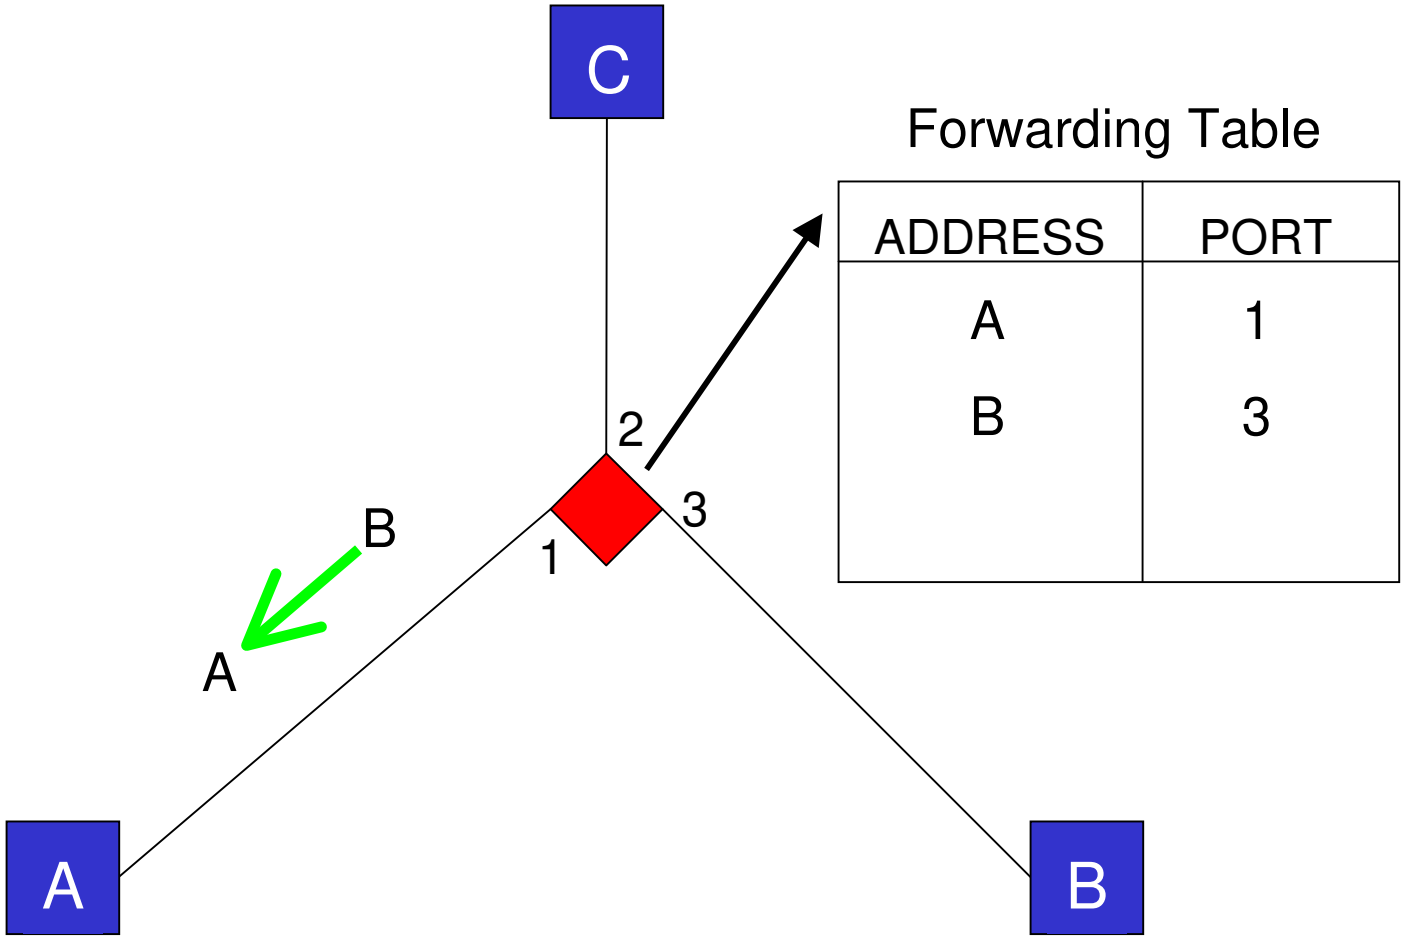

Topology Discovery - Bridging

Transparent bridging is transparent

- hosts don’t need to know tree
- don’t even know what bridges are used

Get list of bridges from network manager

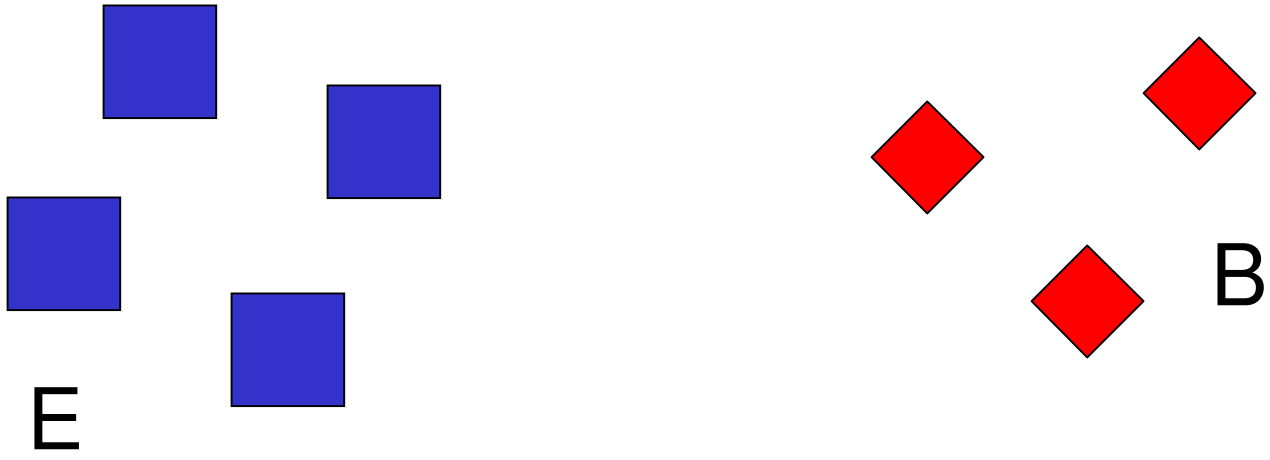

Bridge Discovery Algorithm

E pings E to load forwarding database

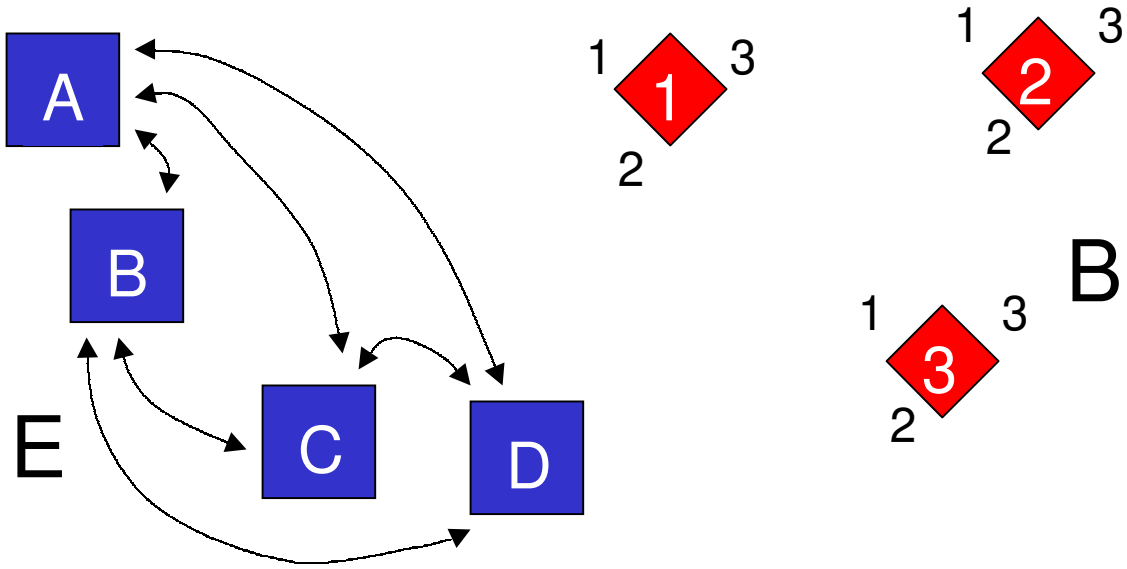

Bridge Discovery Algorithm

Read bridge forwarding databases for E

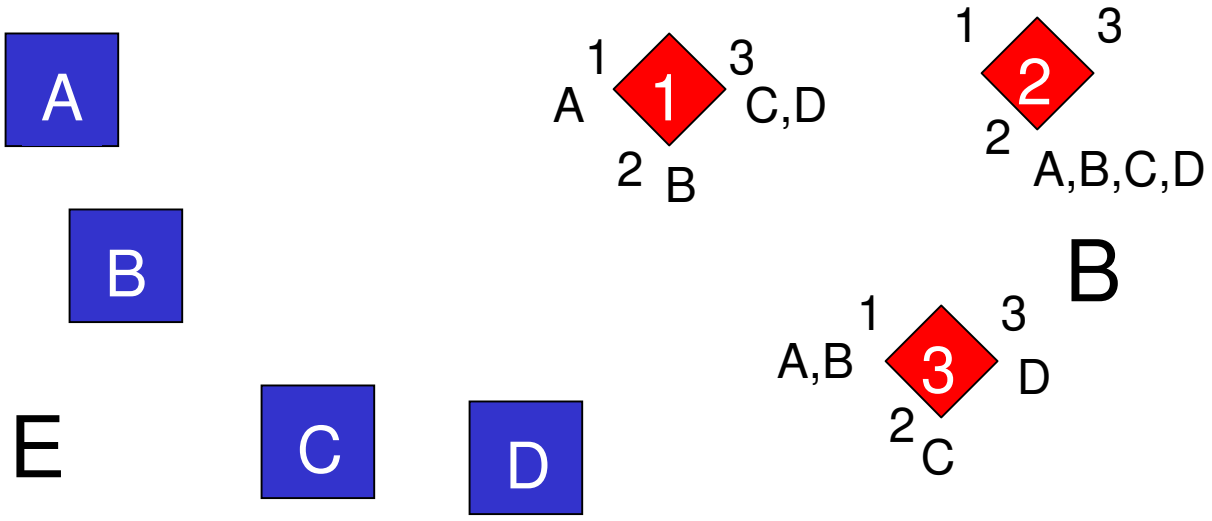

Bridge Discovery Algorithm

Prune unused bridges

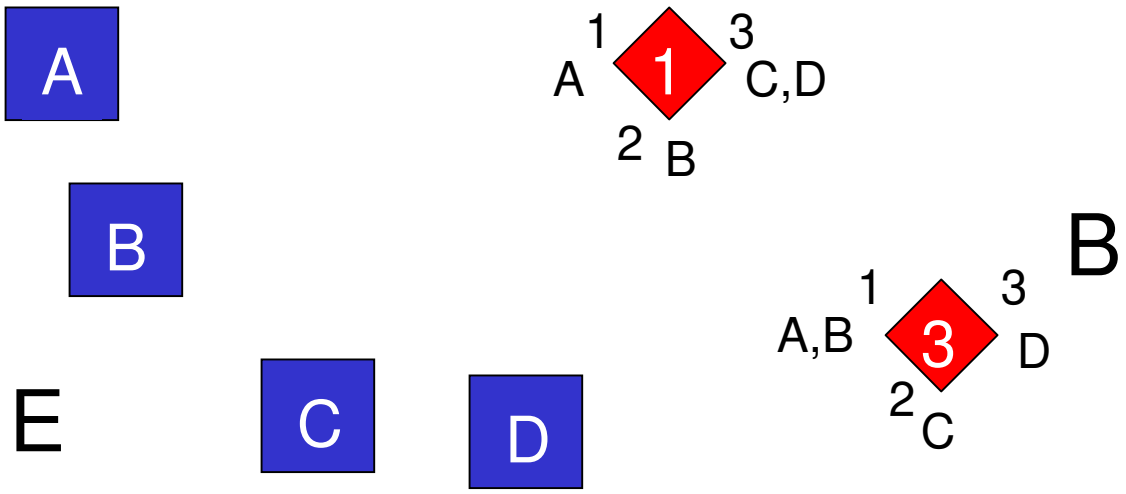

Bridge Discovery Algorithm

E pings B

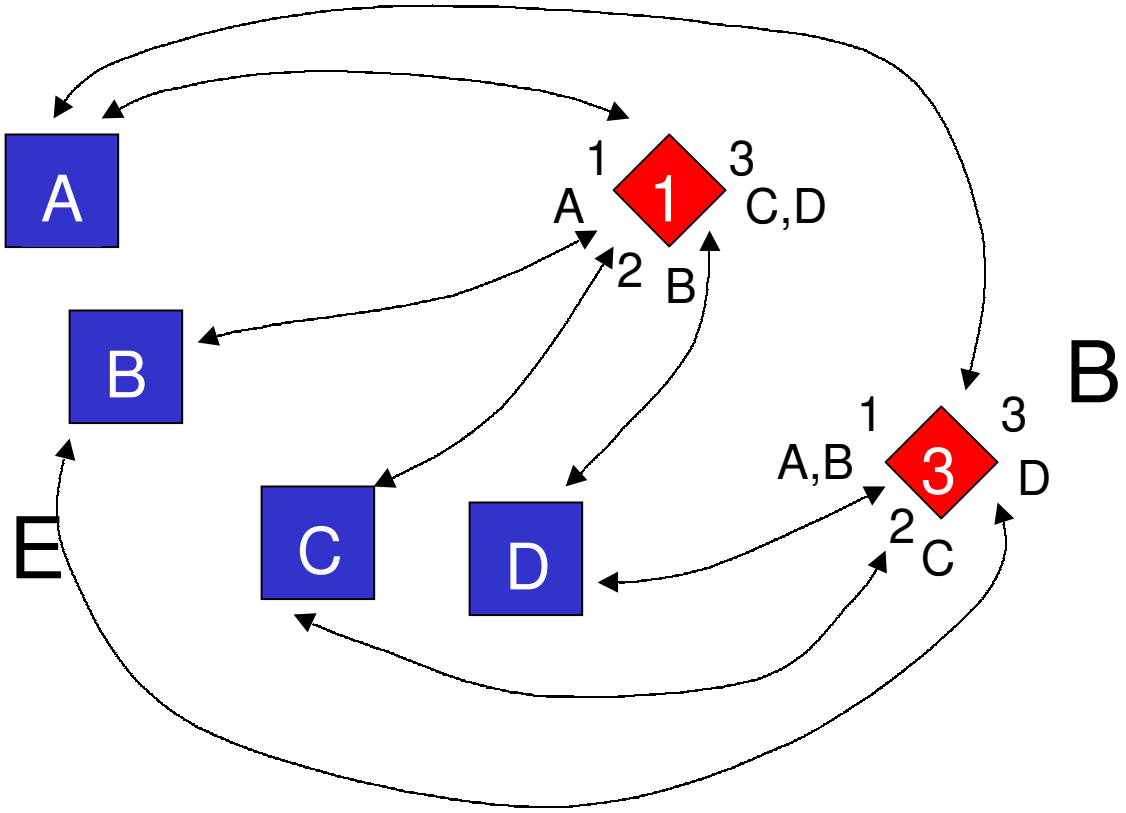

Bridge Discovery Algorithm

Read bridge forwarding databases for B

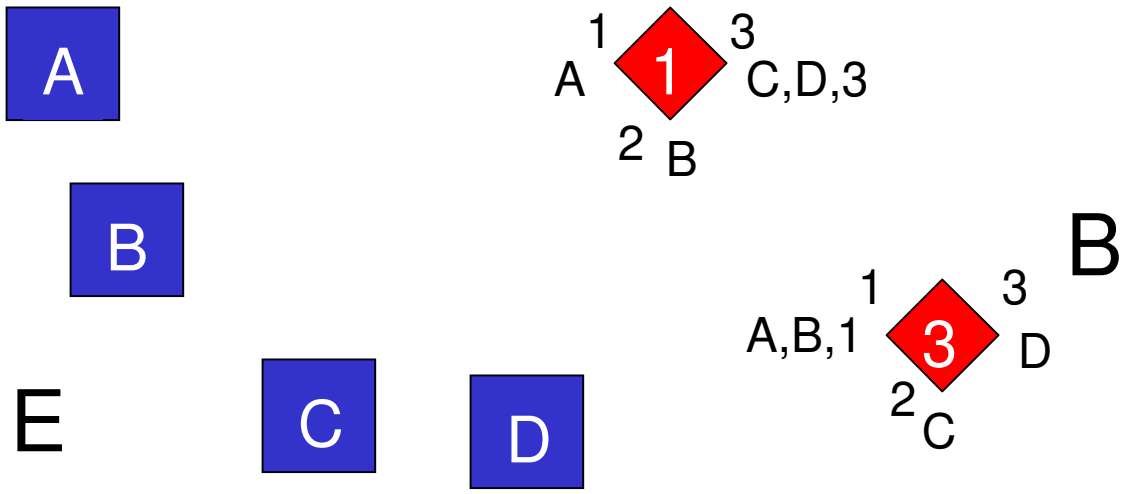

Bridge Discovery Algorithm

Bridges form spanning tree

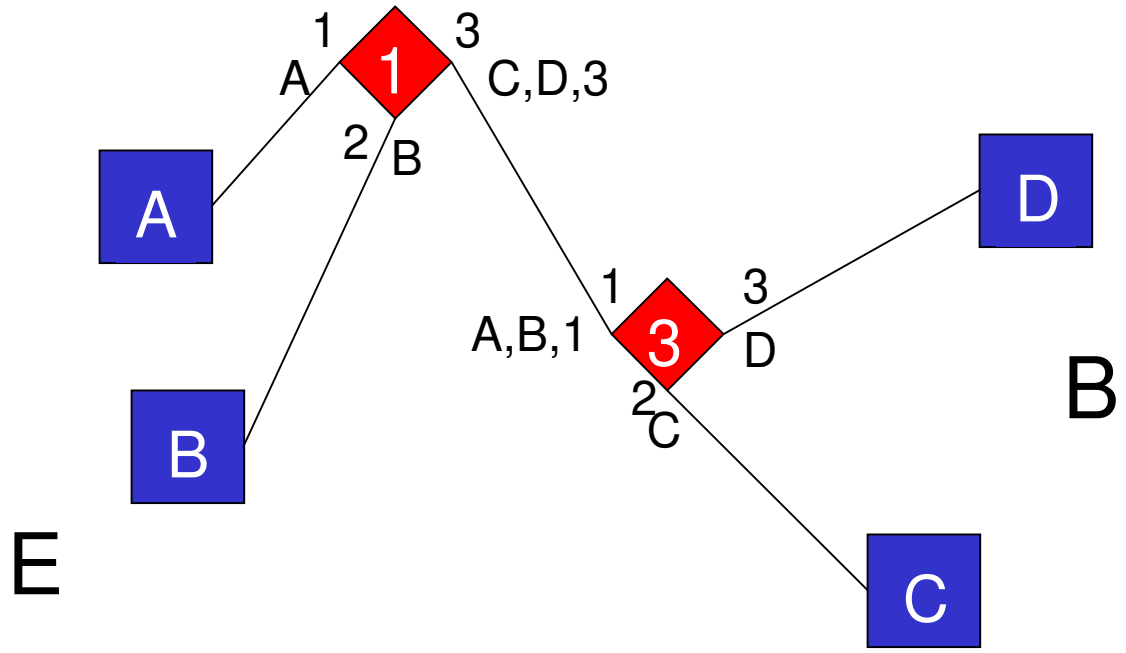

Bridge Example

CMU CS Department Ethernet

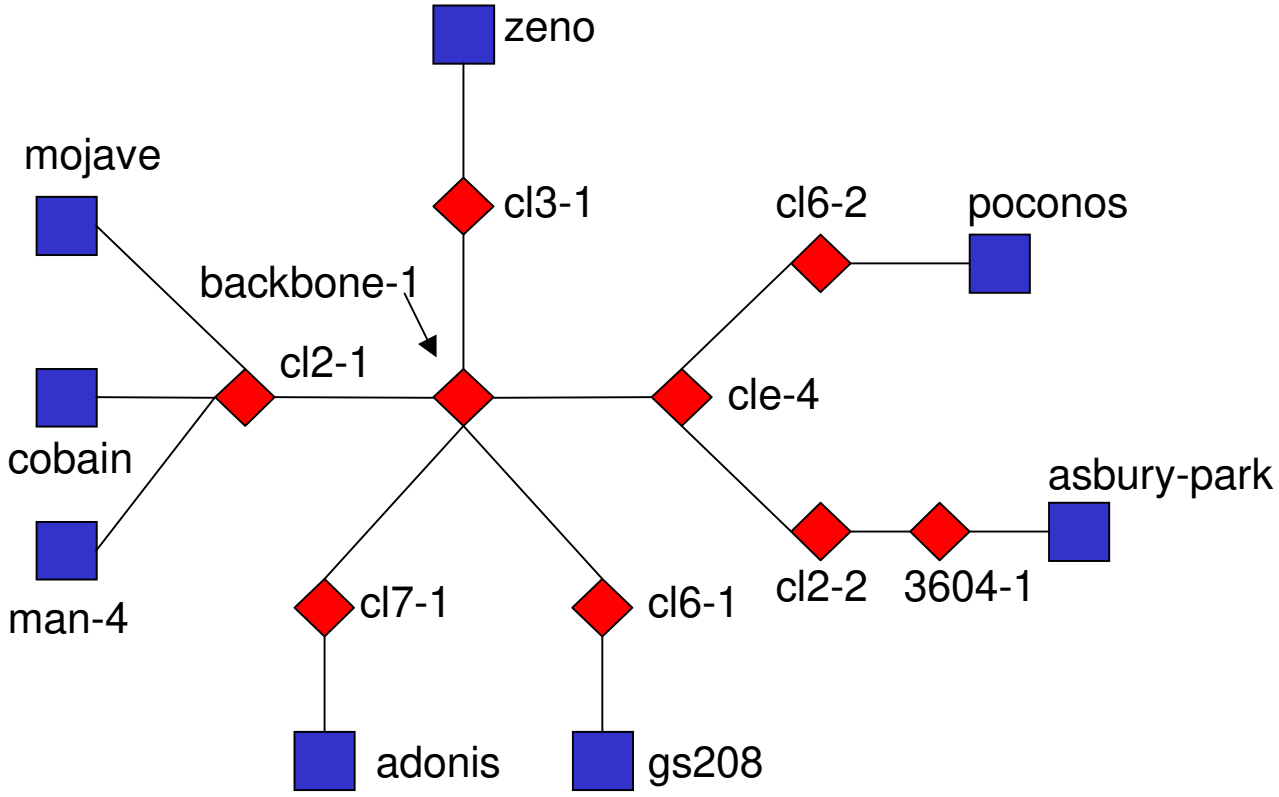

Problems with Bridge Discovery

Complexity is low, O(B (E + B))

Constant factor is high

- Bridge “security”
- Caches can be flushed

Fortunately, bridge topology changes very infrequently.

Discovery can be done offline, performance prediction is still fast.

Practical Considerations

Administration - accessibility

Standards are loose

Implementation is far from ideal

Working with networking community could help improve this situation!

Practical Considerations

Proprietary solutions

End-to-end measurement

Types of competing traffic

Conclusions

Network-based prediction

- scales well
- provides topology information
- accuracy appears promising
- can be implemented today

Being incorporated into Remos

Major challenges remain to make this usable for everyone!

`http://www.cs.cmu.edu/Groups/CMCL/remulac`