

Material for the Networking lab in EITF25 Internet – Techniques and Applications

2013

Preparations

In order to succeed with the lab, you must have understood some important parts of the course. Therefore, before you come to the lab you should have studied the following topics and protocols:

- Local area networks and Ethernet
- Switches and routers
- IPv4 addressing and sub-netting
- ARP and ICMP
- Forwarding and routing

Before you can start the lab, you should answer the following 20 questions:

1. What is the purpose of an ARP request?
2. What layer is ARP operating on?
3. Does ARP use IP headers?
4. When a packet needs to be transmitted to a host on another subnet, what IP address will be used in the destination field of the IP header?
5. When a packet needs to be transmitted to a host on another subnet, what MAC address will be used in the destination field of the MAC header?
6. What MAC destination address is an ARP request ALWAYS sent to?
7. What MAC destination address is an ARP reply sent to?
8. Exactly what information does an ARP request contain?
9. Exactly what information does an ARP reply contain?
10. What is the purpose of a DNS request?
11. When a DNS request is sent to a DNS server on another subnet, what IP destination address will be used in the IP header?
12. When a DNS reply is sent by a DNS server to host on another subnet, what IP destination address will be used in the IP header?
13. When a DNS reply is sent by a DNS server to a host on another subnet, what MAC destination address will be used in the MAC header?
14. What is the difference between a hub, a switch and a router?
15. When a switch relays a packet, how will the destination MAC address be affected?
16. When a router relays a packet, how will the destination MAC address be affected?
17. When a switch relays a packet, how will the destination IP address be affected?
18. When a router relays a packet, how will the destination IP address be affected?
19. Describe the classful addressing method.
20. Describe the classless addressing method.

Networking Lab

This laboratory lesson will give you an idea of how ARP and routing works in an Ethernet based local network. We will also look into network utilisation and DNS.

Apart from network applications and tools included in Ubuntu 10.04 LTS, you will use a sniffer program called WireShark together with an utilisation measuring tool included in the same program, called IO Graphs.

You need to be familiar with the Lab Network Outline, since this is going to be referred to throughout the laboratory session.

🔔 This sign means that you should discuss the question with your lab teacher before you continue.

ARP – Address Resolution Protocol

The addresses you most often talk about in connection with networks for data communication, and especially Internet, are numerical IP addresses. When you for example write the server address with letters (domain name) in your web browser, that address is translated to a numerical address by the Domain Name System (covered in the DNS section). It is quite possible to state the numerical address directly instead of the alphabetic server address.

However, computers that are connected to the same local network must know each other's physical address to be able to communicate with each other. The physical address is, among other things, called MAC address, because this address is defined in the Media Access Control Layer of the OSI model. Another name is Ethernet address, but of course this name can only be relevant if the local network is actually an Ethernet. Yet another name is Hardware address because the address is burnt into the hardware of your computer and in some instances it is even called the Burnt-in address.

There must be a way of correlating the global IP address to the local MAC address. This process is called address mapping and the protocol that specifies how the mapping shall be performed is called ARP, Address Resolution Protocol.

Once a computer has found a MAC address to map to an IP address the computer caches it. The ARP cache holds entries only for a limited time; otherwise the cache would eventually overflow. Each cache entry holds a timer, and when this timer has timed out the entry is deleted from the cache. Another cause for limiting the lifetime of a cached entry is that a computer can change its MAC address. A concrete example of this is when you change network interface card.

Why are previous mappings cached in the first place?

Now we are ready to take a closer look at ARP.

Let us start with examining the IP address and MAC address of your own computer. You need to open a terminal window, by choosing `Applications > Accessories > Terminal` or holding `Ctrl+Alt+T`.

In the terminal window enter the command `ifconfig -a`.

What is your computer's IP address?

What is your computer's MAC address?

In the terminal you can also examine the contents of your computers ARP cache.

Which address pairs are cached (if any)? Hint: Use man arp in terminal. (man shows description of command ARP)

If no addresses were found in the ARP cache, you can force an ARP request by trying to connect to a computer **on the same subnet your computer is connected to**. Check with the network layout document to find a suitable destination. Use the ping program from the terminal: `ping <ip address>`.

Now, do you see any cached address pairs in the ARP cache? Which?

We shall now examine the ARP request and ARP reply in more detail by eavesdropping on the network. We will use WireShark as our sniffer. To make sure that the connection attempt actually triggers an ARP request we first have to clear the ARP cache.

Which command do you use? Hint: Use man arp to find out.

Now start WireShark and start a frame capture. Consult appendix A for information about how to use WireShark.

Try once more to connect to the same computer by using the ping program. Stop the frame capture when you have captured some packets, by hitting the stop button. You can see how many packets that have been captured in the capture window.

You will now see a lot of captured frames. Since all frames on the network are captured, you might see frames that are emanating from your neighbours' connections as well as your own. You can filter out all unwanted frames by using the `eth.addr==<your mac address>` command in the filter field of the capture window.

Find the first frame with the ARP request sent from your computer. Look in the Source column and Protocol column.

What is the destination MAC address of the ARP request? What does this address represent? Which host or hosts is the ARP request aimed for?

The Ethernet frame's length/type field differs for frames containing ARP datagrams and frames containing ICMP datagrams! Check the difference and describe it with corresponding values.

Hint: What layer is ARP operating on? What layer is ICMP operating on?

Now find the corresponding ARP reply.

What is the source MAC address of the ARP reply?

What is the destination MAC address of the ARP reply?

Describe the data contents of the frame! What information is found in the payload?

Use your findings and describe (why not a simple sketch?) how the address resolution process works! What is the function of ARP? Why is ARP needed? Why is the ARP request broadcasted? Why is the ARP reply not broadcasted?

Routing

The purpose of this section is to get an idea of how path selection on the Internet works. The lab network is actually a small Internet. Three separate local networks are connected with routers. In appendix B you will find a drawing of the lab network.

During the lab exercise we will use the ping program extensively, but when we examine packets on the network we will look for ICMP echo packets.

What is the relationship between the ping program and ICMP echo packets?

Segments of an IP-based network divided by routers are often referred to as subnets. Each subnet has its unique network id.

Consider an organization with 4000 hosts. Divide the hosts into two subnets containing 1000 and 3000 hosts. Choose network IDs and define the subnet masks so that the organizations requirements are met.

Not only routers perform routing or path selection. A computer, or host as it is called according to Internet vocabulary, also has to perform simple routing tasks. Accordingly, a computer has a routing table. It can look something like this:

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.0.12.0	10.0.12.10	255.255.255.0	U	1			eth0
10.0.12.10	127.0.0.1	255.255.255.255	U	1	0	0	lo
10.255.255.255	10.0.12.10	255.255.255.255	U	0	0	0	eth0
127.0.0.0	127.0.0.1	255.0.0.0	U	0	0	0	lo
224.0.0.0	10.0.12.10	224.0.0.0	U	0	0	0	eth0
255.255.255.255	10.0.12.10	255.255.255.255	U	1			eth0
0.0.0.0	10.0.12.1	0.0.0.0	UG	0	0	0	eth0

A host must be able to perform routing to find out if the destination is connected to the same subnet as itself. If so, the host can forward the packet directly to the destination. If not, the host forwards the packet to the router that is connecting the subnet to the rest of the network. This router is called the default gateway (denoted 0.0.0.0 in the routing table with gateway 10.0.12.1) and the router's IP address on that subnet is therefore called the default gateway address. All hosts on a subnet must know of the default gateway address to be able to communicate with hosts outside the subnet.

There are some host addresses that have special meanings. The address 127.0.0.1 is the so-called loop back or local host address. It is the address of the computer itself. This address is not accessible from any other host on the network, but all hosts must handle calls to this address from its own applications. A full address with all ones, i.e. both network and host id has all bits set to one, is used as a limited broadcast address. This address can only be used as a destination address, and an IP packet with this destination address may not leave the subnet of the source. An address with the host id set to all ones is a directed broadcast. This broadcast destination address has a valid network identification part pointing out a specified subnet. It is destined to all hosts on this subnet. A directed broadcast packet can be forwarded by routers. A host id of all zeros addresses the subnet itself.

During the routing process the host makes lookups in the routing table to find the next hop for the packet. The next hop might be the destination itself, the default gateway or some other router, being the first hop on the path to the destination.

Using the above routing table, answer the following questions:

What is this computer's limited broadcast IP address?

What is this computer's directed broadcast IP address?

Which of the rows apply if you are sending a packet to IP address 10.0.12.100? Why?

Which of the rows apply if you are sending a packet to IP address 10.1.12.100? Why?

What is the IP address of the computer containing the routing table? Explain how you found your answer!

Use ping and try the address 127.0.0.1! Explain what you actually are testing. Hint: Try to use WireShark and find the transmitted ping packets. Which interface shall be sniffed in this case?

Now let us examine the routing table of **your computer**. Hint: Use `man route` to find out which exact command to use.

What is the IP address of your computer's default gateway?

What is the IP address of your computer?

Which subnet is your computer connected to?


For the next exercise we need WireShark again. Start a frame capture. Use the ping program again to ping the host with IP address 192.168.7.100. Stop the capture and examine the

captured frames. Possibly you see a lot of ICMP echo packets, not only yours. Apply the filter `ip.addr==<your ip address>`.

Examine one of the frames containing an ICMP packet, originating from your computer.

What is the MAC address of the destination in this frame?

*Is this the MAC address of the IP destination? Which host has this MAC address? Explain!
Hint: The destination is not on your subnet.*

 *What is the scope of address fields in the Ethernet frame header and the IP packet header?*

Now take a look on the network drawing.

Which network path do you think your ICMP packets took?

There is a handy tool that can check out network paths for you. It is called trace route. Examples of applications that perform trace route are `tracert` (Windows) and `tracert` (UNIX). Check `man traceroute` for more information. Use `-n` for numerical values.

Trace route the host address you pinged in the exercise above.

Make a note of the result!

Compare it with your findings using the network drawing. A match?

In a working network you normally use a dynamic routing protocol. Your computer's routing information is more or less static, especially when it is equipped with only one network

interface card like the lab computers. However, it is possible to change your computers routing information by using the route command.


First, ping another host than your own, connected to your own subnet.

Which path will it take? Compare your theory with the output of traceroute?

Make a note of the routing table entry regarding the subnet your computer is connected to!

Now remove the routing table entry you found above from your computer's routing table. Use the `route del` command together with `sudo`:

```
sudo route del -net 192.168.X.0 netmask 255.255.255.0 dev eth0.
```

Check the routing table to make sure that you have removed only the entry concerning your subnet! If in doubt, reset your network card and redo. The network card can be reset by right clicking on the network icon  in the upper right corner of your screen, and choosing Enable Networking **twice**.

Capture frames with WireShark, apply the filter `ip.addr==<your ip address> && icmp` and repeat the ping to another host on your own subnet. Stop the ping with `Ctrl + C` immediately after the first ping packet has been transmitted. Then Stop the capture. You should now have a few ICMP packets in WireShark. Examine the ICMP packets.

Which host owns the MAC address of the destination of the first ICMP echo request? Explain?

In this situation you might also see ICMP redirects. ICMP redirects are sent from a router when it receives a packet whose destination has a better path which doesn't involve this router. The router can then tell the sending host to use the better path instead. But host has since long stopped to respond to ICMP redirects since it can be misused in hacker or cracker attacks.

Give an explanation how ICMP redirects could be used in for instance a denial-of-service attack. What network process in the host is affected by an ICMP redirect and how? What would be the effect of a malicious ICMP redirect?

Reinstall the routing table entry you removed before. The command line will look something like this:

```
route add -net 192.168.X.0 netmask 255.255.255.0 dev eth0.
```

Compare with your notes above. Check that the routing table fully restored! If in doubt, reset your network card as described above.

DNS

The Domain Name System is a decentralised database where you, among other information, can look up which IP address corresponds to a computer's name or domain name. Normally an application, like a web browser, performs this lookup for you when you enter the name of a destination. However there are applications that perform only the lookup for you. On Windows NT systems and on many UNIX systems you can try `nslookup`. In newer versions of UNIX distributions this application is replaced by the application `host`.

Start the `nslookup` application in a terminal window.

Use nslookup to find out the domain name of your computer! Hint: Enter your computer's IP address on the command line.

On Windows systems, the local name system of the computer caches resolved name and IP address pairs. However, there is no DNS cache on an UNIX system by default.

Now use WireShark to capture what happens when you use DNS and `nslookup` to resolve the IP address corresponding to your computer's name. Use the filter `ip.addr==<your ip address>`.

What transport layer protocol is used? Why?

Which destination port is used?

A typical network application like Firefox maintains its own temporary DNS cache. Let us examine how. For this exercise we need to make sure that the DNS cache in Firefox is empty, so in case the program is open, simply restart it.

Use WireShark and capture the frames when you try to connect to a neighbour computer with Firefox, using its **domain name** (not the IP address). Apply the filter `ip.addr=<your IP address> && dns`. Restart the capture and repeat the connection attempt without restarting Firefox between the two attempts.

Explain the difference in capture output between the two connection attempts!

Explain what will happen if you use the IP address instead of the domain name in the previous exercise.

Network Utilisation

In this section we will look into sharing of network resources.

The local network you are using for the laboratory work is a 10 Mbps Ethernet.

What are the basic properties of this type of LAN? Hint: What does CSMA/CD stand for?

To be able to observe how an Ethernet works in practise, we have to generate some traffic. We will use Firefox to download a movie from another computer. This computer has the IP address 192.168.7.100 and functions as a web server/FTP server. The observations are performed with the utilisation measuring tool IO Graphs included in WireShark.

Open Firefox and clear the cache with `Ctrl + Shift + Del`.

Now, start WireShark, delete the previous filter, start a capture on interface eth0, and open IO Graphs. Consult appendix A for the proper settings of IO Graphs.

Go to the server home page at `ftp://server.local.lab/pub/` in Firefox and choose a movie to download. Simply click a link and the movie will play during download. Be patient, it can take a while.

Observe the IO Graphs and read your average transmission rate. If the movie has already been downloaded, choose another movie from the server home page.

Which is the theoretical value of the maximum transmission rate? Under which circumstances could this max value be reached?

What is your current reading of the transmission rate?

Give a brief explanation to your observation!

Close IO Graphs and apply the filter `ip.addr==<your ip address>` again. Stop the movie download and the capture.

What data is sent from your computer while downloading?

How many packets are sent back from your computer for each packet received by your computer? How many bytes does a typical return packet contain?

What transport layer protocol is used by the FTP application?

What are the advantages with this protocol?

Observe some FTP-DATA packets sent to your computer.

How many bytes data does a typical packet contain?

Give an estimate of the download overhead in percent. That is how much of the total data on OSI level 2, passed between your computer and the FTP server, containing other information than the actual movie. Include the ACK packets in your calculation!

Appendix A: Wireshark

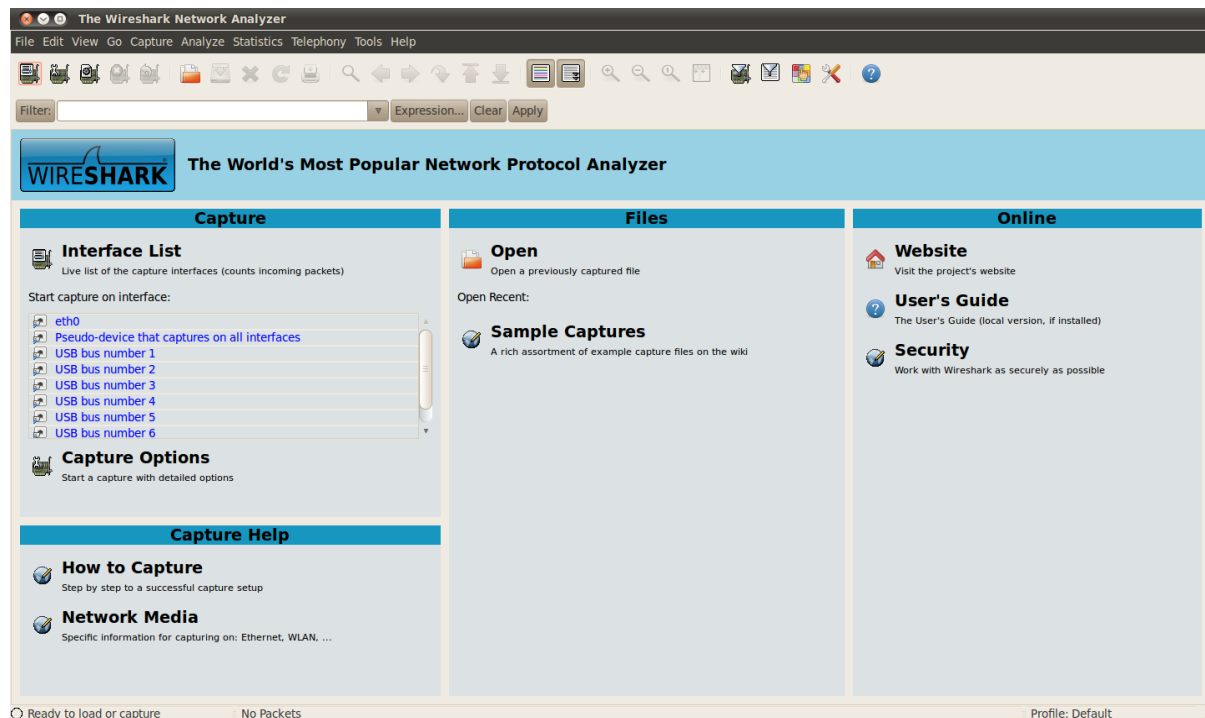
This document applies to Wireshark version 1.2.7.

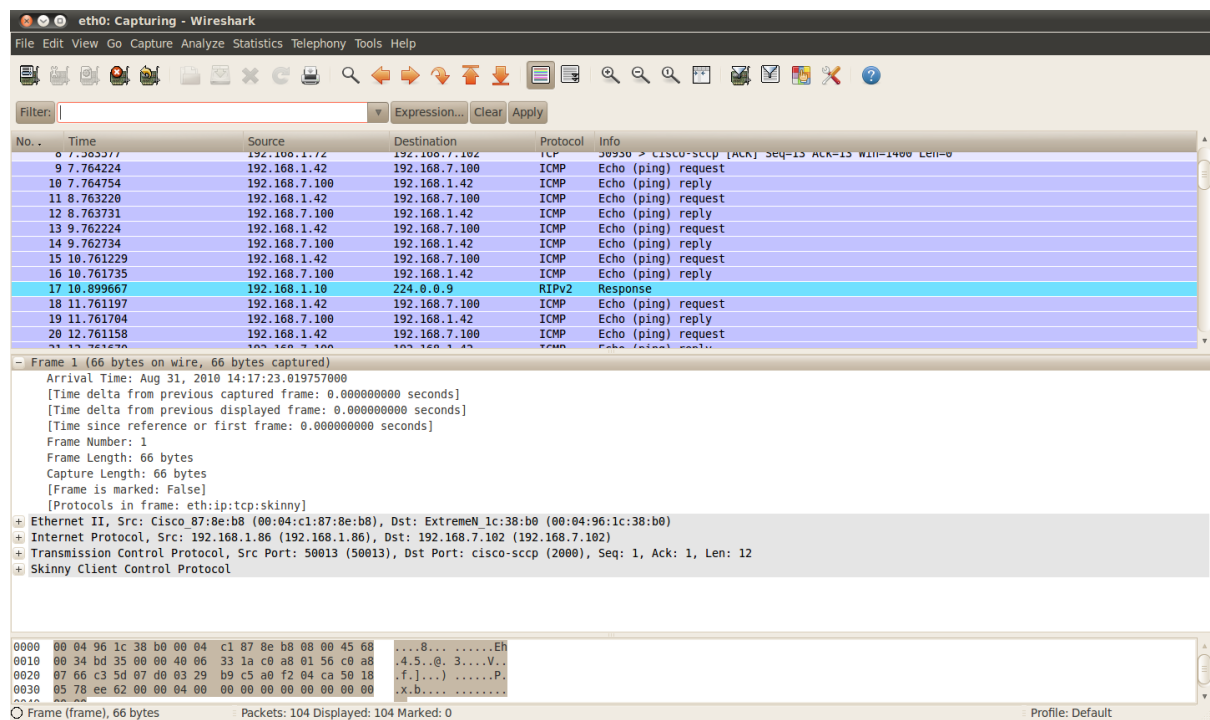
To be able to study data transferred on the local area network you can use the program Wireshark. Wireshark is a so-called sniffer program. It listens and decodes all data frames that are sent on the Ethernet by setting the network interface card, NIC, in so called promiscuous mode. When a NIC is run in promiscuous mode the normal decoding of destination address is not performed. Instead all frames on the local network are received by the NIC and pushed upwards to the application.

To open the program, use the command `sudo wireshark` in a terminal window. Enter the same password as for Telecomuser.


Sudo is as shortcut for *superuser do* or *substitute user do*, which temporarily gives you the super user (root) privileges needed to run the program in sniffer mode.

Start capturing frames by selecting `Capture > Interfaces` in the menu and hitting the `Start` button for interface `eth0`. Alternatively, choose interface `eth0` in the Interface List and the capture will start immediately.



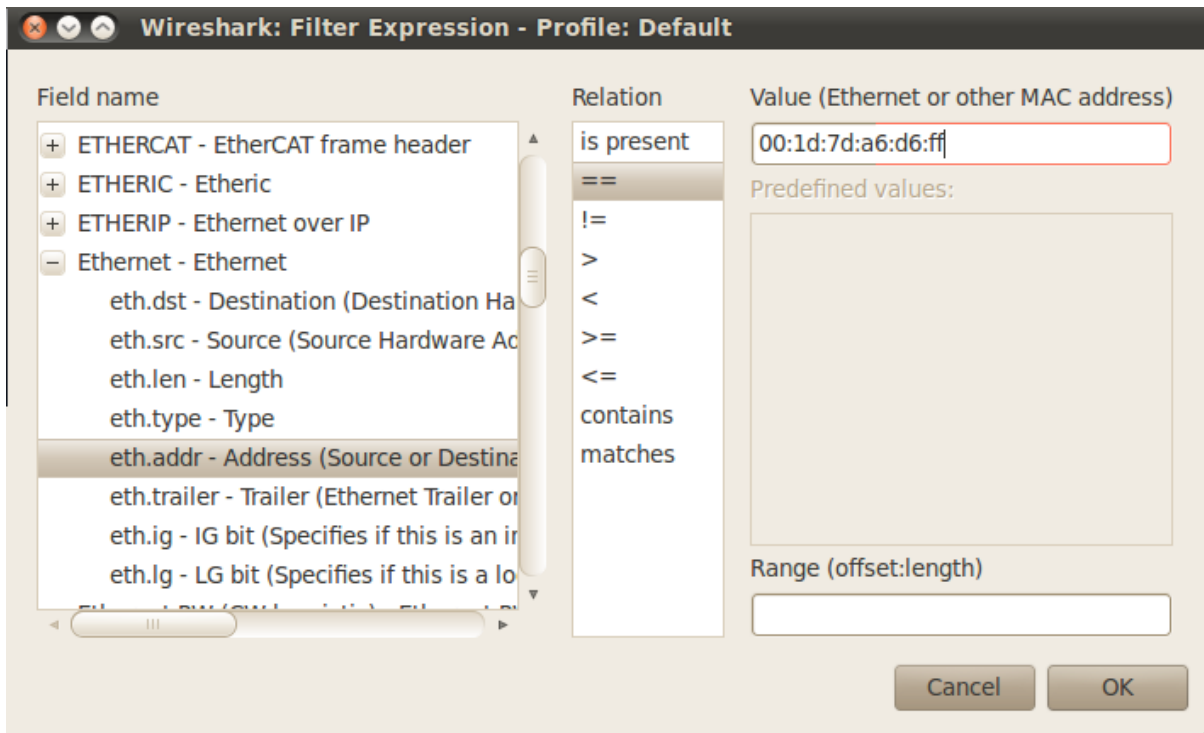


Captured frames are shown in real time in the main WireShark window. You find all captured frames in the upper window, one frame on each row. The frame capturing will continue until you stop it manually by hitting the `Stop capture` button in the upper left corner.



In the middle you can display the properties of the frame selected in the upper window. You can select properties corresponding to each layer in the OSI model by clicking on the + or - knobs.

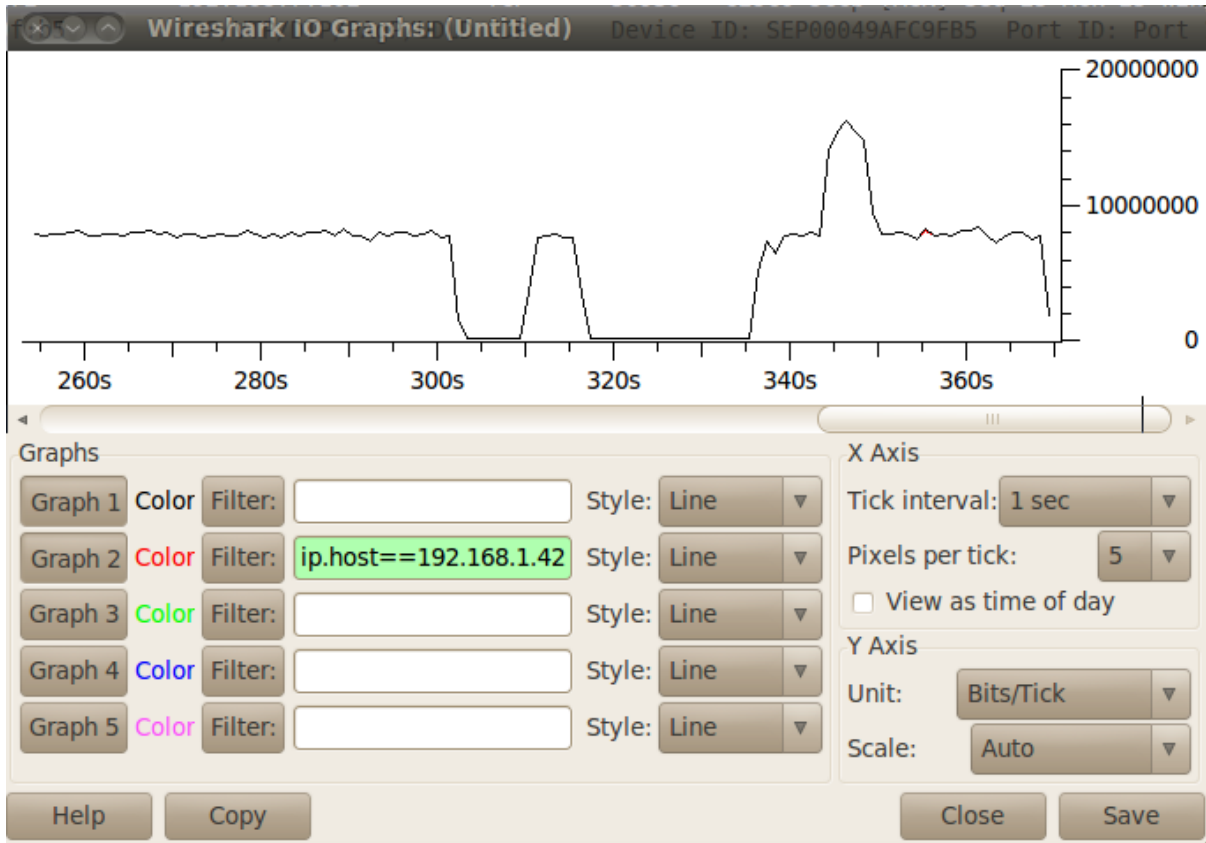
If you want to observe only selected frames you can apply filters. You can get help building filter expressions with the `Expression` button next to the filter text box in the upper left corner. In the Filter Expression window, search for the Ethernet and select field name `eth.addr`. Select the `==` relation and enter your computer's MAC address in the upper right text box. Then hit OK. If the filter expression is valid, it should be displayed with a green background in the filter text box of the main WireShark window. Now hit Apply. Of course, if you already know the syntax for your filter expression, you can enter it directly in the filter text box.



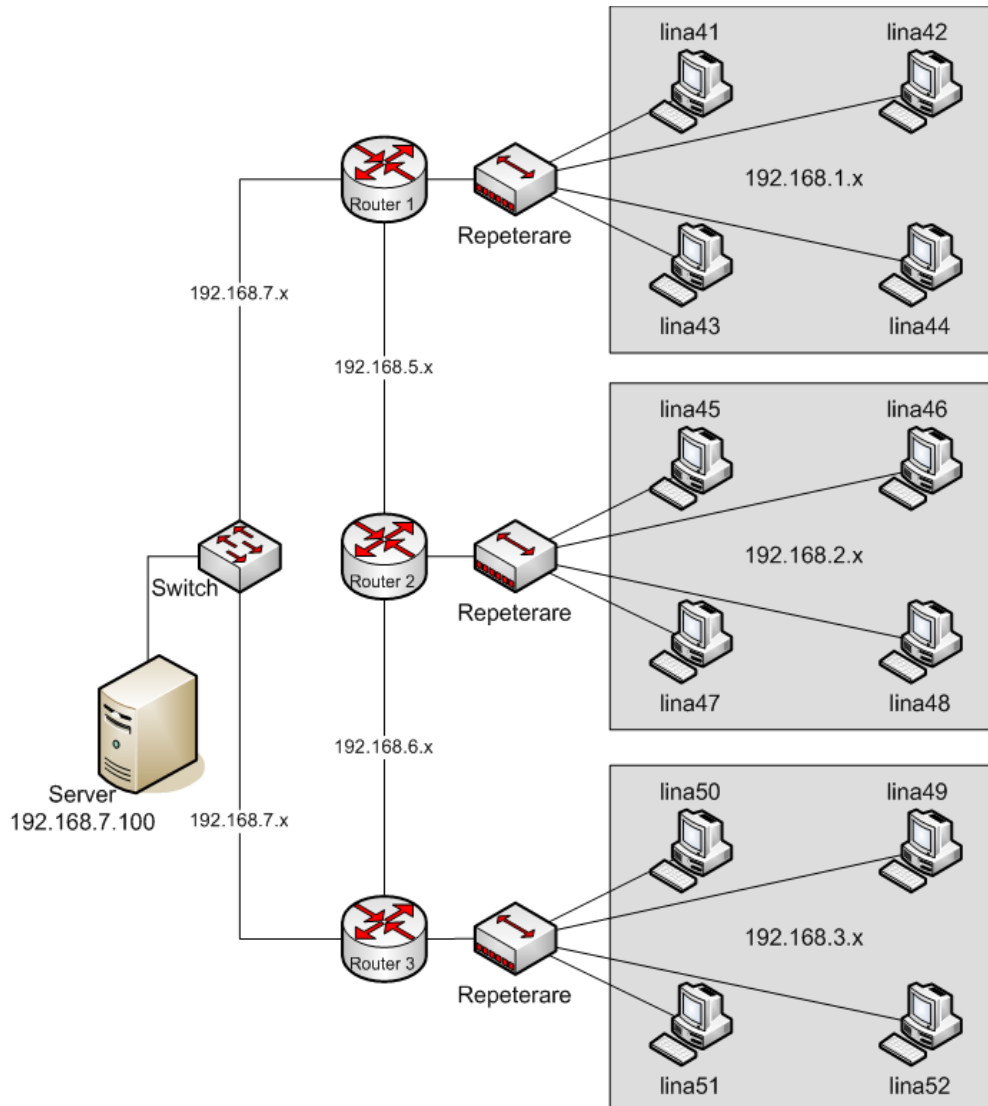
To use the utilisation tool IO Graphs, choose **Statistics > IO Graphs** in WireShark menu.

Set **Units** to **Bits/Tick** and **Tick Interval** to **1 sec**.

The black line, Graph 1, is default and represents the total traffic on the LAN. The red line, Graph 2, can be used to represent the traffic directed to and from your computer. Apply the filter `ip.host==<your ip address>` in the filter box of Graph 2 and press Graph 2. Now, the red line will show up together with the black. If you are alone on the LAN the two graphs might overlap.



Appendix B: Lab Network Outline



IP-adress	MAC-adress	IP-adress	MAC-adress
192.168.1.10	00:04:96:1c:38:b0	192.168.2.45	00:1d:7d:d1:85:3d
192.168.2.20	00:04:96:1c:36:a0	192.168.2.46	00:1d:7d:af:ce:e4
192.168.3.30	00:04:96:1c:3f:30	192.168.2.47	00:1d:7d:d1:83:56
192.168.7.100	00:19:66:47:4e:0c	192.168.2.48	00:1d:7d:af:ce:b9
192.168.3.49	00:1d:7d:d1:85:21	192.168.1.41	00:1d:7d:a6:d5:7d
192.168.3.50	00:1d:7d:af:ce:e0	192.168.1.42	00:1d:7d:a6:d6:ff
192.168.3.51	00:1d:7d:d1:85:27	192.168.1.43	00:1d:7d:af:ce:b5
192.168.3.52	00:1d:7d:d1:85:25	192.168.1.44	00:1d:7d:d1:85:3a

All routers run OSPF.
 DNS server = 192.168.7.100