

Inter-Autonomous-System Routing: Border Gateway Protocol

Antonio Carzaniga

Faculty of Informatics
University of Lugano

December 10, 2014

Outline


- Hierarchical routing
- BGP

Routing

- Goal: each router u must be able to compute, for each other router v , the next-hop neighbor x that is on the least-cost path from u to v

Routing

- Goal: each router u must be able to compute, for each other router v , the next-hop neighbor x that is on the least-cost path from u to v




Network Model

Network Model

- So far we have studied routing over a “flat” network model


Network Model

- So far we have studied routing over a “flat” network model



Network Model


- So far we have studied routing over a “flat” network model



- Also, our objective has been to find the least-cost paths between sources and destinations


More Realistic Topologies

More Realistic Topologies




Even More Realistic

Even More Realistic



An Internet Map



©1999 Lucent Technologies

Higher-Level Objectives

Higher-Level Objectives

- *Scalability*

- ▶ hundreds of millions of hosts in today's Internet

Higher-Level Objectives

■ *Scalability*

- ▶ hundreds of millions of hosts in today's Internet
- ▶ transmitting routing information (e.g., LSAs) would be too expensive

Higher-Level Objectives

■ *Scalability*

- ▶ hundreds of millions of hosts in today's Internet
- ▶ transmitting routing information (e.g., LSAs) would be too expensive
- ▶ forwarding would also be too expensive

Higher-Level Objectives

■ *Scalability*

- ▶ hundreds of millions of hosts in today's Internet
- ▶ transmitting routing information (e.g., LSAs) would be too expensive
- ▶ forwarding would also be too expensive

■ *Administrative autonomy*

Higher-Level Objectives

■ *Scalability*

- ▶ hundreds of millions of hosts in today's Internet
- ▶ transmitting routing information (e.g., LSAs) would be too expensive
- ▶ forwarding would also be too expensive

■ *Administrative autonomy*

- ▶ one organization might want to run a distance-vector routing protocol, while another might want to run a link-state protocol

Higher-Level Objectives

■ *Scalability*

- ▶ hundreds of millions of hosts in today's Internet
- ▶ transmitting routing information (e.g., LSAs) would be too expensive
- ▶ forwarding would also be too expensive

■ *Administrative autonomy*

- ▶ one organization might want to run a distance-vector routing protocol, while another might want to run a link-state protocol
- ▶ an organization might not want to expose its internal network structure

Hierarchical Structure

- Today's Internet is organized in *autonomous systems (ASs)*
 - ▶ independent administrative domains


Hierarchical Structure

- Today's Internet is organized in *autonomous systems (ASs)*
 - ▶ independent administrative domains
- *Gateway routers* connect an autonomous system with other autonomous systems


Hierarchical Structure

- Today's Internet is organized in *autonomous systems (ASs)*
 - ▶ independent administrative domains
- *Gateway routers* connect an autonomous system with other autonomous systems
- An *intra-autonomous system routing protocol* runs within an autonomous system (e.g., OSPF)
 - ▶ this protocol determines internal routes
 - ▶ internal router ↔ internal router
 - ▶ internal router ↔ gateway router
 - ▶ gateway router ↔ gateway router


Hierarchical Structure



Hierarchical Structure



Hierarchical Structure




Inter-AS Routing

- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level

Inter-AS Routing


- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level



At AS3:
AS1 →

Inter-AS Routing


- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level



At AS3:
AS1 → AS1;

Inter-AS Routing

- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level




At AS3:

AS1 → AS1; AS2 →

Inter-AS Routing

- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level




At AS3:

AS1 → AS1; AS2 → AS2;

Inter-AS Routing

- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level




At AS3:

AS1 → AS1; AS2 → AS2; AS4 →

Inter-AS Routing

- An *inter-autonomous system routing protocol* determines routing at the autonomous-system level



At AS3:

AS1 → AS1; AS2 → AS2; AS4 → AS1.

Hierarchical Routing

- All routers within an AS compute their *intra-AS* routing information
 - ▶ using an *intra-domain* routing protocol

Hierarchical Routing

- All routers within an AS compute their *intra-AS* routing information
 - ▶ using an *intra-domain* routing protocol
- Gateway routers figure out *inter-AS* routing information
 - ▶ using an *inter-domain* routing protocol

Hierarchical Routing

- All routers within an AS compute their *intra-AS* routing information
 - ▶ using an *intra-domain* routing protocol
- Gateway routers figure out *inter-AS* routing information
 - ▶ using an *inter-domain* routing protocol
- *inter-AS* routing information is propagated within an AS
 - ▶ using an appropriate protocol

Hierarchical Routing

- All routers within an AS compute their *intra-AS* routing information
 - ▶ using an *intra-domain* routing protocol
- Gateway routers figure out *inter-AS* routing information
 - ▶ using an *inter-domain* routing protocol
- *inter-AS* routing information is propagated within an AS
 - ▶ using an appropriate protocol
- Both *inter-AS* and *intra-AS* routing information is used to compile the forwarding tables

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual
- What about a destination x outside the autonomous system?

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual
- What about a destination x outside the autonomous system?
 - ▶ *inter-AS* information is used to figure out that x is reachable through gateway G_x

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual
- What about a destination x outside the autonomous system?
 - ▶ *inter-AS* information is used to figure out that x is reachable through gateway G_x
 - ▶ *intra-AS* information is used to figure out how to reach G_x within the AS

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual
- What about a destination x outside the autonomous system?
 - ▶ *inter-AS* information is used to figure out that x is reachable through gateway G_x
 - ▶ *intra-AS* information is used to figure out how to reach G_x within the AS
 - ▶ what if x is reachable through multiple gateway routers G_x, G'_x, \dots ?

Hierarchical Routing (2)

- Destinations within the same autonomous system are reached as usual
- What about a destination x outside the autonomous system?
 - ▶ *inter-AS* information is used to figure out that x is reachable through gateway G_x
 - ▶ *intra-AS* information is used to figure out how to reach G_x within the AS
 - ▶ what if x is reachable through multiple gateway routers G_x, G'_x, \dots ?
 - ▶ use *intra-AS* routing information to determine the costs of the (least-cost) paths to G_x, G'_x, \dots
 - ▶ “hot-potato” routing: send it through the closest gateway

Benefits of Hierarchical Routing

- *Administrative autonomy*

Benefits of Hierarchical Routing

- *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use

Benefits of Hierarchical Routing

■ *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use
- ▶ an autonomous system needs to expose only minimal information about the internal structure of its network
 - ▶ essentially only (sub)net addresses

Benefits of Hierarchical Routing

■ *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use
- ▶ an autonomous system needs to expose only minimal information about the internal structure of its network
 - ▶ essentially only (sub)net addresses

■ *Scalability*

Benefits of Hierarchical Routing

■ *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use
- ▶ an autonomous system needs to expose only minimal information about the internal structure of its network
 - ▶ essentially only (sub)net addresses

■ *Scalability*

- ▶ routers within an autonomous system need to know very little about the internal structure of other autonomous systems

Benefits of Hierarchical Routing

■ *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use
- ▶ an autonomous system needs to expose only minimal information about the internal structure of its network
 - ▶ essentially only (sub)net addresses

■ *Scalability*

- ▶ routers within an autonomous system need to know very little about the internal structure of other autonomous systems
 - ▶ essentially only (sub)net addresses

Benefits of Hierarchical Routing

■ *Administrative autonomy*

- ▶ each autonomous system decides what intra-AS routing to use
- ▶ an autonomous system needs to expose only minimal information about the internal structure of its network
 - ▶ essentially only (sub)net addresses

■ *Scalability*

- ▶ routers within an autonomous system need to know very little about the internal structure of other autonomous systems
 - ▶ essentially only (sub)net addresses

■ External subnet addresses are likely to “aggregate” in groups that admit compact representations

- ▶ this process is called *supernetting*

Inter-AS Routing in the Internet

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs
 - ▶ transmits reachability information to all internal routers within an AS

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs
 - ▶ transmits reachability information to all internal routers within an AS
 - ▶ determines good routes to all outside subnets

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs
 - ▶ transmits reachability information to all internal routers within an AS
 - ▶ determines good routes to all outside subnets
 - ▶ based on reachability information

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs
 - ▶ transmits reachability information to all internal routers within an AS
 - ▶ determines good routes to all outside subnets
 - ▶ based on reachability information
 - ▶ based on *policies*

Inter-AS Routing in the Internet

- The *Border Gateway Protocol (BGP)* is the inter-AS routing protocol in today's Internet
 - ▶ provides reachability information from neighbor ASs
 - ▶ transmits reachability information to all internal routers within an AS
 - ▶ determines good routes to all outside subnets
 - ▶ based on reachability information
 - ▶ based on *policies*
 - ▶ BGP is a *path-vector* protocol

BGP Architecture and Terminology

- *BGP session*: a semi-permanent connection between two routers

BGP Architecture and Terminology

- *BGP session*: a semi-permanent connection between two routers
- *BGP peers*: two routers engaged in a BGP session
 - ▶ BGP sessions are established over TCP


BGP Architecture and Terminology

- *BGP session*: a semi-permanent connection between two routers
- *BGP peers*: two routers engaged in a BGP session
 - ▶ BGP sessions are established over TCP
- *BGP external session (eBGP)*: a session across two autonomous systems

BGP Architecture and Terminology

- *BGP session*: a semi-permanent connection between two routers
- *BGP peers*: two routers engaged in a BGP session
 - ▶ BGP sessions are established over TCP
- *BGP external session (eBGP)*: a session across two autonomous systems
- *BGP internal session (iBGP)*: a session within an autonomous system
 - ▶ note that internal sessions carry *inter-AS* information
 - ▶ *intra-AS* routing uses a separate protocol (e.g., OSPF)

Gateway Routers and *eBGP*



BGP Architecture and Terminology (2)

- *BGP advertisement*: a router advertises routes to networks, much like an entry in a distance-vector
 - ▶ destinations are denoted by address *prefixes*

BGP Architecture and Terminology (2)

- *BGP advertisement*: a router advertises routes to networks, much like an entry in a distance-vector
 - ▶ destinations are denoted by address *prefixes*
 - ▶ an AS may or may not forward an advertisement for a foreign network; doing so means being willing to carry traffic for that network

BGP Architecture and Terminology (2)

- *BGP advertisement*: a router advertises routes to networks, much like an entry in a distance-vector
 - ▶ destinations are denoted by address *prefixes*
 - ▶ an AS may or may not forward an advertisement for a foreign network; doing so means being willing to carry traffic for that network
 - ▶ this is where a router may aggregate prefixes (a.k.a., “supernetting”)
E.g.,

$$\left. \begin{array}{l} 128.138.242.0/24 \\ 128.138.243.0/24 \end{array} \right\} \rightarrow 128.138.242.0/23$$

BGP Architecture and Terminology (2)

- *BGP advertisement*: a router advertises routes to networks, much like an entry in a distance-vector
 - ▶ destinations are denoted by address *prefixes*
 - ▶ an AS may or may not forward an advertisement for a foreign network; doing so means being willing to carry traffic for that network
 - ▶ this is where a router may aggregate prefixes (a.k.a., “supernetting”)
E.g.,

$$\left. \begin{array}{l} 128.138.242.0/24 \\ 128.138.243.0/24 \end{array} \right\} \rightarrow 128.138.242.0/23$$

$$\left. \begin{array}{l} 191.224.128.0/22 \\ 191.224.136.0/21 \\ 191.224.132.0/22 \end{array} \right\} \rightarrow$$

BGP Architecture and Terminology (2)

- *BGP advertisement*: a router advertises routes to networks, much like an entry in a distance-vector
 - ▶ destinations are denoted by address *prefixes*
 - ▶ an AS may or may not forward an advertisement for a foreign network; doing so means being willing to carry traffic for that network
 - ▶ this is where a router may aggregate prefixes (a.k.a., “supernetting”)
E.g.,

$$\left. \begin{array}{l} 128.138.242.0/24 \\ 128.138.243.0/24 \end{array} \right\} \rightarrow 128.138.242.0/23$$

$$\left. \begin{array}{l} 191.224.128.0/22 \\ 191.224.136.0/21 \\ 191.224.132.0/22 \end{array} \right\} \rightarrow 191.224.128.0/20$$

BGP Architecture and Terminology (3)

- *Autonomous system number (ASN)*: a unique identifier for each AS (with more than one gateway)

BGP Architecture and Terminology (3)

- *Autonomous system number (ASN)*: a unique identifier for each AS (with more than one gateway)
- *BGP attributes*: a route advertisement includes a number of attributes
 - ▶ *AS-PATH*: sequence of ASNs through which the advertisement has been sent

BGP Architecture and Terminology (3)

- *Autonomous system number (ASN)*: a unique identifier for each AS (with more than one gateway)
- *BGP attributes*: a route advertisement includes a number of attributes
 - ▶ *AS-PATH*: sequence of ASNs through which the advertisement has been sent
 - ▶ *NEXT-HOP*: specifies the interface (IP address) to use to forward packets towards the advertised destination
 - ▶ used to resolve ambiguous cases where an AS can be reached through multiple gateways (interfaces)

BGP Architecture and Terminology (3)

- *Autonomous system number (ASN)*: a unique identifier for each AS (with more than one gateway)
- *BGP attributes*: a route advertisement includes a number of attributes
 - ▶ *AS-PATH*: sequence of ASNs through which the advertisement has been sent
 - ▶ *NEXT-HOP*: specifies the interface (IP address) to use to forward packets towards the advertised destination
 - ▶ used to resolve ambiguous cases where an AS can be reached through multiple gateways (interfaces)
- *BGP import policy*: used to decide whether to accept or reject the route advertisement
 - ▶ e.g., a router may not want to send its traffic through one of the AS listed in *AS-PATH*

Route Selection

1. Router preference: routes are ranked according to a *preference* value
 - ▶ configured at the router
 - ▶ or learned from another router within the same AS
 - ▶ essentially a configuration parameter for the AS

Route Selection

1. Router preference: routes are ranked according to a *preference* value
 - ▶ configured at the router
 - ▶ or learned from another router within the same AS
 - ▶ essentially a configuration parameter for the AS
2. Shortest AS-PATH

Route Selection

1. Router preference: routes are ranked according to a *preference* value
 - ▶ configured at the router
 - ▶ or learned from another router within the same AS
 - ▶ essentially a configuration parameter for the AS
2. Shortest AS-PATH
3. Closest NEXT-HOP router

Route Selection

1. Router preference: routes are ranked according to a *preference* value
 - ▶ configured at the router
 - ▶ or learned from another router within the same AS
 - ▶ essentially a configuration parameter for the AS
2. Shortest AS-PATH
3. Closest NEXT-HOP router
4. ...