Routing part 2

Jens A Andersson

Electrical and Information Technology

Routing

- Introduction
- Inside the Router
- Unicast Routing
  - Intra Domain Routing
  - Inter Domain Routing
- MANET and AdHoc routing
- Multicast Routing

Internet AS Hierarchy

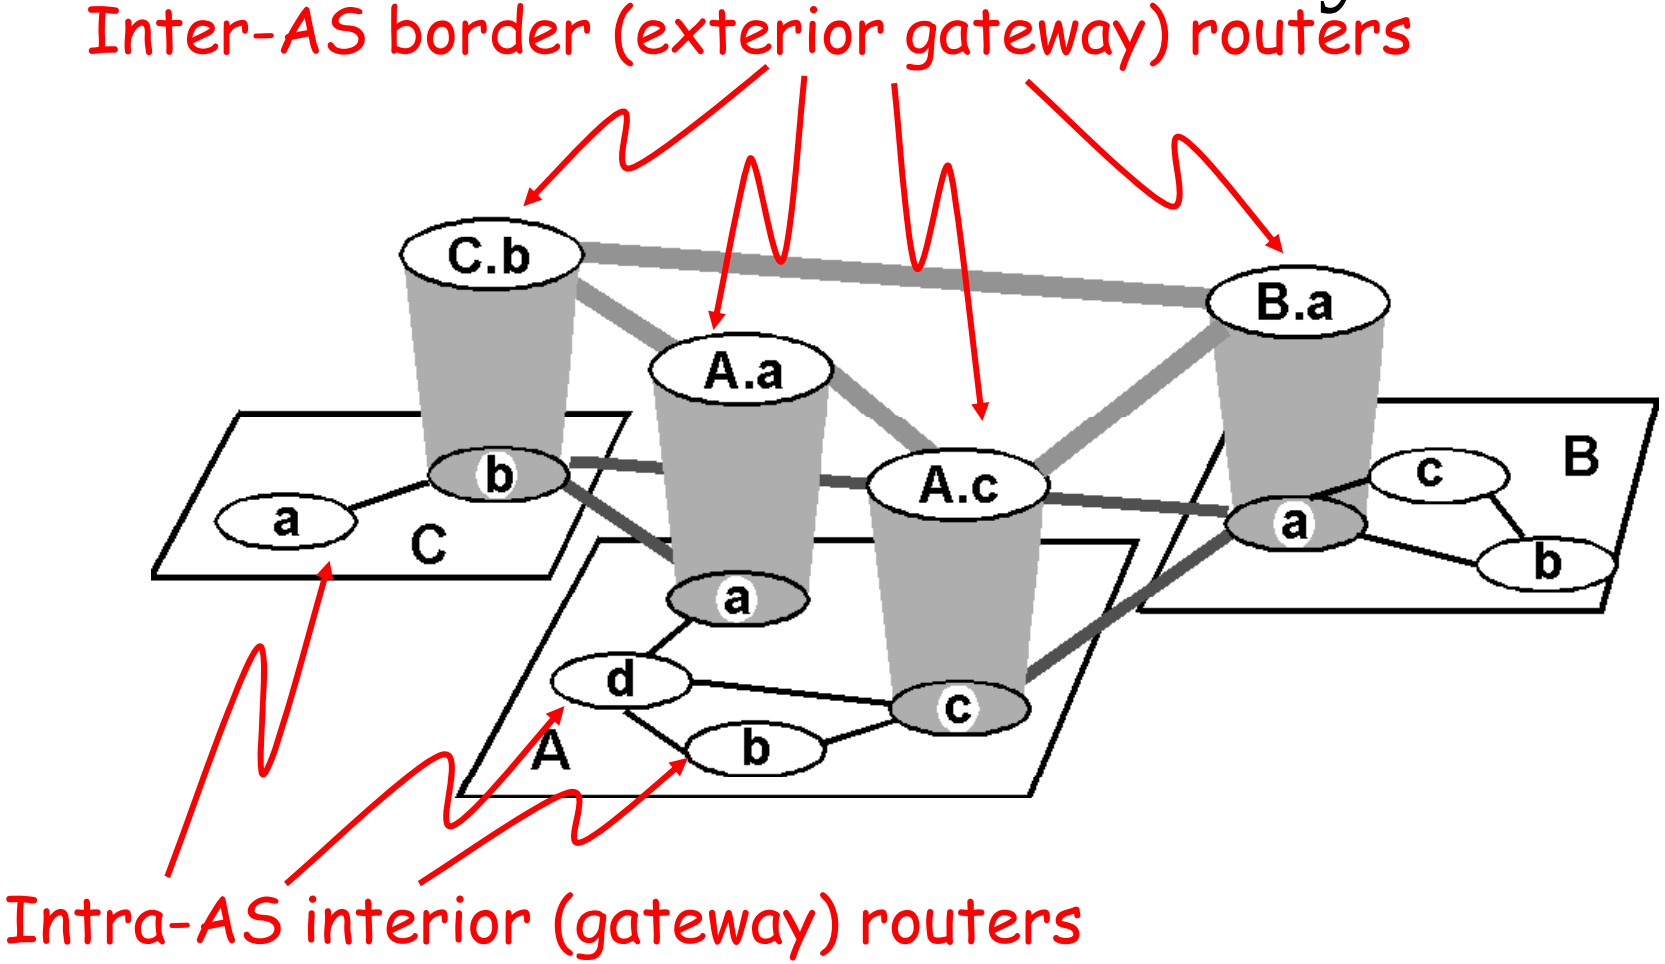

Hierarchical Routing

- aggregate routers into regions, "autonomous systems" (AS)
- routers in same AS run same routing protocol
  - "intra-AS" routing protocol
  - routers in different AS can run different intra-AS routing protocol

gateway routers

- special routers in AS
- run intra-AS routing protocol with all other routers in AS
- *also* responsible for routing to destinations outside AS
  - run *inter-AS routing* protocol with other gateway routers

Why different Intra- and Inter-AS routing ?

Policy:

- Inter-AS: admin wants control over how its traffic routed, who routes through its net.
- Intra-AS: single admin, so no policy decisions needed

Scale:

- hierarchical routing saves table size, reduced update traffic

Performance:

- Intra-AS: can focus on performance
- Inter-AS: policy may dominate over performance

Internet inter-AS routing: BGP

- BGP (Border Gateway Protocol): *the* de facto standard
- **Path Vector** protocol:
  - similar to Distance Vector protocol
  - each Border Gateway broadcast to neighbors (peers) *entire path* (i.e., sequence of AS's) to destination
  - BGP routes to networks (ASs), not individual hosts
  - Gateway X can advertise its path to dest. Z:

    Path (X,Z) = X,Y1,Y2,Y3,…,Z
  - BGP implements policies

Figure 21.21 **Path vector messages**

Path Vector Messages

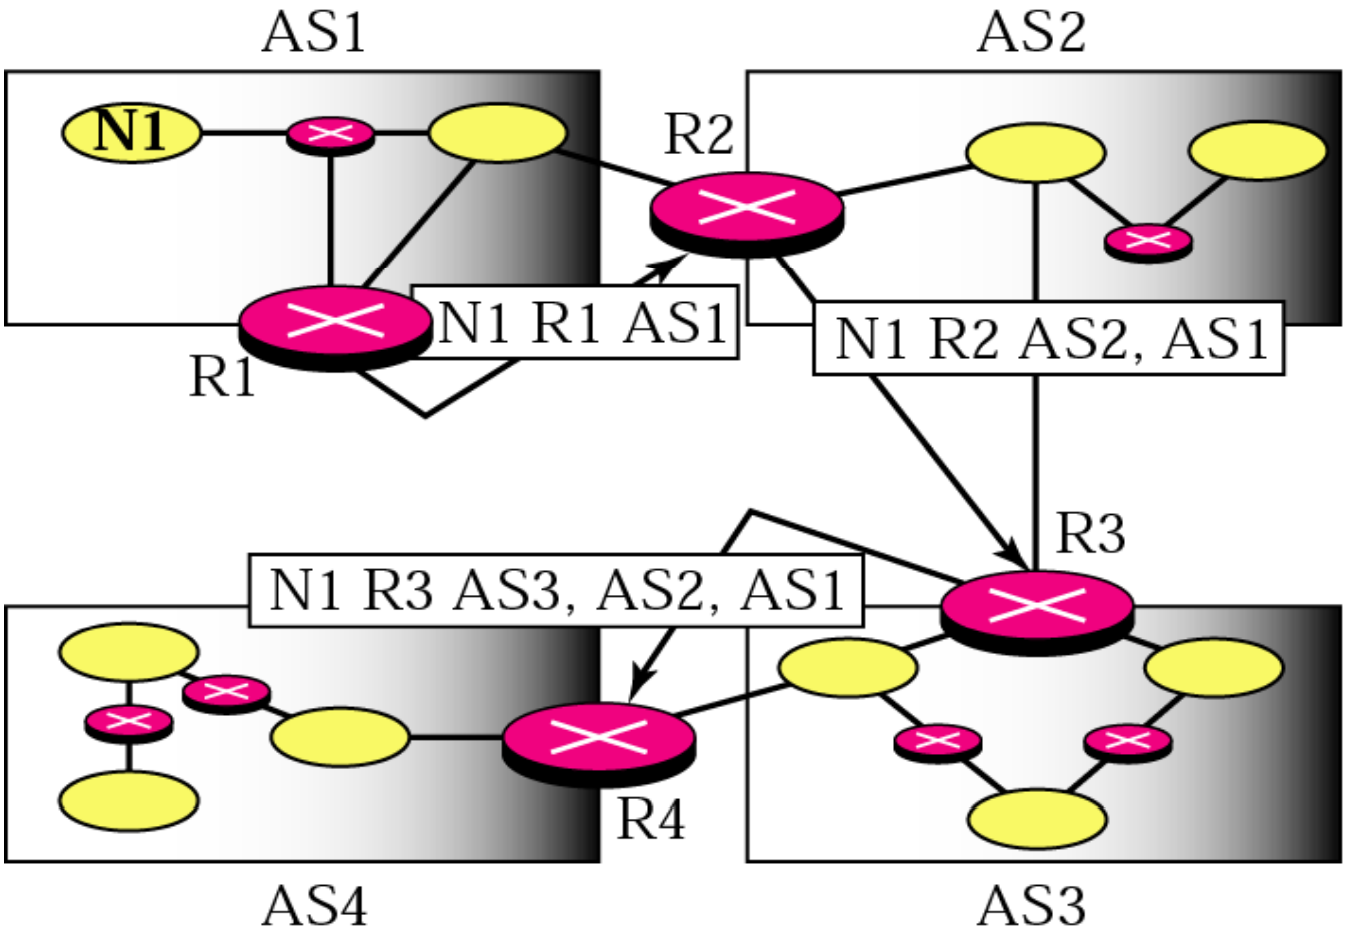

Table 21.3 Path vector routing table

Path Vector Routing Table

Network	Next Router	Path
N01	R01	AS62, AS23, AS67
N02	R05	AS67, AS22, AS05, AS89
N03	R06	AS67, AS89, AS09, AS34
N03	R12	AS62, AS02, AS34

Network id

"Output port"

"Metric"
One of many
ATTRIBUTES

AS = Autonomous System = Organisation

BGP operation

Q: What does a BGP router do?

- Receiving and *filtering* route advertisements from directly attached neighbor(s).
- Route selection.
  - To route to destination X, which path (of several advertised) will be taken?
- Sending route advertisements to neighbours.

BGP Operations (Simplified)

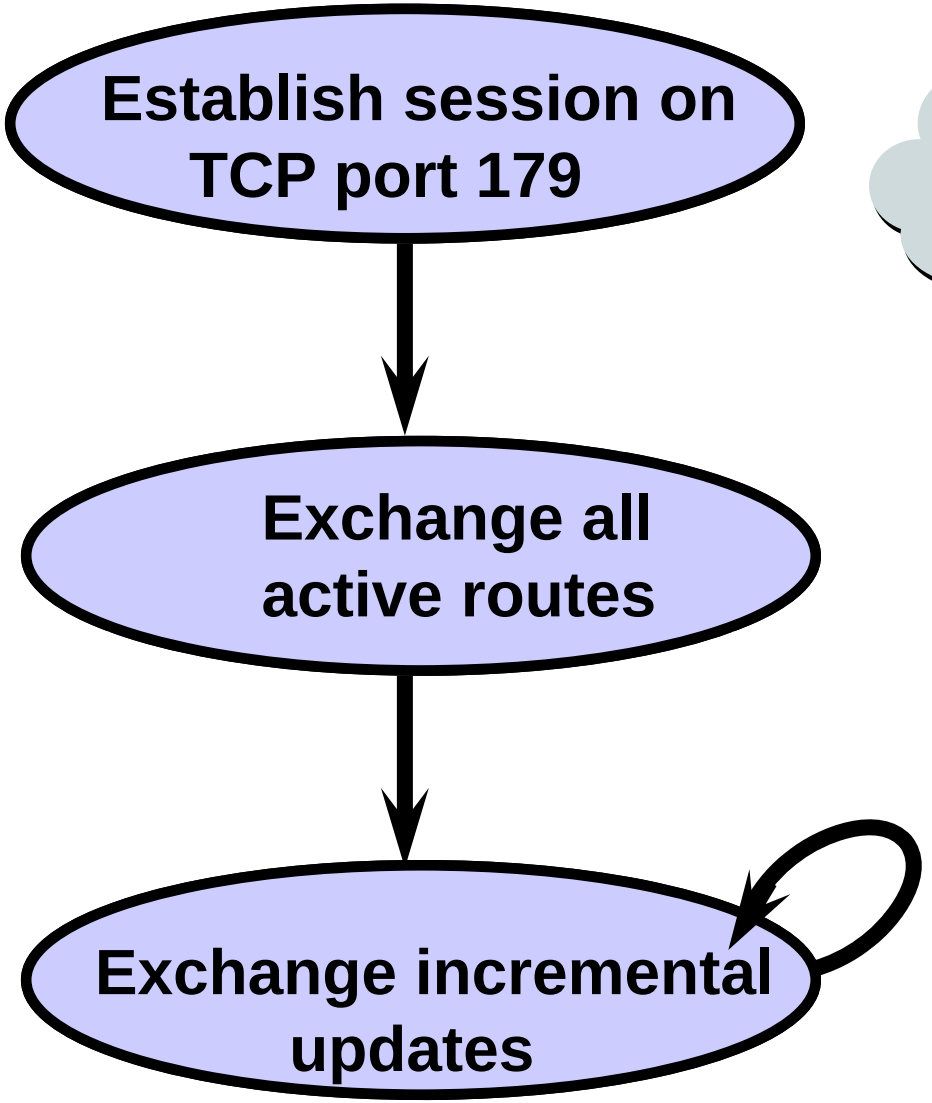

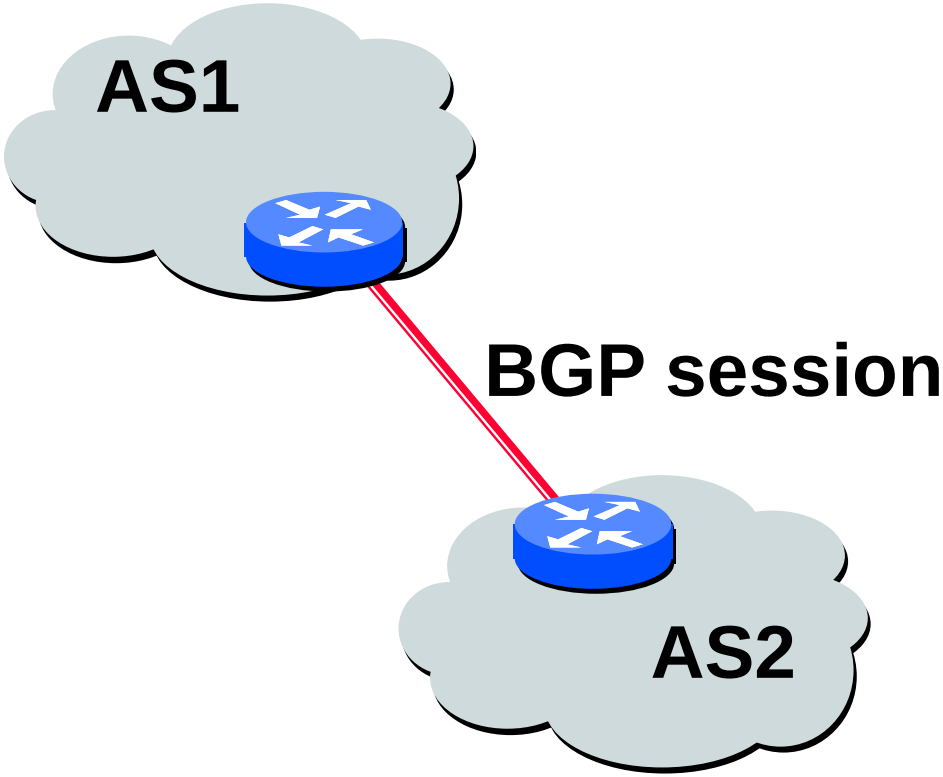

While connection is ALIVE exchange route UPDATE messages

BGP Route Processing

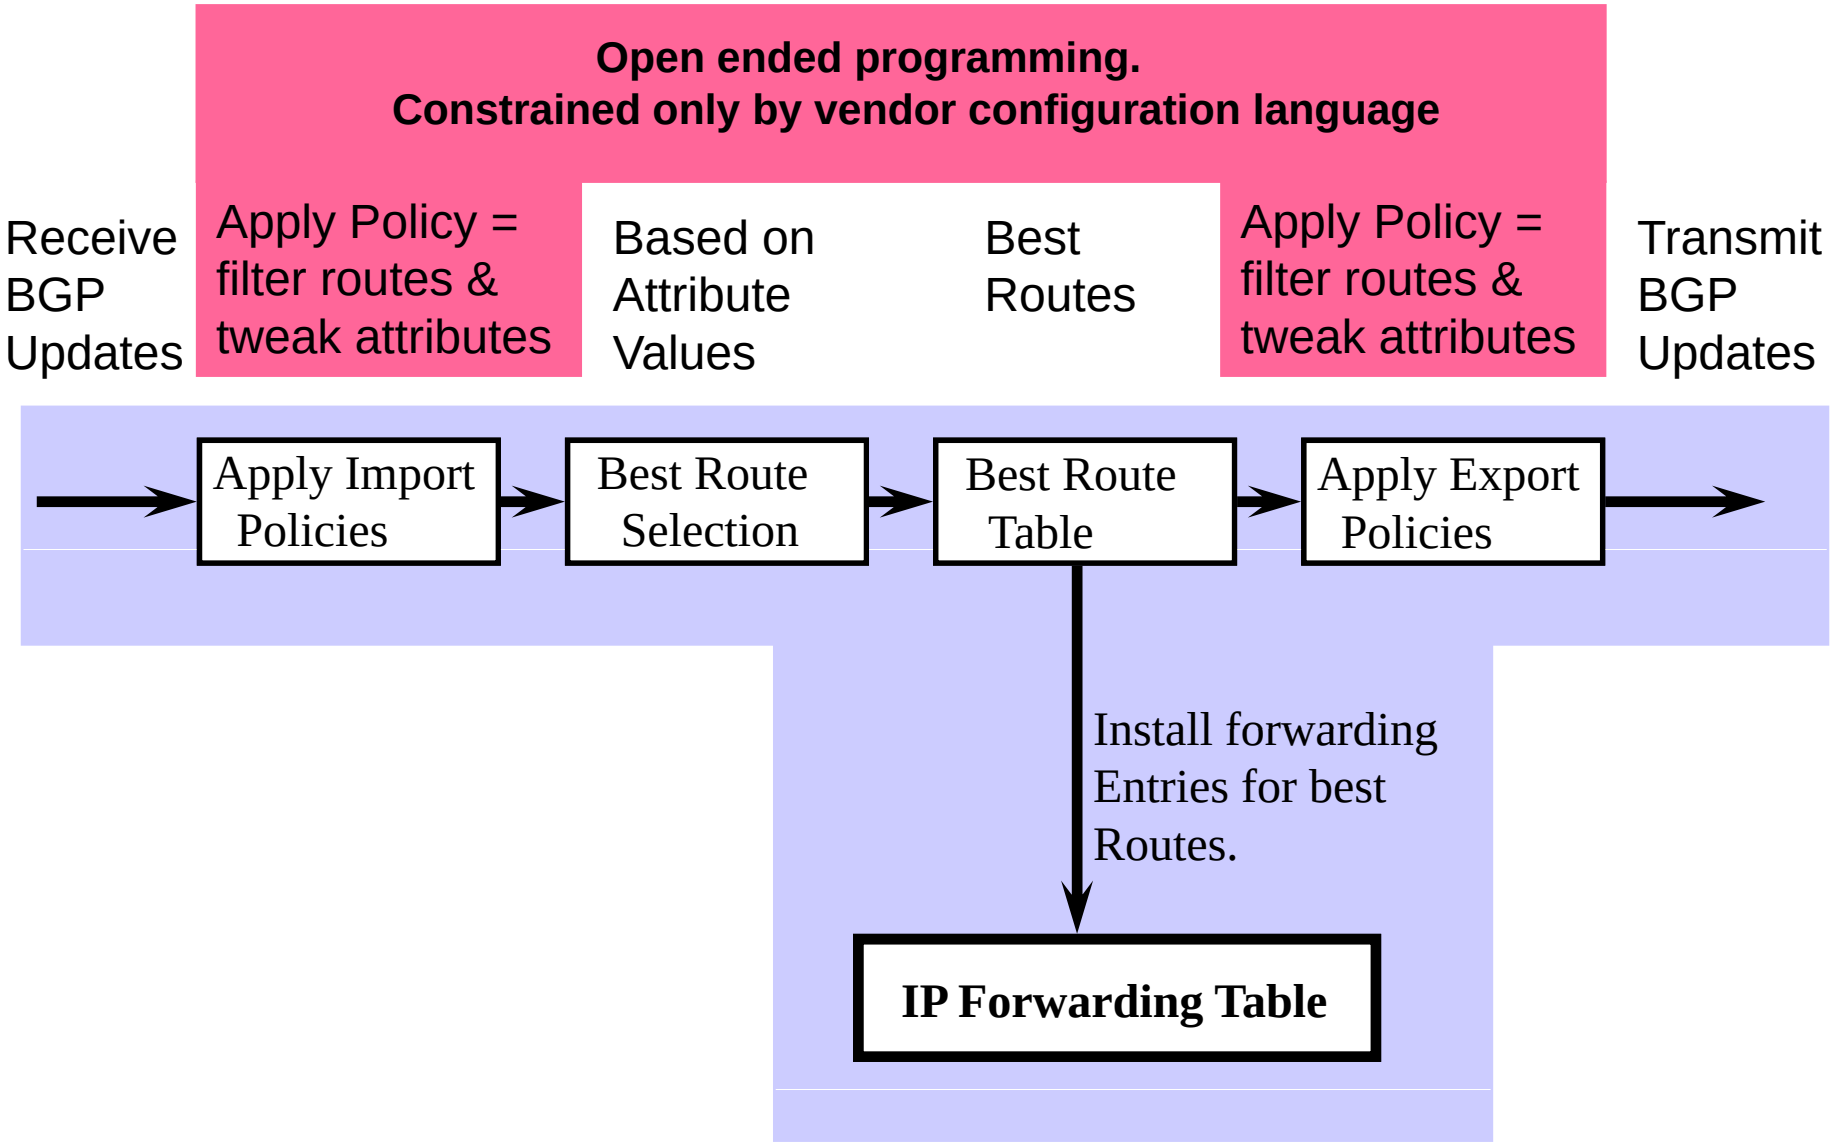

Tweak Tweak Tweak

- For inbound traffic
  - Filter outbound routes
  - Tweak attributes on outbound routes in the hope of influencing your neighbor's best route selection
- For outbound traffic
  - Filter inbound routes
  - Tweak attributes on inbound routes to influence best route selection

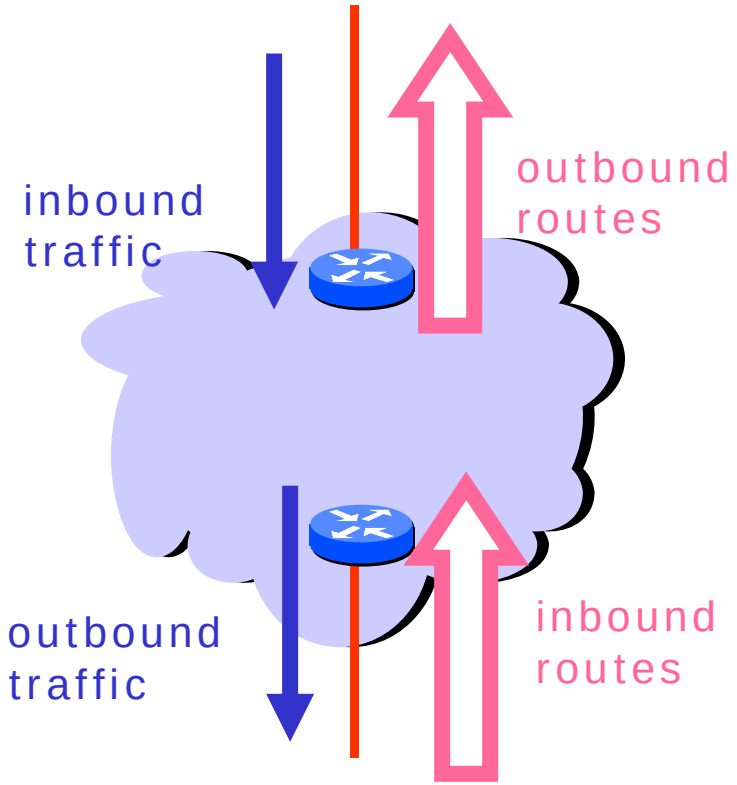

In general, an AS has more control over outbound traffic

BGP messages

- BGP messages exchanged using TCP.
- BGP messages:
  - OPEN: opens TCP connection to peer and authenticates sender
  - UPDATE: advertises new path (or withdraws old)
  - KEEPALIVE keeps connection alive in absence of UPDATES; also ACKs OPEN request
  - NOTIFICATION: reports errors in previous msg; also used to close connection

BGP Attributes

Value	Code	Reference
1	ORIGIN	[RFC1771]
2	AS_PATH	[RFC1771]
3	NEXT_HOP	[RFC1771]
4	MULTI_EXIT_DISC	[RFC1771]
5	LOCAL_PREF	[RFC1771]
6	ATOMIC_AGGREGATE	[RFC1771]
7	AGGREGATOR	[RFC1771]
8	COMMUNITY	[RFC1997]
9	ORIGINATOR_ID	[RFC2796]
10	CLUSTER_LIST	[RFC2796]
11	DPA	[Chen]
12	ADVERTISER	[RFC1863]
13	RCID_PATH / CLUSTER_ID	[RFC1863]
14	MP_REACH_NLRI	[RFC2283]
15	MP_UNREACH_NLRI	[RFC2283]
16	EXTENDED COMMUNITIES	[Rosen]
...		
255	reserved for development	

From IANA: http://www.iana.org/assignments/bgp-parameters

Not all attributes need to be present in every announcement

Route Selection Summary

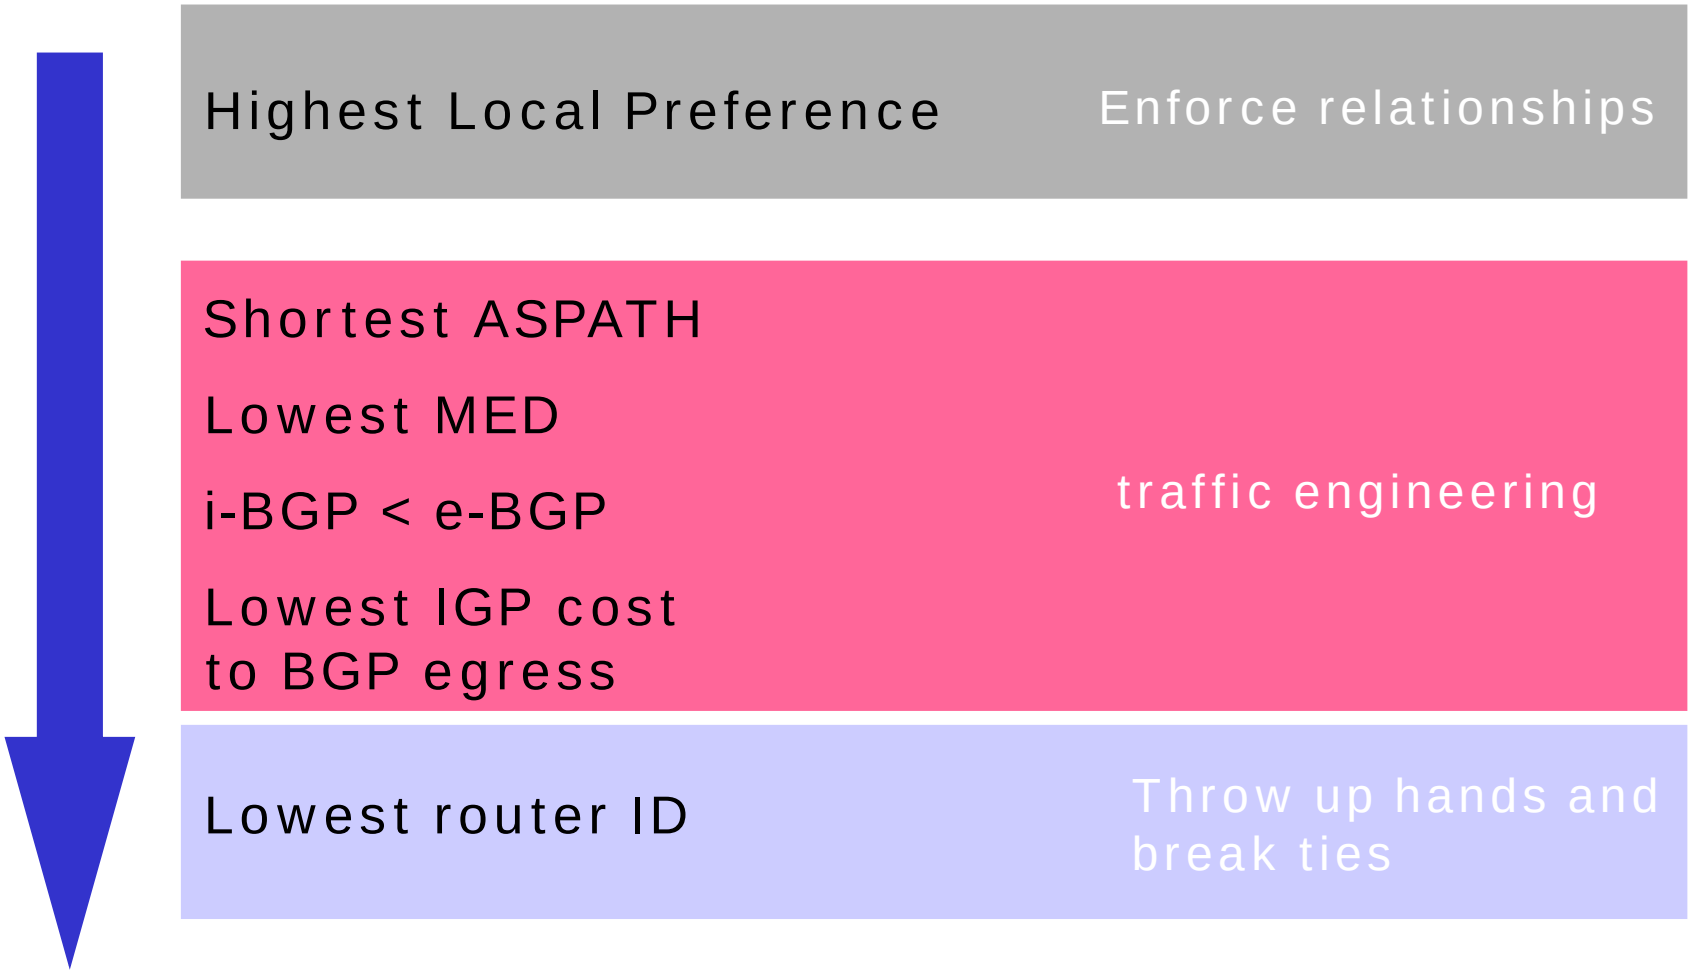

Hot Potato Routing: Go for the Closest Egress Point

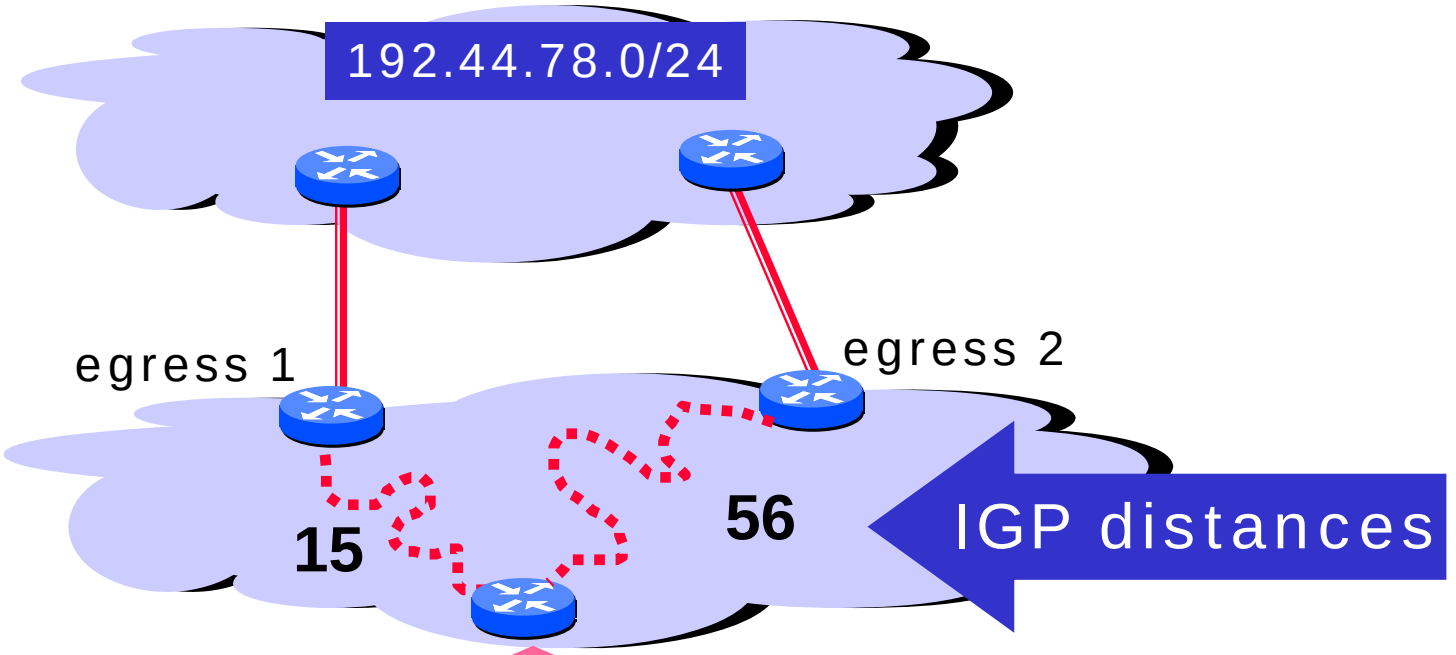

This Router has two BGP routes to 192.44.78.0/24.

Hot potato: get traffic off of your network as Soon as possible. Go for egress 1!

Hot Potato Routing

- Also:

  Routing without queuing

Customer-Provider Hierarchy

Address aggregation!

- Local
- Nation
- Global

Geographical

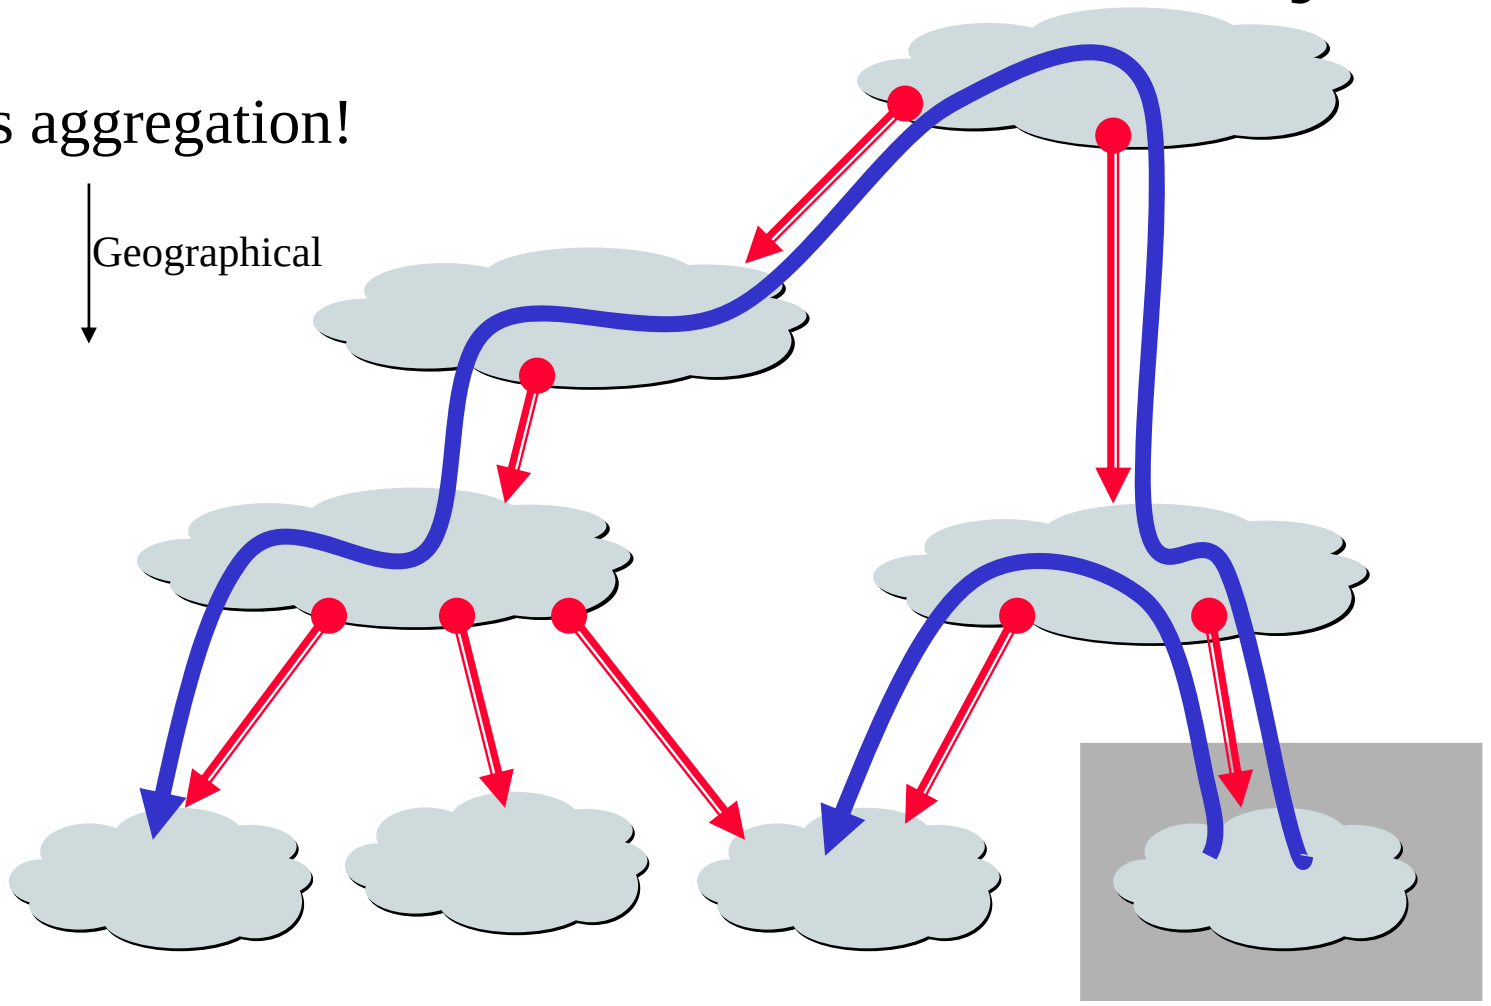

provider → customer

IP traffic

Deaggregation Due to Multihoming

If AS 1 does not announce the more specific prefix, then most traffic to AS 2 will go through AS 3 because it is a longer match

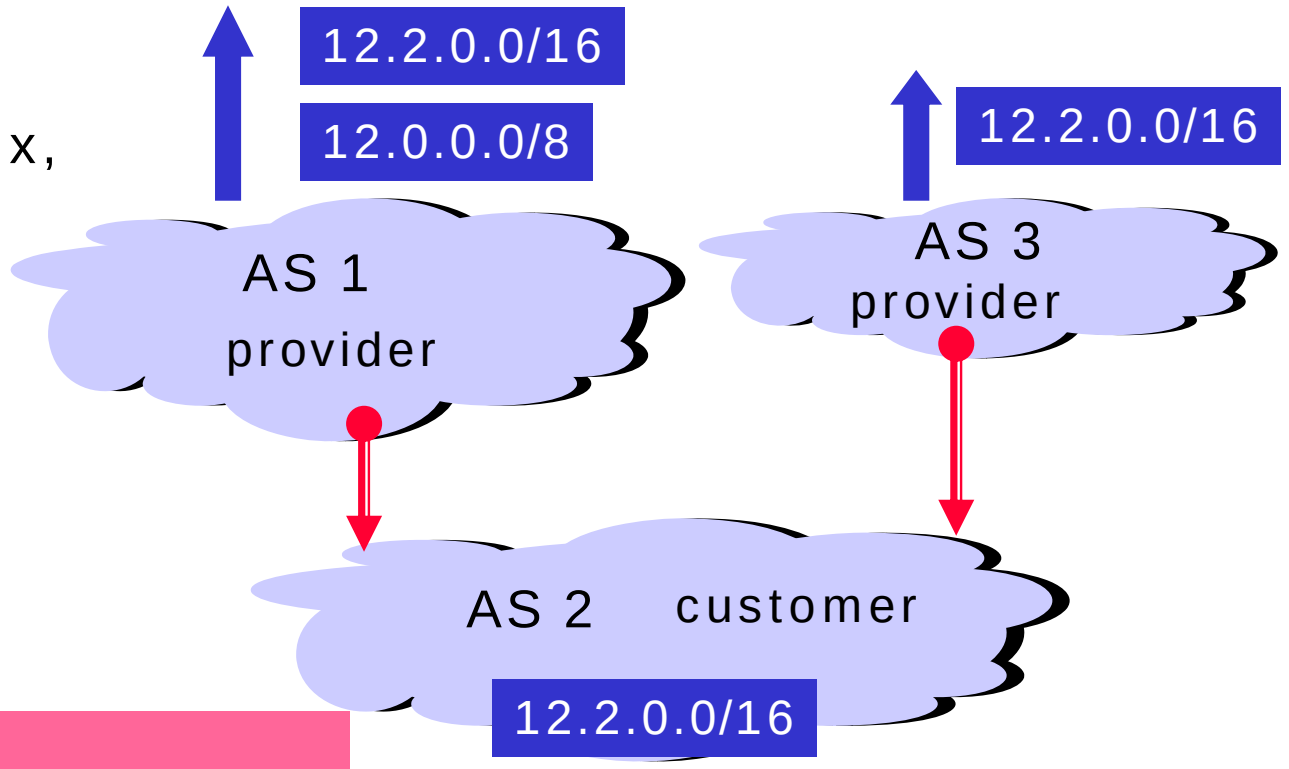

AS 2 is "punching a hole" in The CIDR block of AS 1

Is There A Problem?

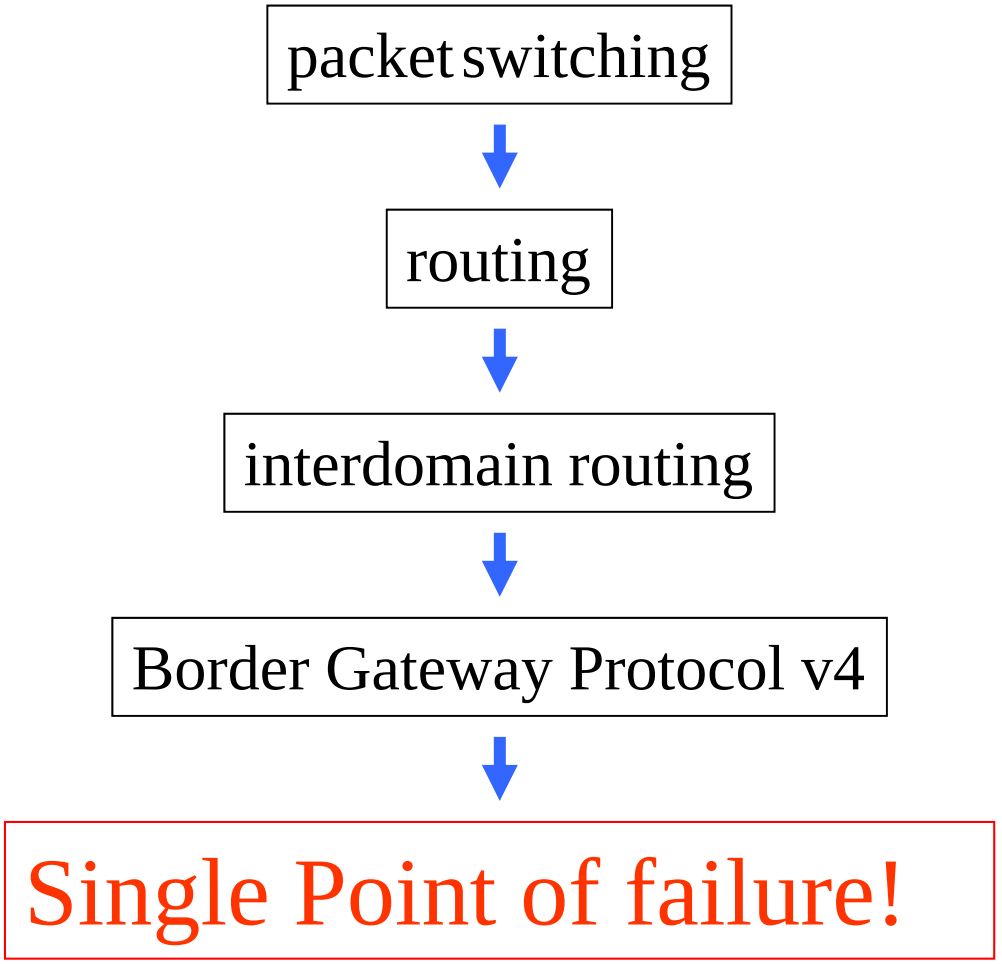

Scarry?

- **BGP is not guaranteed to converge on a stable routing. Policy interactions could lead to "livelock" protocol oscillations.**
  See "Persistent Route Oscillations in Inter-domain Routing" by K. Varadhan, R. Govindan, and D. Estrin. ISI report, 1996
- **Corollary: BGP is not guaranteed to recover from network failures.**

Routing

- Introduction
- Inside the Router
- Unicast Routing
- Multicast Routing
- MANET Routing

MANET routing

- Mobile Ad hoc NETwork
- Ad hoc = dynamic network structure
  - members come and go
- Network nodes move!
- Each host also a routing node

An Ad Hoc Network

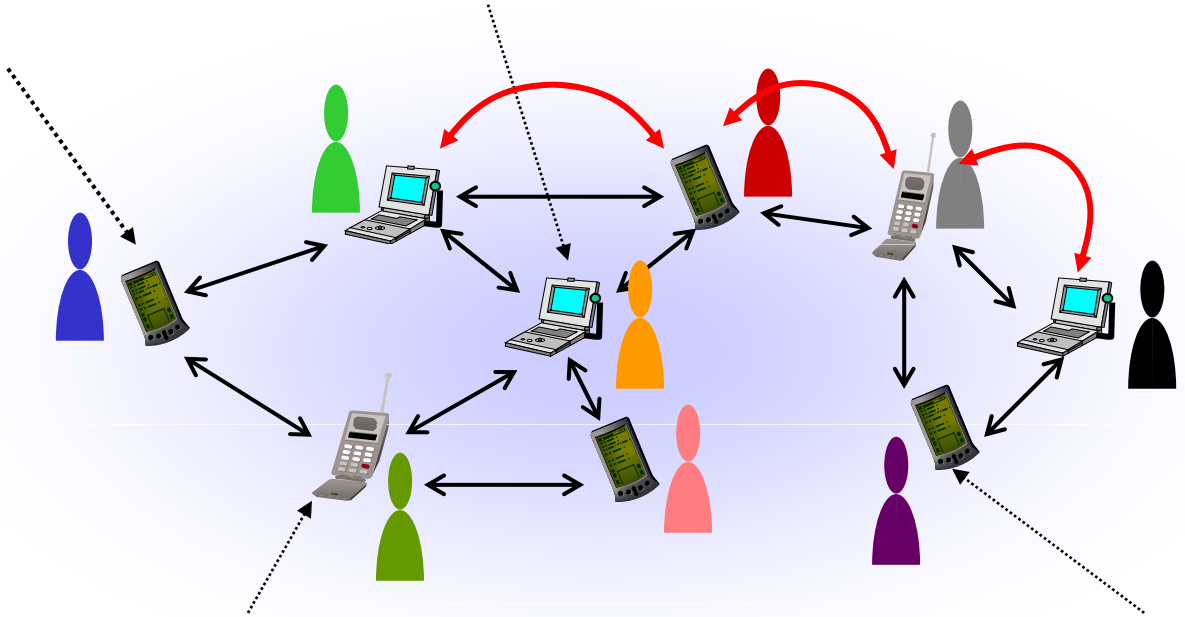

MANET: Special considerations

- Are all stations willing to forward other's packets?
- Do I trust all members in this Ad Hoc net?
- Power consumption is one metric
- Forwarding capacity is one metric

MANET routing

- Proactive
  - complete routing info at hand all the time
  - no special action before sending
  - lots of energy lost in keeping track of paths never used
- Reactive/On demand
  - find best path when connection needed
  - only used paths are exploited
  - delay before connection can be used