La Sapienza
Università degli Studi di Roma
Dipartimento di Informatica e Sistemistica

Computer Networks II
Exercise collection 5 – RIP / OSPF / BGP

Ugo Colesanti
colesanti@dis.uniroma1.it
A.A. 2008/2009

RIP

- **Distance vector routing protocol**
  - **Widespread because simple**
  - **Metric: number of hops**
  - **Implements distributed algorithm (Bellman – Ford)_**
  - **Neighbors: routers having interfaces on the same subnet**
- **Not suitable for large networks**
  - **Slow convergence**
  - **Max. 15 hops**
  - **Possible instability due to persistence of old routes (refresh 30s)_**

Exercise 1

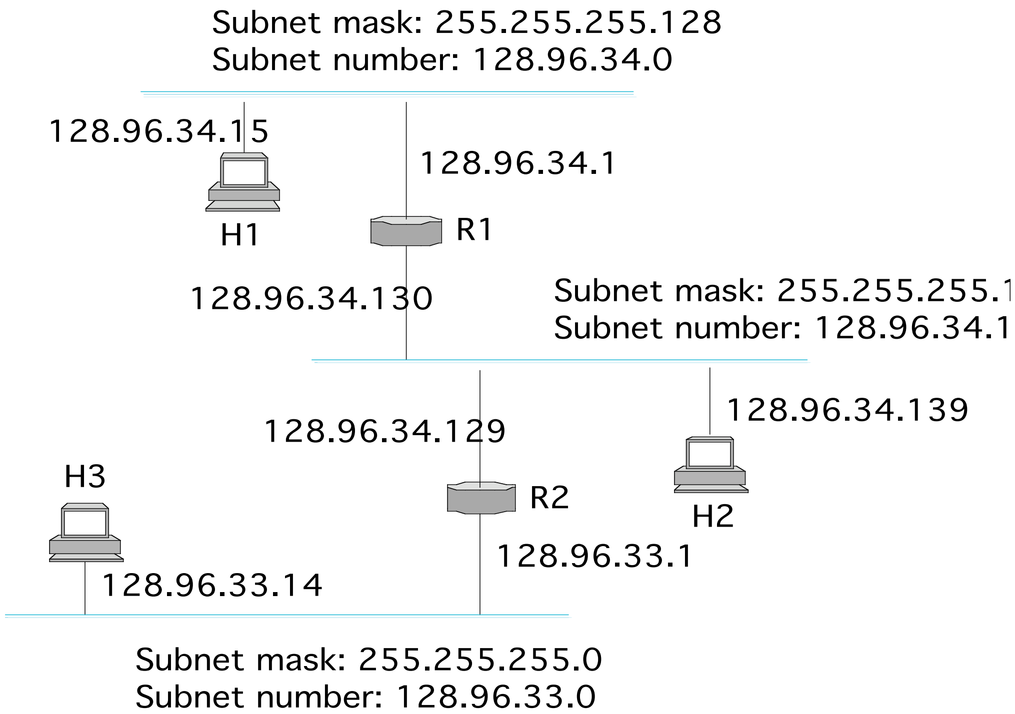

- D1: assume subnetting is used. Write routing table at R2 for network in picture, assuming following entry format:
  - Dest | NextHop | Subnet Mask

Exercise 1 - solution

- **Routing table R2:**

DestNet	NextHop	Subnet
128.96.34.128	Direct	255.255.255.12
128.96.33.0	Direct	255.255,255,0
128.96.34.0	128.96.94.130	255.255.255.12

Exercise 2

- **Given the following topology (see next slide)_**
  - **Link width proportional to link capacities (e.g., 1 and 2 for the 2 kinds of links)_**
  - **Routers use RIP**
- **D1: write routing table for router R1, according to following scheme:**
  - **DestNet | NextHop | NetMask | Interface | Cost**
  - **(note to denote interfaces, use link labels)_**
- **D2: give new routing table, assuming HUB 1 is off.**

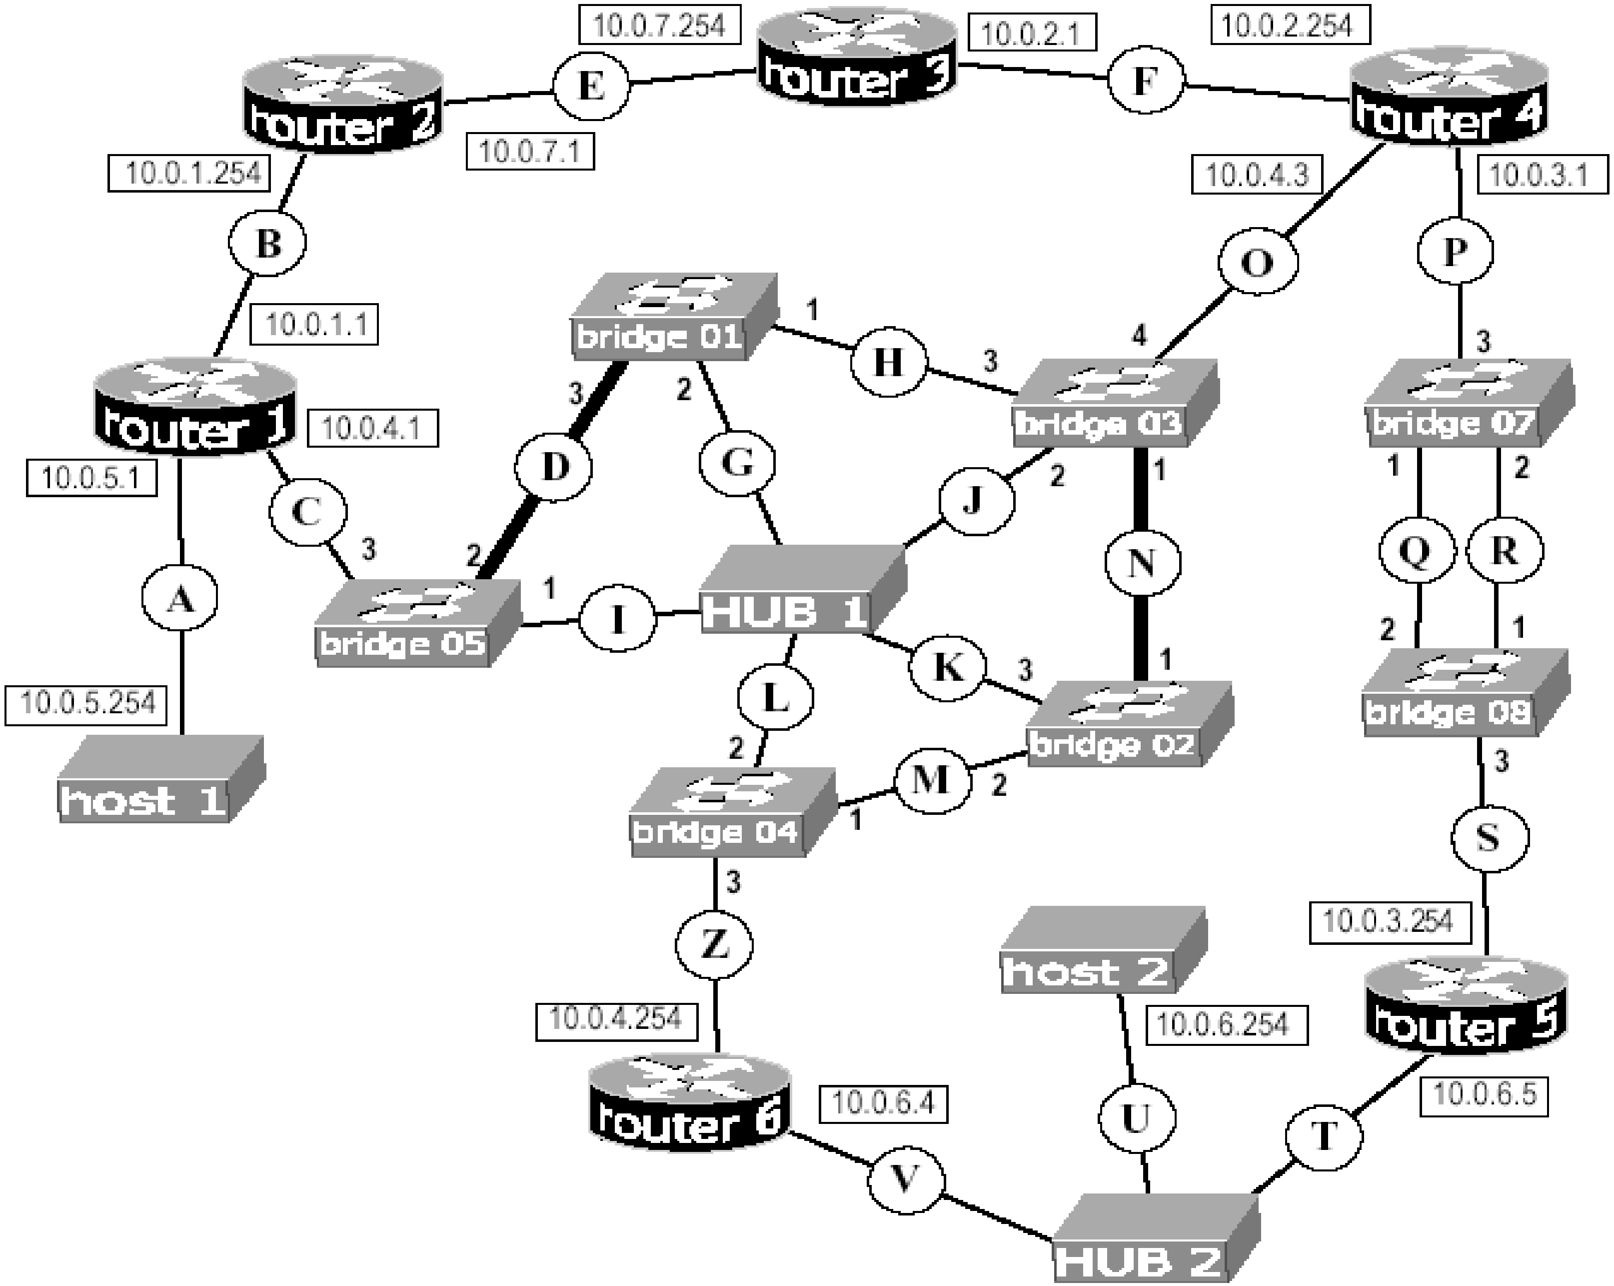
router 1
router 2
router 3
router 4
router 5
router 6
bridge 01
bridge 02
bridge 03
bridge 04
bridge 05
bridge 07
bridge 08
HUB 1
HUB 2
host 1
host 2
10.0.7.254
10.0.7.1
10.0.2.1
10.0.2.254
10.0.1.254
10.0.1.1
10.0.4.1
10.0.5.1
10.0.5.254
10.0.4.3
10.0.3.1
10.0.3.254
10.0.4.254
10.0.6.4
10.0.6.254
10.0.6.5
A
B
C
D
E
F
G
H
I
J
K
L
M
N
O
P
Q
R
S
T
U
V
Z

Exercise 2 - solution

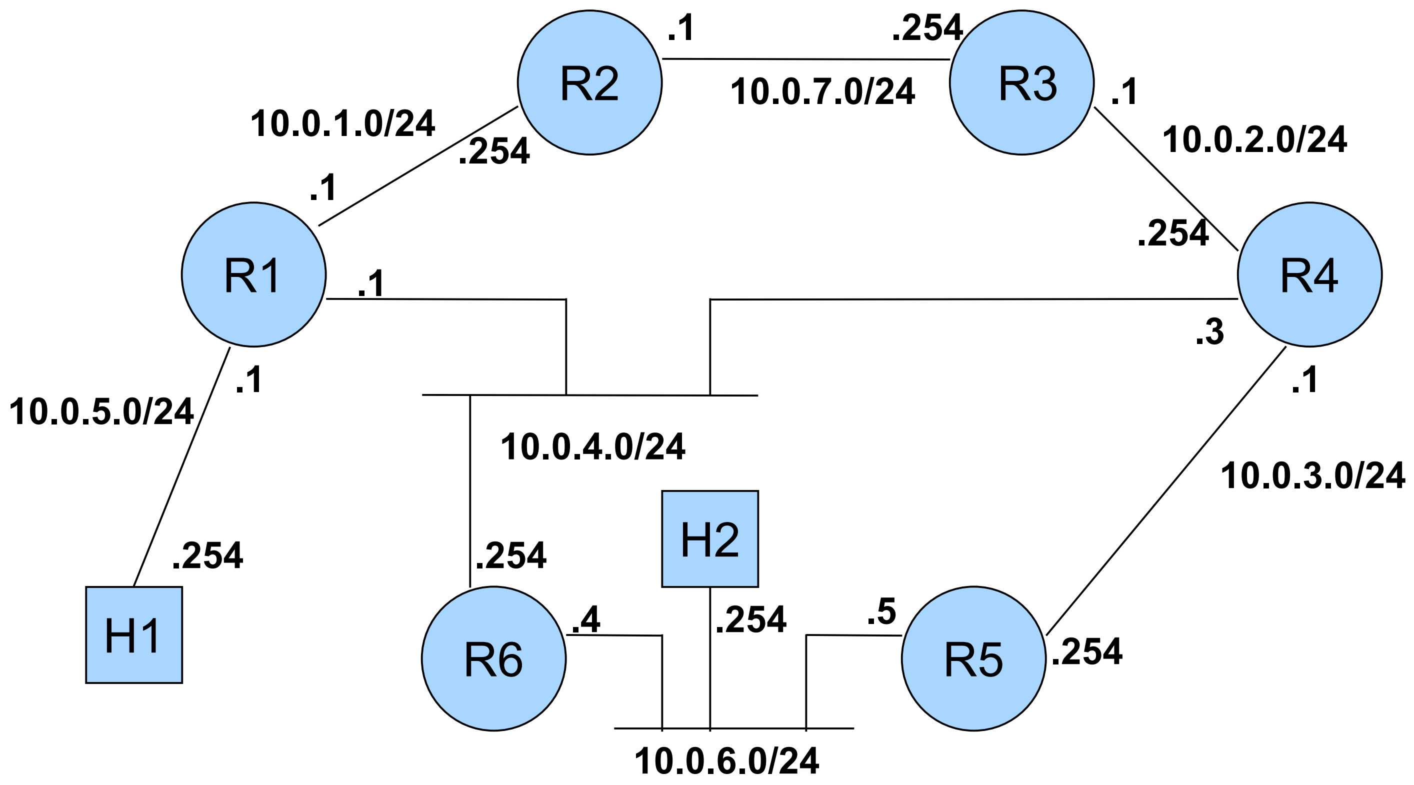

Exercise 2 - solution/cont.

DestNet	NextHop	Subnet	Interface	HopCount
10.0.5.0	direct	255.255.255.0	A	0
10.0.1.0	direct	255.255.255.0	B	0
10.0.4.0	direct	255.255.255.0	C	0
10.0.7.0	10.0.1.254	255.255.255.0	B	1
10.0.2.0	10.0.4.3	255.255.255.0	C	1
10.0.6.0	10.0.4.254	255.255.255.0	C	1
10.0.3.0	10.0.4.3	255.255.255.0	C	1

CIDR (1/3)

- **Dual of subnet addressing:**
  - **Subnet addressing: different network prefixes for different physical subnets**
  - **CIDR addressing: a block of contiguous C class addresses assigned to the same organization (e.g., instead of a class B address, 256 class C ones).**
- **Hierarchical organization:**
  - **ISP give network connectivity to organization -> ISP have large address spaces**
  - **ISPs redistribute blocks of their address spaces to organization, under a contract**
- **Problem: if many addresses to the same organization -> routing tables may become large**

CIDR (2/3)

- **Class Inter-Domain Routing: compresses blocks of contiguous addresses into one pair <network address, subnet mask>**
- **In practice:**
  - **Network address: lowest address in block**
- **Compact notation: address/x , x=1..32**
  - **e.g. 128.134.168.0/21**
- **Aggregation:**
  - **Two /20 contiguous addresses with common 19 bit prefix can be aggregated into a /19 address**
  - **Etc.**

CIDR (3/3)

- **Example:**

  ISP manages block 128.211.0.0/16

  Assigns to organization 2048 CIDR address block starting at 128.211.168.0:

  CIDR block: 128.211.168.0/21 -> address space: from 128.211.168.0 to 128.211.175.255

Exercise 3

- **D1: given routing table below, show how it can be compressed using CIDR**

Networks address	NETMASK	INTERFACCIA
194.100.0.0	255.255.255.0	int 1
194.100.1.0	255.255.255.0	int 1
194.100.2.0	255.255.254.0	int 1
194.100.4.0	255.255.252.0	int 1
194.100.8.0	255.255.248.0	int 1
194.100.48.0	255.255.240.0	int 1
194.100.64.0	255.255.240.0	int 1

Exercise 3 - solution

- **194.100.0.0/24 -> from 194.100.0.0 to 194.100.0.255 (256 addresses)_**

  194.100.1.0/24 -> from 194.100.1.0 to 194.100.1.255 (256 addresses)_

  194.100.2.0/23 -> from 194.100.2.0 to 194.100.3.255 (512 addresses)_

  194.100.4.0/22 -> from 194.100.4.0 to 194.100.7.255 (1024 addresses)_

  194.100.8.0/21 -> from 194.100.8.0 to 194.100.15.255 (2048 addresses)_

  _ all contiguous with same configuration of highest 20 bits, 4096 addresses _ 194.100.0.0/20 in CIDR

- **194.100.48.0/20 -> from 194.100.48.0 to 194.100.63.255 (4096 addresses)_**

  194.100.64.0/20 -> from 194.100.64.0 to 194.100.79.255 (4096 addresses)_

  _ contiguous, but 48 -> 00110000 while 64 -> 01000000, no common 19 bits prefix, no aggregation possible

Exercise 3 - solution

- **Compressed routing table:**

DestNet	Interface
194.100.0.0/20	int1
194.100.48.0/20	int1
194.100.64.0/20	int1

Exercise 4

- **D1: given routing table below, show how it can be compressed using CIDR**

network	netmask	interface
200.0.0.0	255.255.192.0	A
200.0.64.0	255.255.192.0	A
200.0.128.0	255.255.128.0	A
200.1.0.0	255.255.0.0	A
193.0.2.0	255.255.255.0	B
193.0.3.0	255.255.255.0	B
193.0.4.0	255.255.255.0	B
193.0.5.0	255.255.255.0	B

Exercise 4 - solution

- **200.0.0.0/18 -> from 200.0.0.0 to 200.0.63.255 (16384 addresses)_**
  200.0.64.0/18 -> from 200.0.64.0 to 200.127.255 (16384 adrresses)_
  200.0.128.0/17 -> from 200.0.128.0 to 200.0.255.255 (32768 addresses)_
  200.1.0.0/16 -> from 200.1.0.0 to 200.1.255.255 (65536 addresses)_
  _ all contiguous with same 15 bit prefix, 131072 addresses _ 200.0.0.0/15
- **193.0.2.0/24 -> from 193.0.2.0 to 193.0.2.255 (256 addresses)_**
  193.0.3.0/24 -> da 193.0.3.0 a 193.0.3.255 (256 addresses)_
  193.0.4.0/24 -> da 193.0.4.0 a 193.0.4.255 (256 addresses)_
  193.0.5.0/24 -> da 193.0.5.0 a 193.0.5.255 (256 addresses)_
  -> 193.0.2.0 and 193.0.3.0 contiguous with 23 bit common prefix -> 193.0.2.0/23
  -> 193.0.4.0 and 193.0.5.0 contiguous with 23 bit common prefix -> 193.0.4.0/23

Exercise 4 - solution

- **Compressed routing table:**

DestNet	Interface
200.0.0.0/15	A
193.0.2.0/23	B
193.0.4.0/23	B

Exercise 5

- **ISP manages address space 194.48.0.0/16.**
- **3 organizations A,B e C respectively need 2048, 8192 and 4096 addresses and respective NETMASKs.**
- **Organizations A and C reachable over interface 1, while B is reachable over interface 2**
- **D1: optimaze address allocation using CIDR, so as to compress local routing table. Describe tesulting routing table entries**

Exercise 5 - solution

- **A and C on same interface, B on a different one -> at least 2 entries in RT**
- **A -> 2048 addresses -> /21**

  C -> 4096 addresses -> /20

  _ no aggregation possible: no common 19 bit prefixes
- **A -> 194.48.0.0, subnet mask 255.255.248.0 , interface 1**

  C -> 194.48.16.0 subnet mask 255.255. 240.0 , interface 1

  (2048 address block between A and C è left for possible future aggregation)_
- **Why not 194.48.8.0/19 to aggregate A and C?**
- **B -> 8192 indirizzi -> /19**

  B -> 194.48.32.0, subnet 255.255.224.0 , interface 2

OSPF (1/2)_

- **Link-state protocol**
  - **Suitable for large networks**
  - **Hello protocol to build adjacency relationships**
  - **Use of Link State Packest -> list of networks and metric costs**
  - **All routers in an area acquire same routing info in LSP (DB)_**
  - **LSP propagated via flooding**
- **Areas**
  - **Leduce traffic control (full flooding only within an area -> summary info outside)_**
  - **Backbone Area: interconnects different areas**
  - **ABR: connects backbone to different areas**
  - **IR: a router having all interfaces internal to a same area**
  - **Route summarization: info on LS in one area collected by ABRs and injected into backbone to other ABRs**

OSPF (2/2)

- **ASBR: interconnects two different ASes**
- **Route**
  - **Inter-area: a route belonging to a same AS and spanning different areas**
  - **External: route spanning different ASes**
- **Summary routes:**
  - **Limit amount of control traffic**
  - **Summary info collected from an ABR and injected into backbone**
  - **Also: info collected fron an ASBR and forwarded to a different AS**

Exercise 6

- **Assume institution X is assigned IP address block 131.108.64/20.**
- **These addresses must be assigned among departments A, B and C of the organization. A needs at most 2000 addresses, B and C both at most 1000**
- **D1: propose a suitable network architecture to meet these requirements, assuming OSPF is the interior gateway protocol of choice. Design a network organization into areas and propose a network address partition for the proposed network.**

Exercise 6 - solution

- **Assigned block: 4096 addresses -> OK**
- **3 areas (+ area 0), one for each department.**
- **Address block assignment:**
  - **A: 2048 addresses -> /21 -> 131.108.64.0/21 , area1**
  - **B: 1024 addresses -> /22 -> 131.108.72.0/22 , area2**
  - **C: 1024 addresses -> /22 -> 131.108.76.0/22 , area3**
- **Architecture:**

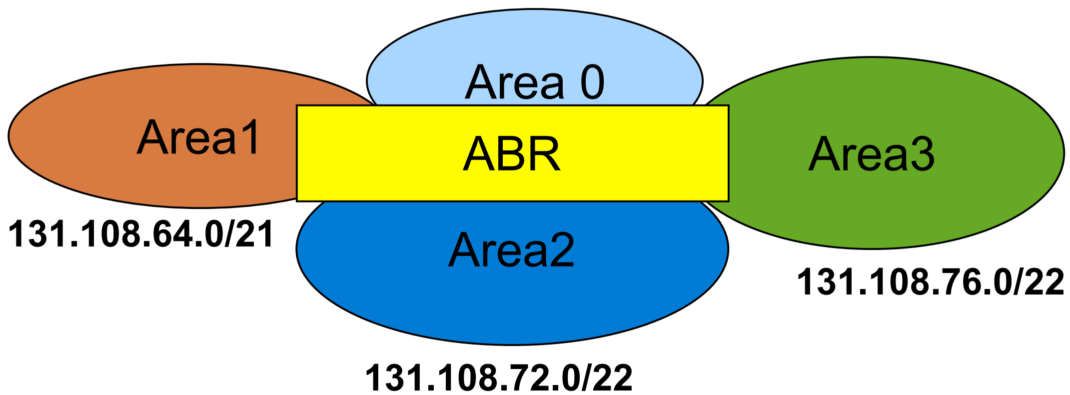

Exercise 7

- **A network administrator is assigned a class B address block. She decides to design an OSPF architecture having 3 areas (including area 0) . To this purpose, the administrator decides to use a single ABR**
- **The network administrator decides to equally partition available addresses among the 2 areas, assigning even addresses (those terminating with an even decimal in dotted decimal notation) to area 1 and odd ones to area 2**
- **D1: determine the order of magnitude of the number of entries in the ABR's routing table**
- **D2: give an approximation of the amount of memory necessary to store the routing table under this assignment if every RT entry is 80 bytes**
- **D3: propose (if possible) a more efficient address assignment**

Exercise 7 - solution

- **Architecture:**

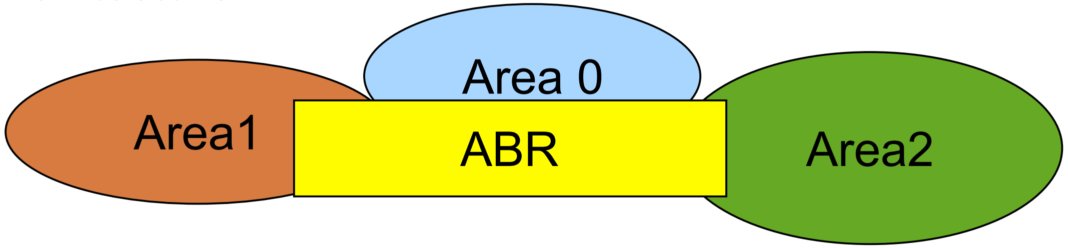

ABR has 2 interfaces, one towards area 1 and the other towards area 2
Even IP addresses: end with 0
Odd IP numbers end with a 1
_ no aggregation possible _ 2X2^15 = 2^16 entries -> 2^16X80 ~ 5MB

Autonomous systems - BGP

- **Ogni organizzazione è composta da un insieme di router e LAN sotto una singola amministrazione**
- **Un AS definisce in maniere coerente le politiche di instradamento all’interno della sua oganizzazione**
- **In una inter-rete comprendente più AS necessita di punti di collegamento**
- **Instradamento inter-domain -> EGP -> BGP**
- **Router che parlano BGP: gateway o router di confine**
  - **Nasconde parte interna all'AS**
  - **Mantiene zone di demarcazione e router di frontiera**
  - **Ogni router di frontiera mostra le reti interne come se fossero locali**
- **Path vector routing:**
  - **quali reti possono essere raggiunte attraverso un router**
  - **quali AS sono attraversati lungo il cammino**

Exercise 8

- **Given the figure below:**

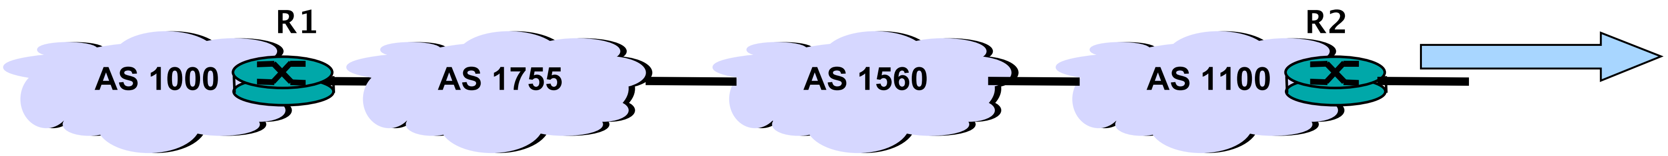

- **Assume that AS 1000 aggregates address block 124.100.0.0/16.**

- **D1: describe synthetically the path announced by BGP speaker R2 in the direction indicated by the arrow for this address block**

- **D2: assume that speaker R1's interface towards AS 1755 has IP address 124.102.12.1. Explain which BGP update messages have this IP address as NEXT HOP attribute value**

Exercise 8 - solution

- **The path announced by R2 is:**
  - **(124.100/16; 1100, 1560, 1755, 1000)_**
- **124.102.12.1 is the NEXT_HOP attribute value in the update message sent from R1 to its BGP peer belonging to AS 1755. This is also the value of this attribute for update message possibly exchanged via IBGP within AS 1755 for this path**

Exercise 9

– **Given the picture below:**

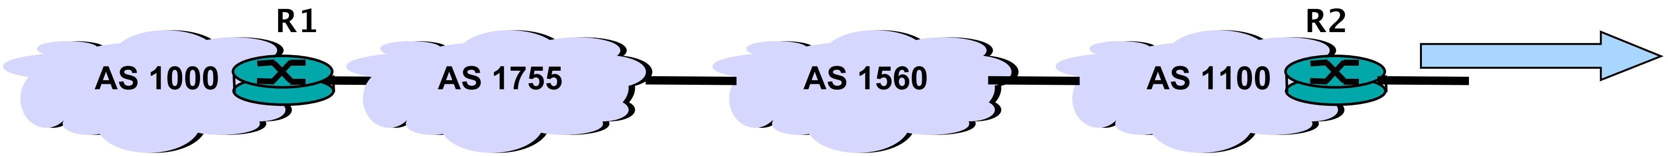

- **assume the system administrator of AS 1100 intends to prevent traffic originating from AS 1560 within her domain**
- **D1: which basic BGP update policy can achieve this result?**
- **D2: how is it possible to prevent transit traffic towards AS 1560 from traversing AS 1100?**

Exercise 9 - solution

- **D1: No BGP routes are announced to AS 1560**
- **D2: no BGP updates are sent towards networks that are right of AS 1100 for routes that traverse AS 1560**

Exercise 10

- **Given the following picture:**

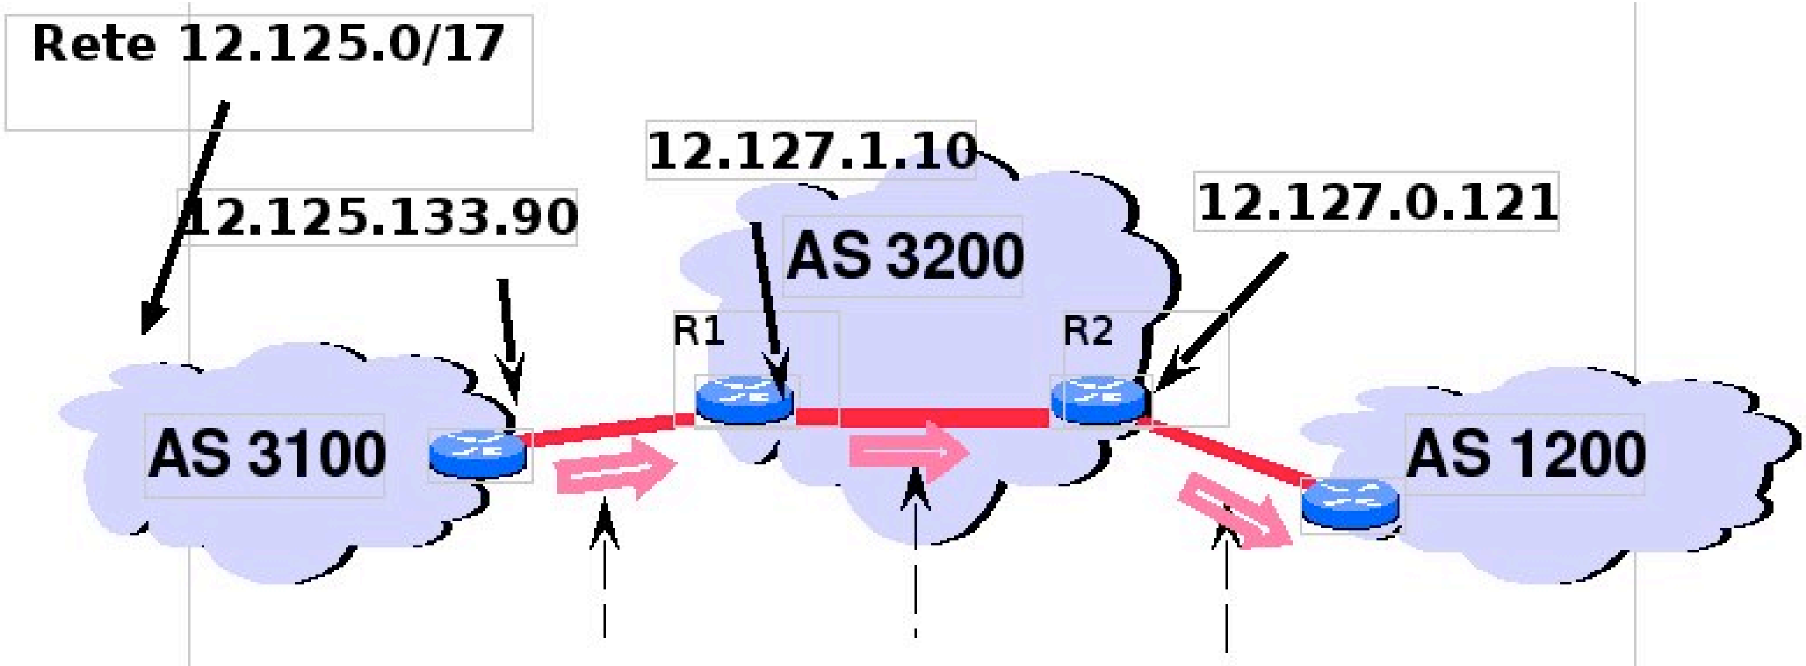

D1: Give the values of the NEXT_HOP attribute for the BGP update messages indicated in the picture

Esercizio 10

- **D2: briefly discuss the constraints on the contents of the routing tables of routers that are internal to the ASes in the picture, deriving from the use and meaning of the NEXT_HOP attribute in BGP.**

Exercise 10 - solution

- **D1: the values of the attribute for the messages in the picture are 12.125.133.9, 12.125.133.9 and 12.127.0.12 respectively.**
- **Within the As, it is necessary to maintain explicit routes towards border BGP speakers**
  - **->routing tables of routers internal to AS 3200 must be such as to allow routing an IP packet from right BGP speaker to left BGP speaker and viceversa, in order to reach destination**
  - **For example, consider the simple case in which R1 and R2 are directly connected. R2's internal routing table must have 1 line for destination 12.125.0/17 (or a destination including this block) with Next Hop value (IGP protocol for the AS) equal to 12.127.1.10.**

Exercise 11

- **Consider the AS configuration in the picture below:**

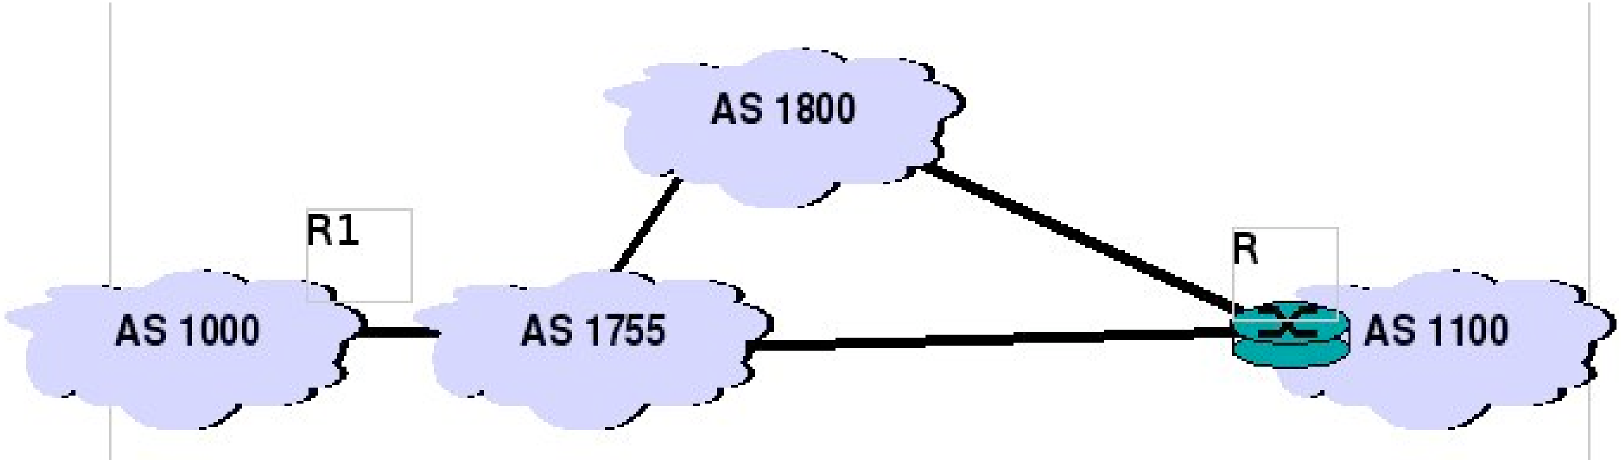

Assume AS 1000 aggregates address block 124.100.0.0/16. Is it possible for R to receive more BGP updates (i.e., with different attribute values) for this block?

- **D1: which announcement is selected by AS 1100 when all the rest is the same?**

Exercise 11 - solution

- **If all other attributes are the same, the announcement with shortest path vector (number of ASes crossed by route) is preferred. In our case, two routes are announced (124.100/16, 1800, 1755, 1000) and (124.100/16, 1755, 1000), hence the latter is preferred when all other attributes are the same**

Exercise 12

- **Consider the following picture:**

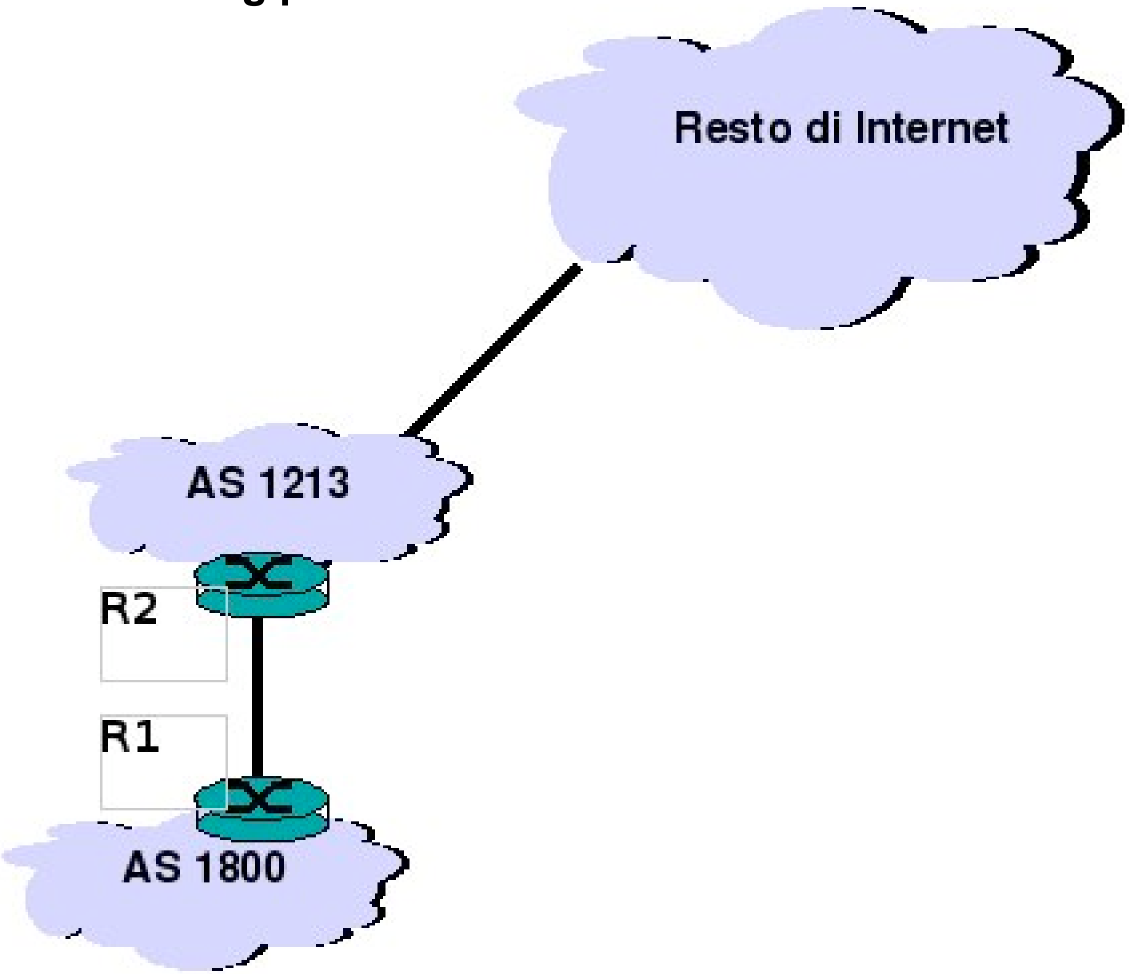

Esercizio 12

- **Assume AS 1800 is a stub of AS 1213 (its provider). AS 1213 is responsible for address block 112.0.0.0/8. Assume AS 1800 receives address block 112.64.0.0/16.**
- **D1: propose an organization for the network of AS 1800, defining its internal organization and its connection to the provider, so as to meet the following requirements: i) the internal network is organized into 3 subnets A, B and C, each approximately receiving 32000, 16000 and 16000 addresses respectively; also propose a solution for the following point:**
- **D2: how is AS 1800 connected to its provider?**

Exercise 12 - solution

- **Large networks organized into subnets -> reasonable to use OSPF for AS 1800.**
- **R1 is an ABR interconnecting 3 areas and at the same time serving as ASBR for external routing**
- **AS 1800 organization:**

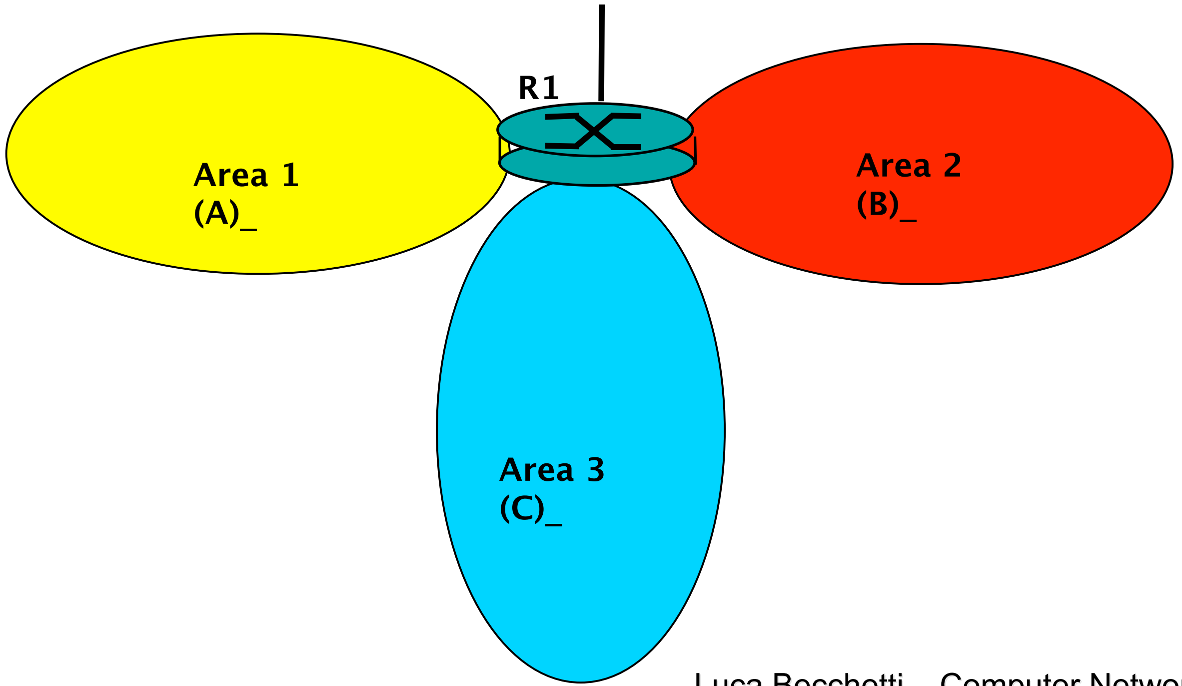

Exercise 12 - solution

- **Area 1 -> /17 -> 112.64.0.0/17 ,**

  Areas 2 and 3 -> /18 -> 112.64.128.0/18 and 112.64.192.0/18
- **Since AS 1800 is stub with a unique router connected to its provider, we can use a static routing, without resorting to BGP**