Redes de Computadores (RCOMP)

Theoretical-Practical (TP) Lesson 05

2017/2018

- IPv4 routeing and static routeing tables.
- IPv4 dynamic routeing.

Routeing

Routeing is the activity of analysing a packet's **destination address** and, in conformity, decide **where to send it**, thus every node executes this task. Nevertheless, we usually only call router to a node that performs this tasks for packets that have been sent by other nodes and not by the node itself.

To make decisions about to where packets should be sent, some knowledge about other nodes location is required. On a layer two Ethernet switch, that information is stored in the MAC table where each node location is present.

By convention, the expression **routeing** is reserved for layer three, whereas in layer two, the equivalent activity is called **switching**.

Unlike switching in layer two that handles node addresses, routeing in layer three uses network addresses, this reduces dramatically the amount of information to be handled. Instead of specifying every node address on a network, specifying the network address is enough. This information is stored in the routeing table, a routeing table can be manually or automatically built.

We call end-nodes, nodes that only send packets of their own and don't retransmit other nodes' packets (packet forwarding is disabled). Still, an end-node is also required to have a routeing table.

IPv4 routeing decisions

Every IPv4 node is required to have a routeing table, it's based on the routeing table that the node takes routeing decisions.

A routeing table is made of a column with a sequence of **known networks**, each specified by a network address and a network prefix-length (network mask).

The routeing activity consists on confronting a packet's IPv4 destination address with each of these **known networks**.

This confrontation aims at determining if the packet's **destination address** belongs to that **known network**. Routers do this by:

1st – Performing a bit-by-bit logical **and** operation between the packet destination IP address and the network mask.

2nd – Check if the result is the network address. If so, this means the destination address belongs to the network and the remaining routeing table is ignored. Otherwise, proceed to the confrontation with the next known network in the routeing table (next line).

In essence, the router first assumes the packet destination address may belong to the network and performs the test to check if that is so.

The routeing table

When a node wants to send a packet, and the packet's destination address belong to a known network, the routeing table also has instructions about how to transmit that packet, that's the role of the routeing table's second column, usually called the **next-hop**. There are two options for how to transmit a packet: direct transmission (local interface) or using a neighbour router (local IP address).

Direct transmission means the network to be reached is directly connected to the node, thus no third parties are required to intervene.

Using a neighbour router means the network to be reached is not directly connected to the node, thus the packet must be sent to a router. That router will, in turn, confront the packet destination address with its own routeing table and act accordingly, but that's another story and not a concern for this node anymore.

The details about how the packet is transmitted in each of these two cases differ. On **direct transmission**, ARP is used to obtain the MAC address matching the IPv4 destination address, then a layer two frame is created and the MAC address is used to fill the frame destination address, the IPv4 packet is placed on the payload and the frame is directly sent to the destination node.

Using a neighbour router

In contrast, if transmitting a packet requires the use of a neighbour router (next-hop), then ARP is used to obtain the **MAC address of the neighbour router**. A layer two frame is created and the router MAC address is used to fill the frame destination address, the IPv4 packet is placed on the payload and the frame is sent to the router.

This is why the next-hop must always be an address of a directly connected network, otherwise, ARP cannot work.

Here is a sample of a routeing table:

We can see, this node is aware of five networks, two of which are directly connected to it. In addition, three other remote networks are available by using neighbour routers.

Network	Next-hop
192.160.0.0/16	Direct – Interface 1
192.150.0.0/16	Direct – Interface 2
192.140.0.0/24	192.160.0.1
195.10.9.0/27	192.160.0.1
15.20.50.0/24	192.150.0.1

As stated before, neighbour routers must belong to a directly connected network. This routeing table tells nothing about how far the remote networks are, they could be connected to the mentioned neighbour router or on the other side of the world.

The default route

Imagine a routeing table for an IPv4 node connected to the internet, it can't hold all networks existing on the internet, however, **if a packet doesn't belong to any known network in the routeing table, it won't get transmitted**.

The issue can be solved by adding a default route to the table, it's represented by network 0.0.0.0 with prefix length zero. This means, any destination IPv4 address belongs to this network (if the network mask is zero, then the bit-by-bit and logical operation with any address will always result in 0.0.0.0).

Of course, the default route always ought to be the last row in the routeing table, if fact, subsequent lines will never be reached.

The next-hop entry for the default route is usually called the default-router or default-gateway, if there's an internet connection this should be the neighbour router that is suited to send packets toward the internet.

If other networks have the same next-hop as the default-route, they can normally be removed from the routeing table. If those lines are removed, then packets previously matching them will fall to the last line of the routeing table (default-route), and thus, will end up being sent to the same next-hop.

End-nodes' routeing table

Every IPv4 node has a routeing table, in the simplest case, it's derived from its own network configuration. The most basic scenario is an end-node with a single network interface, the configuration data is: the local IPv4 address, the network mask, and also the default gateway. The routeing table is based on this data. Example:

Interface IPv4 address: 193.100.10.50

Interface network mask: 255.255.255.0

Default gateway: 193.100.10.1

Then, the routeing tables is:

Network	Next-hop
193.100.10.0/24	Direct – Interface
0.0.0.0/0	193.100.10.1

Usually, routers have extensive routeing tables, although for better performance they should be kept as small as possible. Often, routeing tables can be simplified, this means, reducing the number of lines, but keeping the same behaviour. One already seen simplification case encompasses the aggregation to the default-route.

Aggregation to the default route

Simplifying routeing tables is important, this consists of combining several lines of a routeing table into a single line, and yet, keeping the same behaviour. There are several rules to obey, the first one is:

Two lines can only be combined in one if they have the same next-hop

Only after checking this condition, further analysis should be carried.

The easiest simplification is aggregation to the default-route, in principle, any line with the same next-hop of the default-route can be removed. There are exceptions to this rule, here is one example:

The second line has the same next-hop of the default-route, however, it can't be removed because that would change the behaviour.

Network	Next-hop
193.100.10.0/24	Direct – Interface
120.5.10.0/24	193.100.10.1
120.5.0.0/16	193.100.10.5
0.0.0.0/0	193.100.10.1

This happens because the third line's network includes the second line's network, as it is, packets with destination address belonging to network 120.5.10.0/24 are sent to 193.100.10.1. If the second line was removed they would match the third line and would be sent to 193.100.10.5. **Order matters.**

Instituto Superior de Engenharia do Porto – Departamento de Engenharia Informática – Redes de Computadores (RCOMP) – André Moreira

8

Networks aggregation

One other chance to simplify routeing tables is by combining two networks into a single (double sized) equivalent network. This should be tried only after aggregation to the default route. Once more, the next-hop must be the same.

Though, another criterion must be met: **both networks to be aggregated must have the same prefix length** (network mask).

Aggregating two networks is achieved by reducing both networks prefix-length in one bit. As result, one bit that was previously a network bit will become a host bit. This may or not lead to a change in each network address, if both result in the same address, then they can be aggregated. Examples:

170.10.1.0/24 and 170.10.0.0/24: If we reduce one bit to the prefix-length, the first network will become 170.10.0.0/23, while the second will keep the same address 170.10.0.0/23, we can therefore say 170.10.0.0/23 = 170.10.1.0/24 + 170.10.0.0/24.
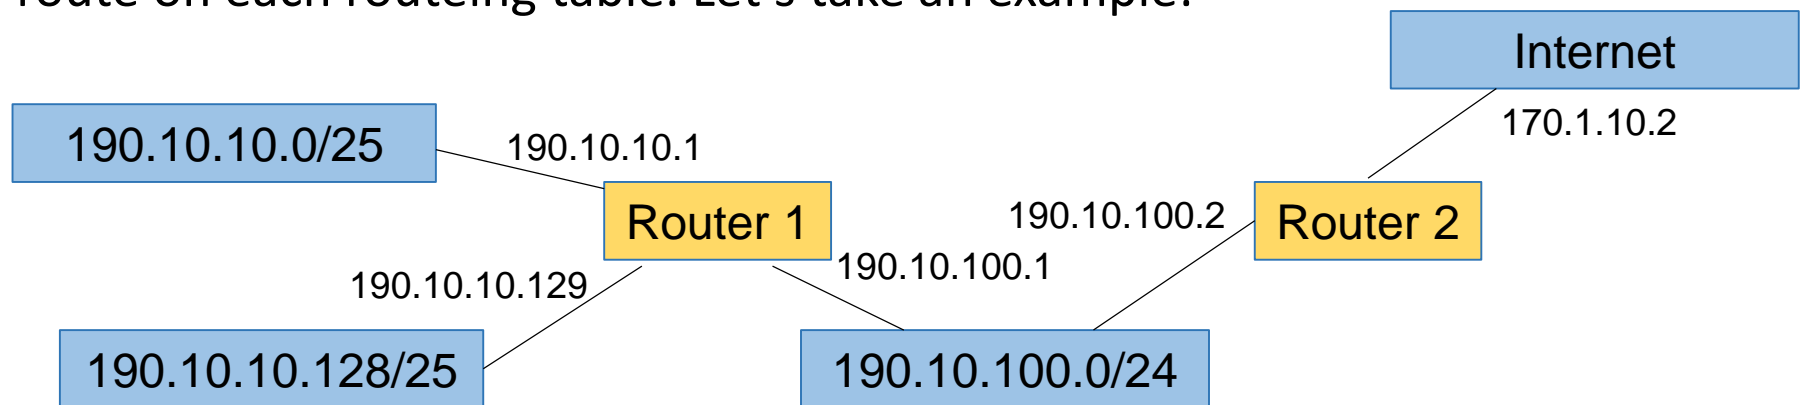
170.10.1.0/24 and 170.10.2.0/24: If we reduce one bit to the prefix-length, the first network becomes 170.10.0.0/23, while the second keeps the same address 170.10.2.0/23, unlike the previous example, **they cannot be aggregated**.

Building a routeing table

Routeing tables can be manually built on each router by the network administrator (static routeing) or they can be automatically built using routeing protocols (dynamic routeing). To build a routeing table we must know the routers relative positions and networks that interconnect them.

Each router will have its own routeing table, to build it, we must focus on that router's position and, from its point of view, **list all remote networks** (not directly connected to the router). Directly connected networks are already known to the router and there is no point in adding them to the routeing table.

If the infrastructure has a internet connection we must also add a default route on each routeing table. Let's take an example:

Instituto Superior de Engenharia do Porto – Departamento de Engenharia Informática – Redes de Computadores (RCOMP) – André Moreira

10

Building a routeing table

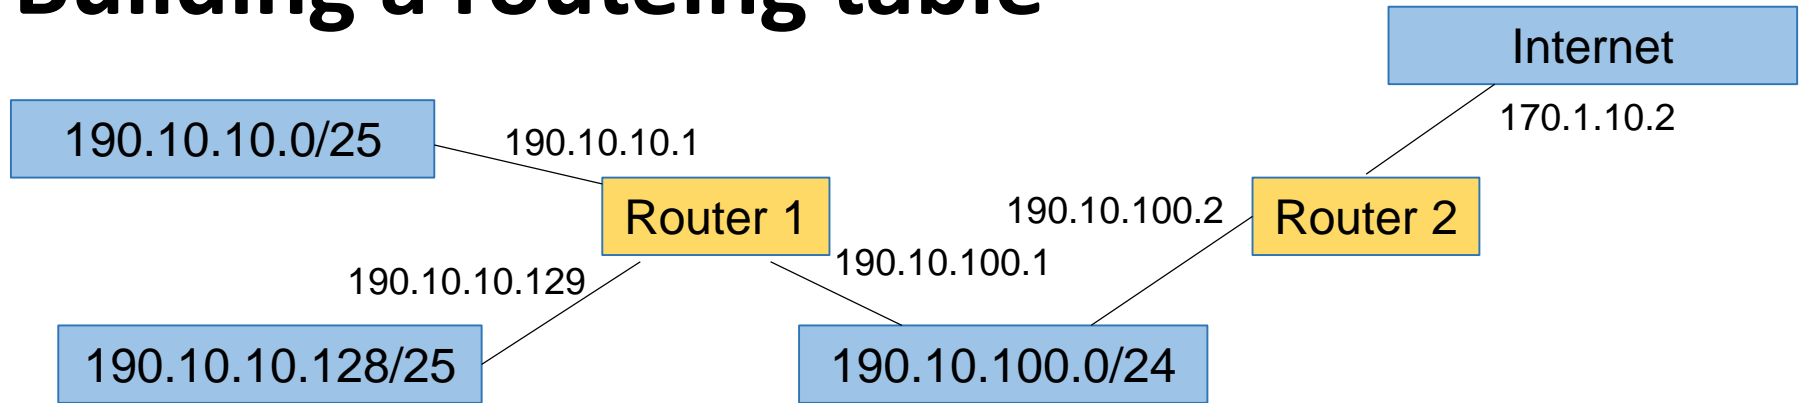| 190.10.10.0/25 | | |
Router 1	190.10.100.2	Router 2
190.10.10.1	190.10.100.1	170.1.10.2
190.10.10.129		Internet
190.10.10.128/25	190.10.100.0/24	

In this layout we have two routers and three local networks, router 2 has an internet connection, near each router connection the router's IPv4 address is shown in the diagram.

Routeing table for router 1 is simple, it knows all local networks, so there is no need to add any of then to its routeing table, so the routeing table for Router 1 required only the default route:

Network	Next-hop
0.0.0.0/0	190.10.100.2

The meaning of this table is: any packet with a destination address not belonging to a directly connected network (190.10.10.0/25, 190.10.10.128/25 or 190.10.100.0/24) must be sent to 193.10.100.2 (Router 2).

Instituto Superior de Engenharia do Porto – Departamento de Engenharia Informática – Redes de Computadores (RCOMP) – André Moreira

11

Building a routeing table

190.10.10.0/25 — 190.10.10.1 — **Router 1**

Internet — 170.1.10.2

190.10.100.2 — **Router 2**

190.10.10.129 — 190.10.100.1

190.10.10.128/25 — 190.10.100.0/24

Router 2 is not aware networks 190.10.10.0/25 and 190.10.10.128/25 do exist (because it's not directly connected to them). So we must add then to the routeing table, also saying that to reach them it should send packets to 190.10.100.1 (Router 1).

Router 2 is connected to network 190.10.100.0/24, so there is no need to add it. The default route for router 2 will send packets to the internet router (170.1.10.2).

Network	Next-hop
190.10.10.0/25	190.10.100.1
190.10.10.128/25	190.10.100.1
0.0.0.0/0	170.1.10.2

Building a routeing table

We can try a simplification for **Router 2** routeing table.

Aggregation to the default route is not possible because the next-hop is different. However, networks 190.10.10.0/25 and 190.10.10.128/25 have the same mask and addresses are very similar, they may become the same if we reduce one bit to the network mask. In fact, representing the network bits in red:

Network	Next-hop
190.10.10.0/25	190.10.100.1
190.10.10.128/25	190.10.100.1
0.0.0.0/0	170.1.10.2

190.10.10.0/25   -> **190.10.10.(0**000 0000$)_2$

190.10.10.128/25 -> **190.10.10.(1**000 0000$)_2$

This means if we reduce one bit to the mask, both address will become the same: **190.10.10.0/24**. The simplified version of the routeing table is therefore:

Network	Next-hop
190.10.10.0/24	190.10.100.1
0.0.0.0/0	170.1.10.2

Dynamic routeing

An IPv4 networking infrastructure may be redundant, this means there are several paths to reach the same destination network. This can be clearly seen in the following routeing table:

This table shows us, on this router, there are three alternative next-hops to reach network 155.0.0.0/24. The question is: what's the best choice?

Network	Next-hop
155.0.0.0/24	190.10.100.1
155.0.0.0/24	170.1.10.2
155.0.0.0/24	190.10.100.5

To be able to answer that, an additional column must be added representing the cost of each alternative, the lowest cost alternative should then be elected.

In this new version we can now see the best option is 190.10.100.5. However, networks are dynamic, later in the day router 190.10.100.5 may crash, and then what?

Network	Next-hop	Cost/Metric
155.0.0.0/24	190.10.100.1	4
155.0.0.0/24	170.1.10.2	6
155.0.0.0/24	190.10.100.5	1

With static routeing, the network administrator would have to rush and go to every router removing all routes passing through the crashed router. This is where dynamic routeing becomes a must.

Dynamic routeing

Dynamic routeing uses a routeing protocol. The routeing protocol establishes a dialogue between routers, leading them to cooperate and automatically build the routeing table in all of them. Furthermore important than building the tables once, they are permanently updated to reflect the network status.

Dynamic routeing enforces an effective load balancing when there are several alternatives. Each router always selects the lowest cost alternative it can find on the routeing table, but cost values in routeing tables are permanently updated by the routeing protocol. For instance, if one router is heavily loaded with traffic, the routeing protocol will increase the cost of using that alternative on all routers tables.

Also, if one router is simply unavailable, the routeing protocol removes that alternative from all routeing tables.

There are similarities with redundant layer two infrastructures and STP failover, however, unlike with STP, all alternatives are permanently active and in use. With dynamic routeing the traffic is distributed by all available alternatives with load balancing and not just failover.

Cost/metric

Each router only handles with neighbour routers (connected to shared networks), these neighbour routers represent all options available to retransmit packets. This retransmission represents only a single step (hop) in the path of a packet to reach the destination network. Along the path, all routers must make the correct decision on transmitting the packet to the next router, otherwise, the packet will not reach the destination network.

If there are alternative paths to reach a destination network we are interested in using the best available, however, routers don't handle with paths, they handle with next-hops (neighbour routers).

The cost/metric is inserted into the routeing table by the routeing protocol and represents **the path cost** of using each alternative next-hop. Thus, the only thing the router is required to do is select the lowest cost/metric alternative in the routeing table. Of course, this only applies when more than one line has the network address to which the packet belongs.

In fact, when a router has alternative paths to reach a network, it will automatically remove those with higher cost, leaving only the one with lower cost. So, when analysing a packet's destination address the best alternative has already been selected.

The path cost/metric

Different routeing protocols use different metrics, in the first place, there is a **hop or link metric**, this represents the cost of sending to a neighbour router and its calculation can take in account values related to the link used, like transmission rate, error rate, MTU and values related to the next hop like its performance, latency or current traffic load.

We are though interested in the **path metric**, and this is what routeing protocols calculate and insert in routeing tables.

The **path metric** represents the cost of, at a given router, using one particular next-hop, that ultimately represents following one path to reach the destination network.

Routeing protocols calculate the path metric by taking into account all hop or link metrics along the path, the straightest way to do this is simply calculating the sum of all hop or link metrics along the path, but other methods can be used.

RIP (routeing Information Protocol)

RIP is a very simple distance-vector routeing protocol. Distance-vector means knowledge about existing networks is received from neighbour routers, incorporated in the local routeing tables and retransmitted to other neighbour routers. In this process, each router updates the path cost by adding to it the cost of the interface from which the information was received.

In RIP the cost of each interface is typically 1, greater integer values can be used, however, when a path cost reaches 16, it will be accounted as unreachable. If all interfaces have cost 1, then the path metric will represent the number of hops (routers) to reach the destination network.

As an example, lets consider the following network infrastructure:
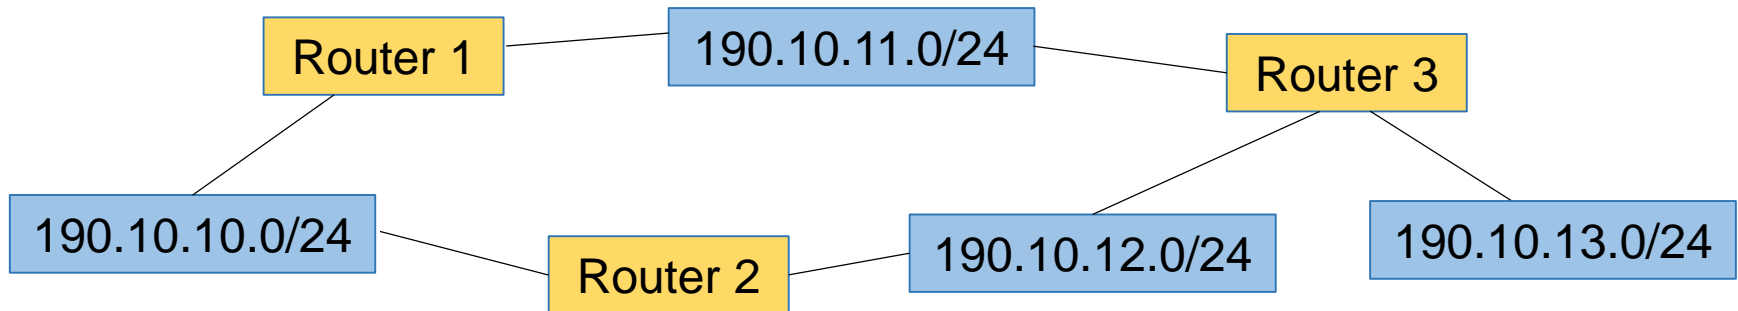
RIP example – initial knowledge

As usual in RIP, we assume every interface has cost one. In the next image we can also see the rightmost octet of each router address in each network.
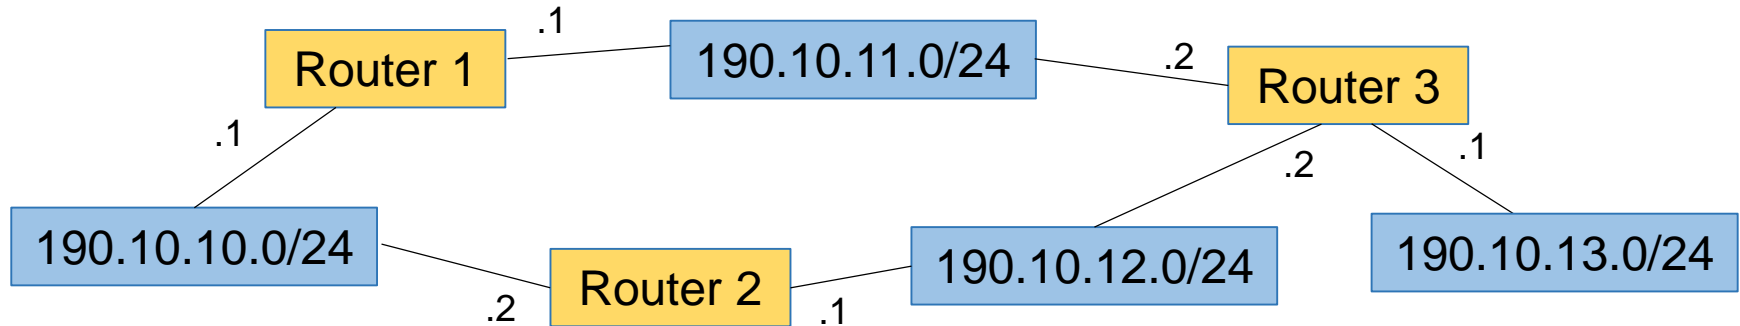
Before RIP starts to broadcast information about known networks, each router already knows the networks it's directly connected to:

Router 1		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.11.0/24	Direct	1

Router 2		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.12.0/24	Direct	1

Router 3		
Network	Next-hop	Cost
190.10.13.0/24	Direct	1
190.10.11.0/24	Direct	1
190.10.12.0/24	Direct	1

Instituto Superior de Engenharia do Porto – Departamento de Engenharia Informática – Redes de Computadores (RCOMP) – André Moreira

19

RIP example – receiving RIP tables

Now RIP starts broadcasting, each router broadcasts a **table with the known networks and corresponding metric**.

Router 1		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.11.0/24	Direct	1
190.10.10.0/24	190.10.10.2	2
190.10.12.0/24	190.10.10.2	2
190.10.11.0/24	190.10.11.2	2
190.10.12.0/24	190.10.11.2	2
190.10.13.0/24	190.10.11.2	2
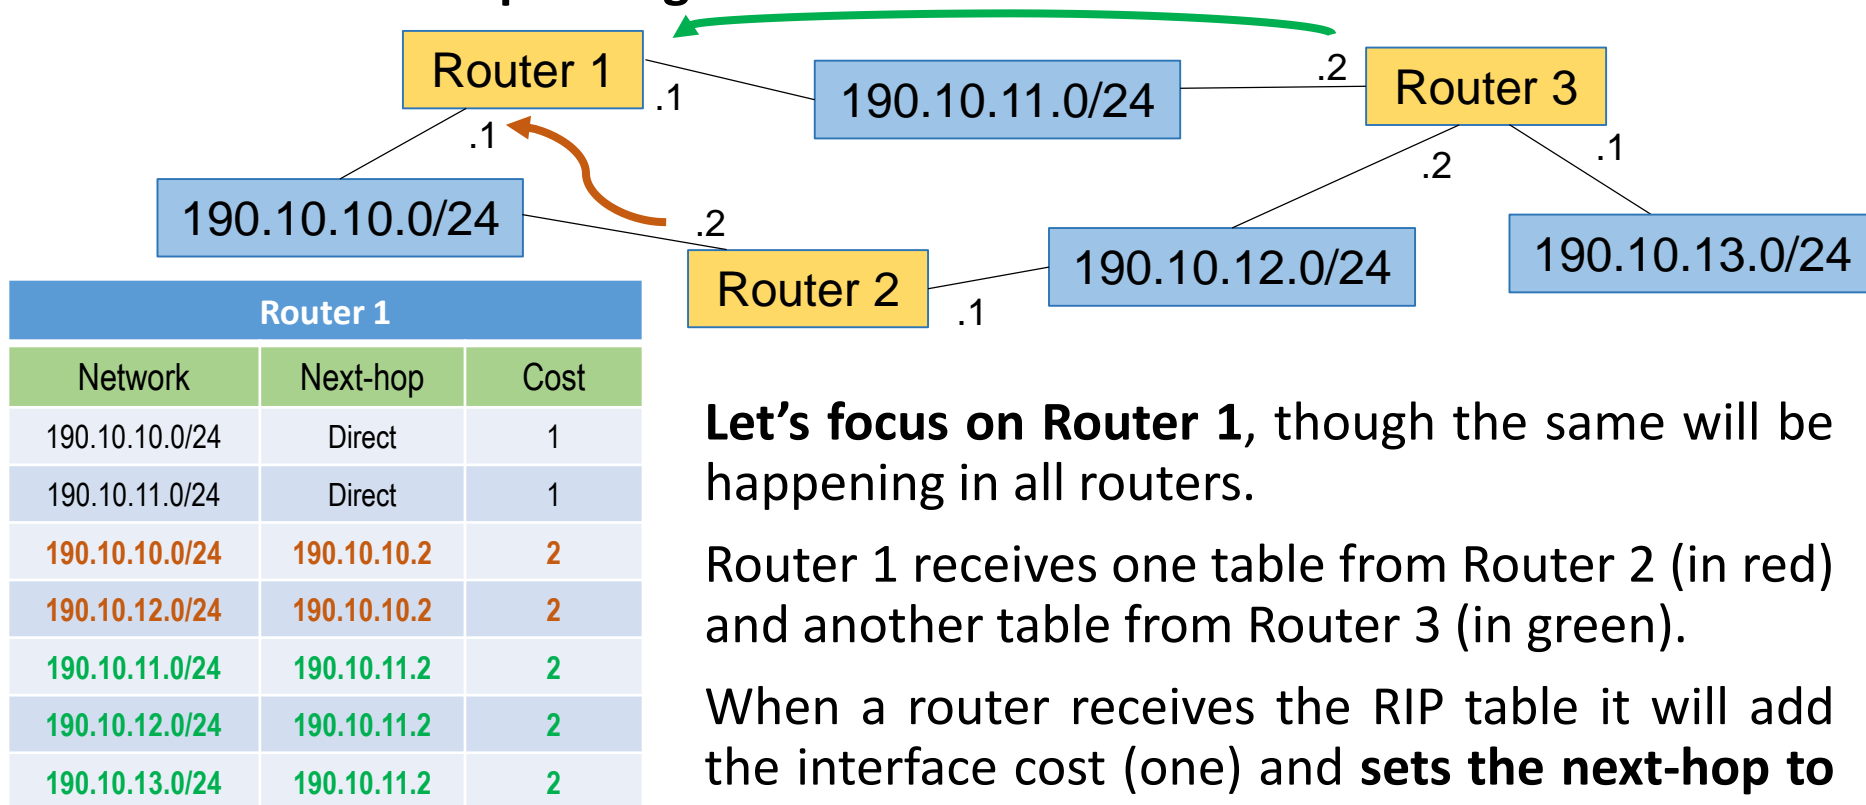
Let's focus on Router 1, though the same will be happening in all routers.

Router 1 receives one table from Router 2 (in red) and another table from Router 3 (in green).

When a router receives the RIP table it will add the interface cost (one) and **sets the next-hop to the source address from where the table came from**.

The resulting routeing table in Router 1 is shown on the left.

RIP example – tables cleanup

Next step is tables cleanup, every 120 seconds each router will remove repeated networks with higher metric.

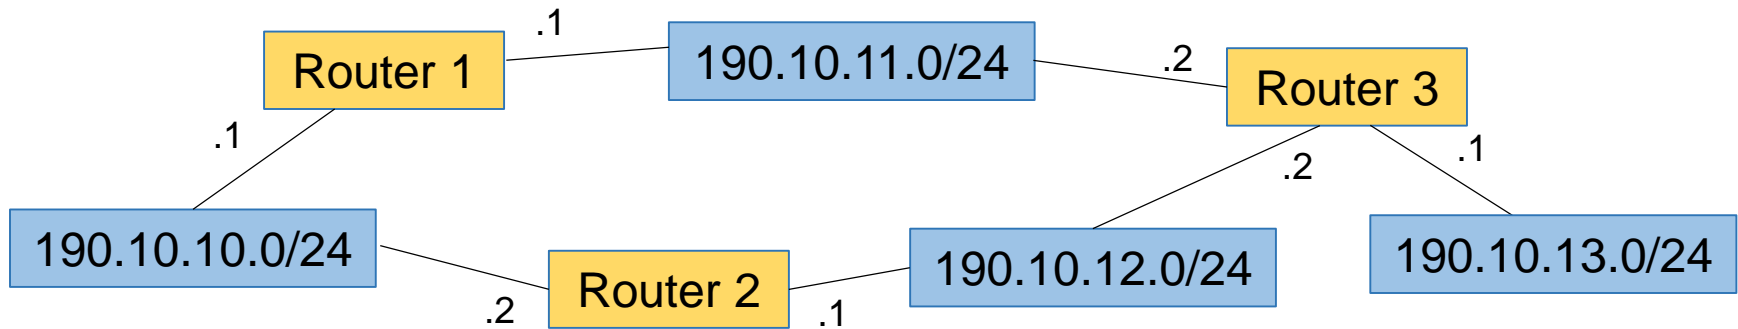| Router 1		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.11.0/24	Direct	1
~~190.10.10.0/24~~	~~190.10.10.2~~	~~2~~
190.10.12.0/24	190.10.10.2	2
~~190.10.11.0/24~~	~~190.10.11.2~~	~~2~~
190.10.12.0/24	190.10.11.2	2
190.10.13.0/24	190.10.11.2	2

Keeping the focus on Router 1, two lines are eliminated because for each there's an alternative with a lower path cost. The remaining lines will be kept for now because alternatives have the same cost.

If a line is not refreshed within 180 seconds is removed. Information is sent between routers every 30 seconds, ultimately routeing tables converge and no further changes occur when the network is stable.

RIP convergence time is rather long, it takes 3 minutes to detect a connection is down. Then, for each router the change is propagated through, up to 30 additional seconds must be added.

RIP example – convergence

If network connections are stable, routeing tables will in due course converge. Routeing tables will then look like this:

Router 1		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.11.0/24	Direct	1
190.10.12.0/24	190.10.10.2	2
190.10.12.0/24	190.10.11.2	2
190.10.13.0/24	190.10.11.2	2

Router 2		
Network	Next-hop	Cost
190.10.10.0/24	Direct	1
190.10.12.0/24	Direct	1
190.10.11.0/24	190.10.10.1	2
190.10.11.0/24	190.10.12.2	2
190.10.13.0/24	190.10.12.2	2

Router 3		
Network	Next-hop	Cost
190.10.13.0/24	Direct	1
190.10.11.0/24	Direct	1
190.10.12.0/24	Direct	1
190.10.10.0/24	190.10.12.1	2
190.10.10.0/24	190.10.11.1	2

Instituto Superior de Engenharia do Porto – Departamento de Engenharia Informática – Redes de Computadores (RCOMP) – André Moreira

22