

Intra-domain Routing (2)

Networked Systems (H)

Lecture 10

Lecture Outline

- Intra-domain unicast routing
 - ...
 - Link state routing

Limitations of Distance Vector Routing

- Distance vector routing tries to minimise state at nodes – as a consequence, is slow to converge
- An alternative is *link state* routing


Link State Routing

- Nodes know the links to their neighbours, and the cost of using those links
 - The *link state* information
- Reliably flood this information, giving all nodes complete map of the network
- Each node then directly calculates shortest path to every other node, uses this as routing table


Link State Information

- Link state information updates are flooded on start-up, and when the topology changes
- Each update contains:
 - The address of node that sent the update
 - List of directly connected neighbours of that node
 - With the cost of the link to each neighbour
 - With the range of addresses assigned to each link, if it connects to more than one host
 - A sequence number

Flooding Link State Updates




Flooding Link State Updates



Node C sends an update to each of its neighbours


Flooding Link State Updates



Node C sends an update to each of its neighbours

Each receiver compares the sequence number with that of the last update from C, if greater it forwards the update on all links except the link on which it was received.

Flooding Link State Updates




Node C sends an update to each of its neighbours

Each receiver compares the sequence number with that of the last update from C, if greater it forwards the update on all links except the link on which it was received.

Each receiver compares the sequence number with that of the last update from C, if greater it forwards the update on all links except the link on which it was received.

Flooding Link State Updates



Node C sends an update to each of its neighbours

Each receiver compares the sequence number with that of the last update from C, if greater it forwards the update on all links except the link on which it was received.

Each receiver compares the sequence number with that of the last update from C, if greater it forwards the update on all links except the link on which it was received.

Eventually, the entire network has received the update

Calculate Shortest Paths

- Flooding link state data from all nodes ensures all nodes know the entire topology
- Each node uses Dijkstra's shortest-path algorithm to calculate optimal route to every other node
 - Optimal is assumed to be the shortest path, by weight

Shortest Path Algorithm

Definitions:

N set of all nodes in the graph
 $l(i, j)$ weight of link from i to j (∞ if no link, 0 if $i = j$)
 s source node from which we're calculating shortest paths

Dijkstra's Algorithm for an undirected connected graph:

```
 $M = \{s\}$                                 The set of nodes that have been checked  
foreach  $n$  in  $N - \{s\}$ :                    The distance to directly connected neighbouring nodes  
     $C(n) = l(s, n)$   
  
while ( $N \neq M$ ):  
     $c = \infty$   
    foreach  $n$  in  $(N-M)$   
        if  $C(w) < c$  then  $w = n$           Find node  $w$  such that  $C(w)$  is the minimum for all nodes in  $(N-M)$   
  
     $M += \{w\}$                             Add one node at a time, starting with the closest  
    foreach  $n$  in  $(N-M)$ :  
        if  $C(n) > C(w) + l(w, n)$  then  $C(n) = C(w) + l(w, n)$     Best route to  $n$  is via  $w$ 
```

Result:

$C(x)$ cost of the shortest path from s to x

Forwarding and Route Updates

- Forward packets based on calculated shortest path
 - Static forwarding decision based on weights distributed by the routing protocol
 - Does not take into account network congestion
- Recalculate shortest paths on every routing update
 - Updates occur if a link fails, or a new link is added

Distance Vector vs Link State

Distance vector routing:

- Simple to implement
- Routers only store the distance to each other node: $O(n)$
- Suffers from slow convergence

Link State routing:

- More complex
- Requires each router to store complete network map: $O(n^2)$
- Much faster convergence

Slow convergence times make distance vector routing unsuitable for large networks

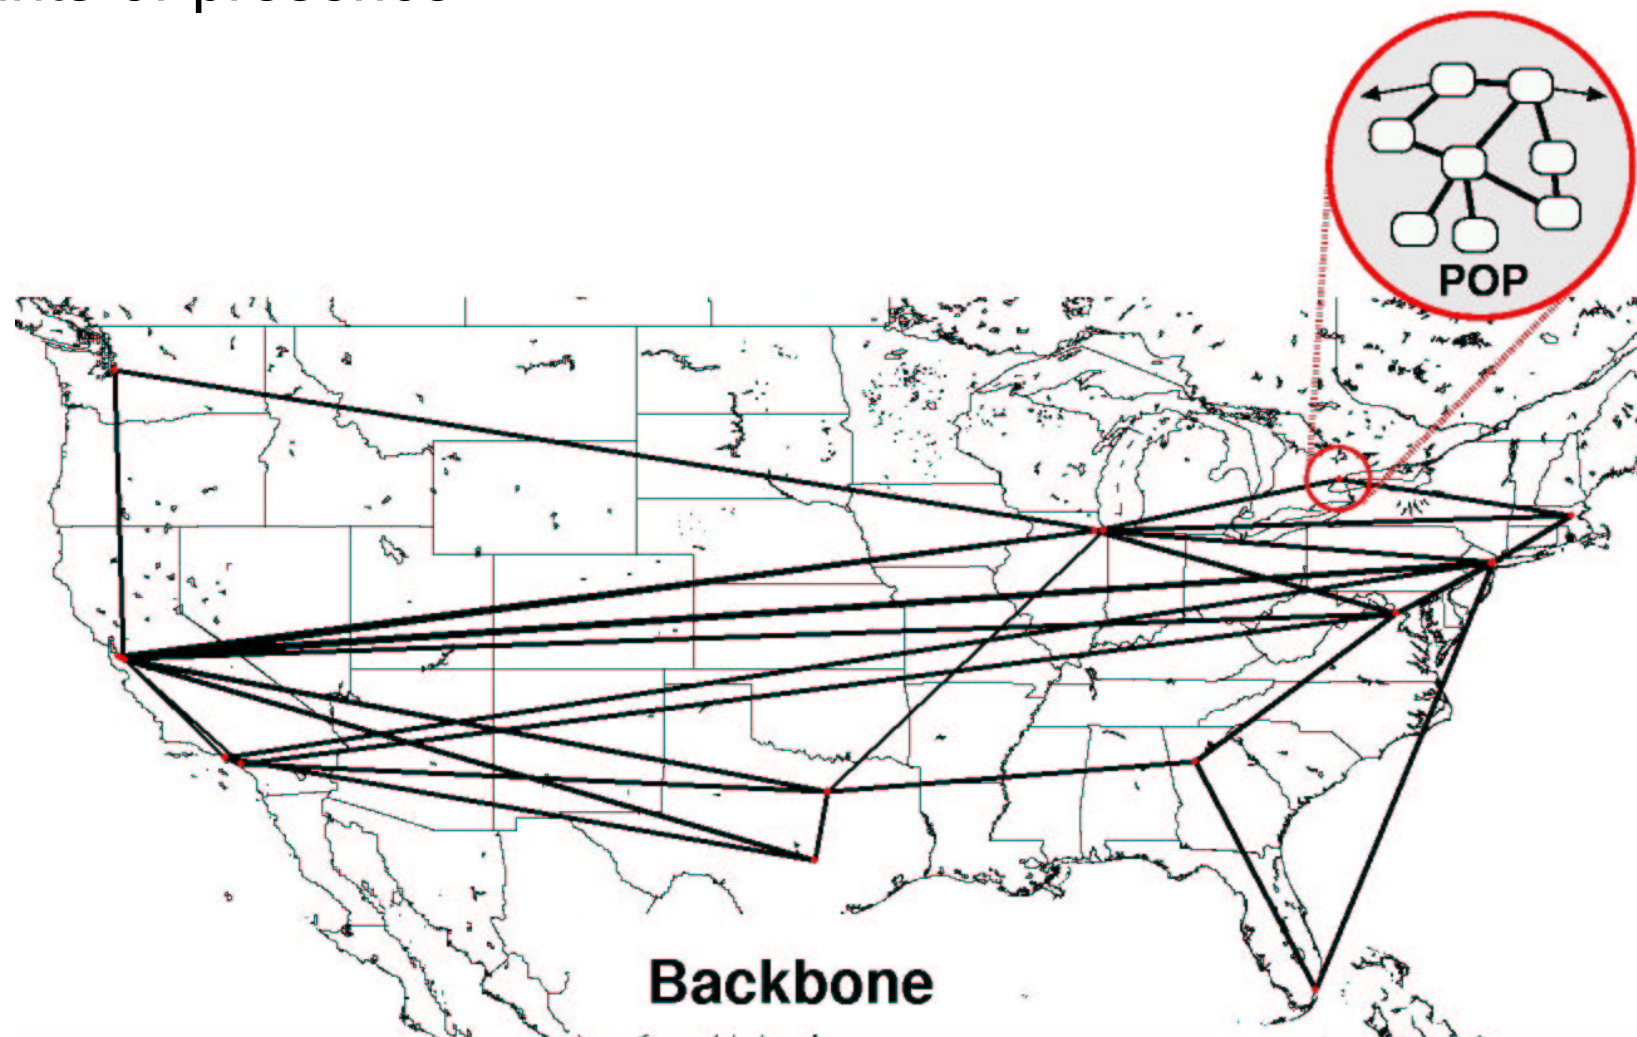
Where is intra-domain routing used?

- Intra-domain routing operates within a network
 - Any network operated as a single entity – autonomous system – could be local area, nationwide, or even worldwide
 - Operates a single routing protocol – typically the OSPF link-state protocol exchanging routes to IP address prefixes
 - Running on IP routers within an autonomous system, typically with fibre connections wide area and ethernet local area
 - Exchange routes to IP prefixes, representing regions in the network topology



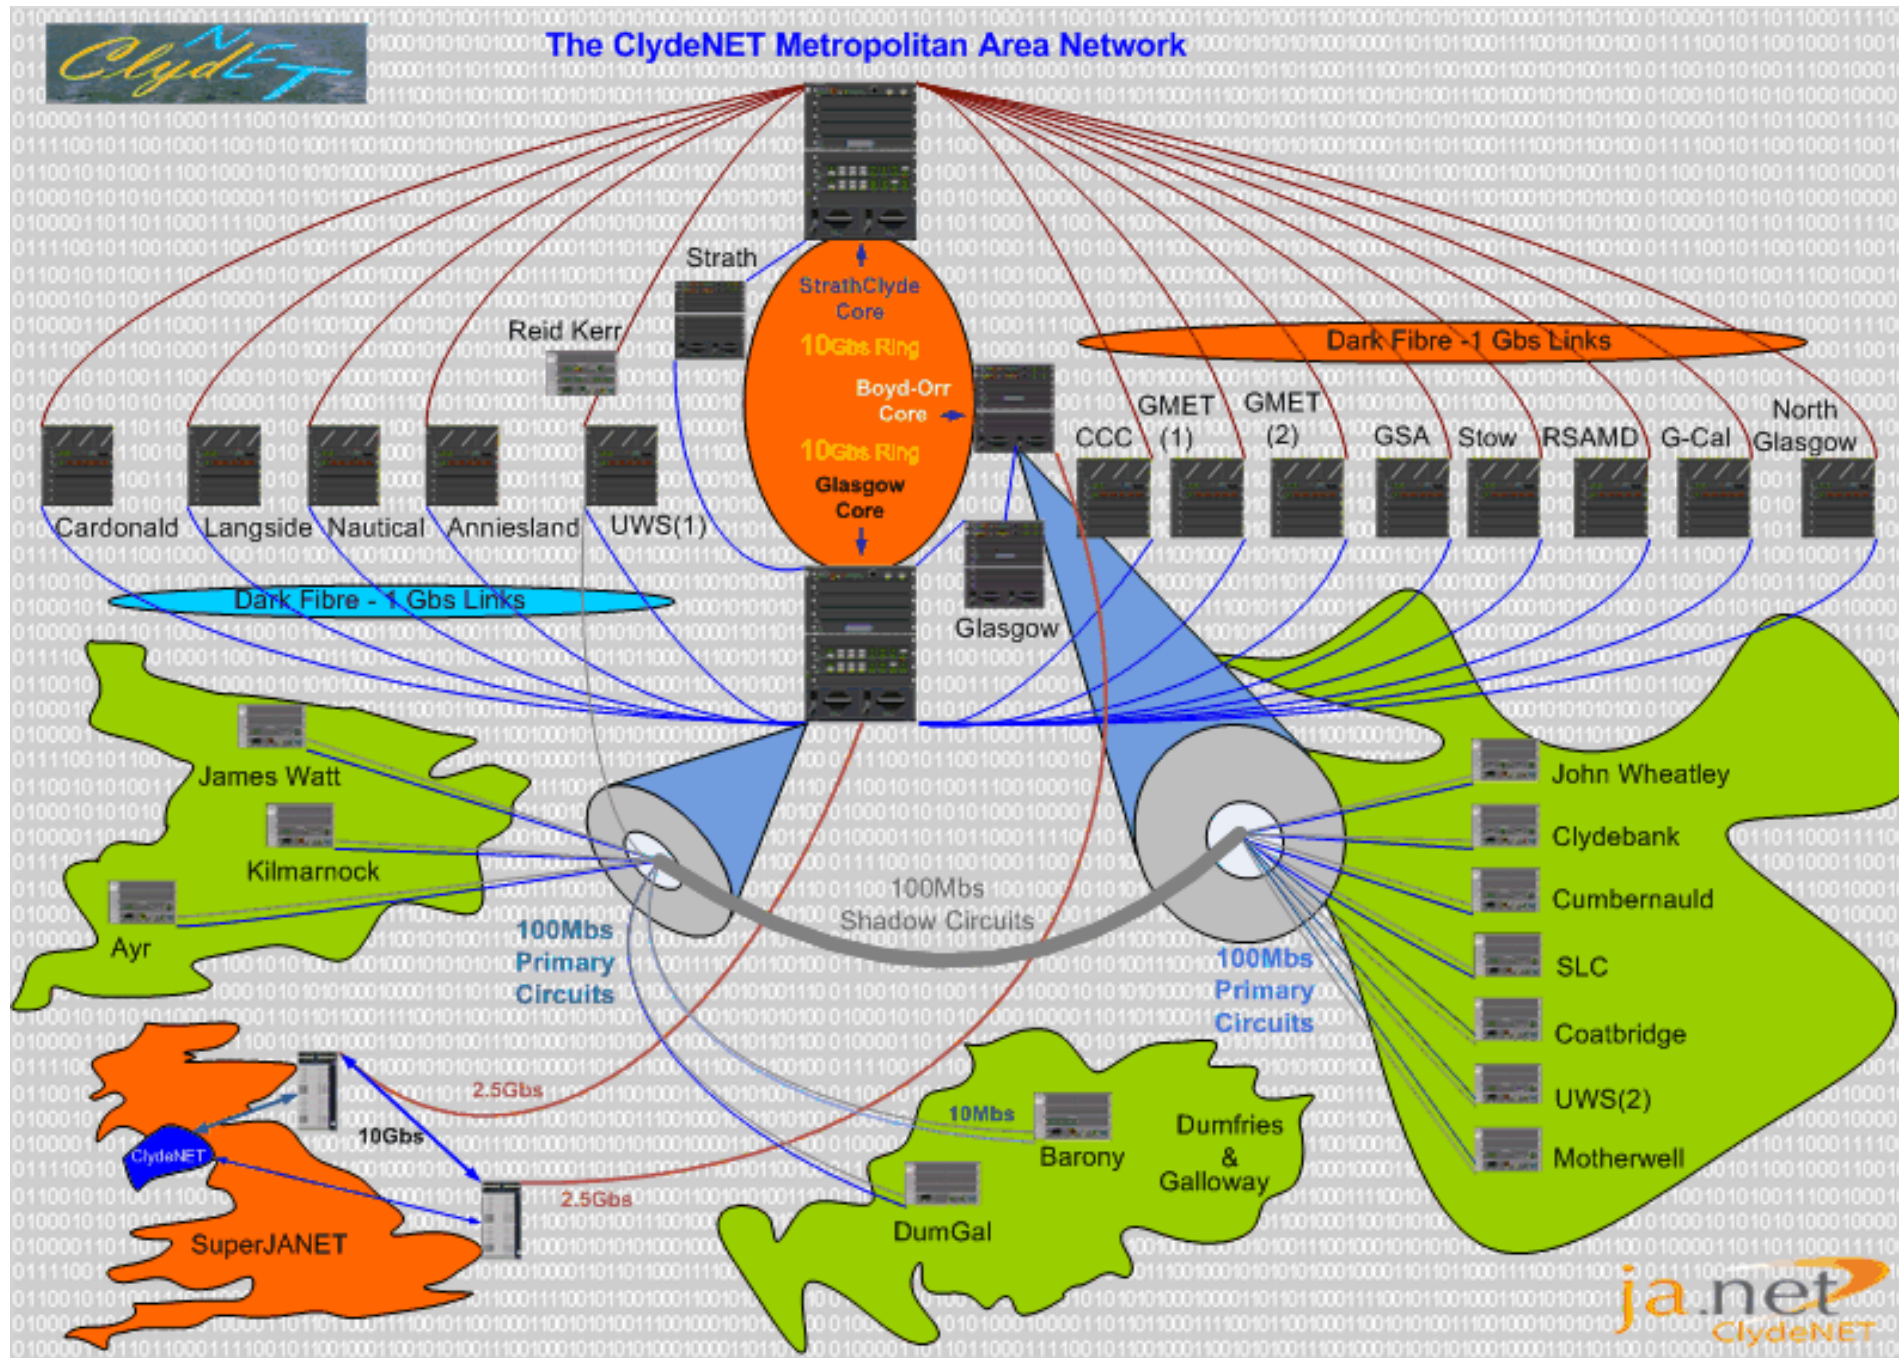
What is the Topology of Real Networks?

- Wide-area backbone network connecting local points-of-presence



N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Network Topologies with Rocketfuel", Presentation at ACM SIGCOMM conference, 2002

Example: ClydeNet

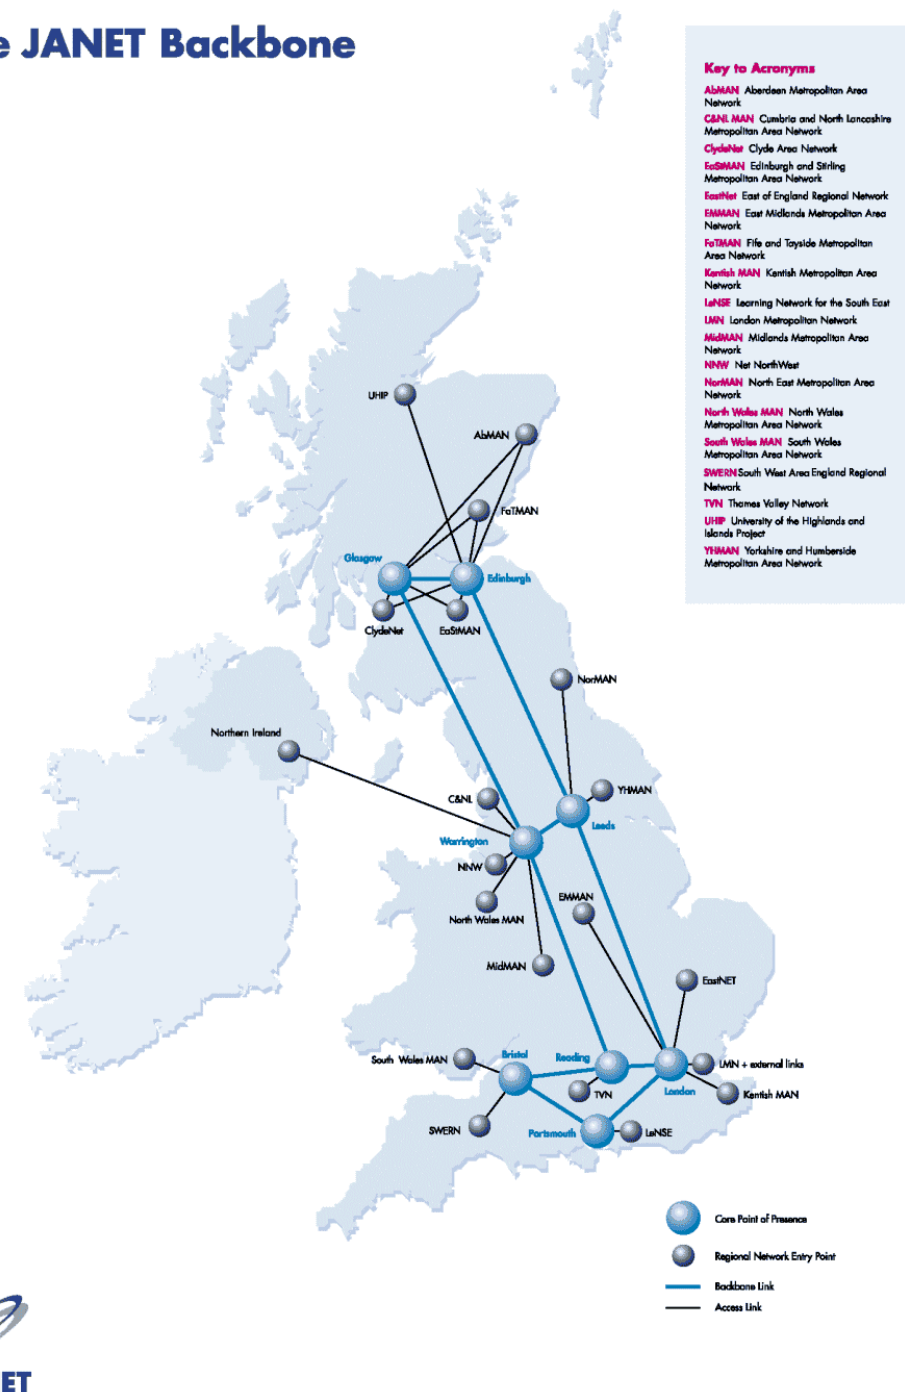


ClydeNet is the metropolitan area network for the Glasgow region – the regional JANET PoP

Example: The JANET Backbone (2001)

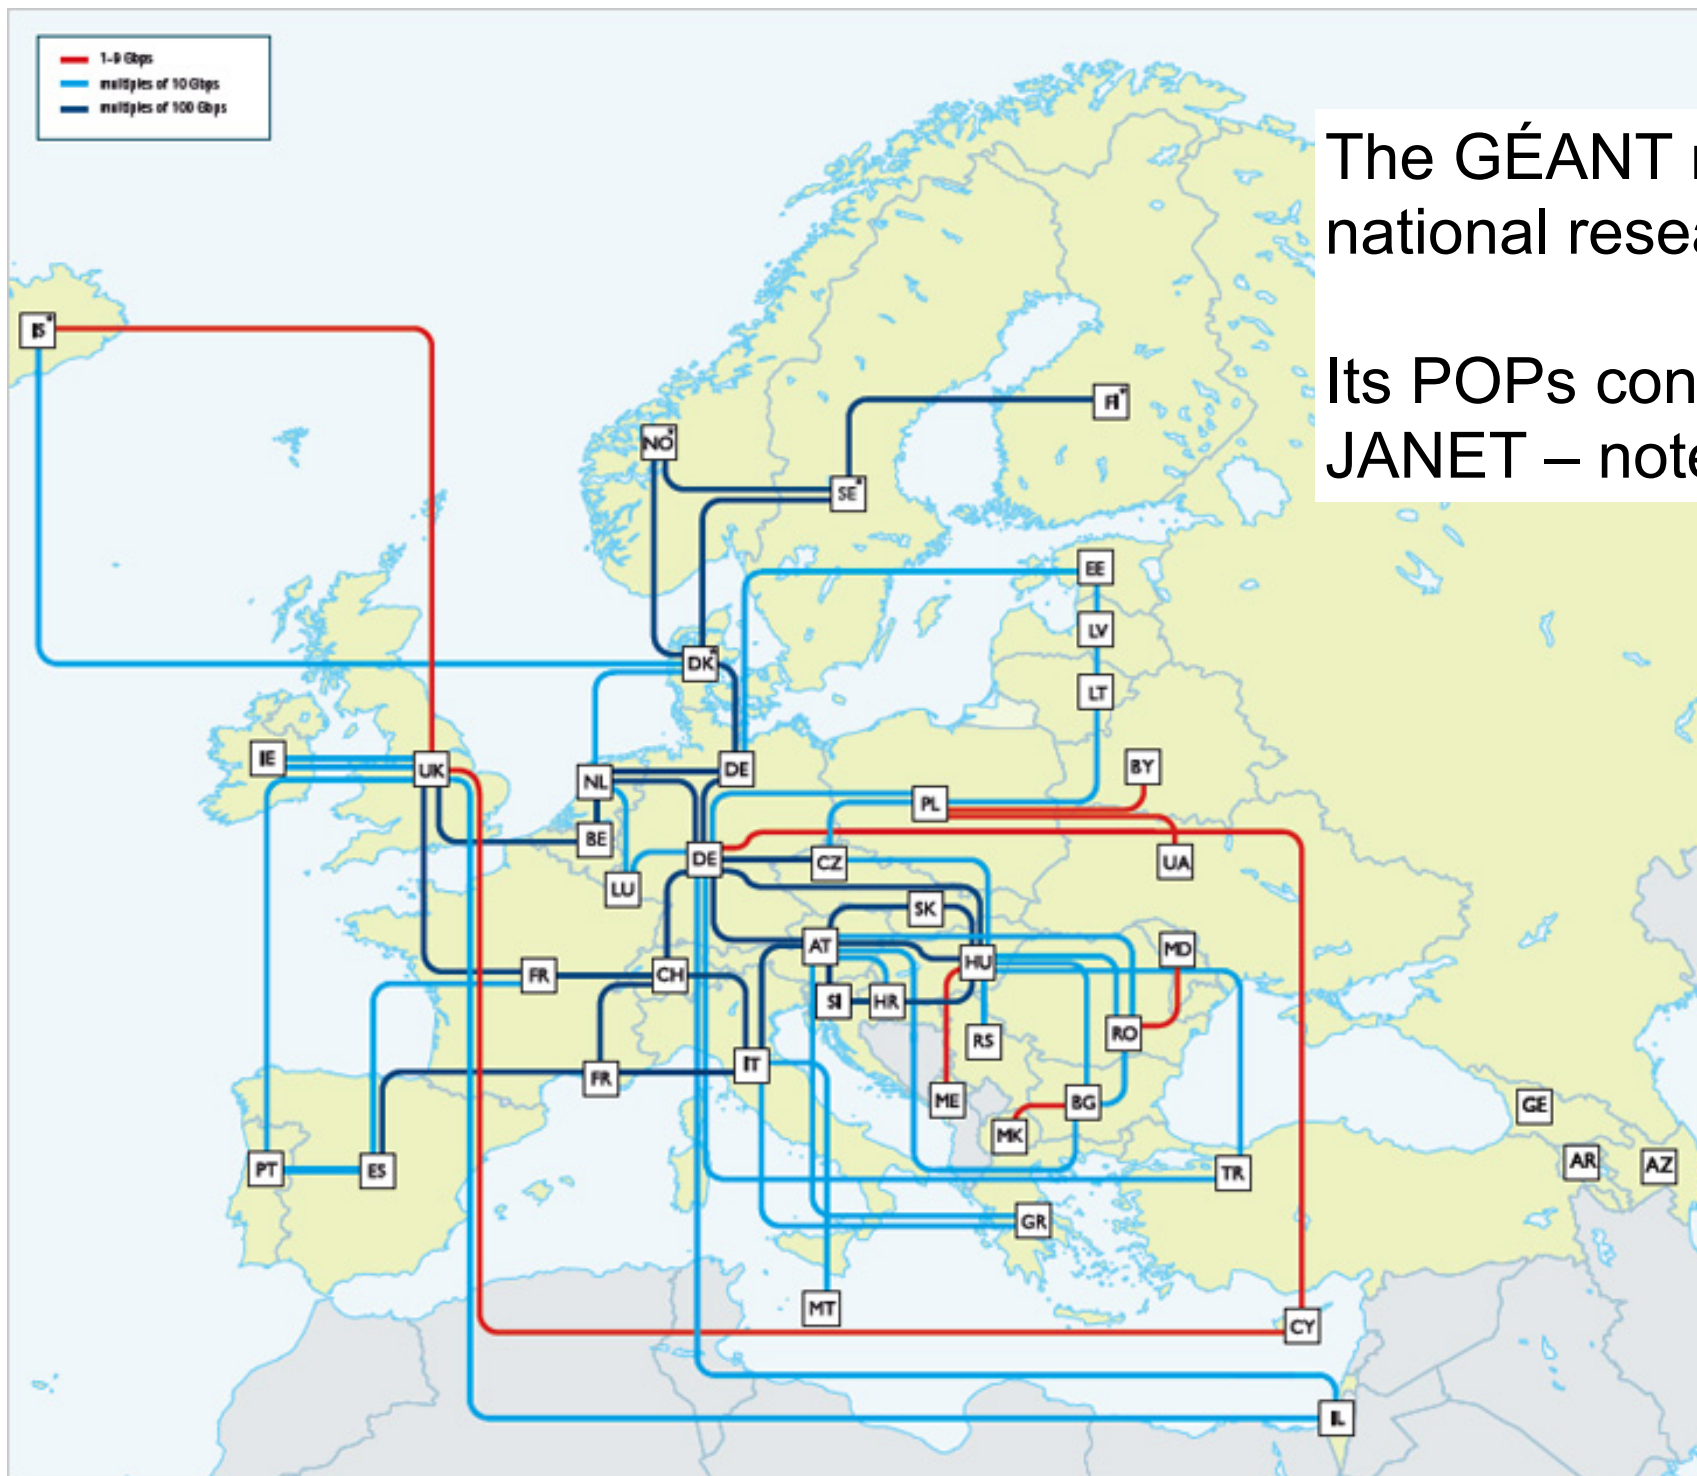
- JANET is the UK national research network
- It interconnects major universities and metropolitan area networks – its customers are a mix of end sites and other networks

The JANET Backbone



© The JNT Association 2000 JNS/NAF/004

Example: GÉANT



The GÉANT network inter-connects the national research networks in Europe

Its POPs connect to networks such as JANET – note almost fractal structure

Summary

- Distance vector vs link state routing
- Local, metropolitan, or wide area – key is routing within a single autonomous system