Unmanaged Internet Protocol

Taming the Edge Network Management Crisis

Bryan Ford
Massachusetts Institute of Technology
HotNets II – November 21, 2003

“Ubituitous Networking”

- What is it?
- Why isn't it here yet?
- How can we make it work?

A Ubiquitous Networking Scenario

A Ubiquitous Networking Scenario

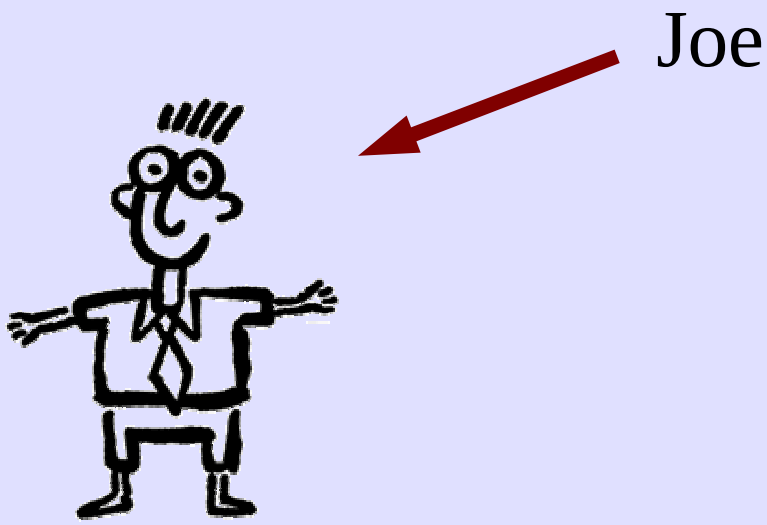

A Ubiquitous Networking Scenario

A Ubiquitous Networking Scenario

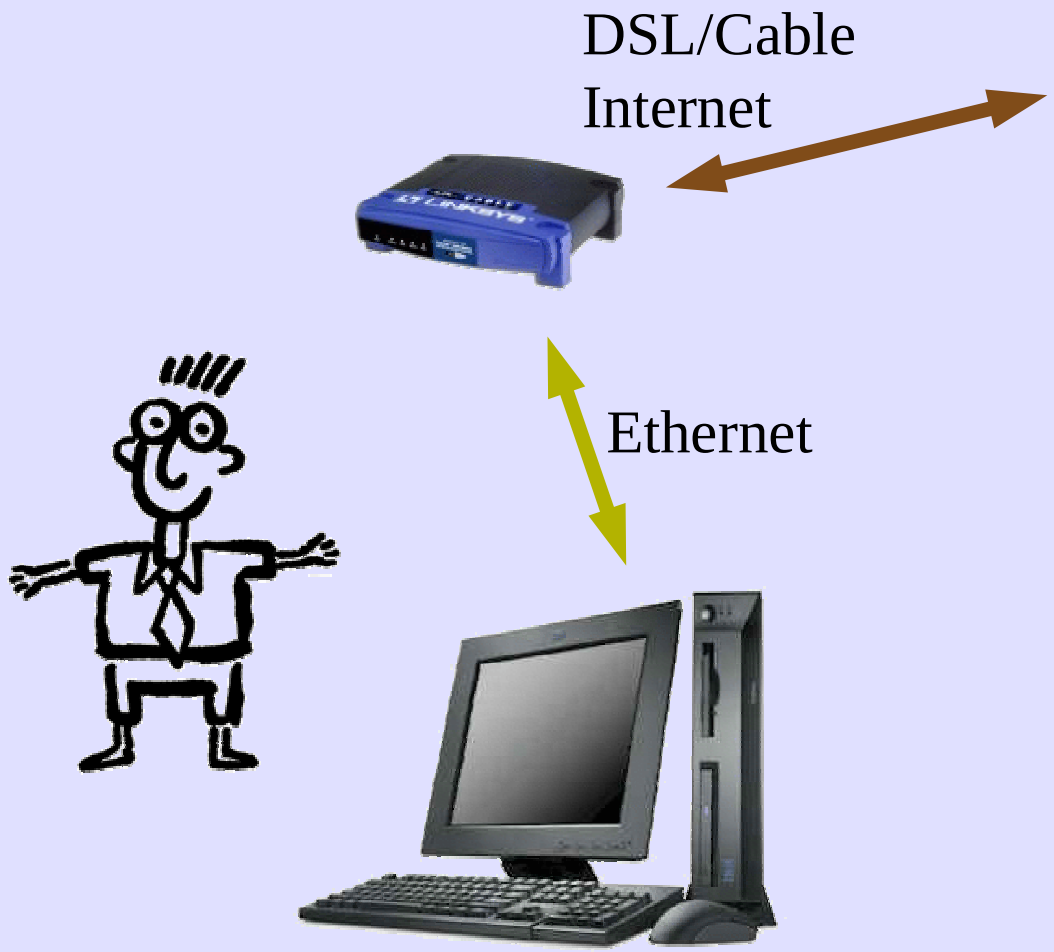

A Ubiquitous Networking Scenario

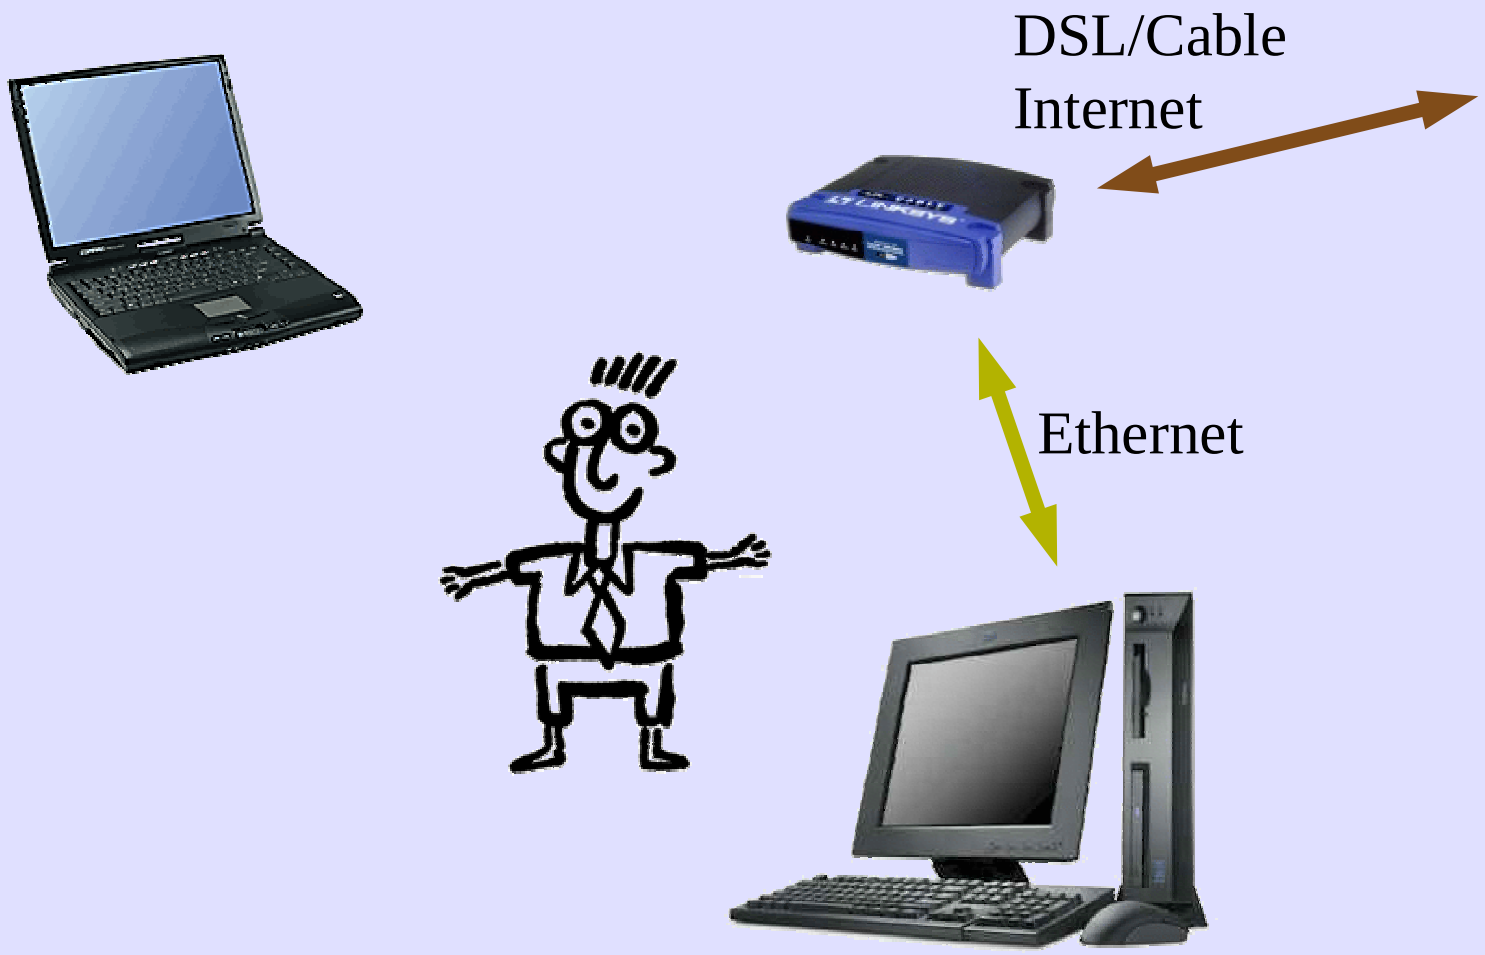

A Ubiquitous Networking Scenario

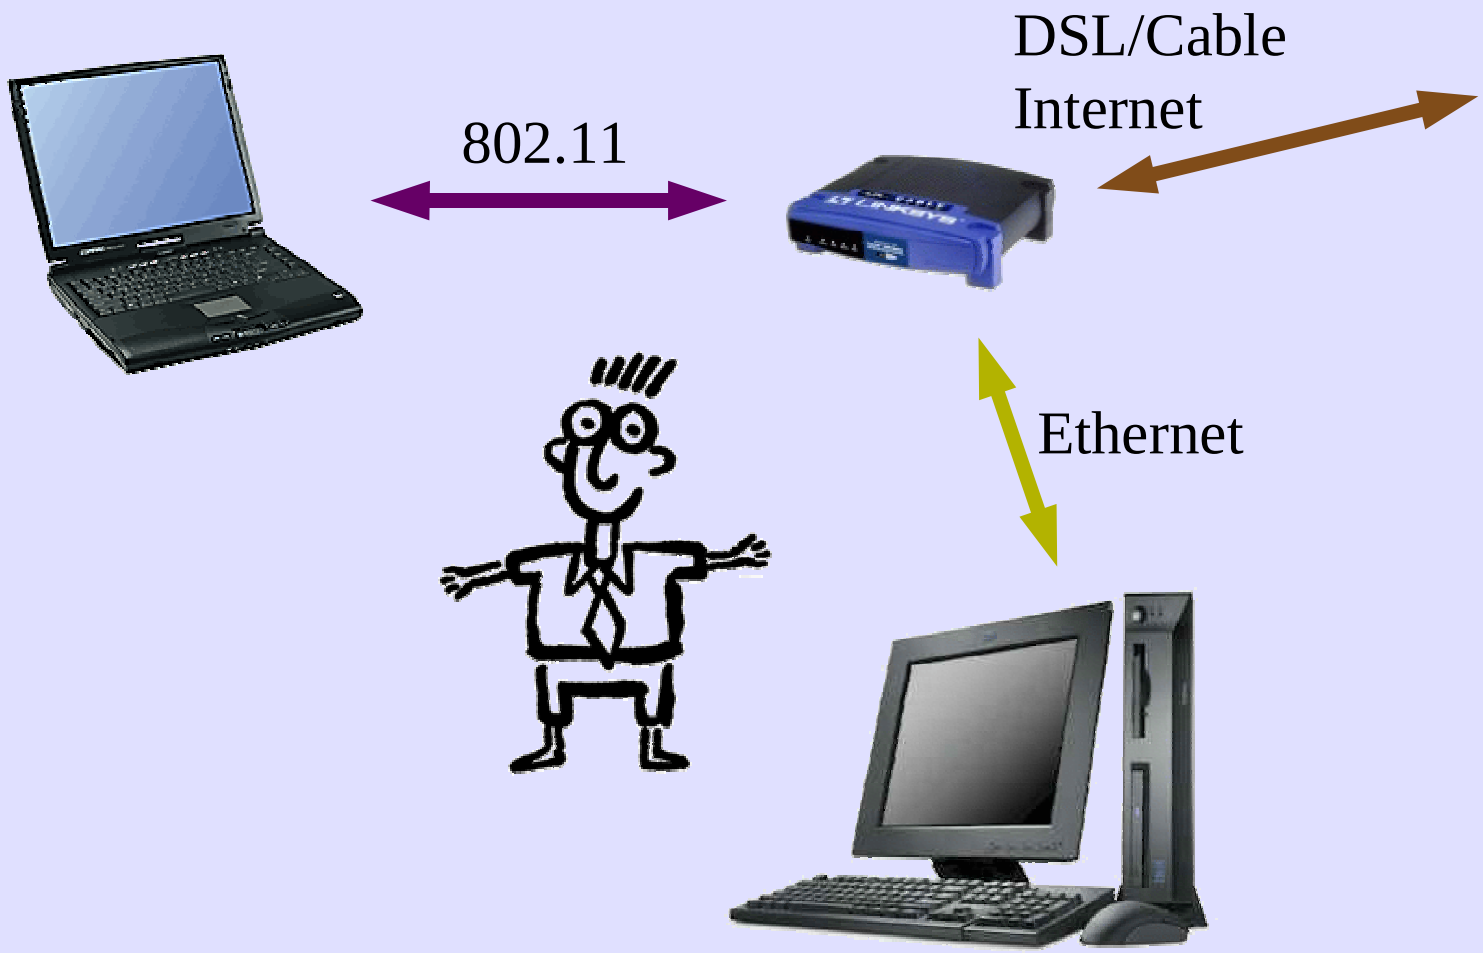

A Ubiquitous Networking Scenario

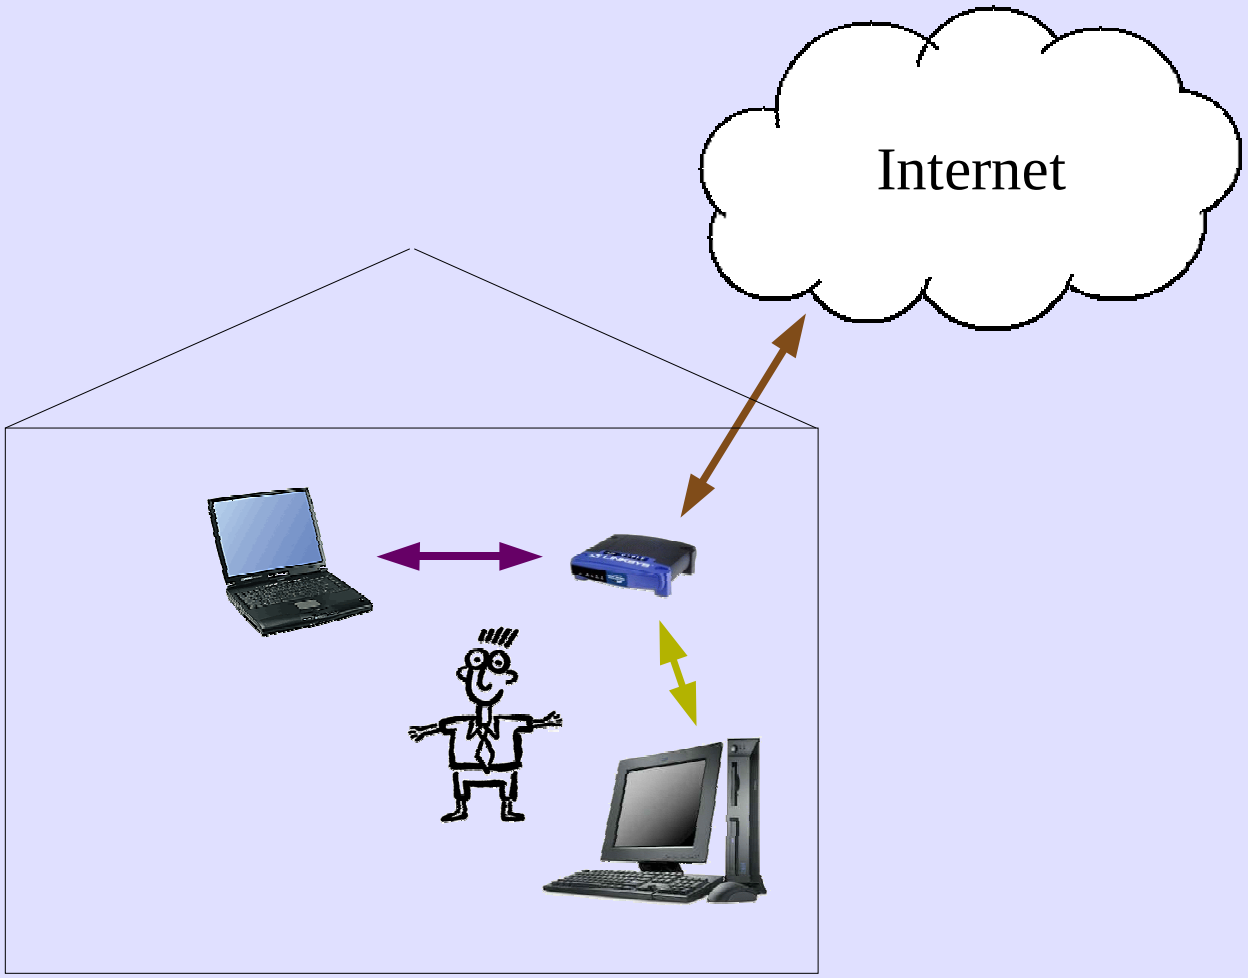

A Ubiquitous Networking Scenario

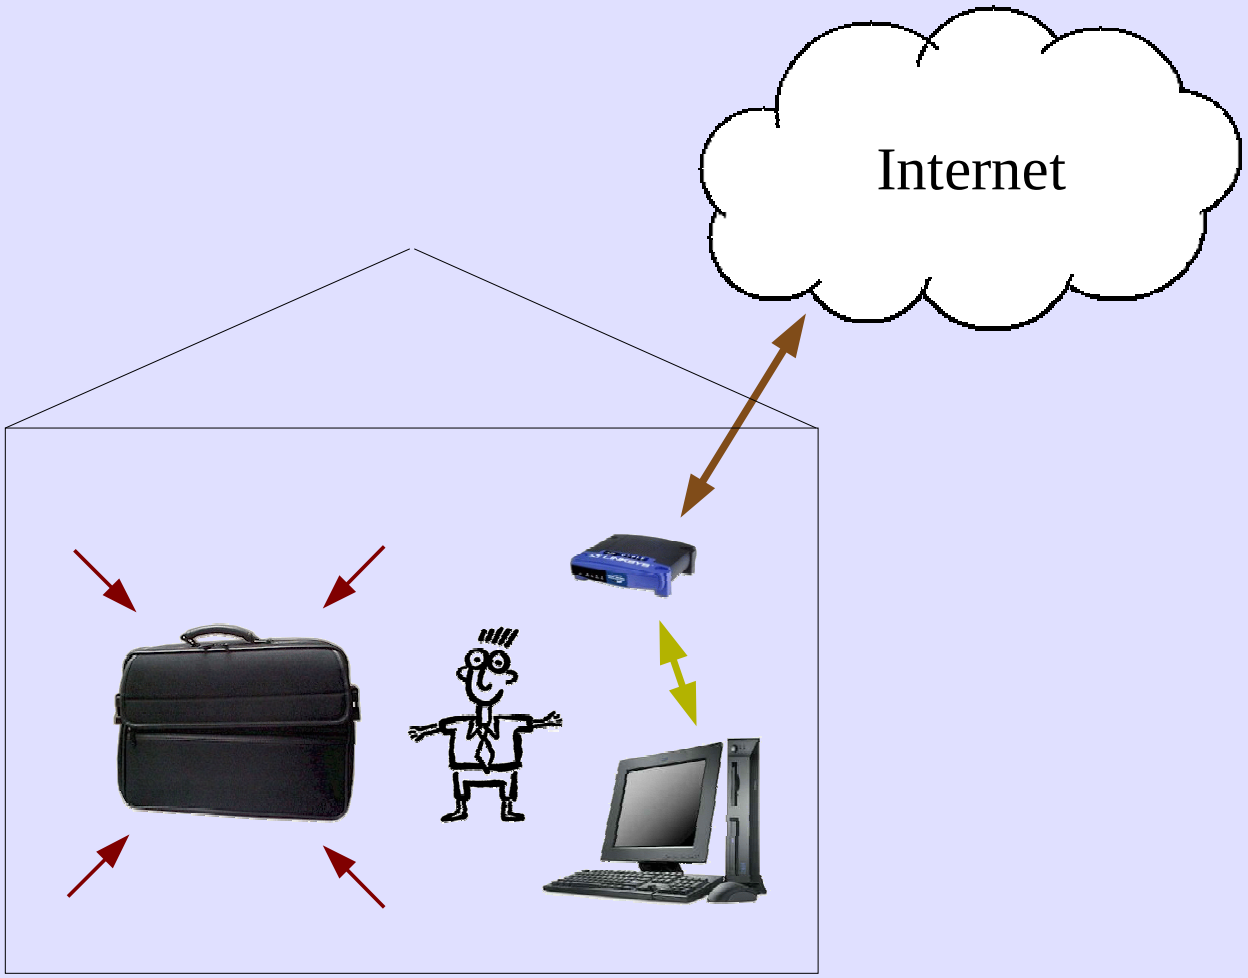

A Ubiquitous Networking Scenario

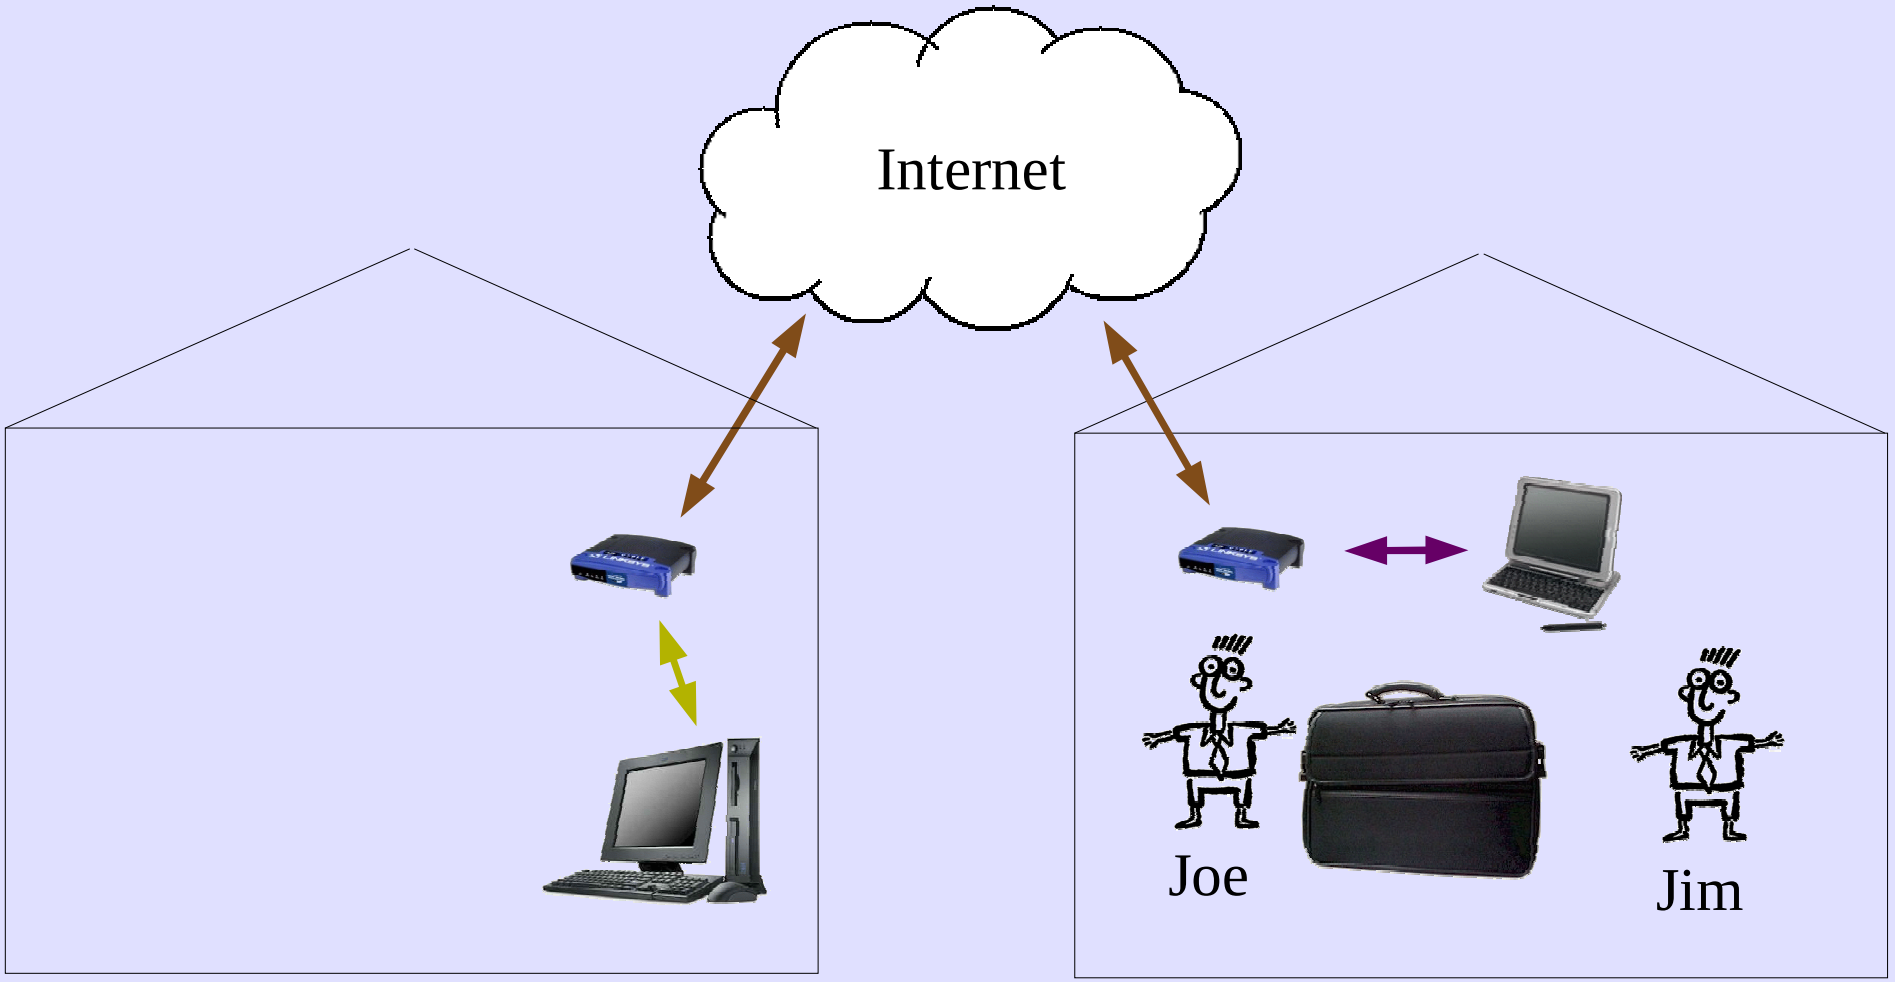

A Ubiquitous Networking Scenario

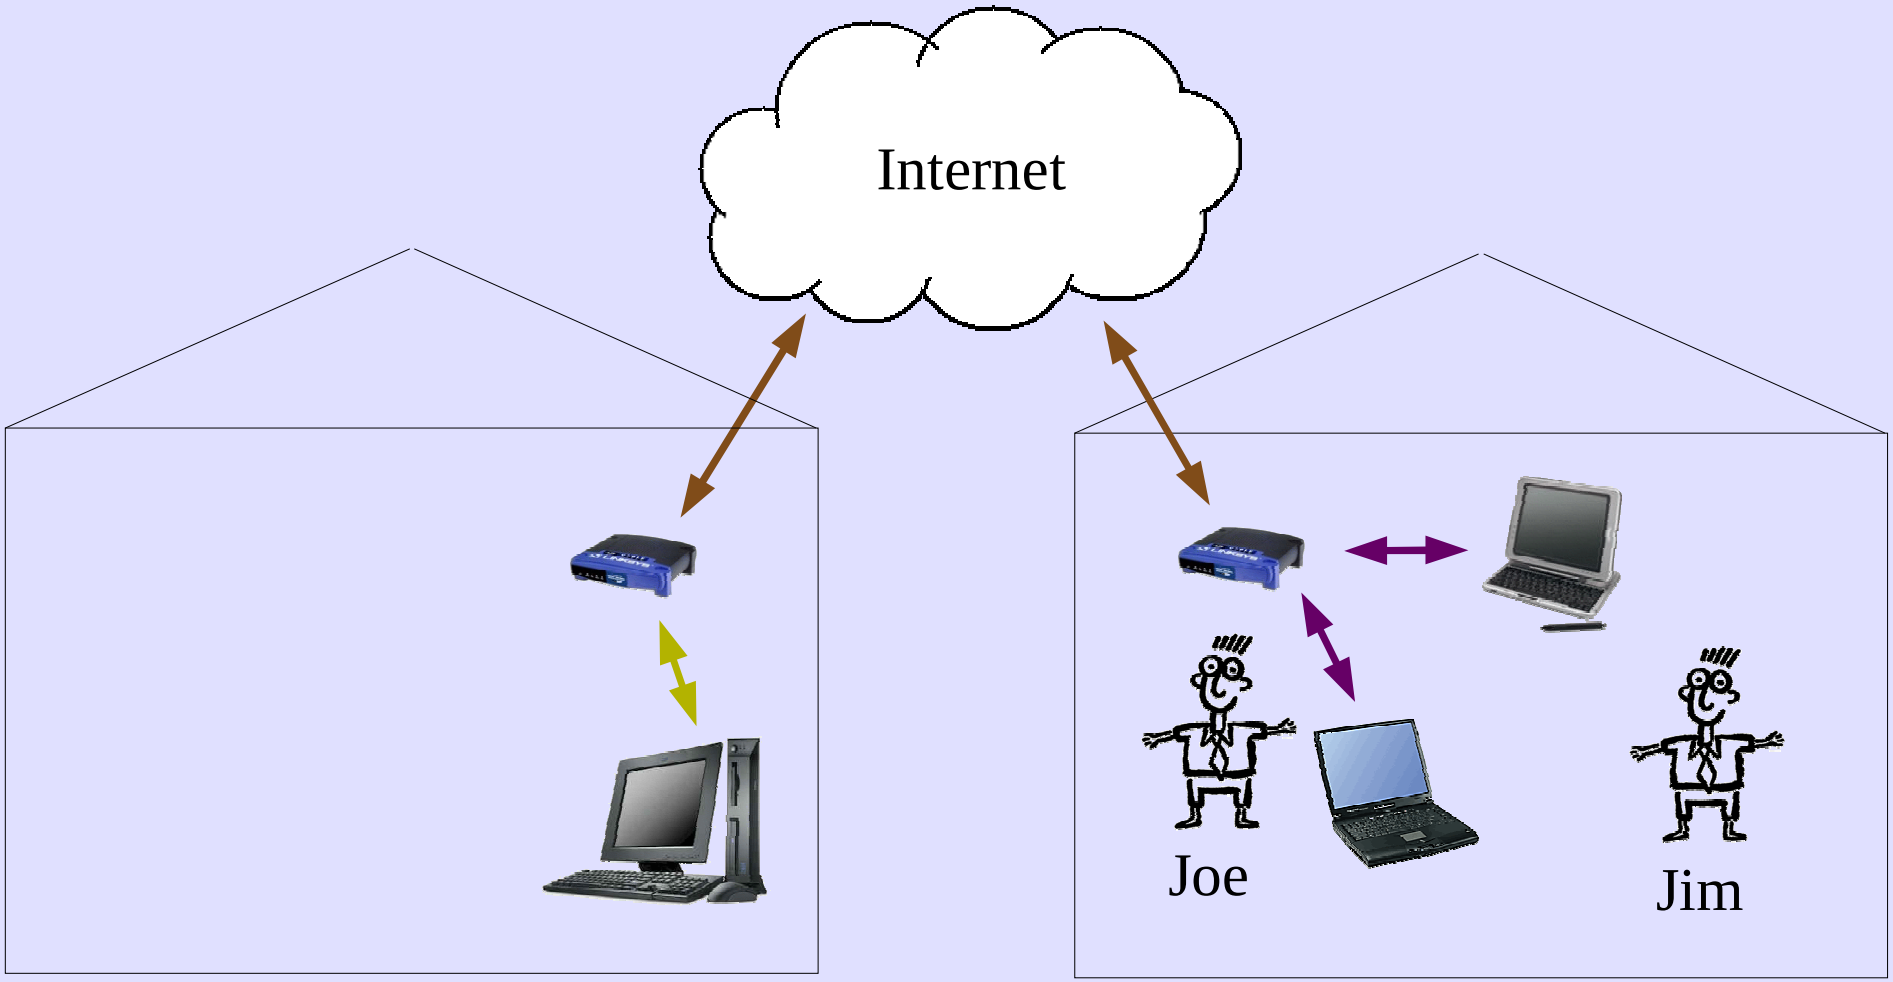

A Ubiquitous Networking Scenario

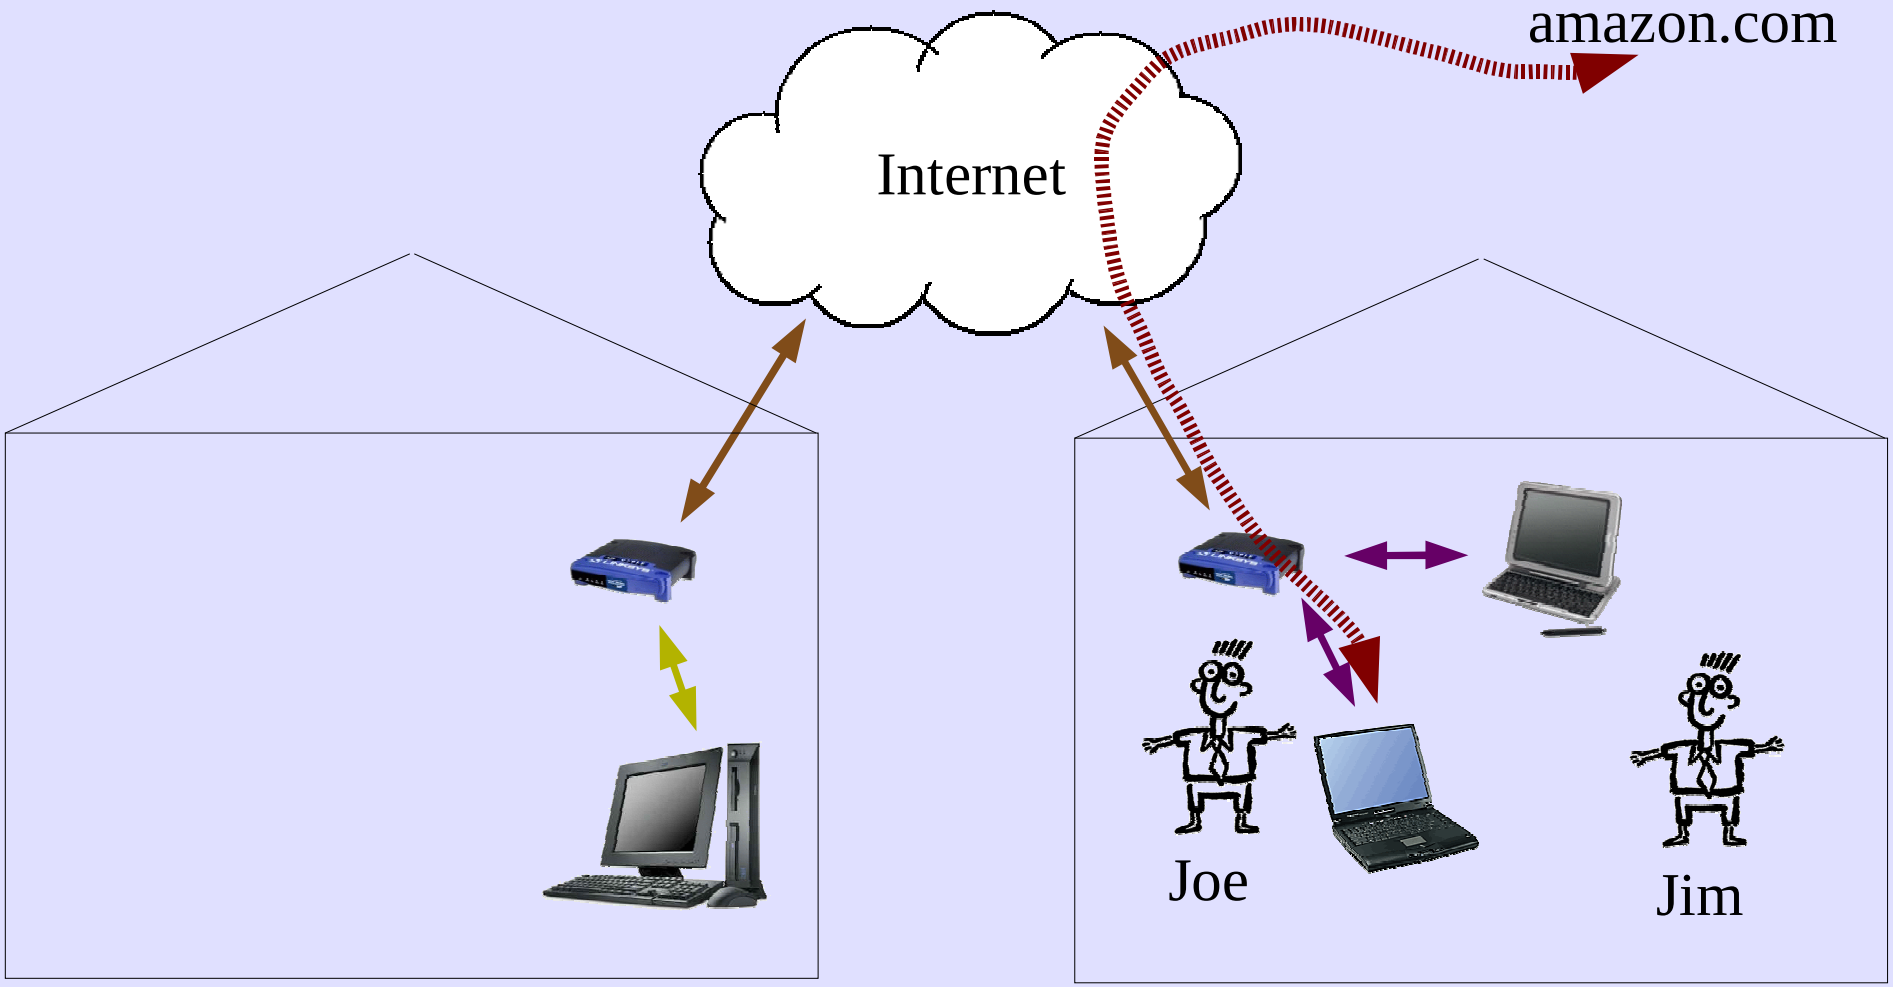

A Ubiquitous Networking Scenario

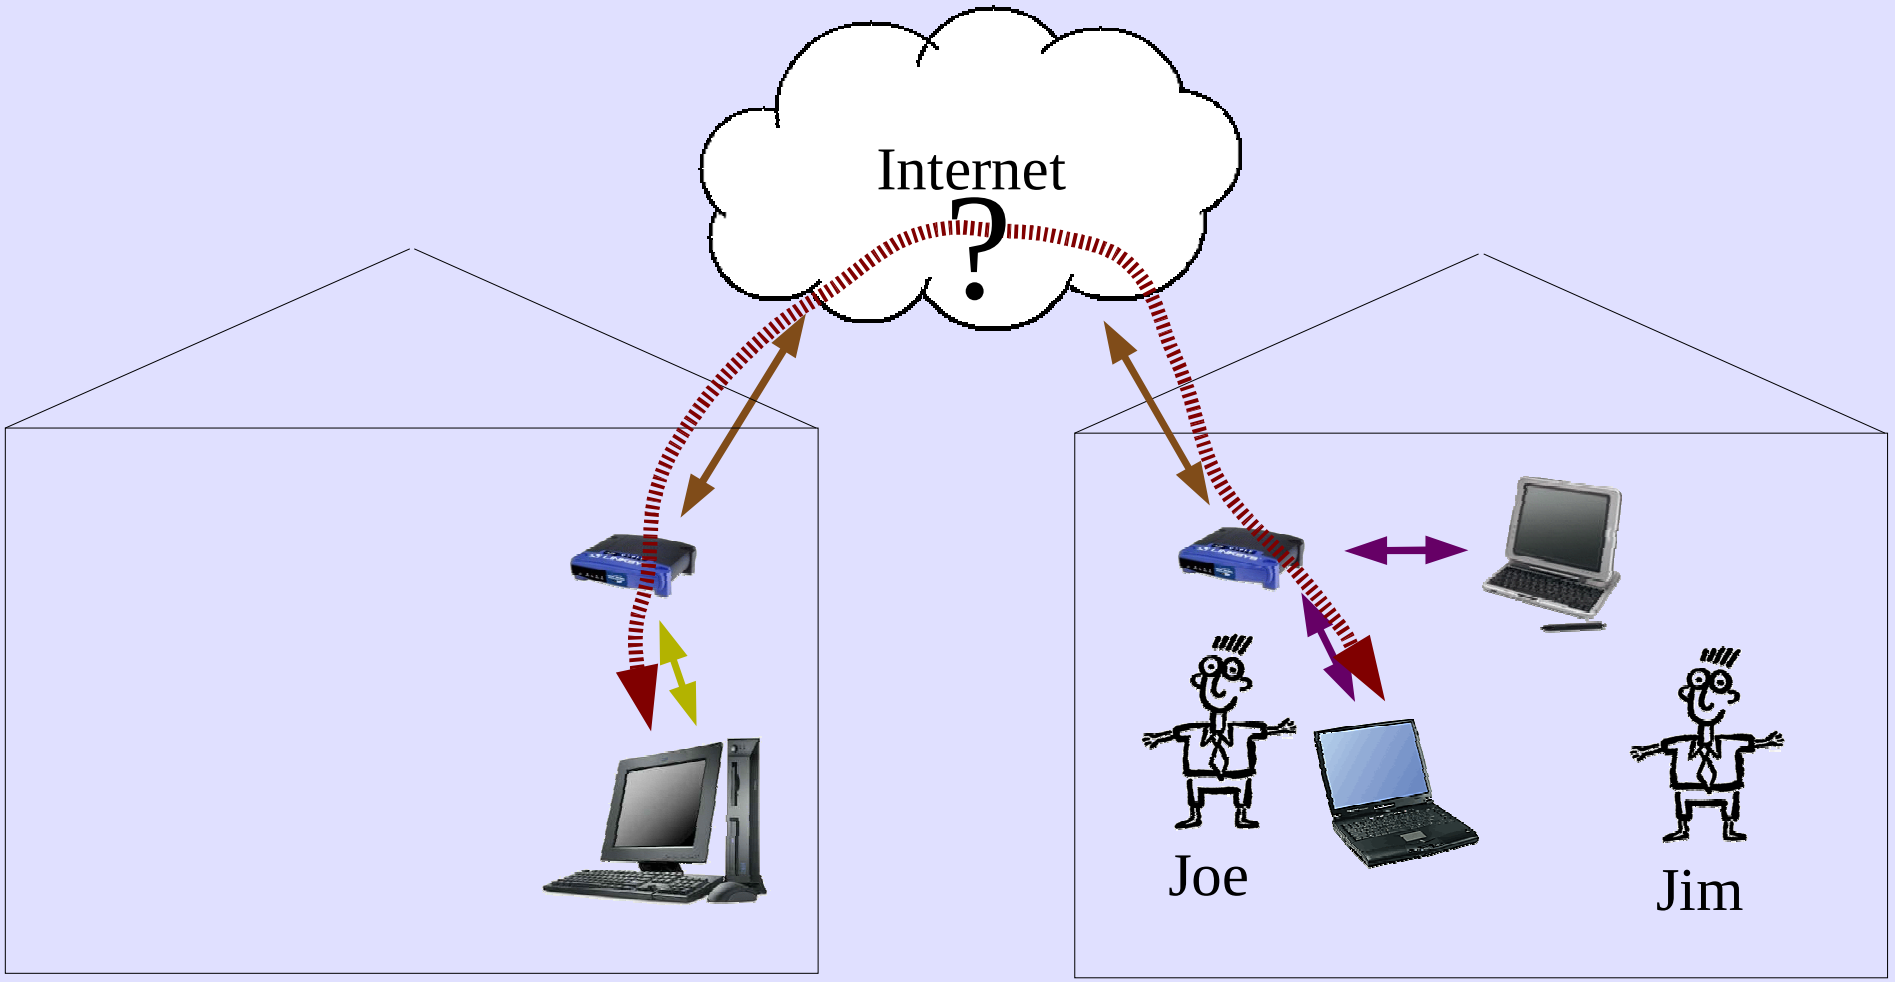

A Ubiquitous Networking Scenario

“NAT?”

A Ubiquitous Networking Scenario

“NAT?”

“Dynamic DNS?”

A Ubiquitous Networking Scenario

“NAT?”

“Dynamic DNS?”

“Mobile IP?”

A Ubiquitous Networking Scenario

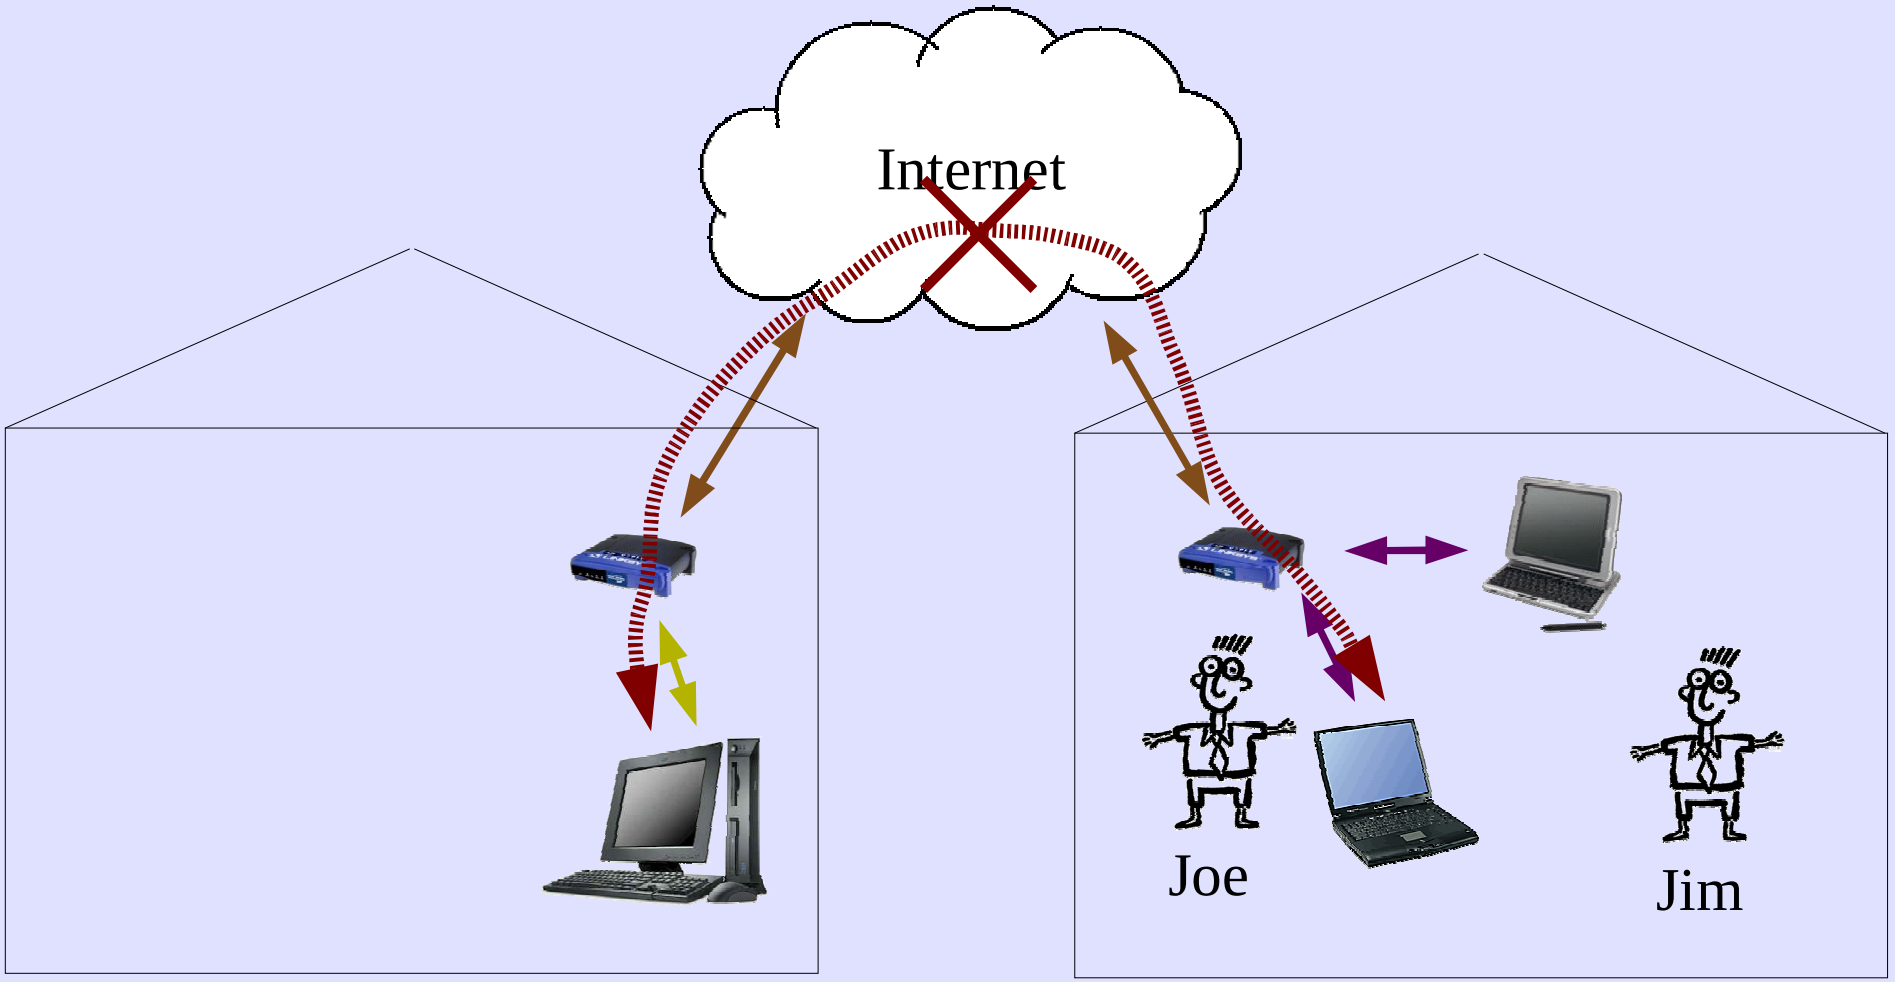

A Ubiquitous Networking Scenario

Joe

Jim

A Ubiquitous Networking Scenario

Joe

Jim

A Ubiquitous Networking Scenario

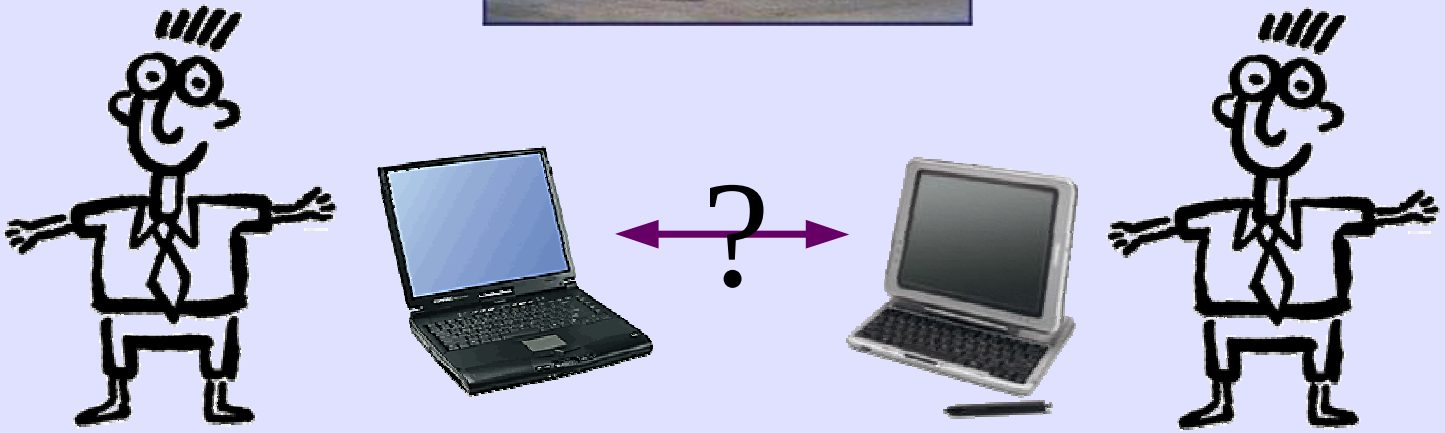

A Ubiquitous Networking Scenario

"Ad-hoc mode?"

Joe

Jim

A Ubiquitous Networking Scenario

“Ad-hoc mode?”

“DHCP?”

Joe

Jim

A Ubiquitous Networking Scenario

“Ad-hoc mode?”

“DHCP?”

“Static IP addresses?”

Joe

Jim

A Ubiquitous Networking Scenario

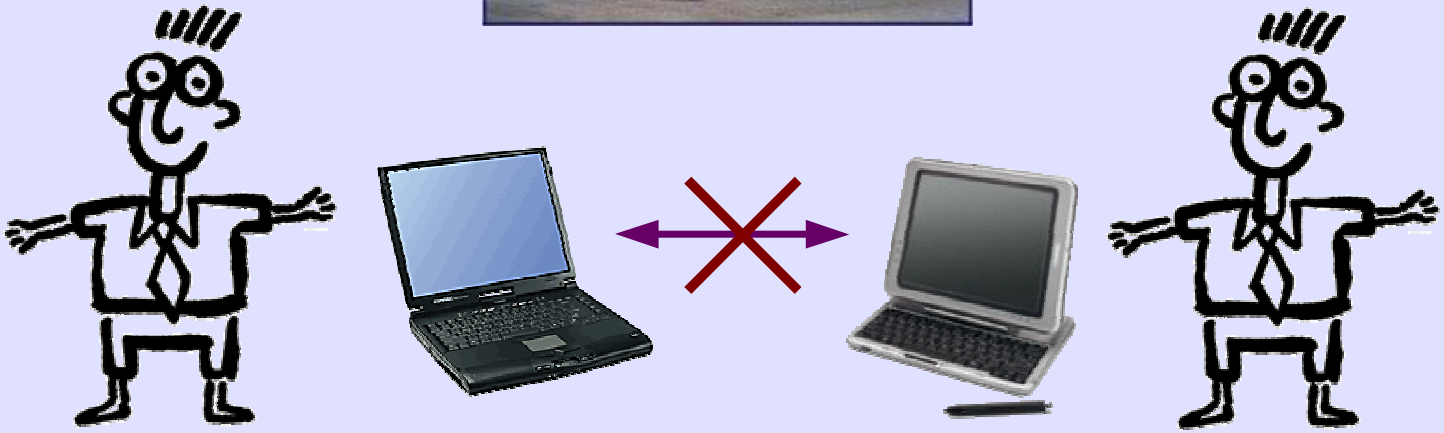

Joe

Jim

The Problem

Getting “ubiquitous networking” devices to

ubiquitously network

is way too complicated,

even when the technology is available.

Outline

- ✔ Motivation: What's wrong?
- Why doesn't ubiquitous networking work?
  - *Answer:* hierarchical address-based routing (ABR).
- How do we fix it?
  - *Answer:* scalable identity-based routing (IBR).
- A proposed identity-based routing architecture
- Conclusion

Why IP is Wrong for Edge Networks

- Hierarchical address architecture
  - Routable addresses must be allocated from central administrative authorities
  - Each node must be assigned an address:
    - Static assignment ⇒ inconvenient, requires knowledge
    - DHCP ⇒ nodes can't talk at all without DHCP server
  - Address hierarchy must reflect topology
    - Node mobility ⇒ address instability, broken connections
  - Good for scalability, bad for useability

What about ad-hoc routing protocols?

- Landmark, DSR, DSDV, AODV, etc.
- A big step in the right direction, *but:*
  - Not scalable beyond local area (≈hundreds of nodes)
- Good for outdoor geek parties
- Useless for Joe and Jim

We need ad-hoc routing

at Internet-Wide Scale

We need ad-hoc routing
at Internet-Wide Scale

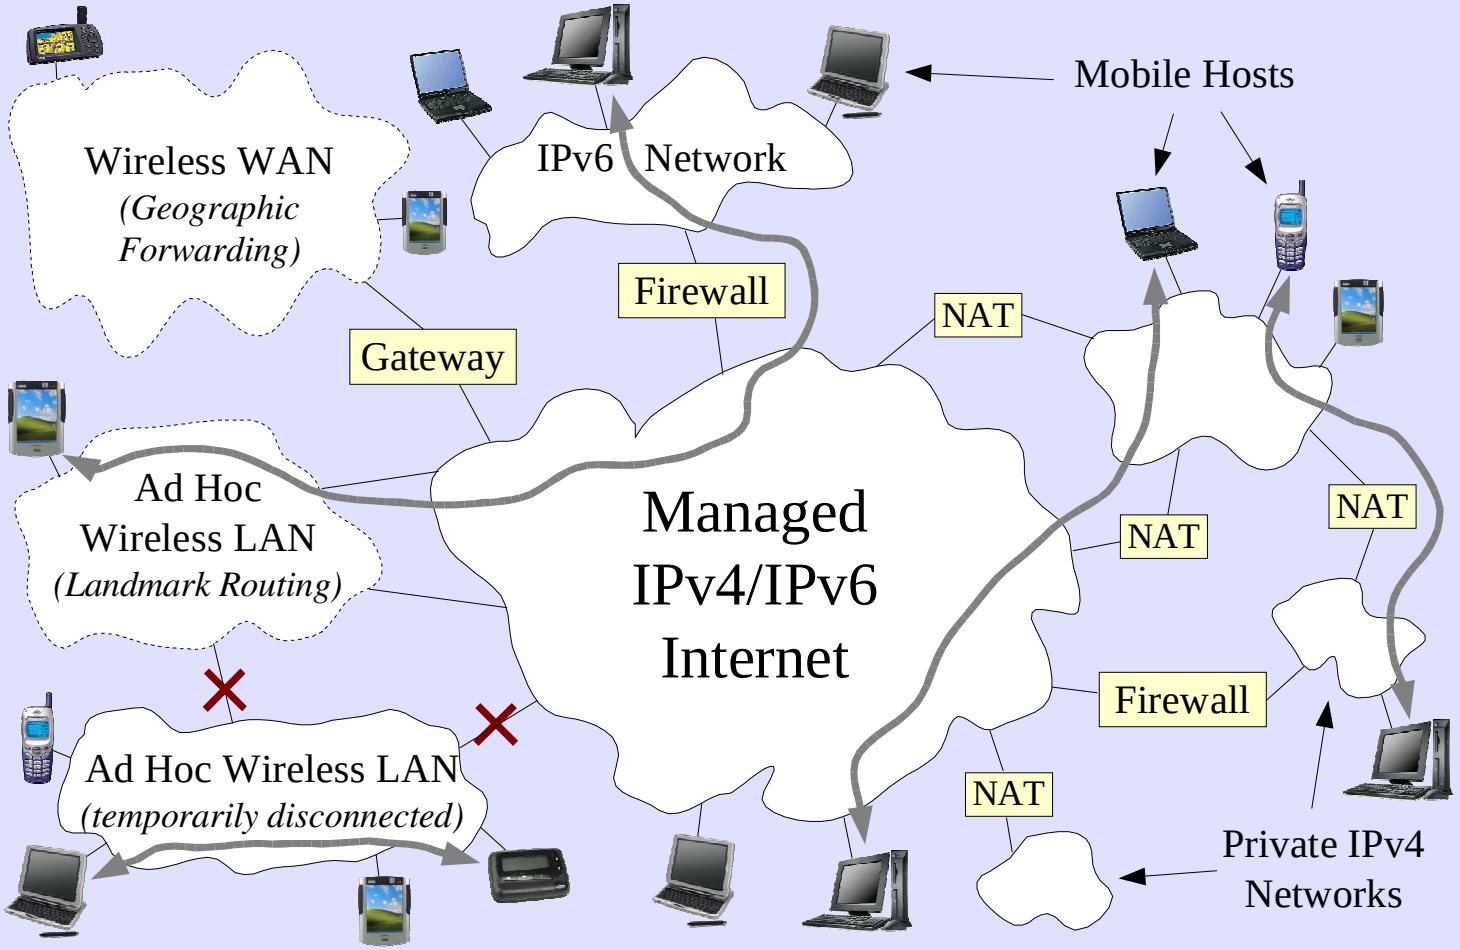

A Proposed Identity-Based Routing Protocol Architecture

UIP: “Unmanaged Internet Protocol”

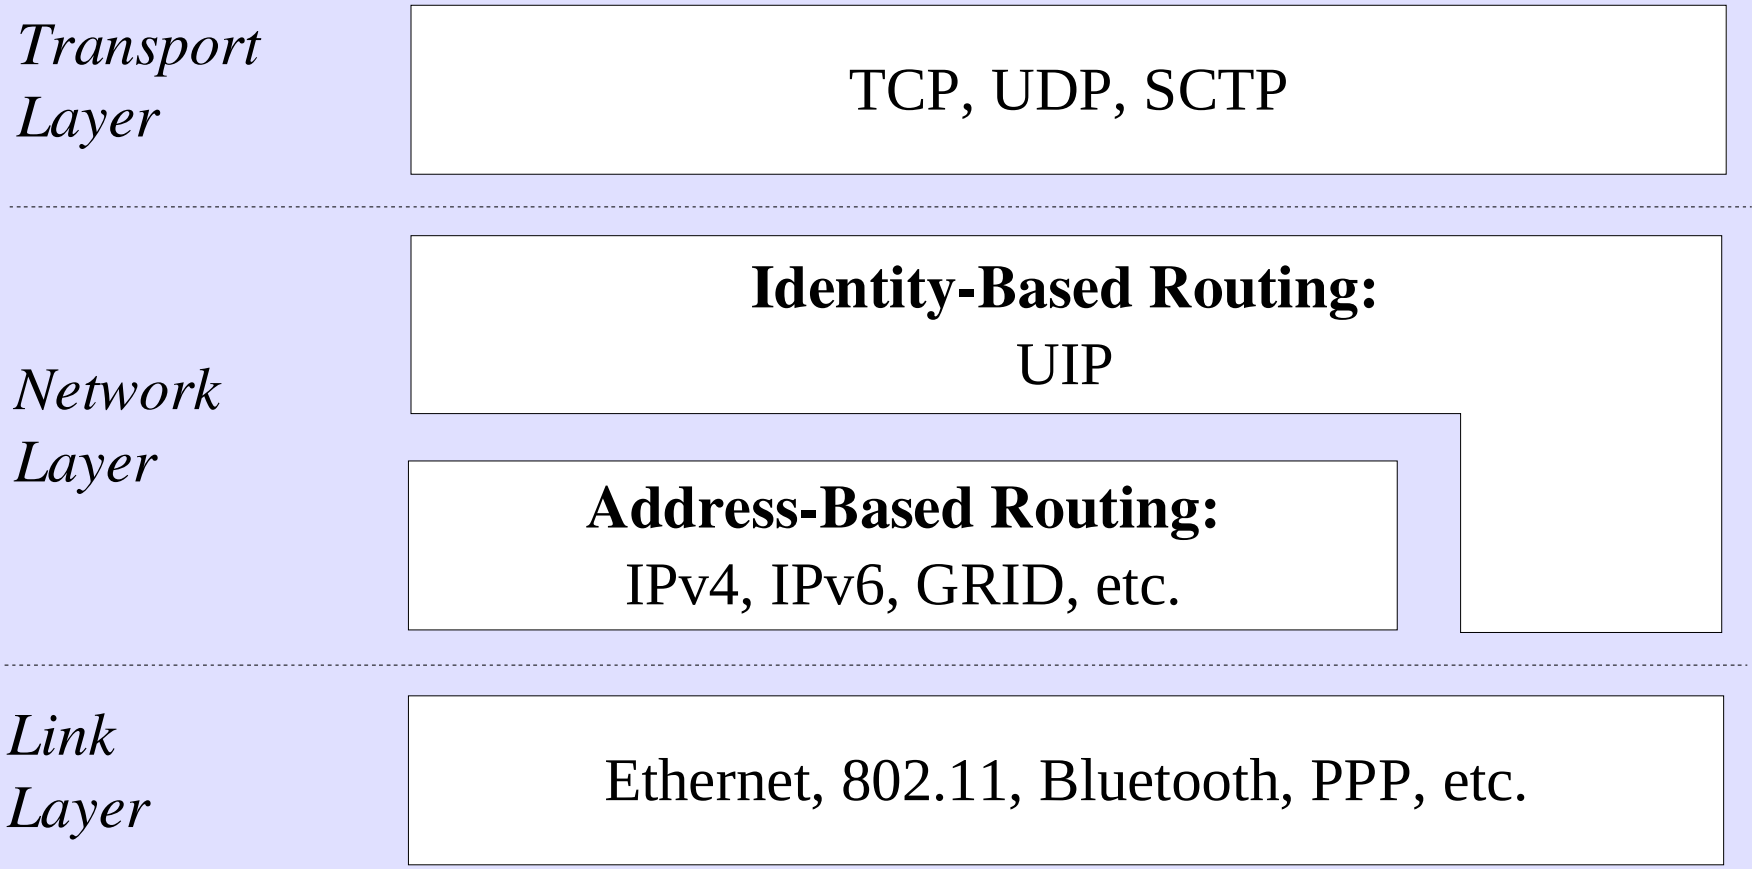

Key Properties of UIP

- “Unmanaged” = “Manages Itself”
  - No central authority required to hand out addresses
  - No explicit maintenance of routing and forwarding
  - No futzing or broken connections when nodes move
- Operates both:
  - Over IPv4/IPv6 as a scalable overlay network
  - Directly over Ethernet and other link layers

UIP Node Identifiers

Cryptographic hash of node's public key (ala HIP):

- **Automatically generated** by node itself
- **Stable** for as long as owner of node desires
- **Self-authenticating** for privacy and integrity
- **Topology-independent** for host mobility
- **Globally unique,** cryptographically unforgeable

Why This Is Hard

- Must give up hierarchical address architecture, but still get scalability to millions of nodes!
- Can't require each node to maintain and propagate state about every other node
- ...*But theoretically feasible:*
  Arias et al. “Compact Routing with Name Independence,” SPAA 2003

Idea!

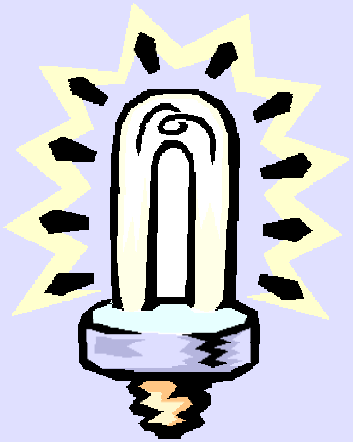

What about adapting Peer-to-Peer Distributed Hash Table (DHT) lookup algorithms?

The Intuition

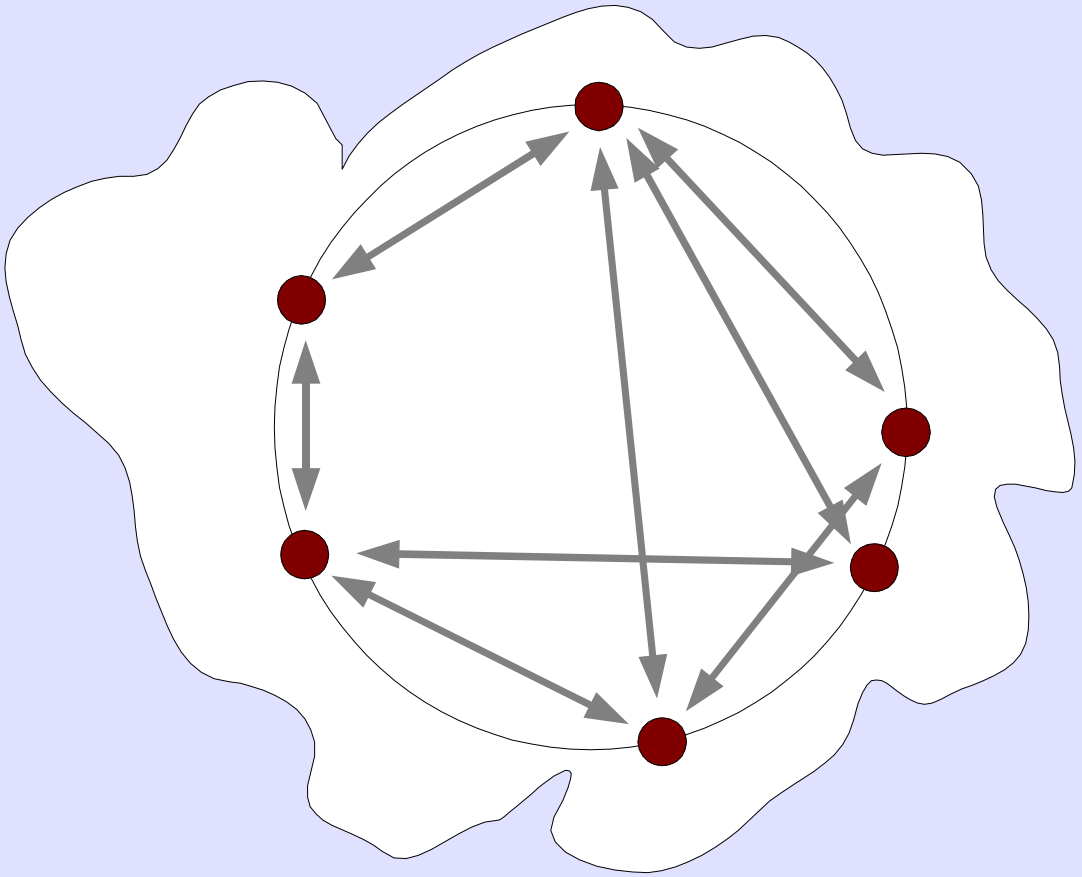

- DHTs provide:
- Lookup on topology-independent keys
- *O(log n)* state, maint. traffic per node

The Intuition

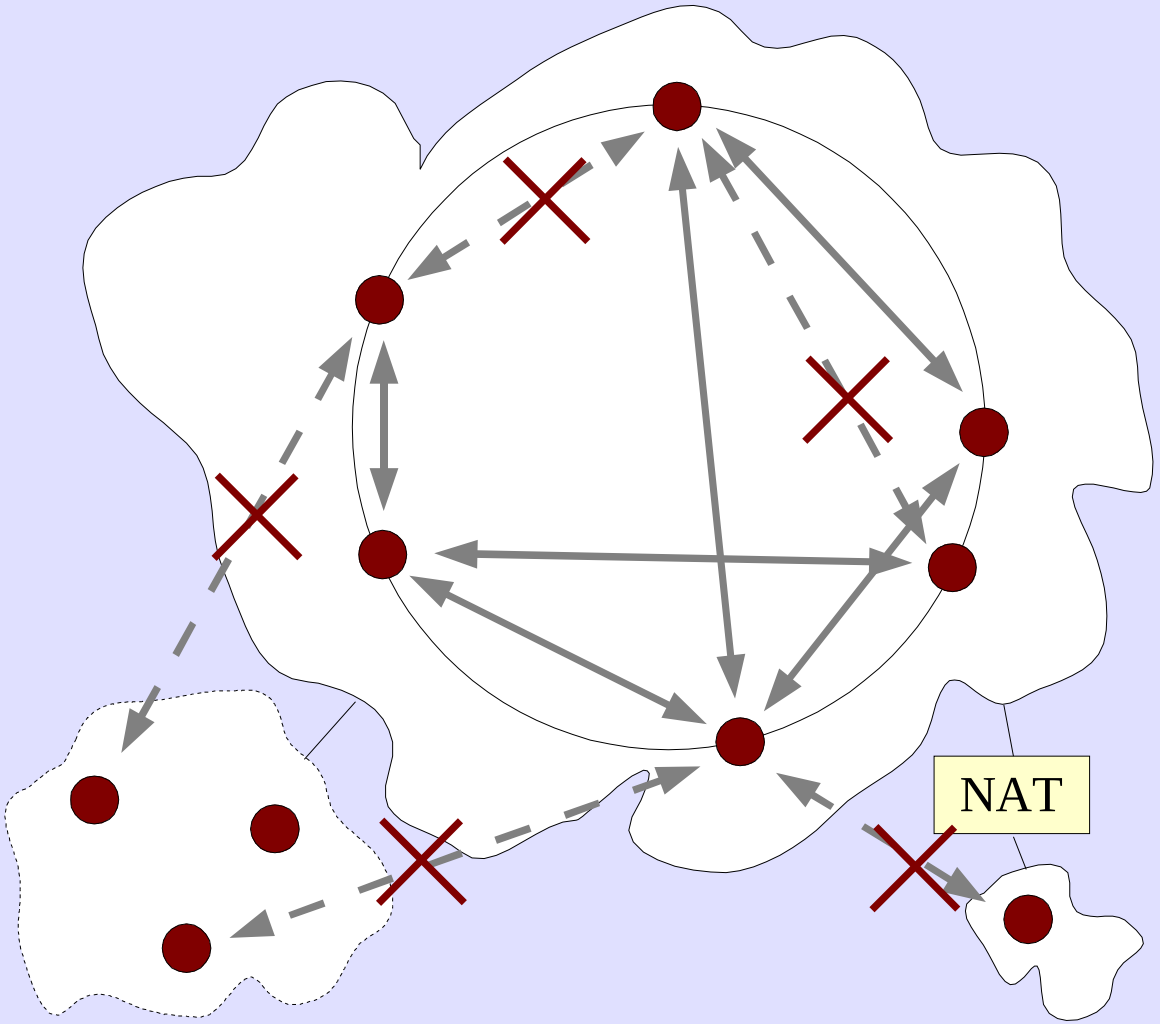

- DHTs *don't*:
- Forward around discontinuities
- Traverse NATs (usually)
- Route between Internet & Ad-hoc Networks

A First Approximation

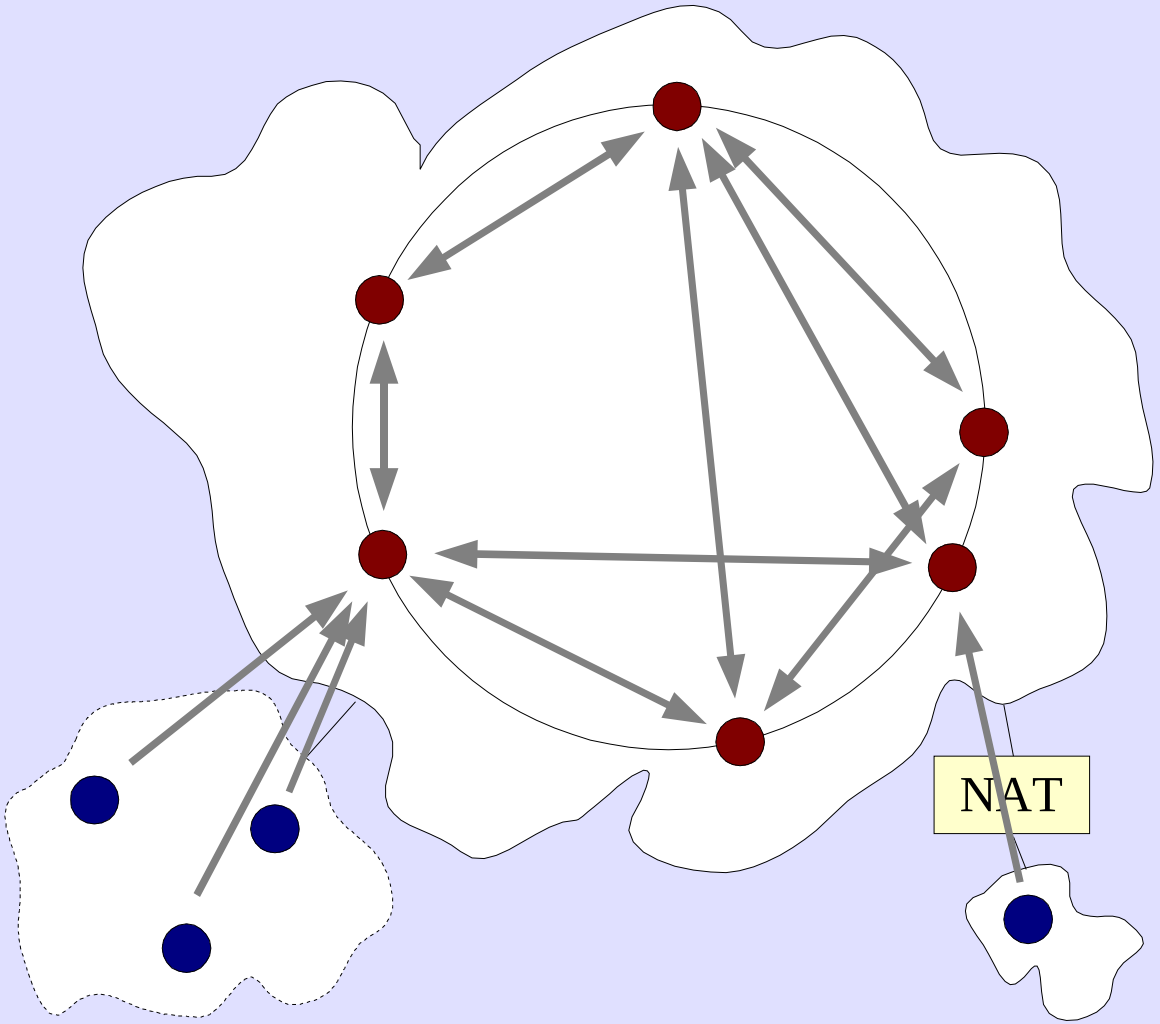

- Two-level stratification
- “Core” nodes maintain DHT
- “Edge” nodes reachable thru core nodes
- Example: *i3*

A First Approximation

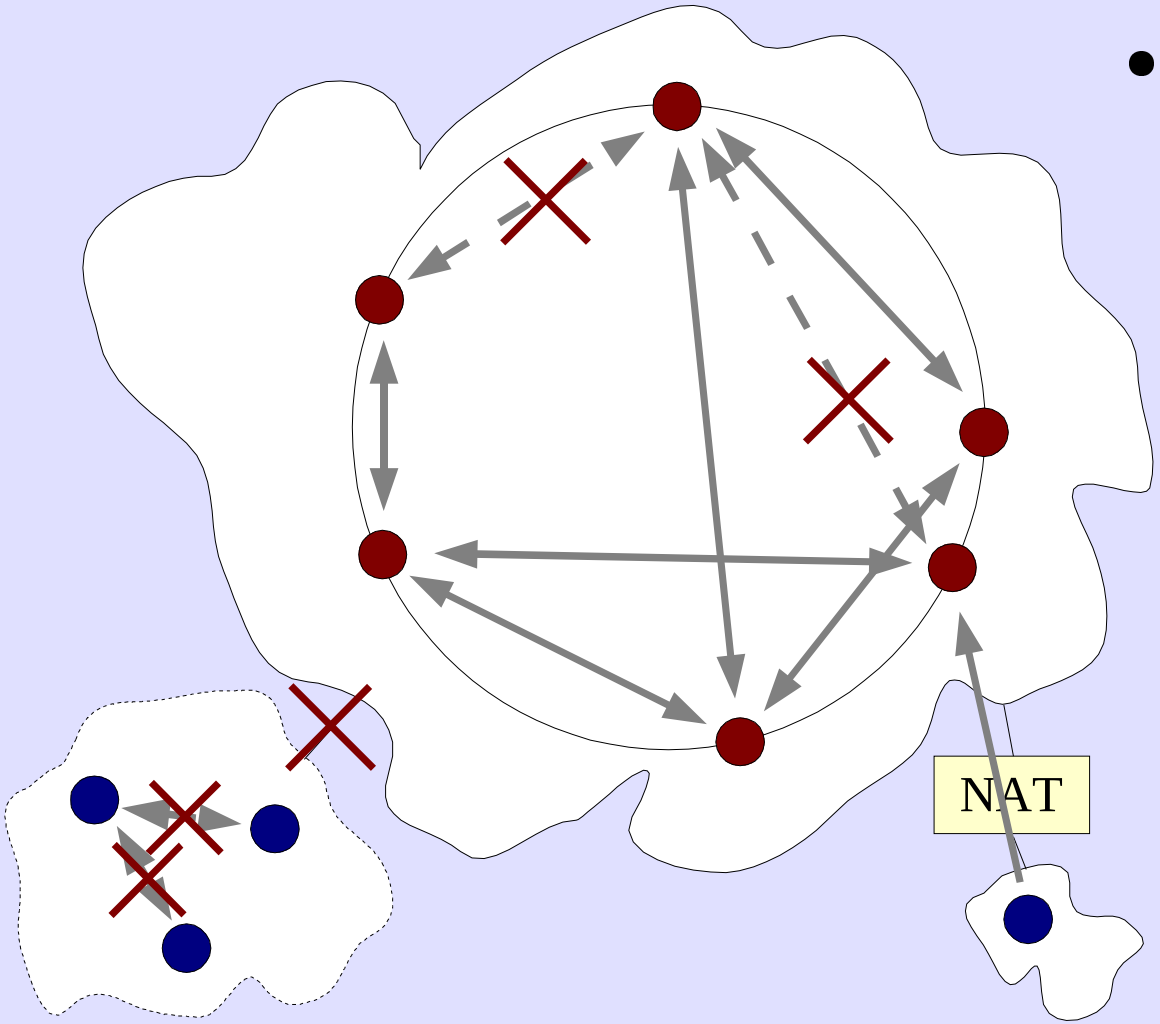

- Limitations:
  - Must configure whether node is “core” or “edge”
  - Discontinuities in “core” network
  - Disconnected edge nodes can't talk

What We Want

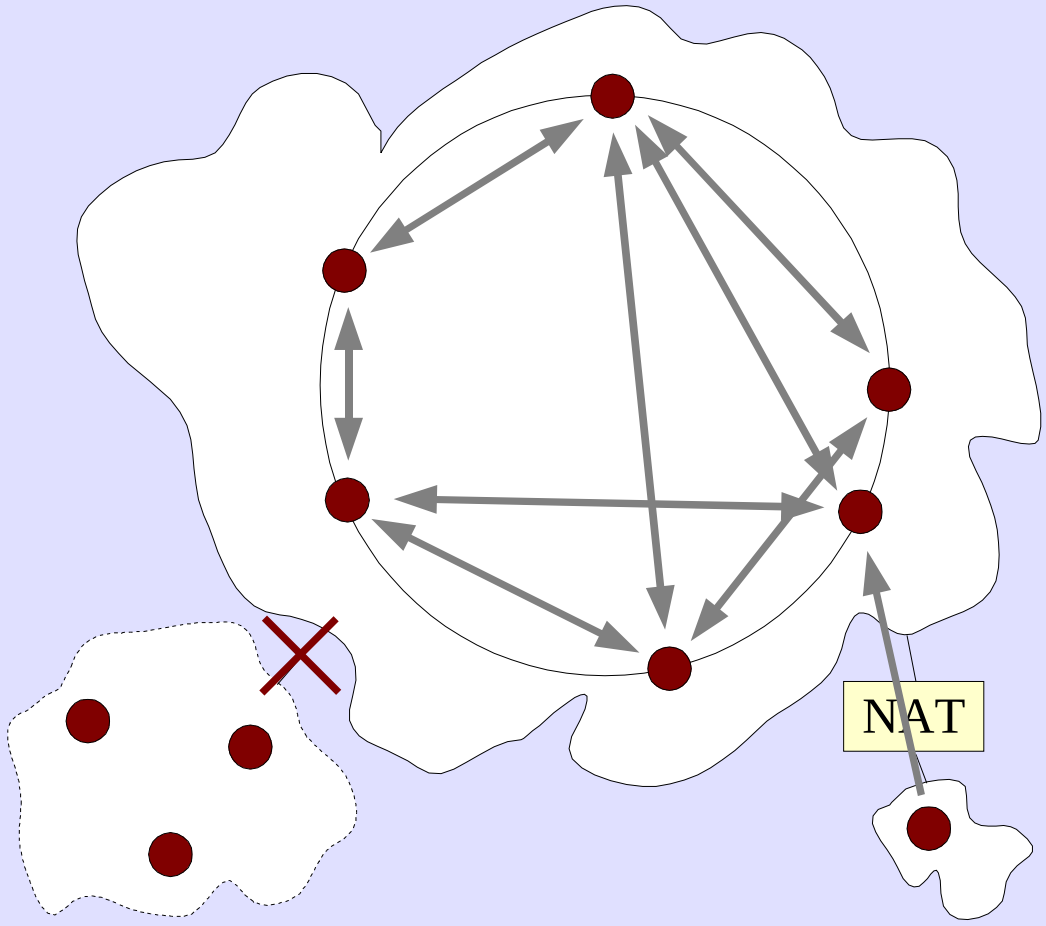

- Unstratified

What We Want

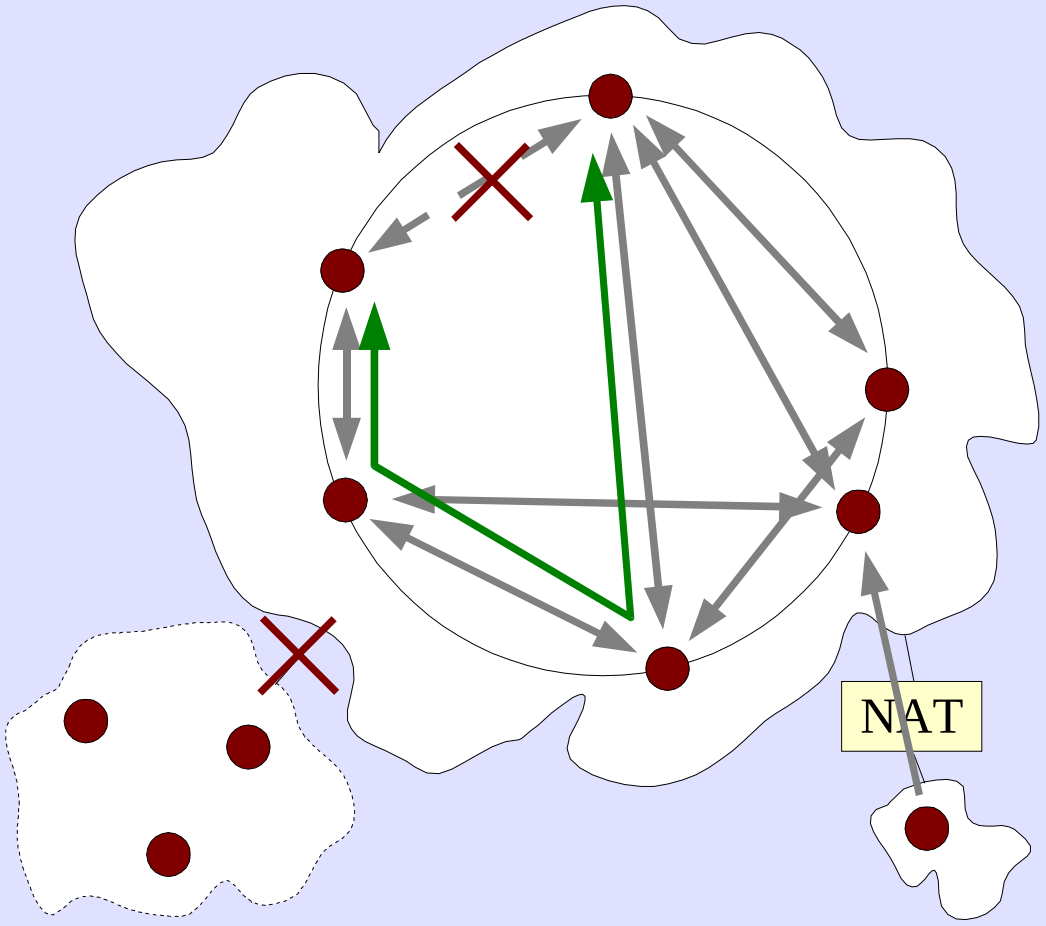

- Unstratified
- Forwarding around holes (RON)

What We Want

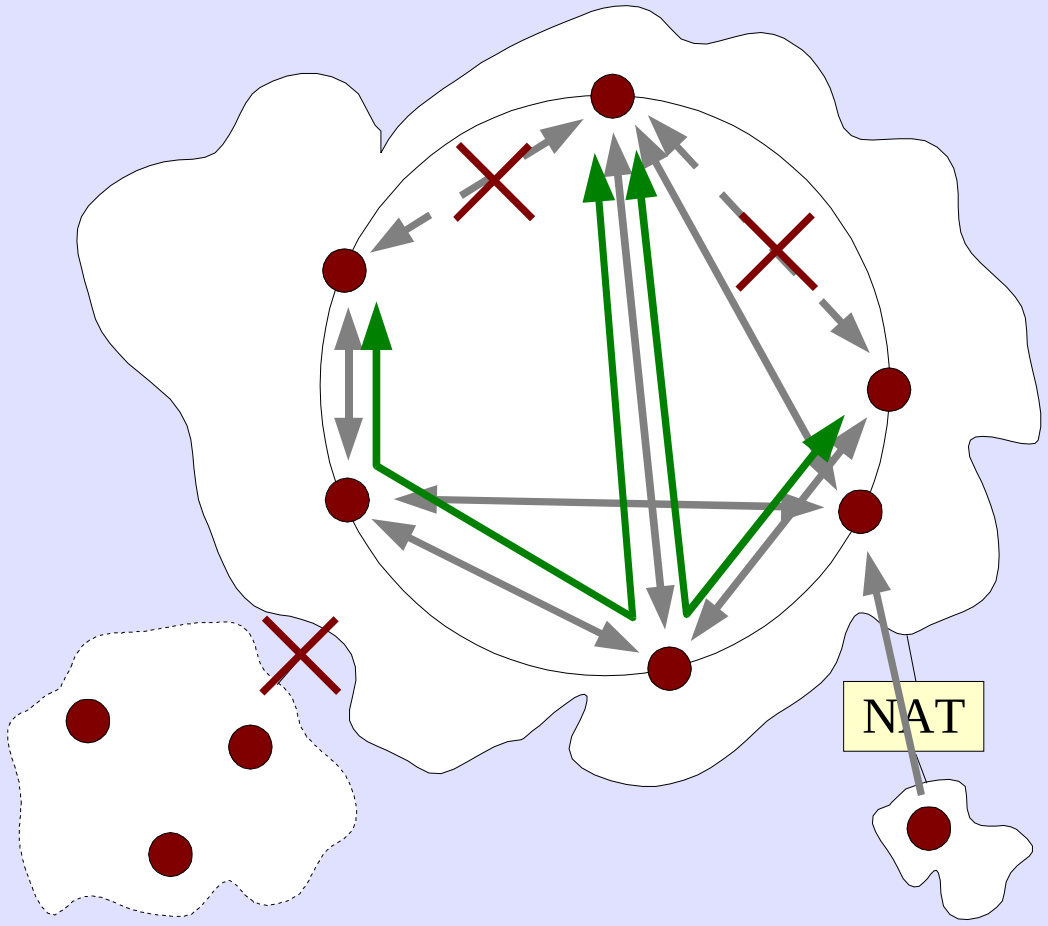

- Unstratified
- Forwarding around holes (RON)

What We Want

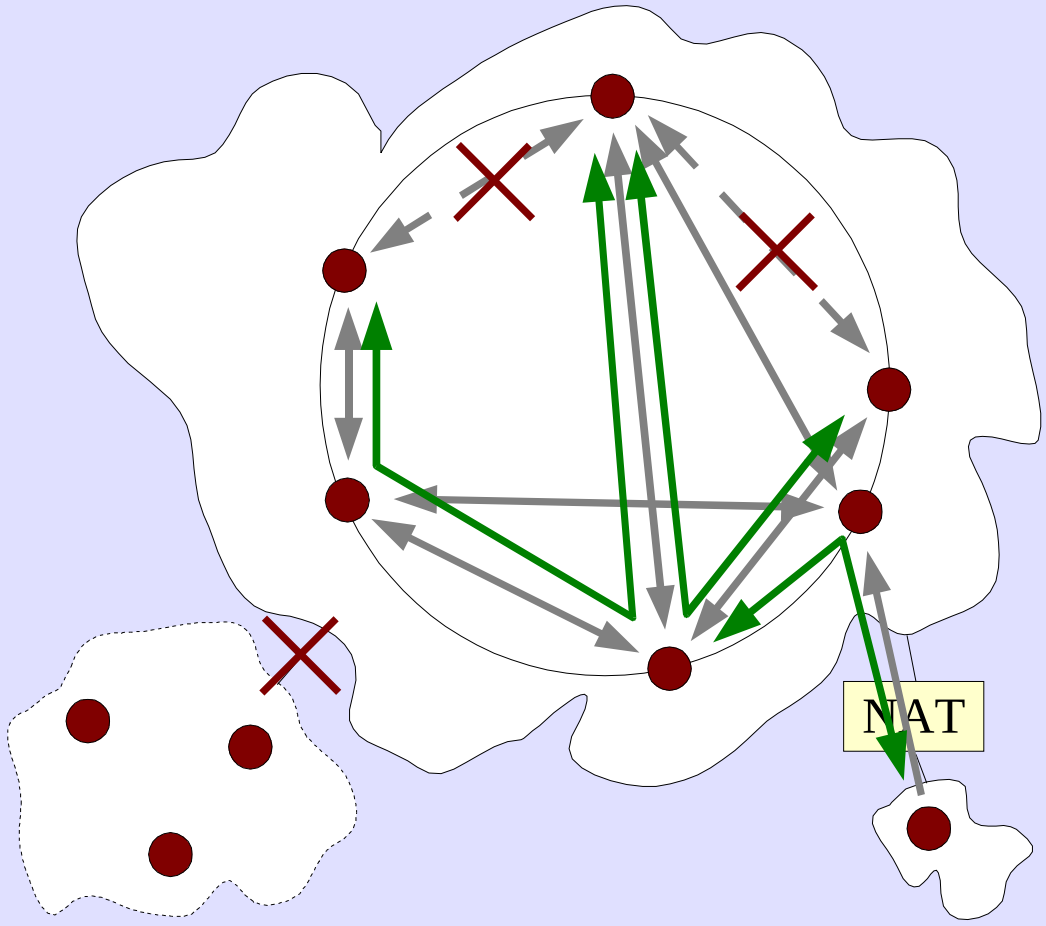

- Unstratified
- Forwarding around holes (RON)
- ...thru NATs

What We Want

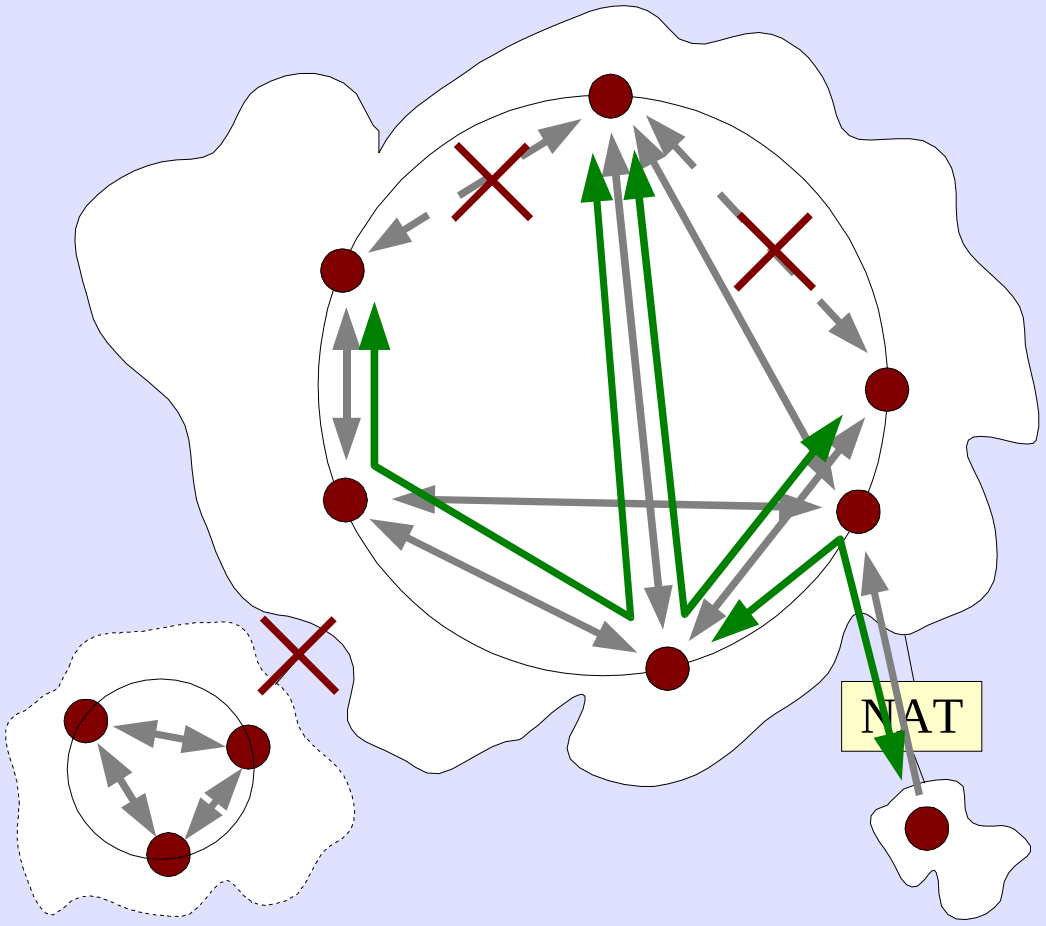

- Unstratified
- Forwarding around holes (RON)
- ...thru NATs
- Autonomous ad-hoc rings

What We Want

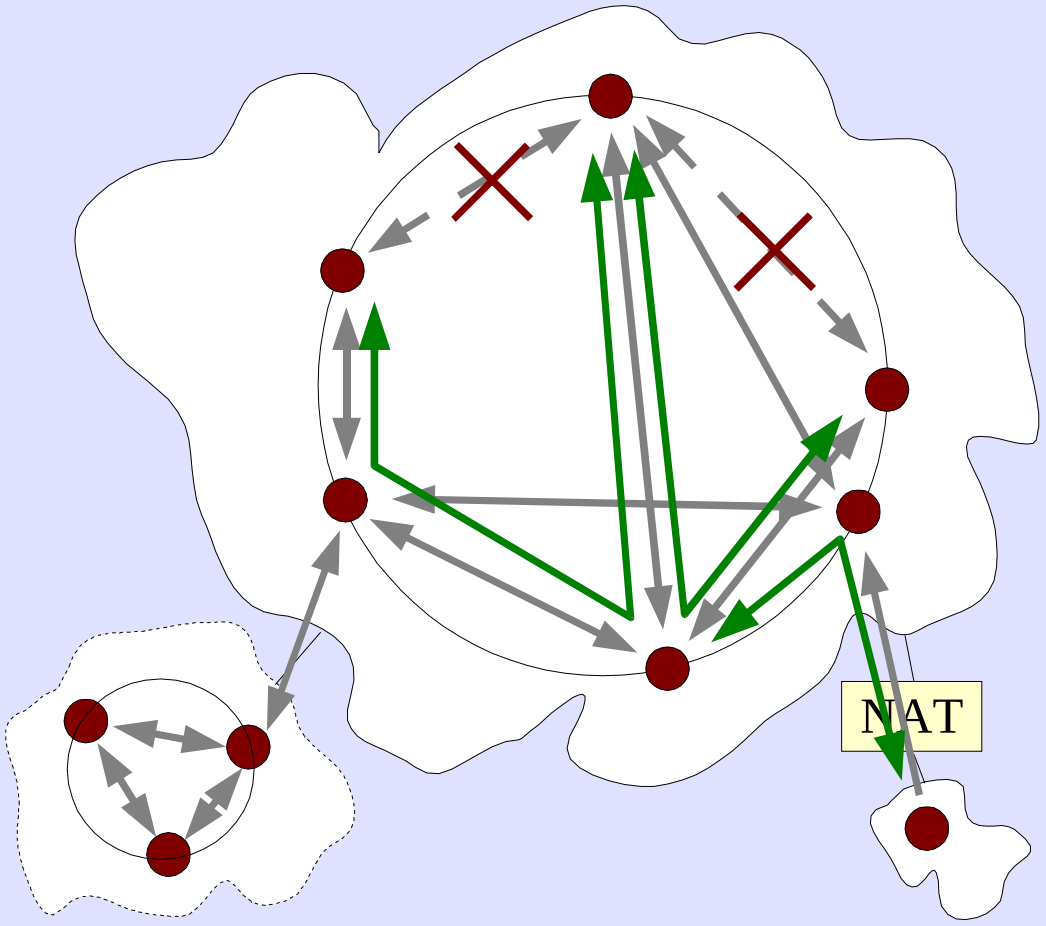

- Unstratified
- Forwarding around holes (RON)
- ...thru NATs
- Autonomous ad-hoc rings

What We Want

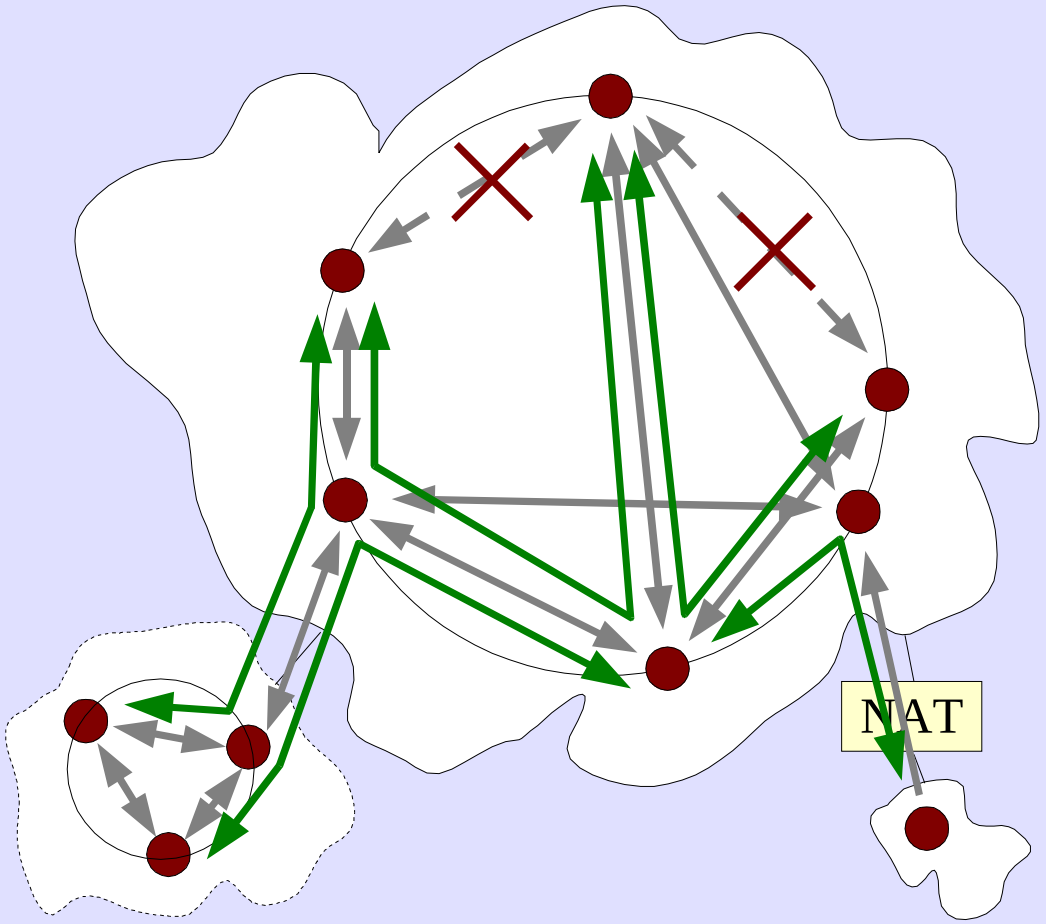

- Unstratified
- Forwarding around holes (RON)
- ...thru NATs
- Autonomous ad-hoc rings
- Inter-domain routing

Forwarding Mechanisms

- Source Routing
  - Nodes can store source routes, not just IP addresses, in their DHT neighbor tables.
  - Source routes not usually very long, because UIP sees Internet as “one big link.”
- Virtual Link Forwarding
  - Source routes restricted to two hops, but recursively composable
  - Distributes routing information throughout path

Source Routing

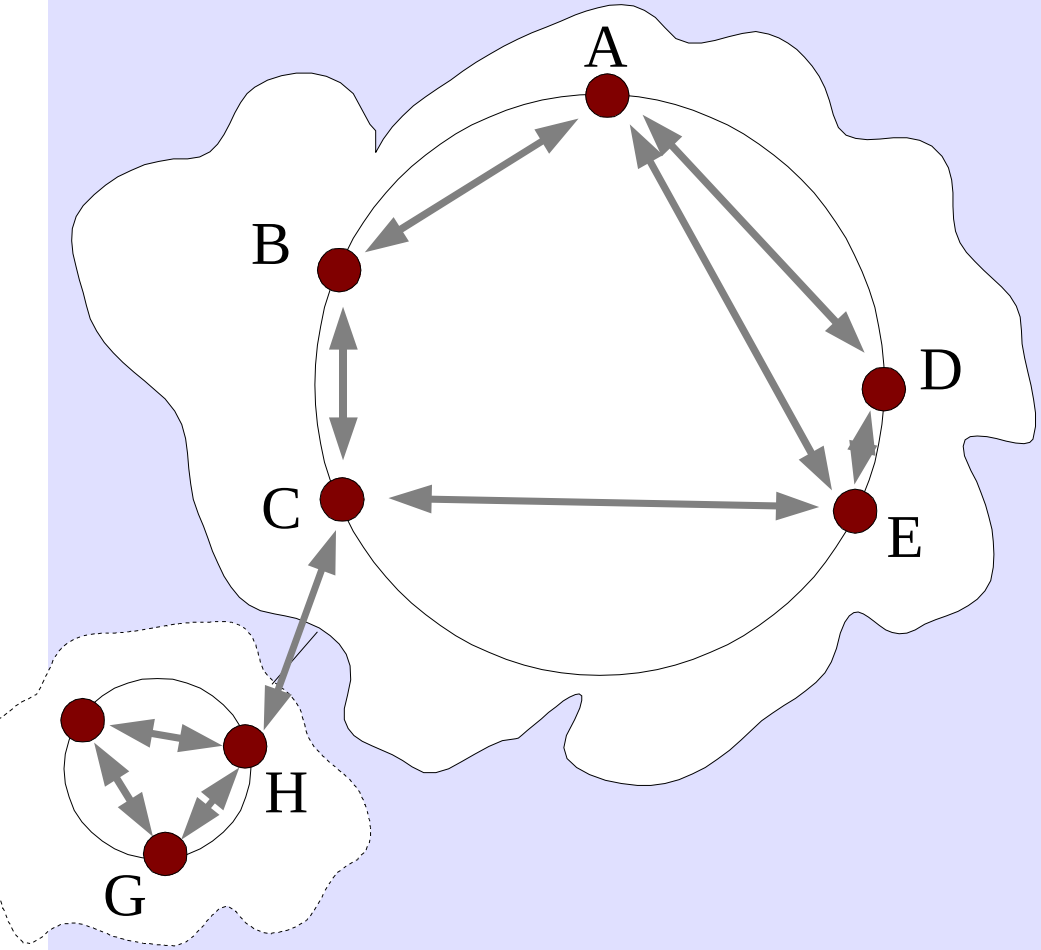

Source Routing

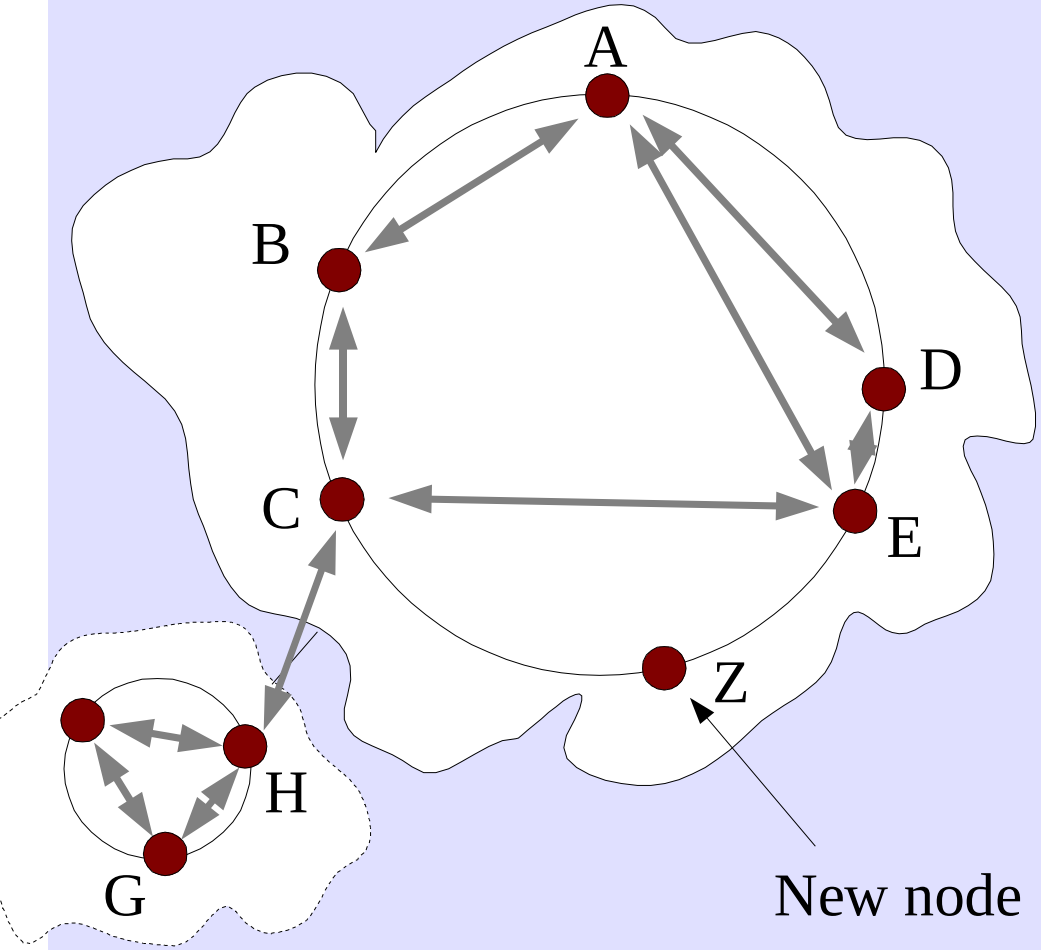

Source Routing

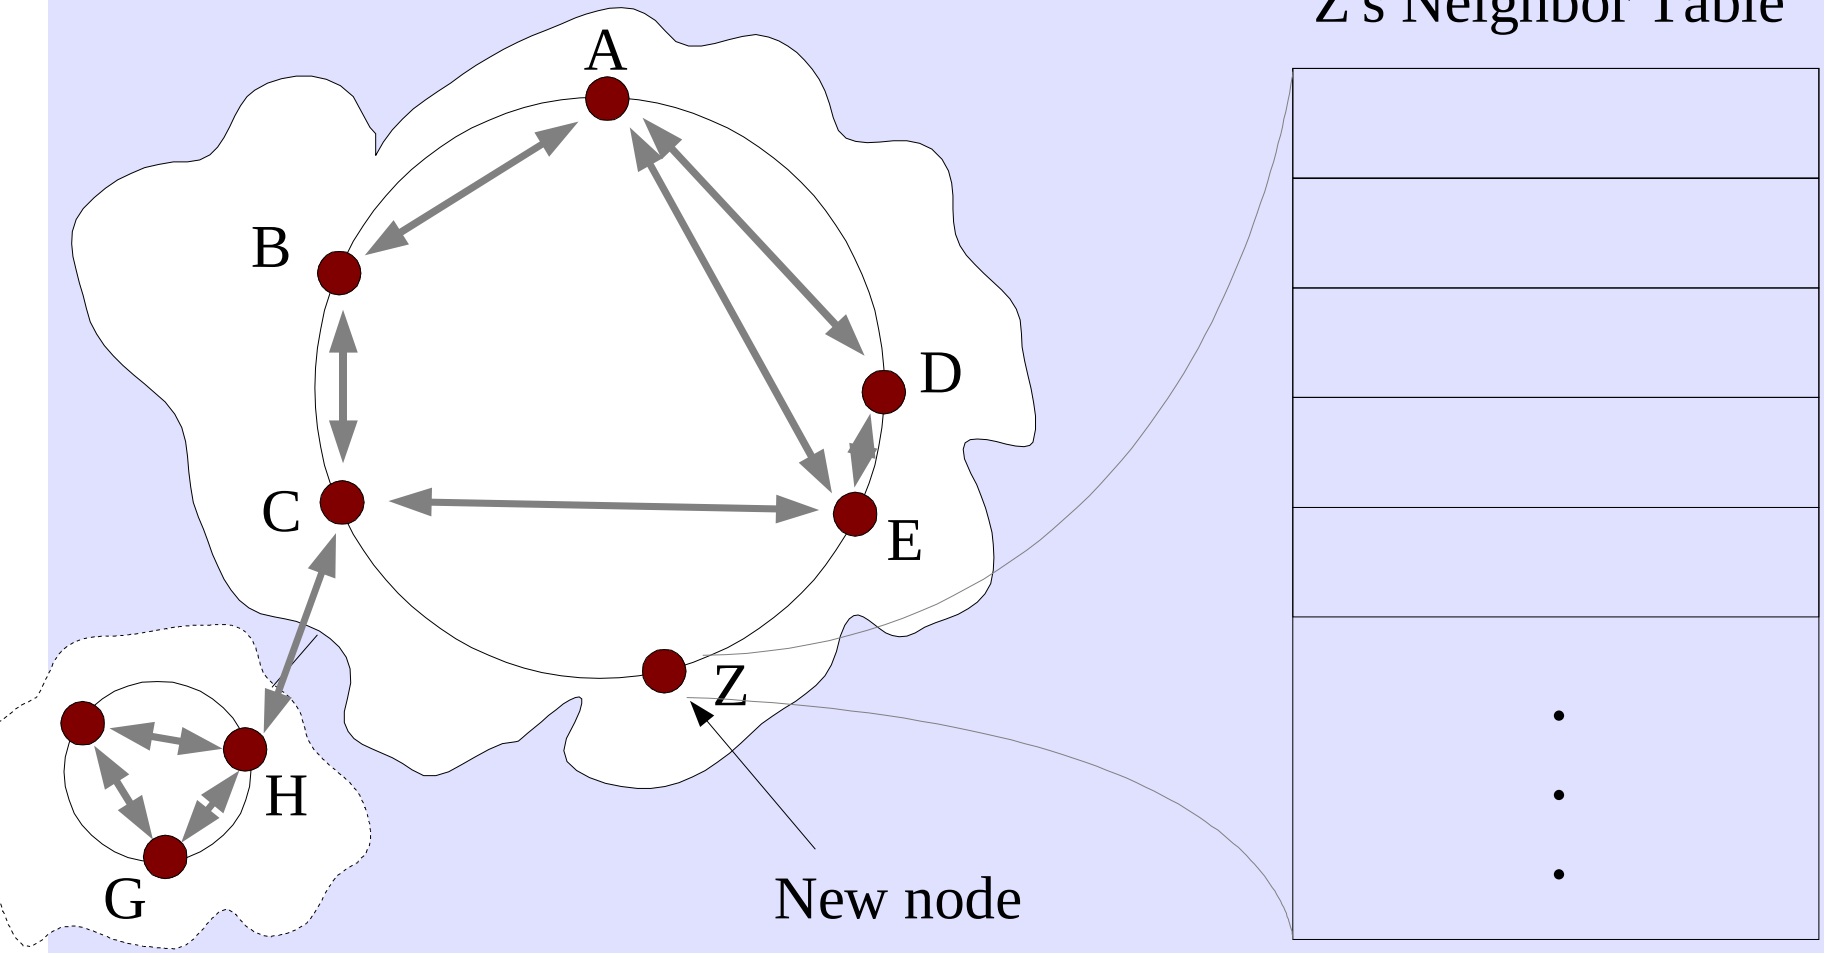

Source Routing

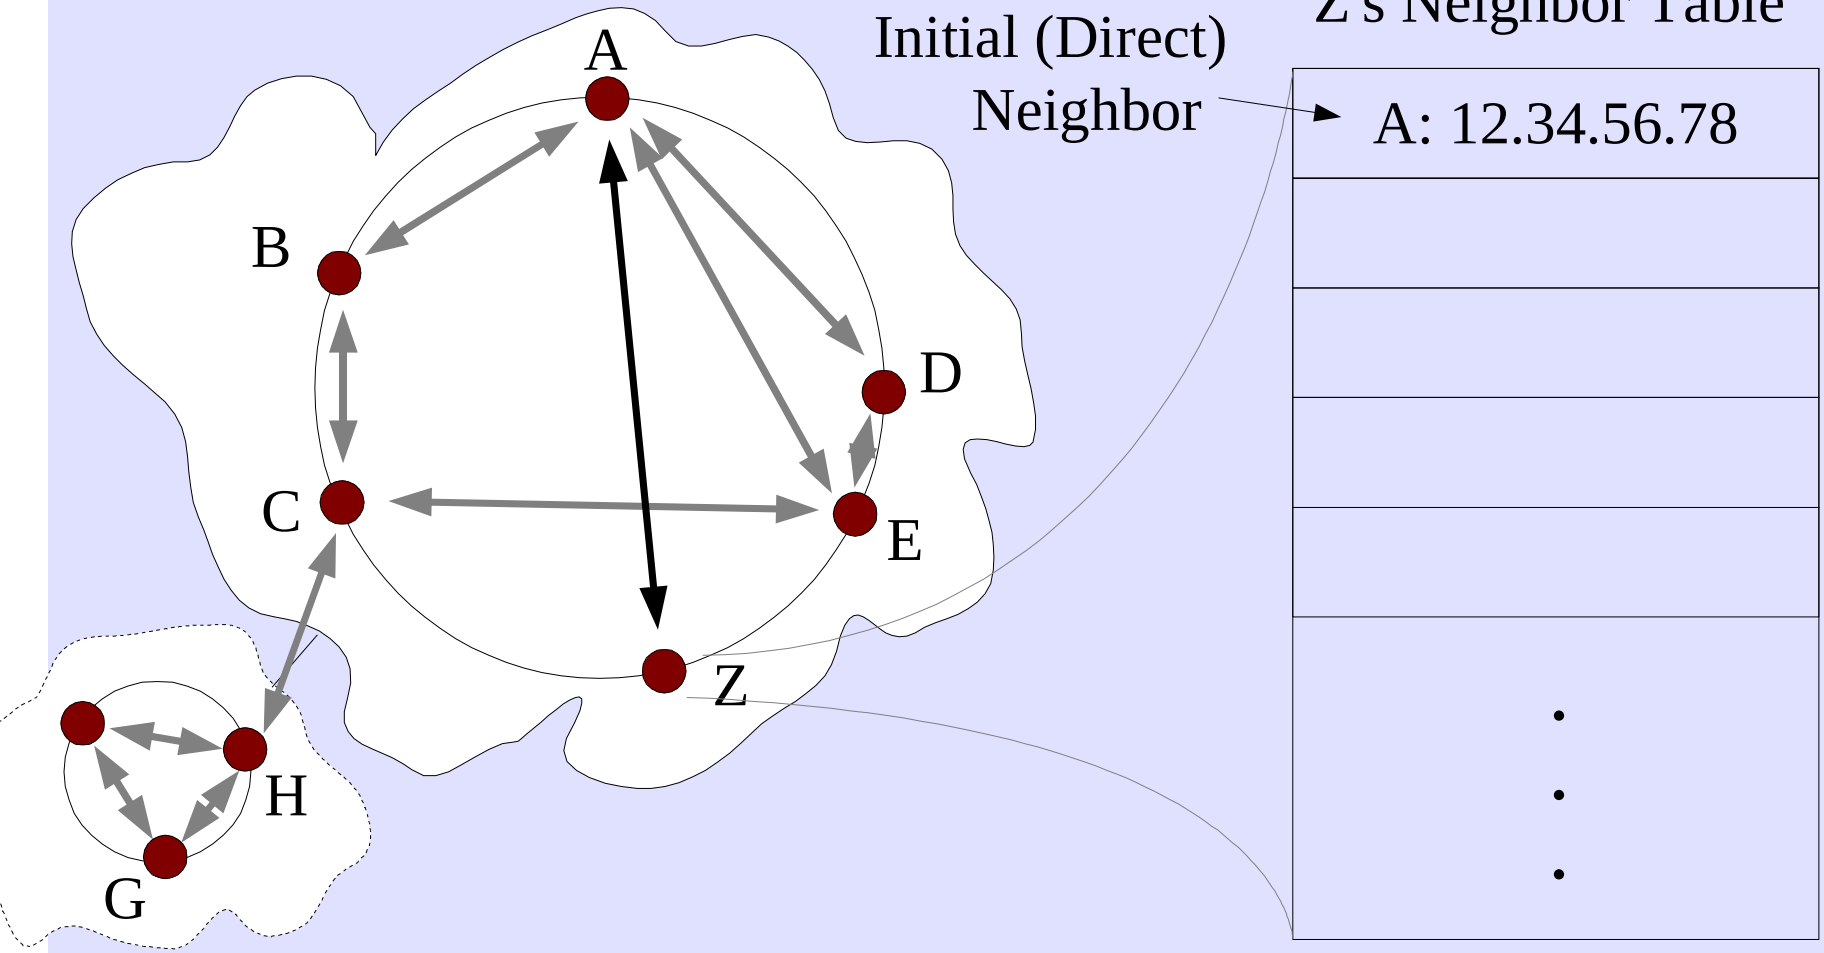

Source Routing

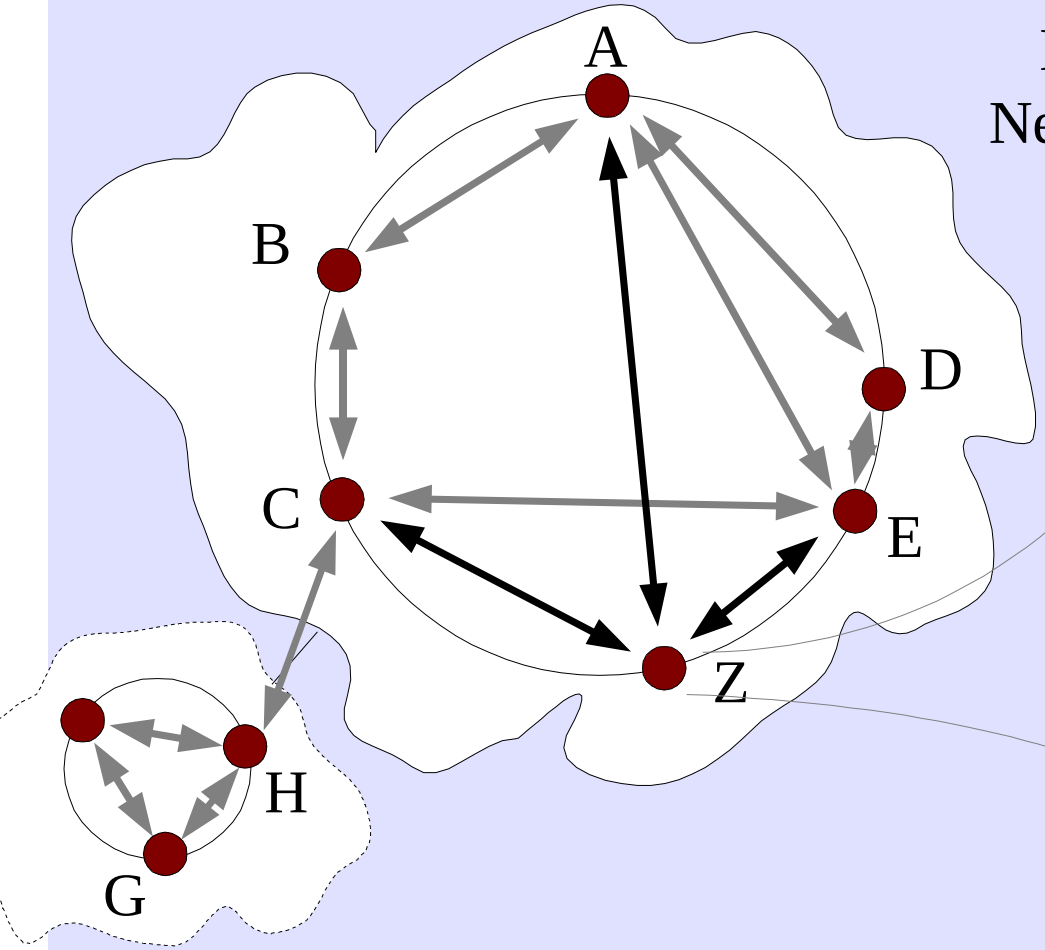

Direct Neighbors

Z's Neighbor Table

A: 12.34.56.78
C: 23.45.67.89
E: 34.56.78.90
. . .

Source Routing

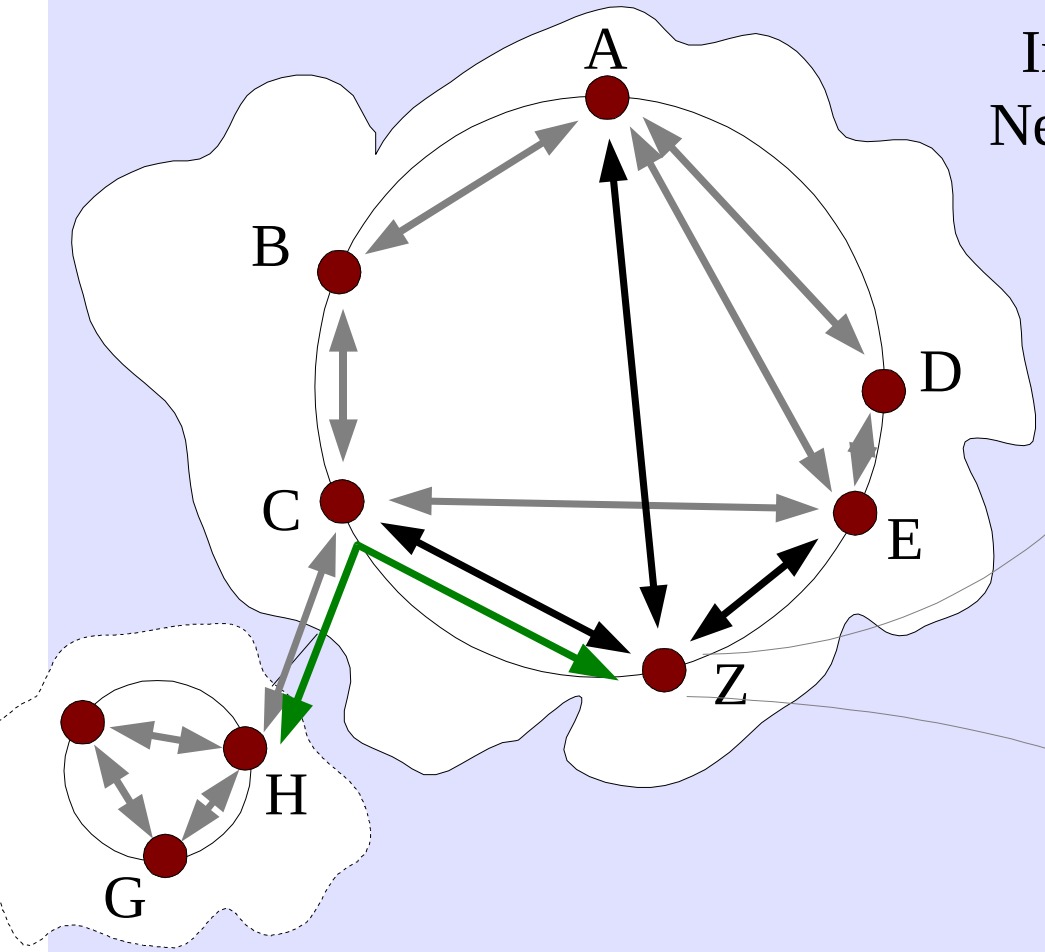

Indirect Neighbors

Z's Neighbor Table

A: 12.34.56.78
C: 23.45.67.89
E: 34.56.78.90
H: [C → H]
. . .

Source Routing

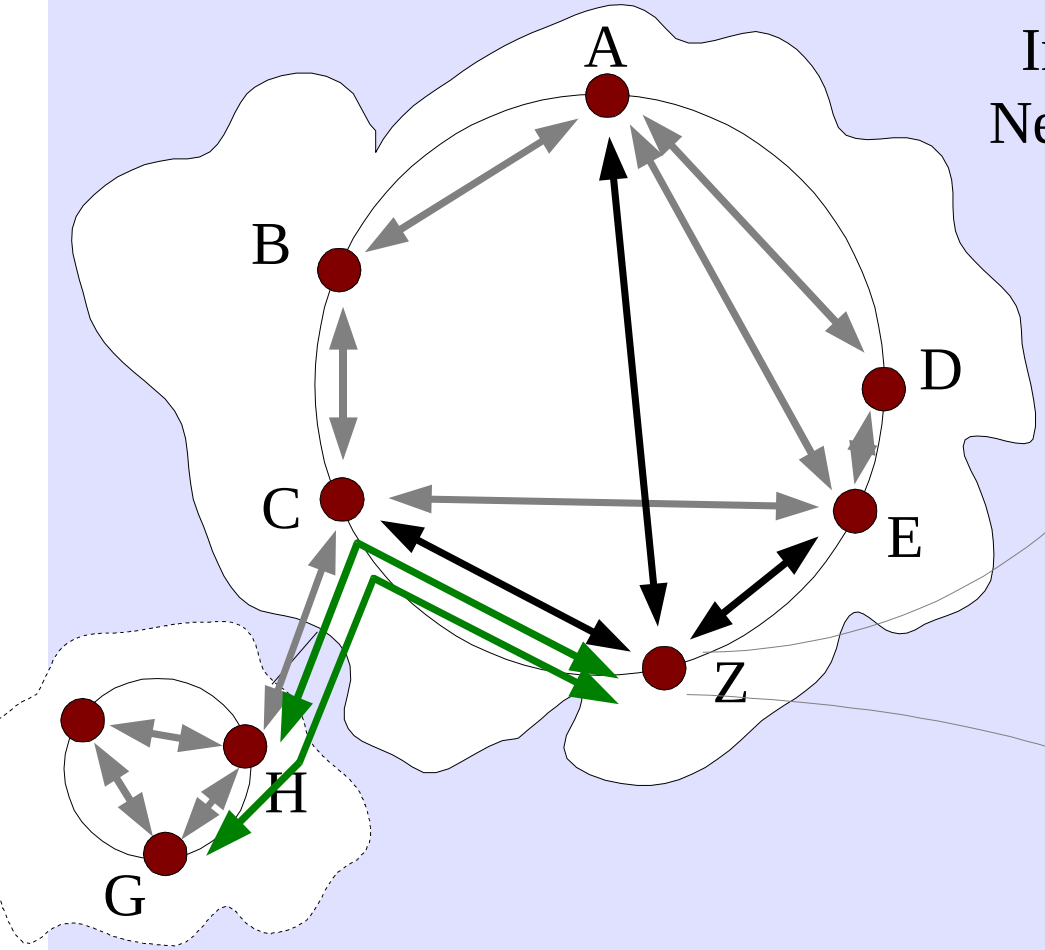

Indirect Neighbors

Z's Neighbor Table

A: 12.34.56.78
C: 23.45.67.89
E: 34.56.78.90
H: [C → H]
G: [C → H → G]
. . .

What We Have

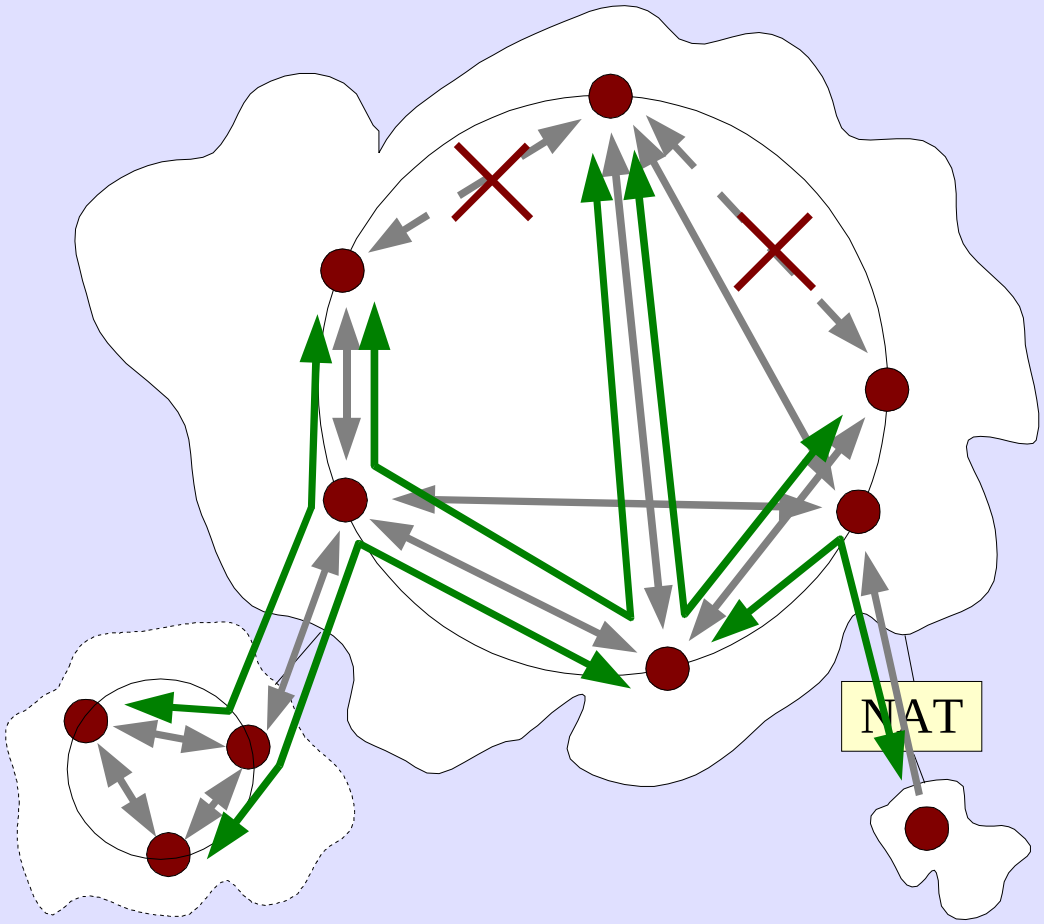

What We Have

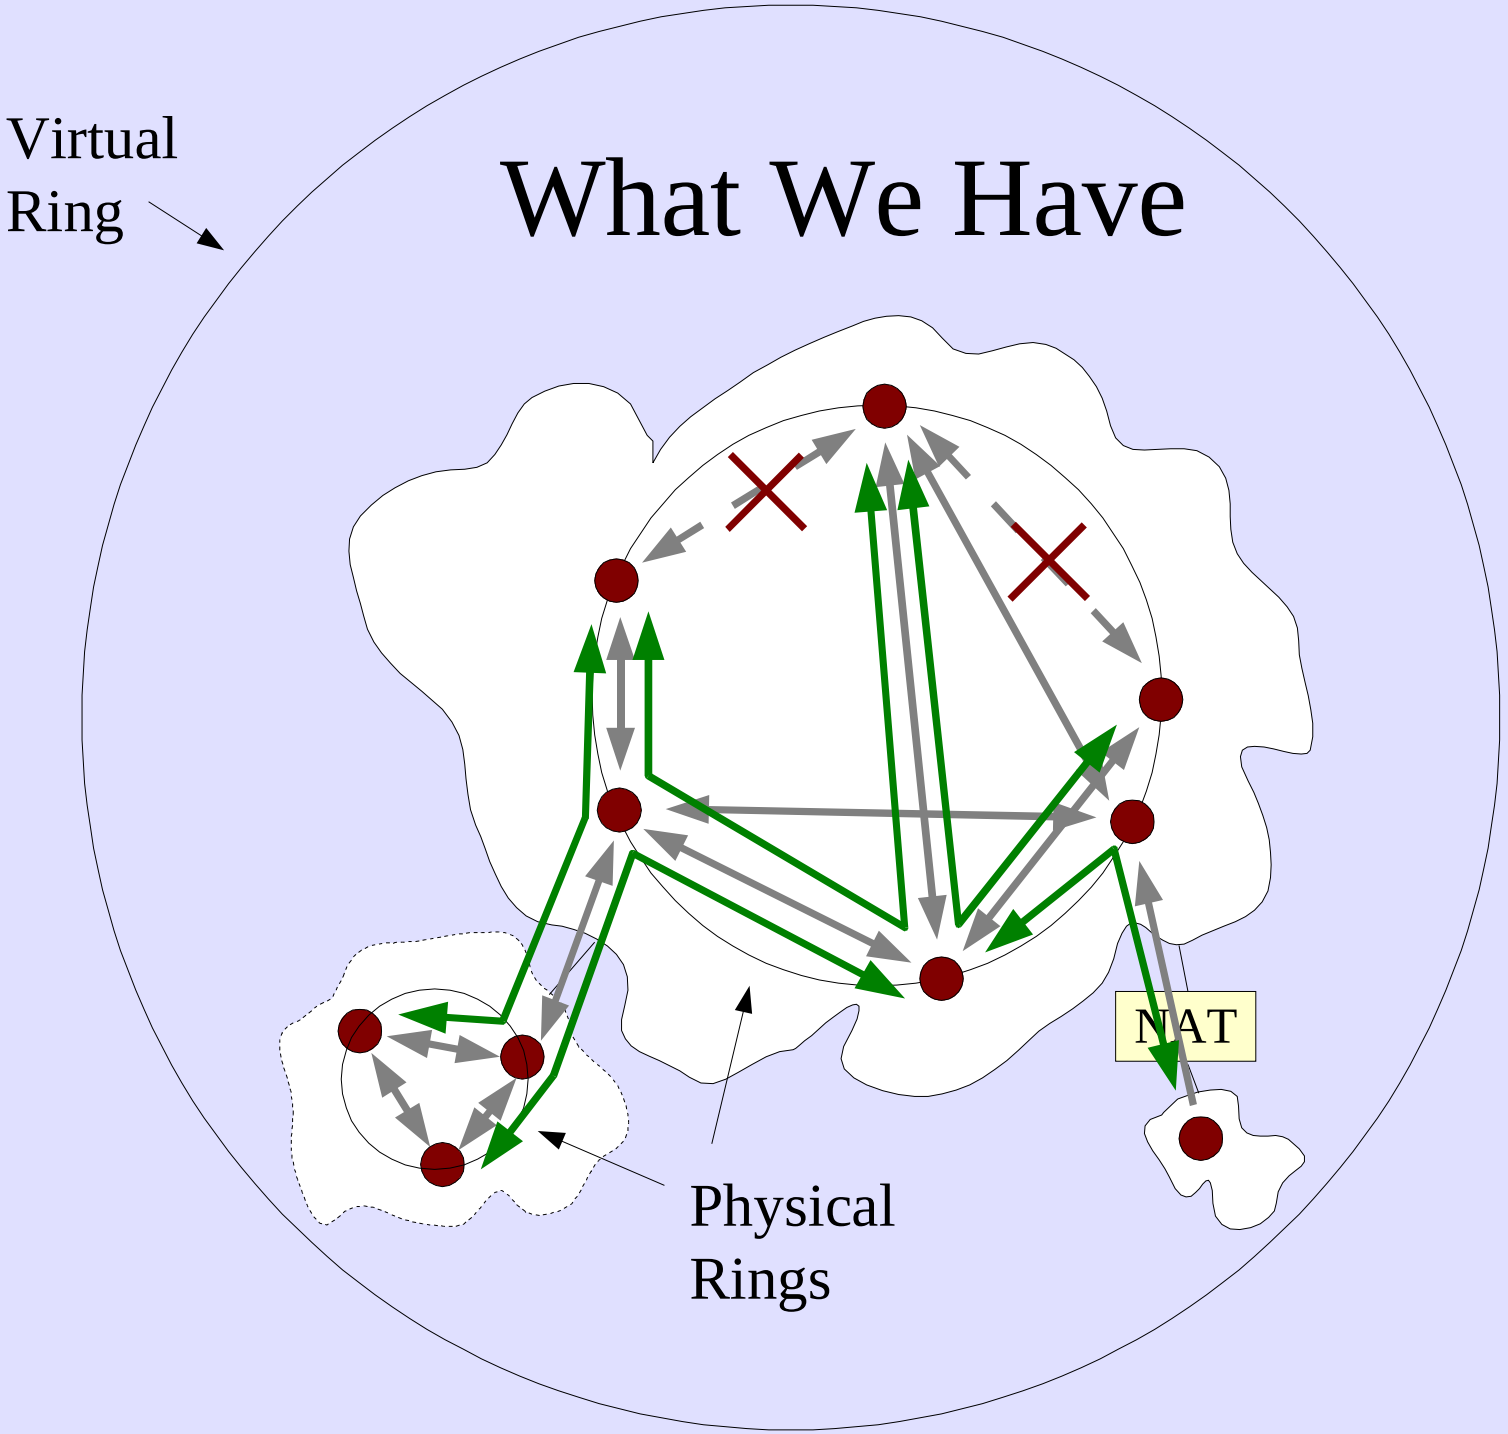

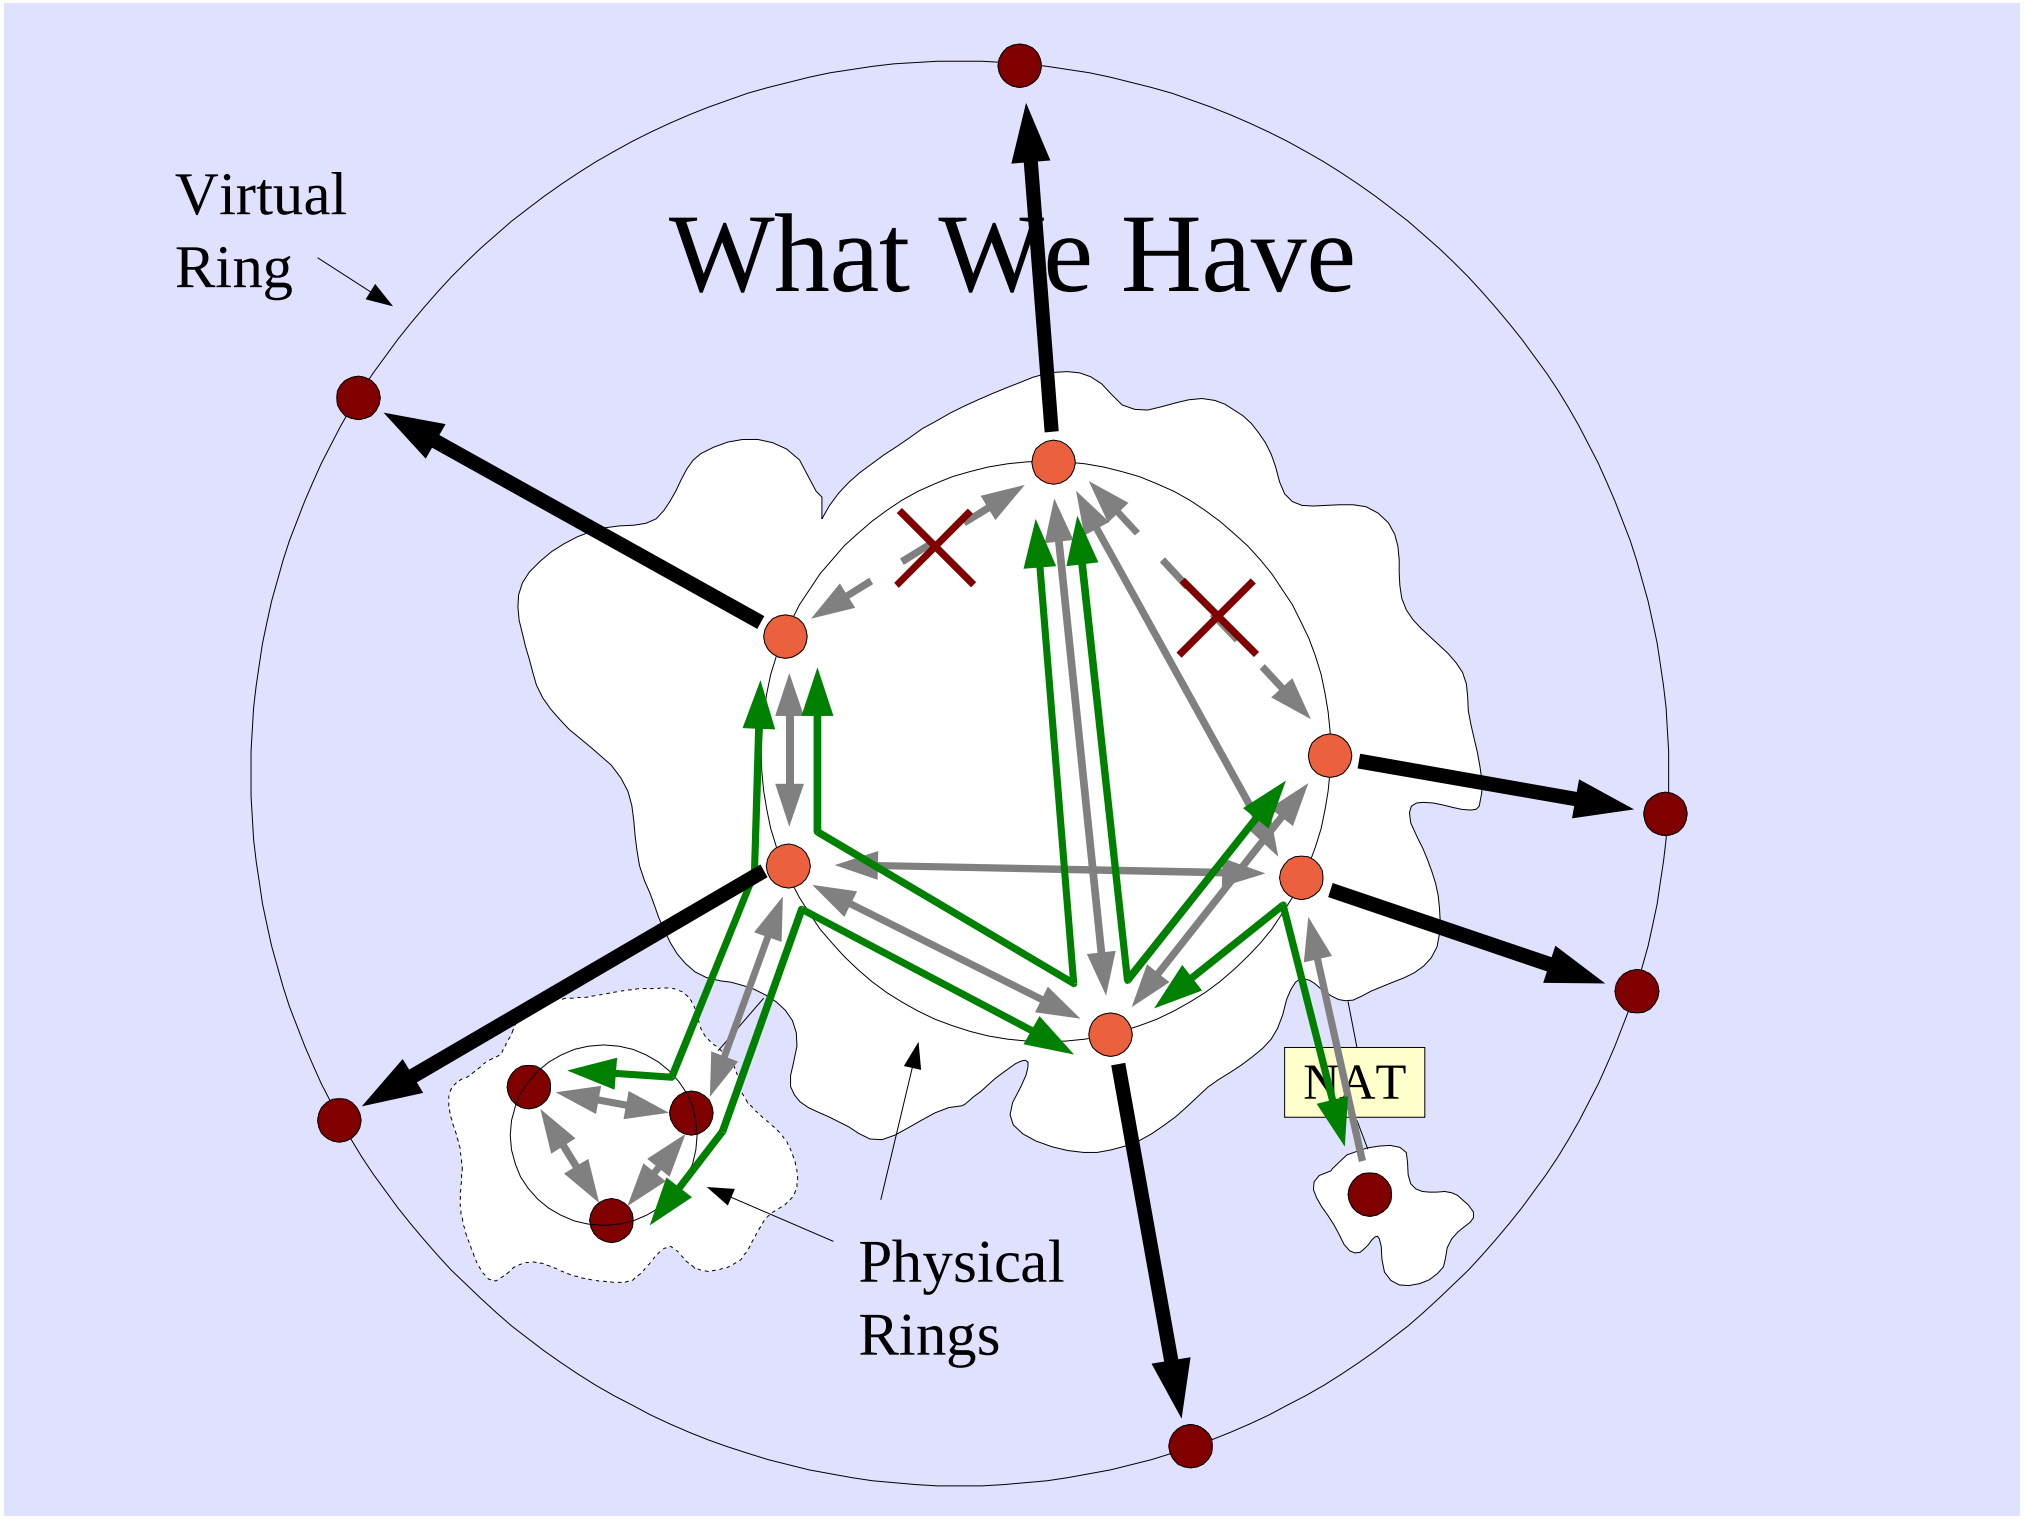
Virtual
Ring
What We Have
NAT
Physical
Rings

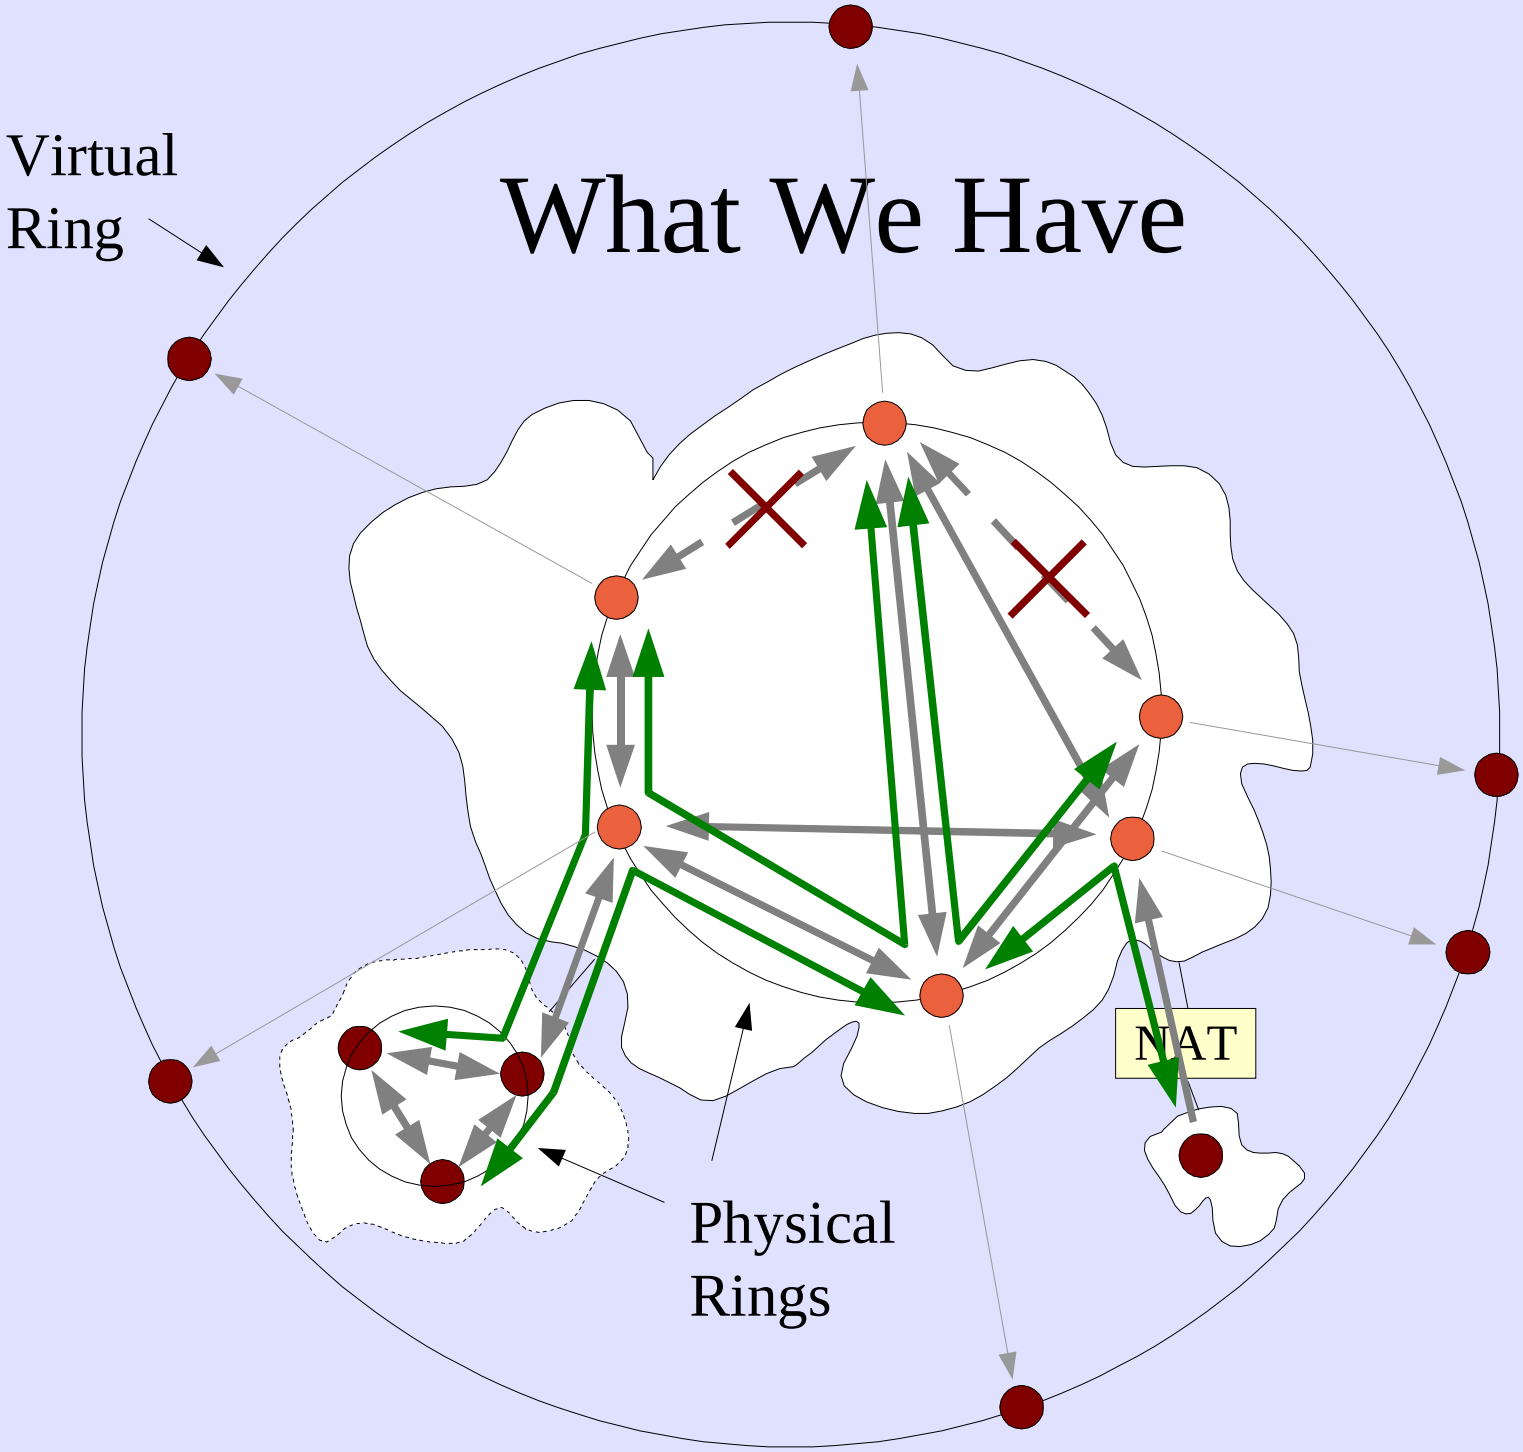
Virtual
Ring
What We Have
NAT
Physical
Rings

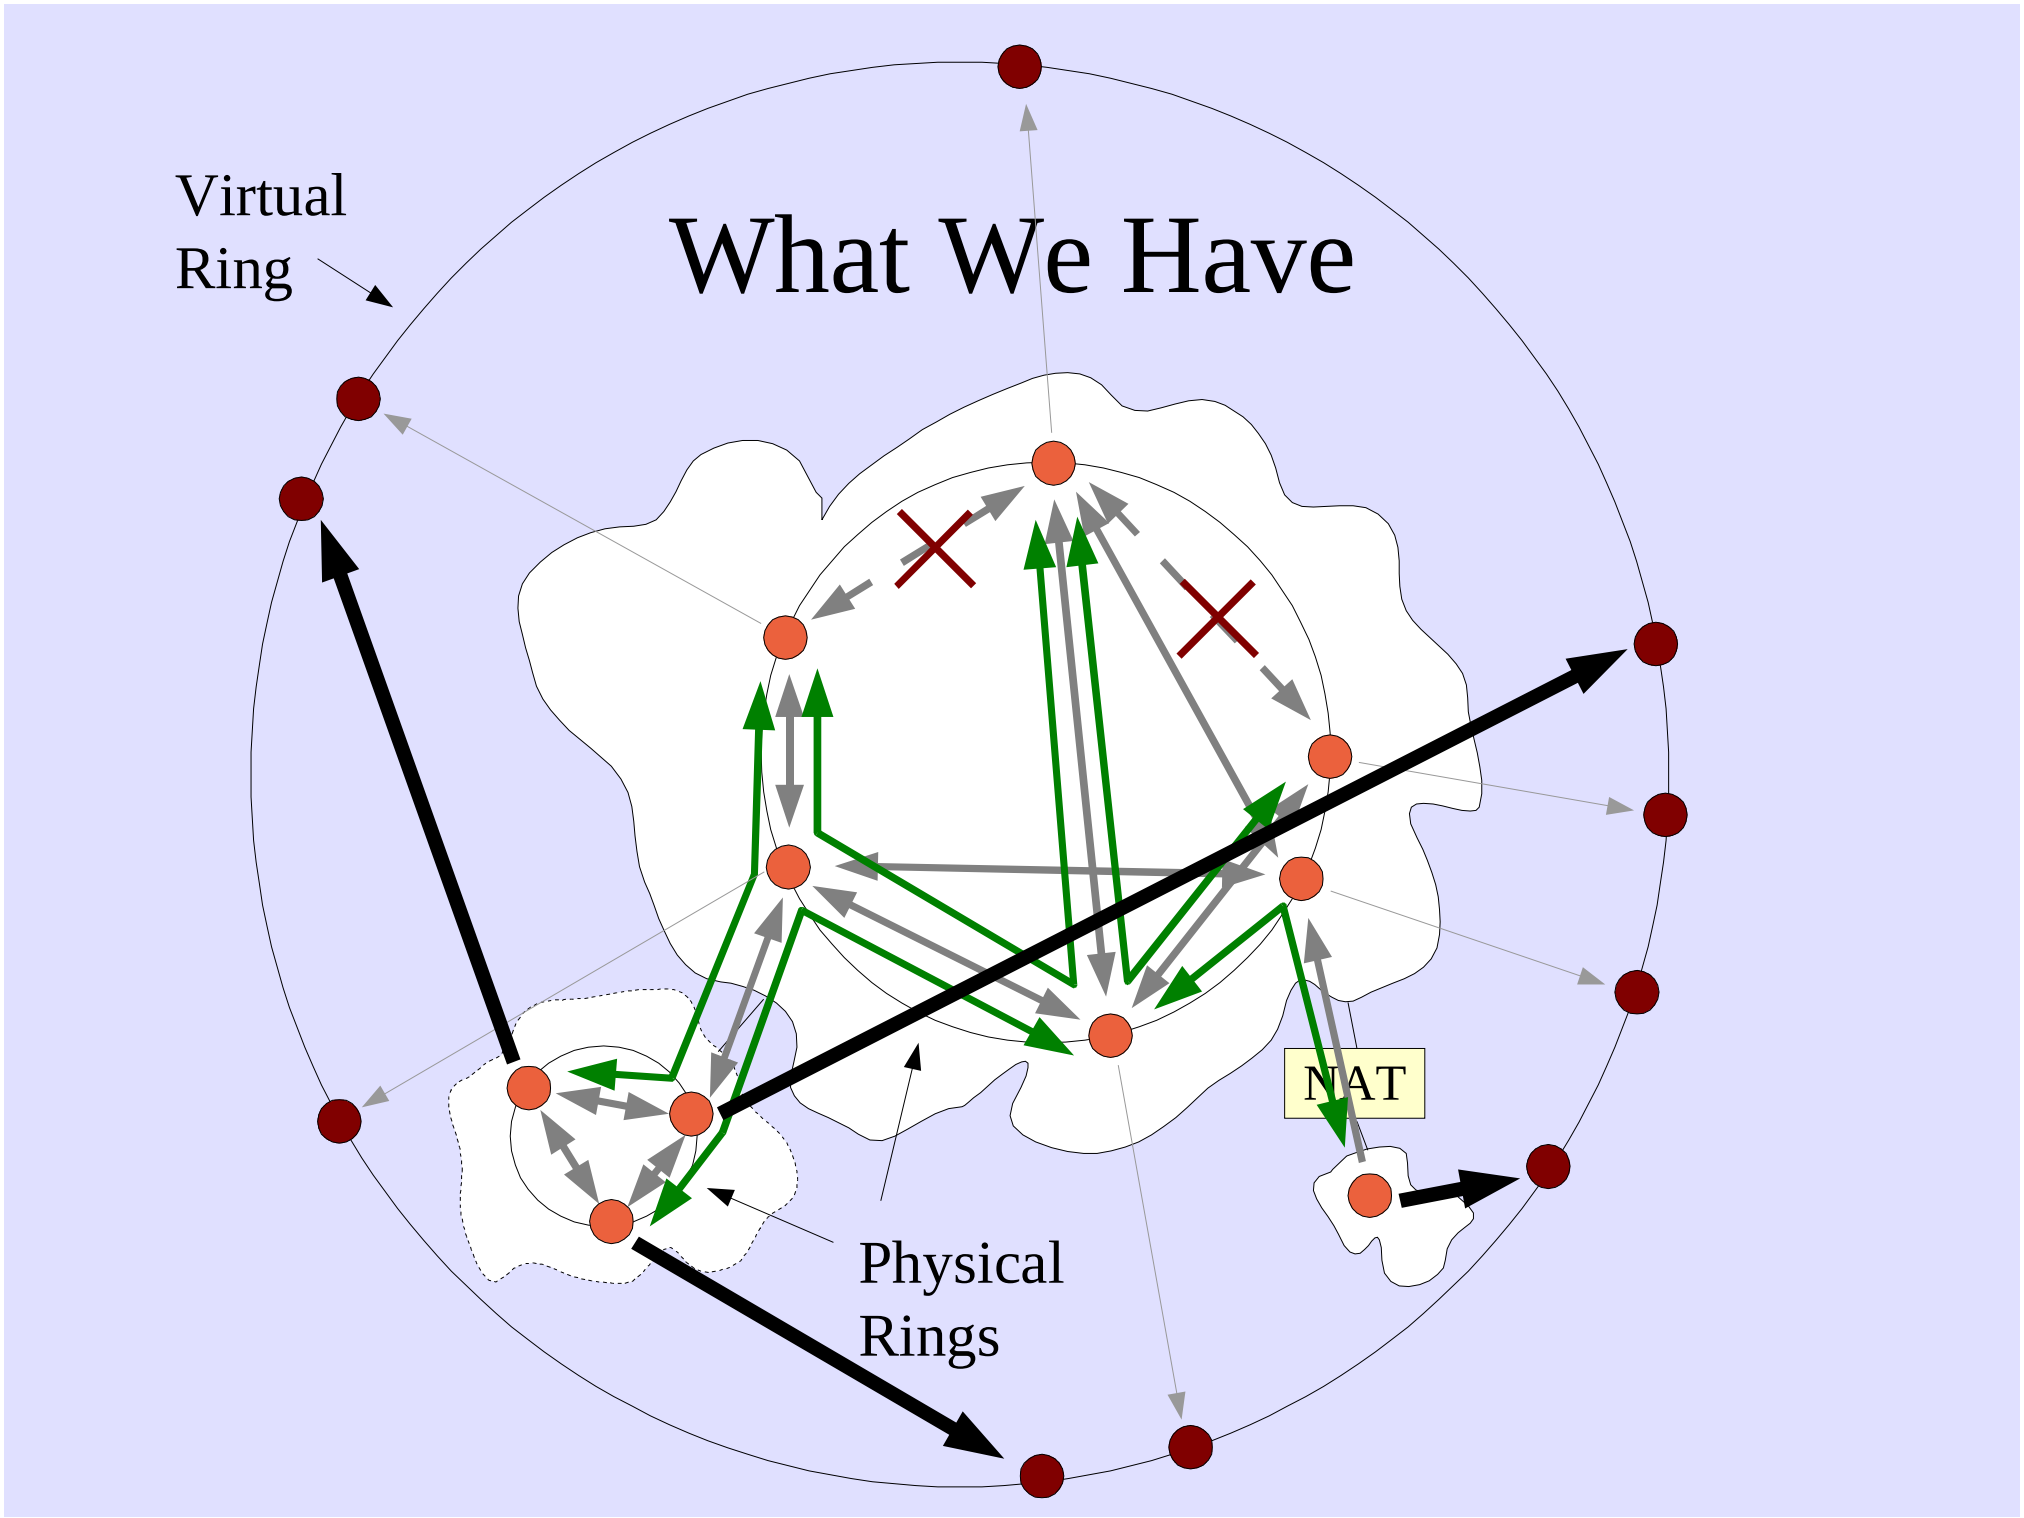
Virtual
Ring
What We Have
NAT
Physical
Rings

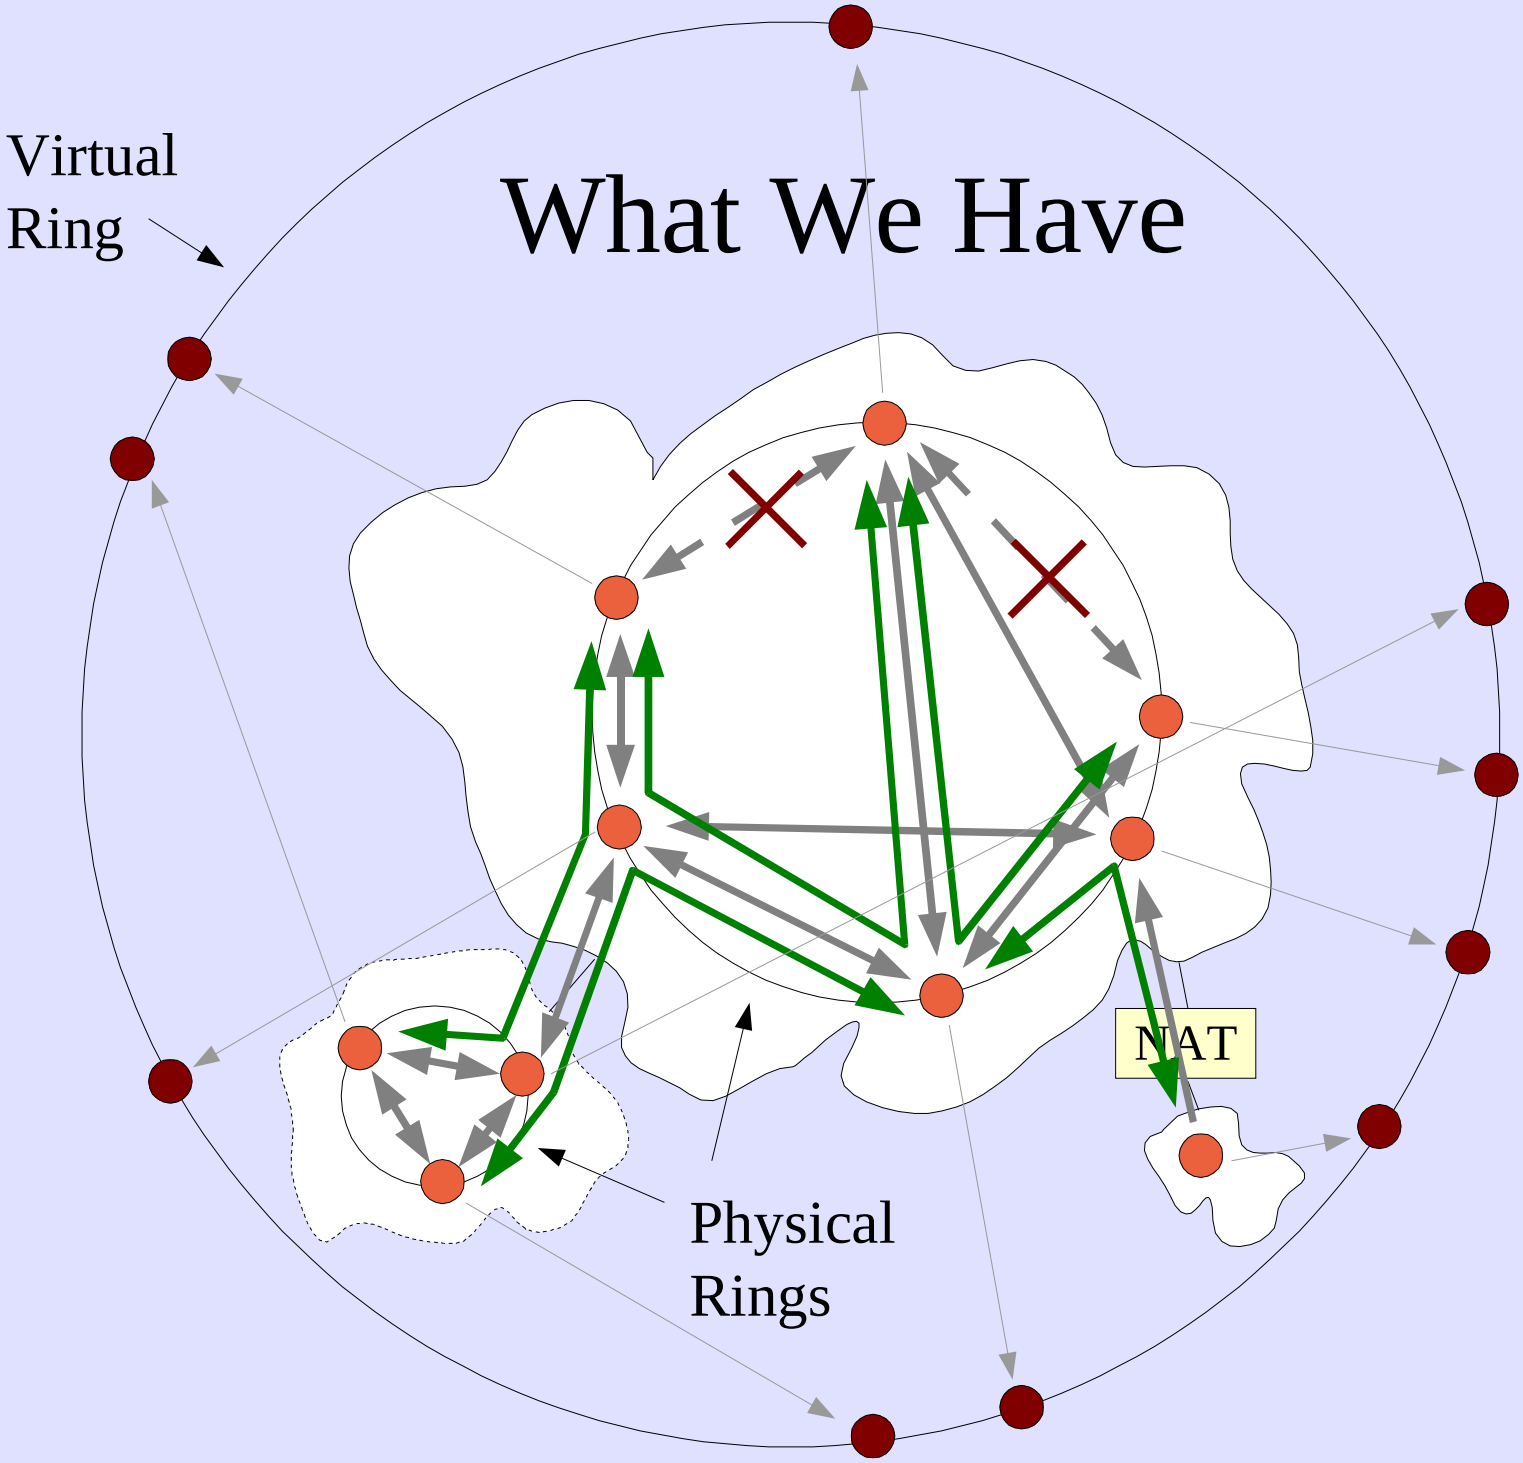
Virtual
Ring
What We Have
NAT
Physical
Rings

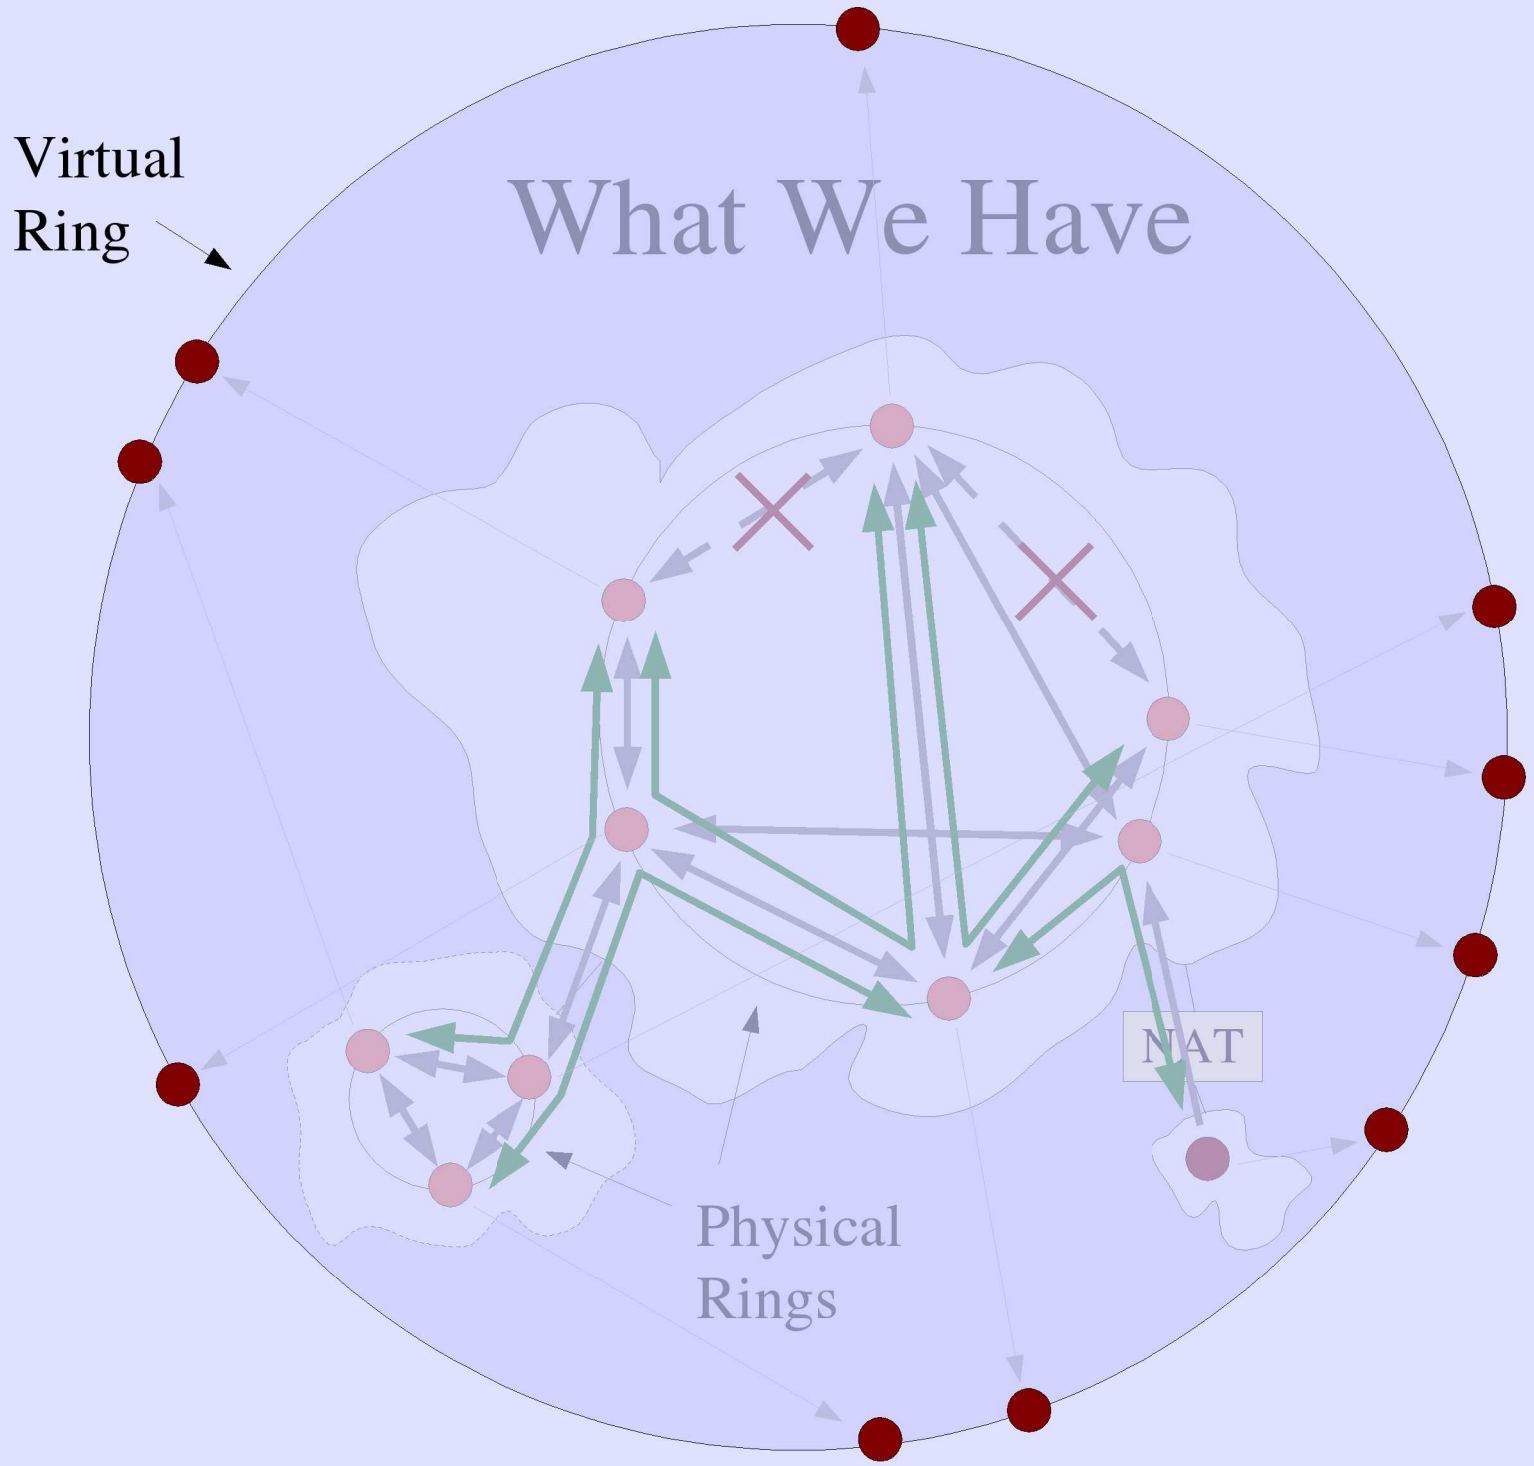
Virtual
Ring
What We Have
NAT
Physical
Rings

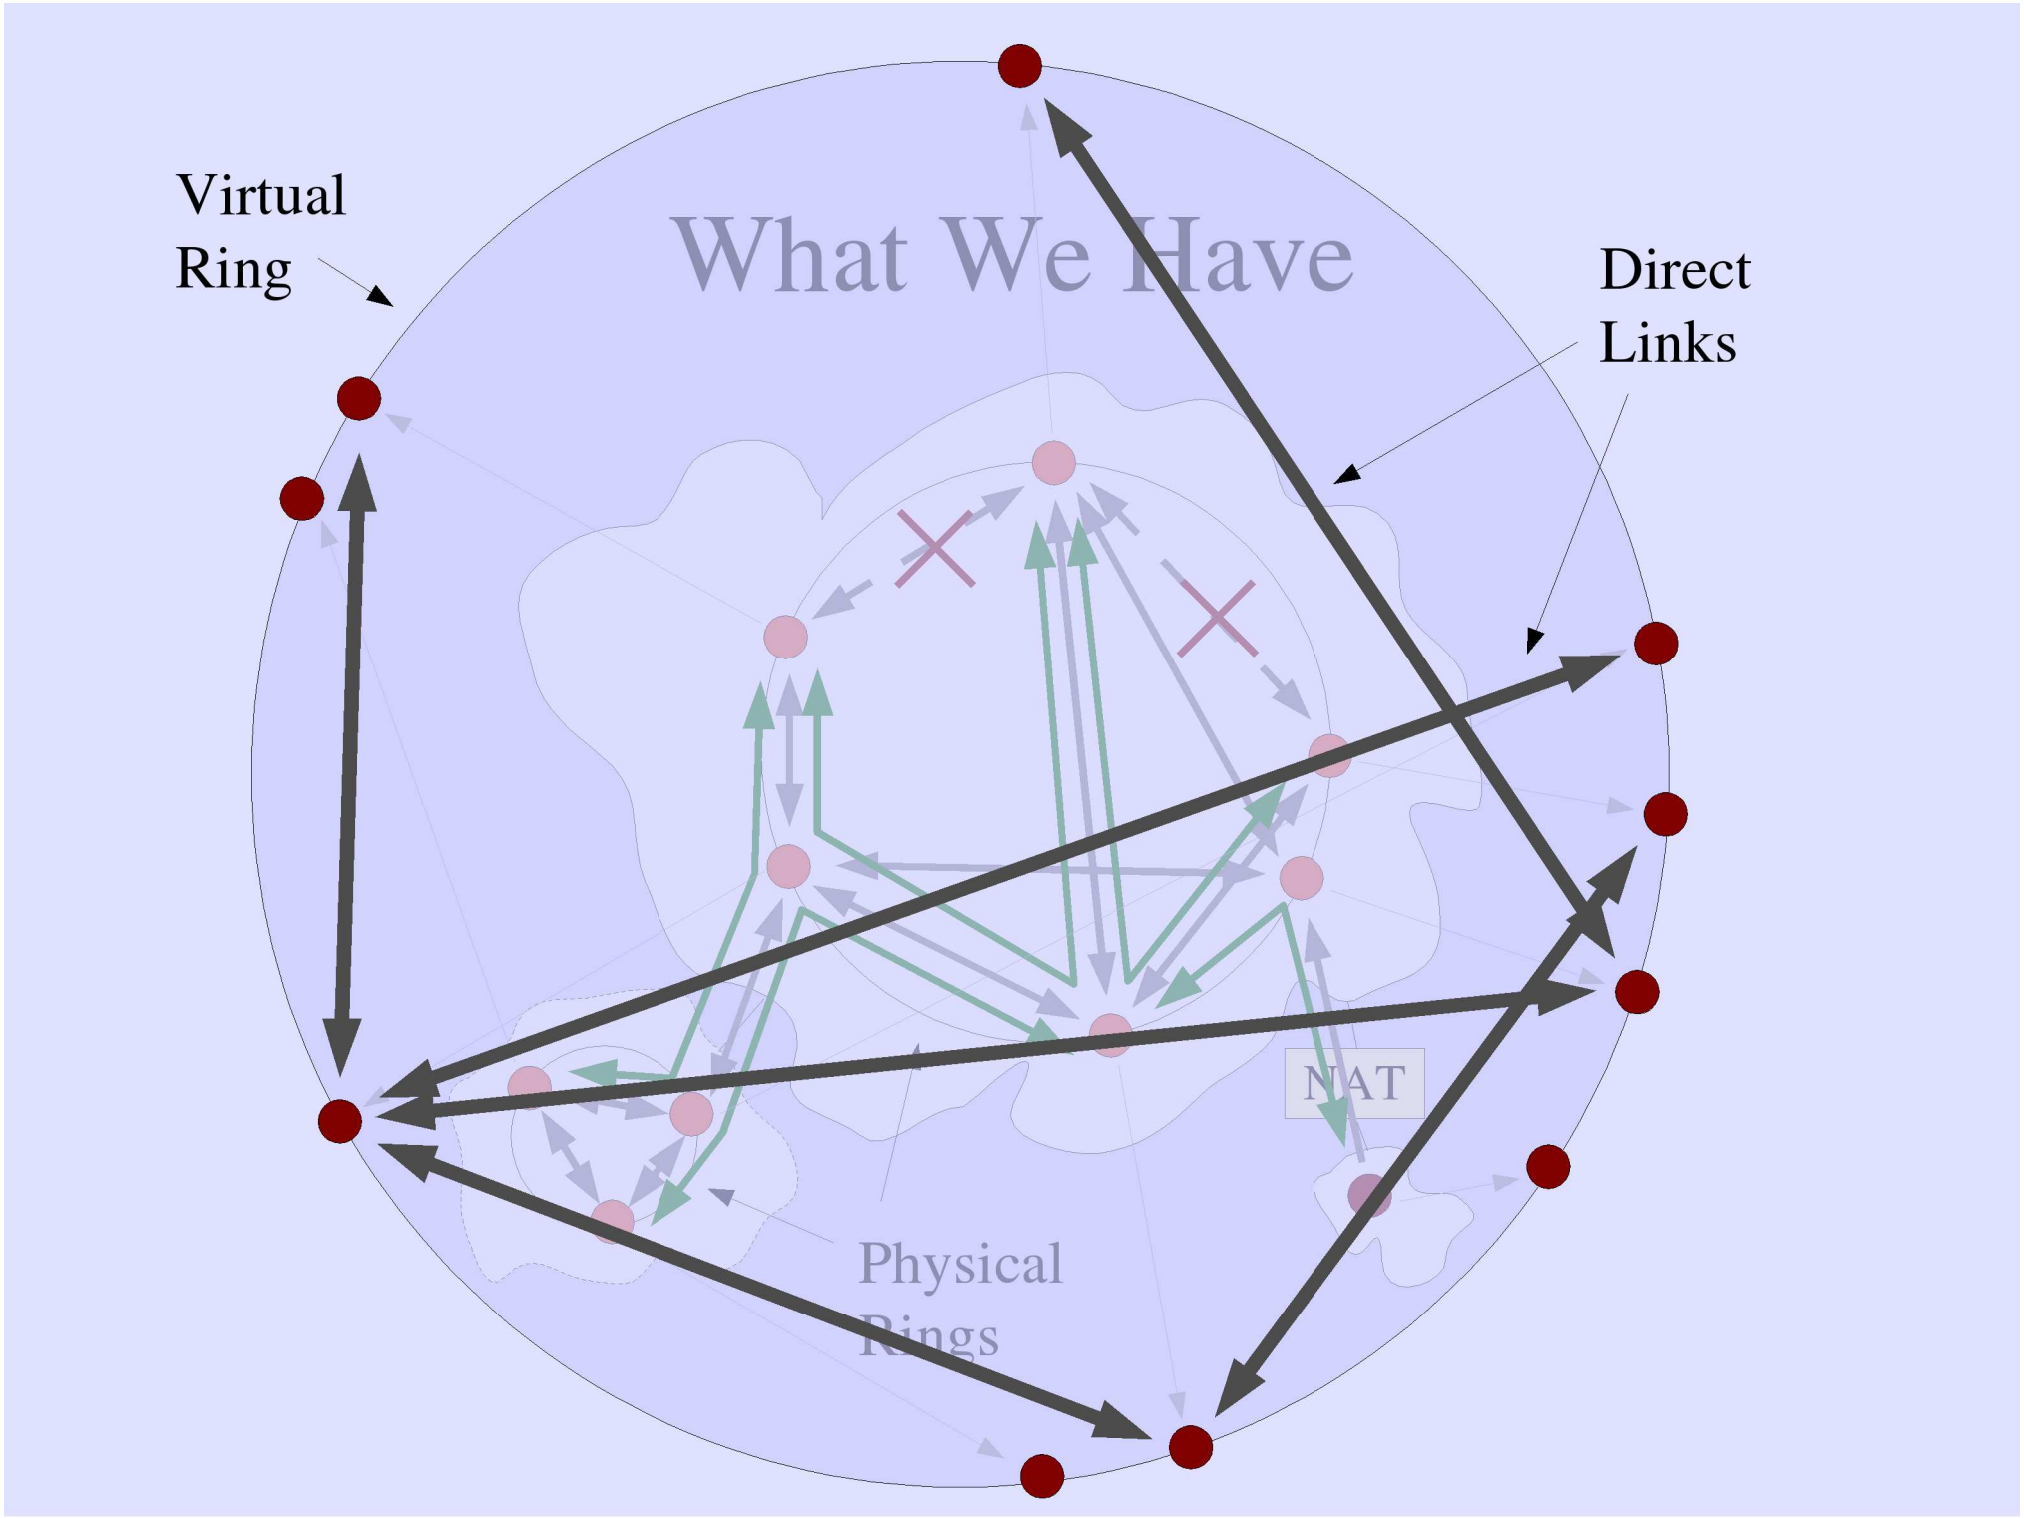
Virtual
Ring
What We Have
Direct
Links
NAT
Physical
Rings

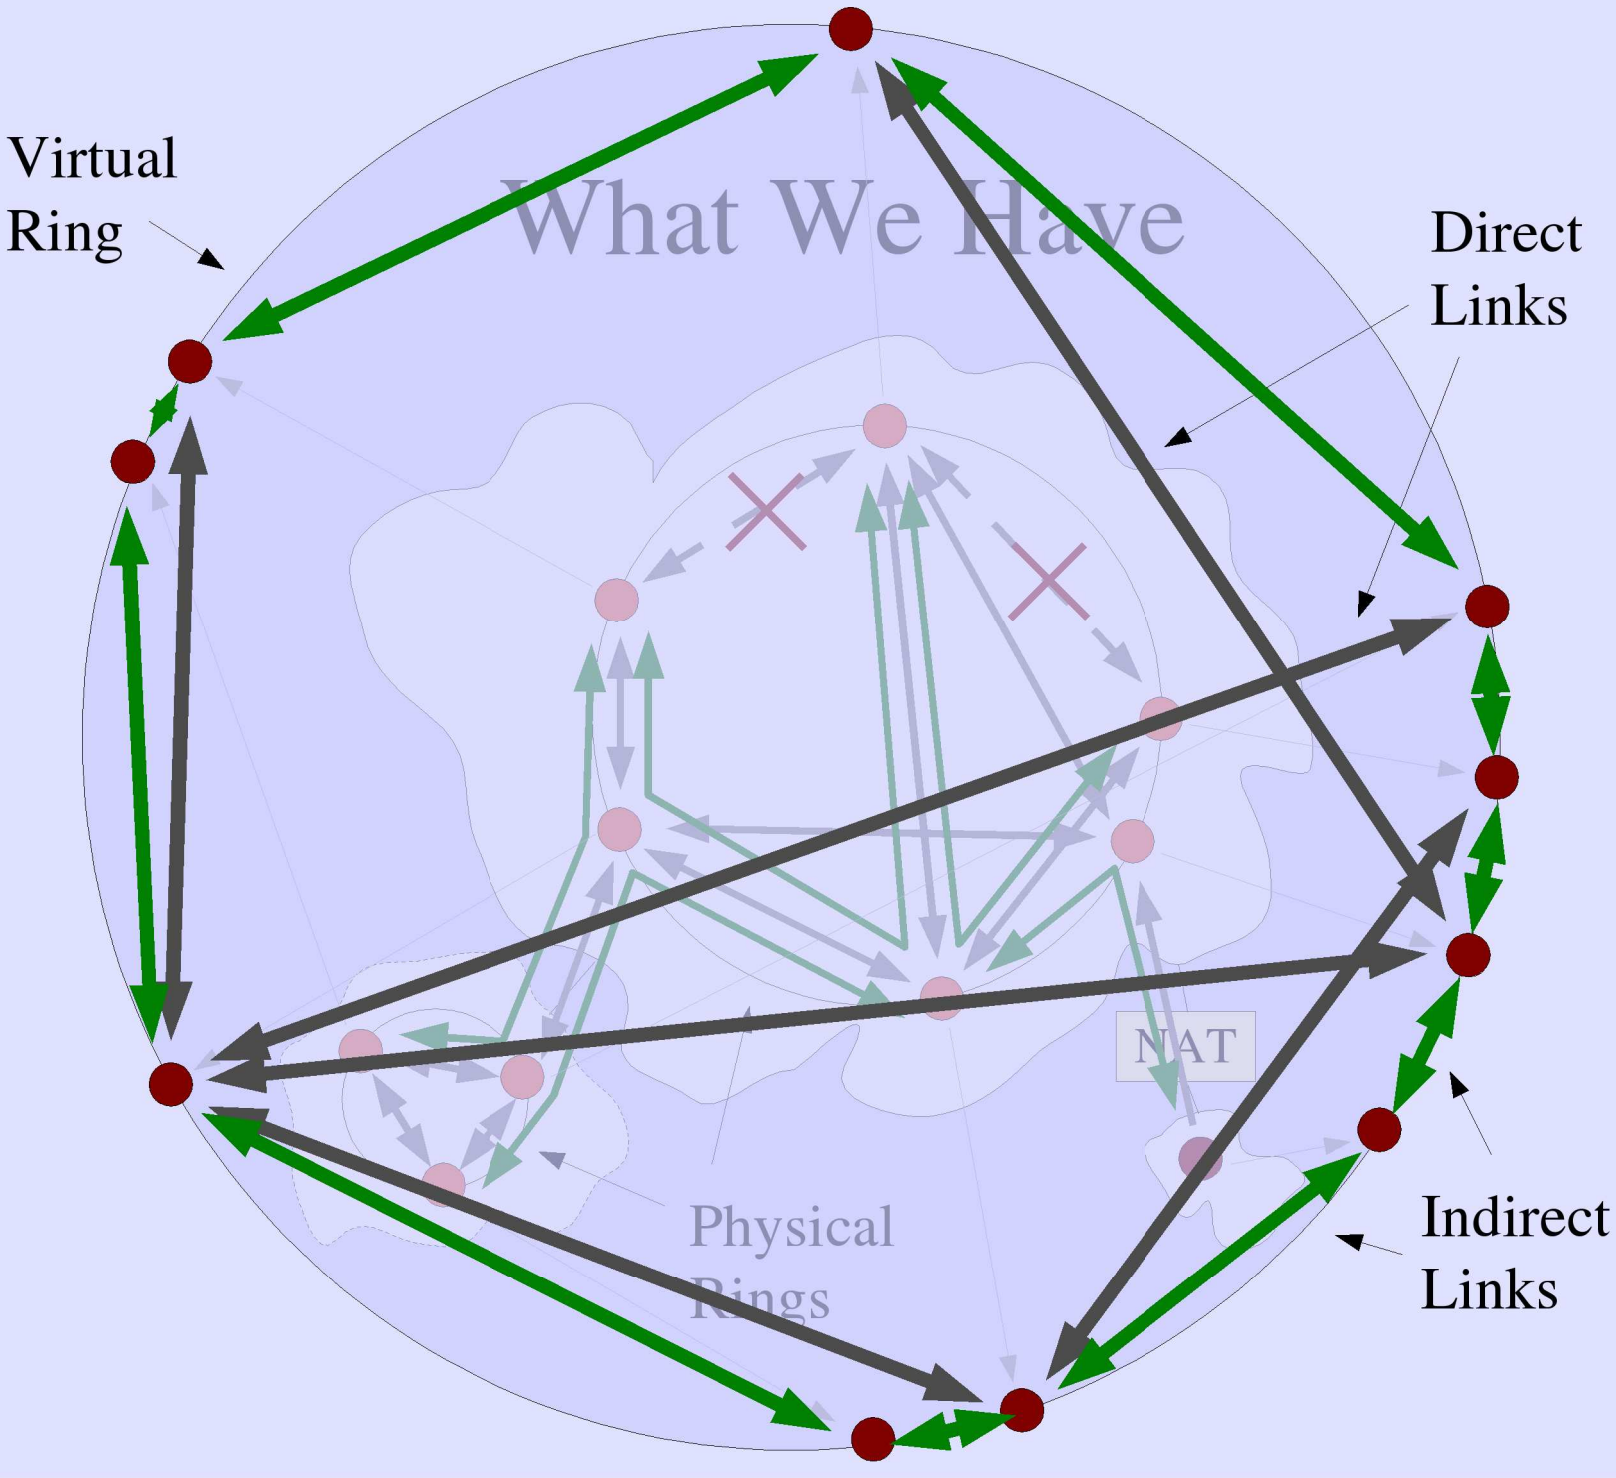
What We Have
Virtual
Ring
Direct
Links
NAT
Physical
Rings
Indirect
Links

Source Routing

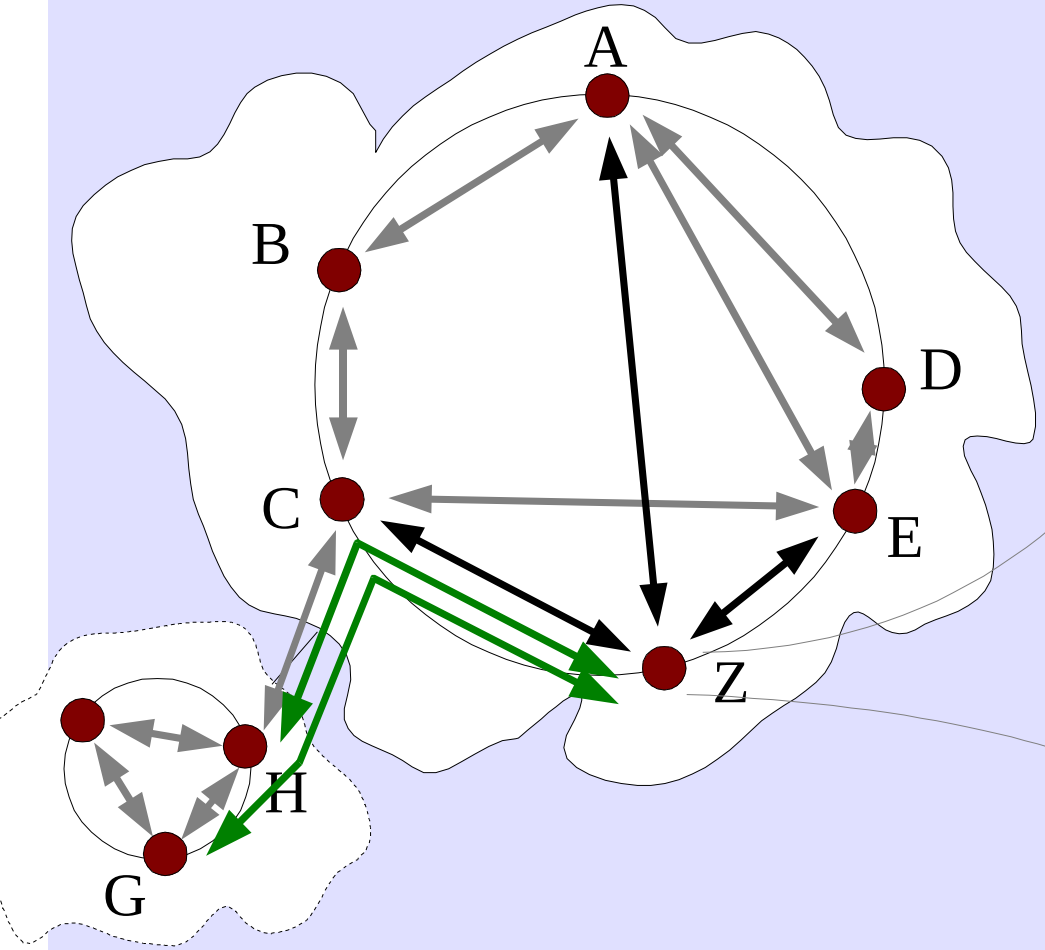

Z's Neighbor Table

A: 12.34.56.78
C: 23.45.67.89
E: 34.56.78.90
H: [C → H]
G: [C → H → G]
. . .

Source Routing

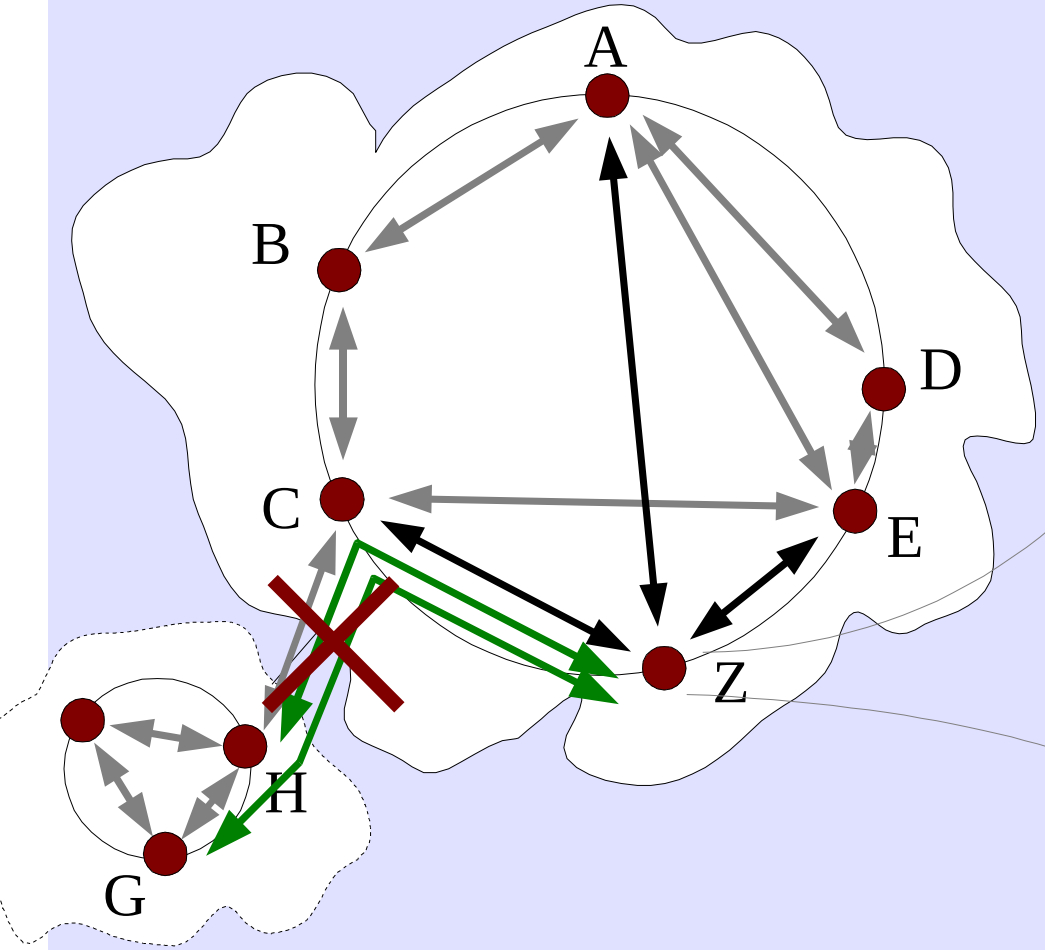

Z's Neighbor Table

A: 12.34.56.78
C: 23.45.67.89
E: 34.56.78.90
~~H: [C → H]~~
~~G: [C → H → G]~~
. . .

Challenges

- Forwarding path optimization
- Healing efficiently after arbitrary partitions
- Incentives for good behavior,
  resistance to denial-of-service attacks

Implementation Status

- Algorithm works under simulation
  - Up to 10,000 nodes, “Internet-like” networks
  - $\approx O(\log n)$ state and maintenance traffic observed
  - Heals quickly after partitions
- In progress:
  - Further algorithm refinement
  - Real-world prototype

Conclusion

- To get ubiquitous networking:
  - Edge nodes *must* be able to operate without centralized address assignment:
    Address-Based Routing ⇨ *Identity-Based Routing*
  - Edge routing protocols must be self-managing
    at global Internet-wide scales, not just locally
- Scalable IBR is hard, but should be feasible