# More on Ethernet

Michael Tsai
2017/03/27

# MAC address

- AB:CD:EF:01:23:45 or ab-cd-ef-01-23-45
  —> 6 byte, each byte represented by HEX

- Usually hard-coded on the network interface card

- The first 3 bytes correspond to a specific vendor
  Example: a MAC address lookup tool:
  http://aruljohn.com/mac.pl

# Address Resolution Protocol (ARP)

hsinmu.csie.ntu.edu.tw

mvnl-1.csie.ntu.edu.tw

**mvnl-1** wants to send a frame to **hsinmu**

1. Get IP address from the domain name (DNS, Application Layer)
2. Assume IP address is on the same subnet (more on this next time)
3. Get dest. MAC address (卡號) from the IP (ARP, IP/Link Layer)
4. Send the frame with “dest. address” in the header filled with the learned address

# Address Resolution Protocol (ARP)

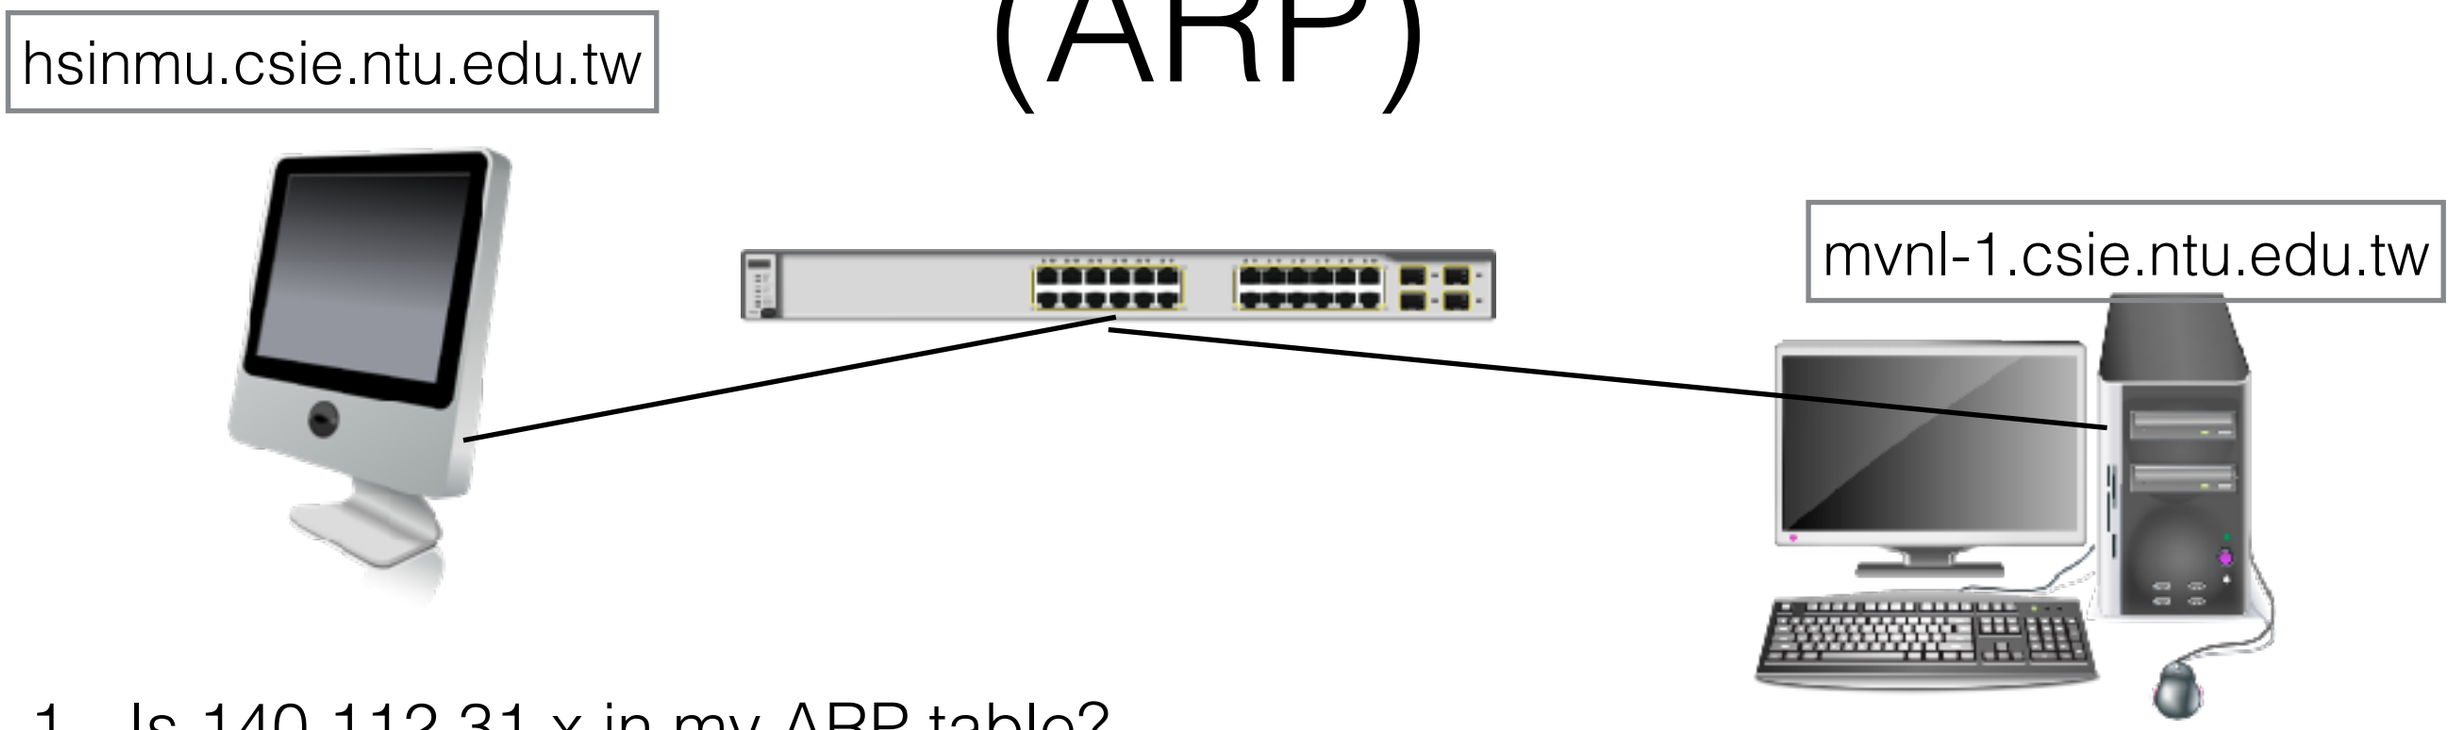


1. Is 140.112.31.x in my ARP table?
2. If yes, send the frame to the MAC address associated with this IP. (put it in the header)
3. If no, send a ARP request (who is 140.112.31.x?)
   this request is destined to FF:FF:FF:FF:FF:FF (broadcast address, received by all hosts on the network)
   (all switches along the way put mvnl-1's IP and MAC address in the table)
4. hsinmu responds with its MAC address and IP
5. mvnl-1 heard the response and put it in the table
   (and the switches along the way also do the same)

# In-Class Exercise

- Spend a few minutes to look for information about "ARP Spoofing"
  - What are possible ways to defend against this attack?

- (Do it when you have time)
  Read about the arping tool and try it out.
  http://www.habets.pp.se/synscan/programs.php?prog=arping

# Unicast, Broadcast, and Multicast

- Unicast = send to one destination
  (send to the address of the destination host)

- Broadcast = send to all hosts on the network
  (send to broadcast address)

- Multicast = send to a group of hosts
  (send to multicast address)

- Broadcast = at a switch, an incoming broadcast packet is copied and sent through **all other ports**.

# Switching Loop

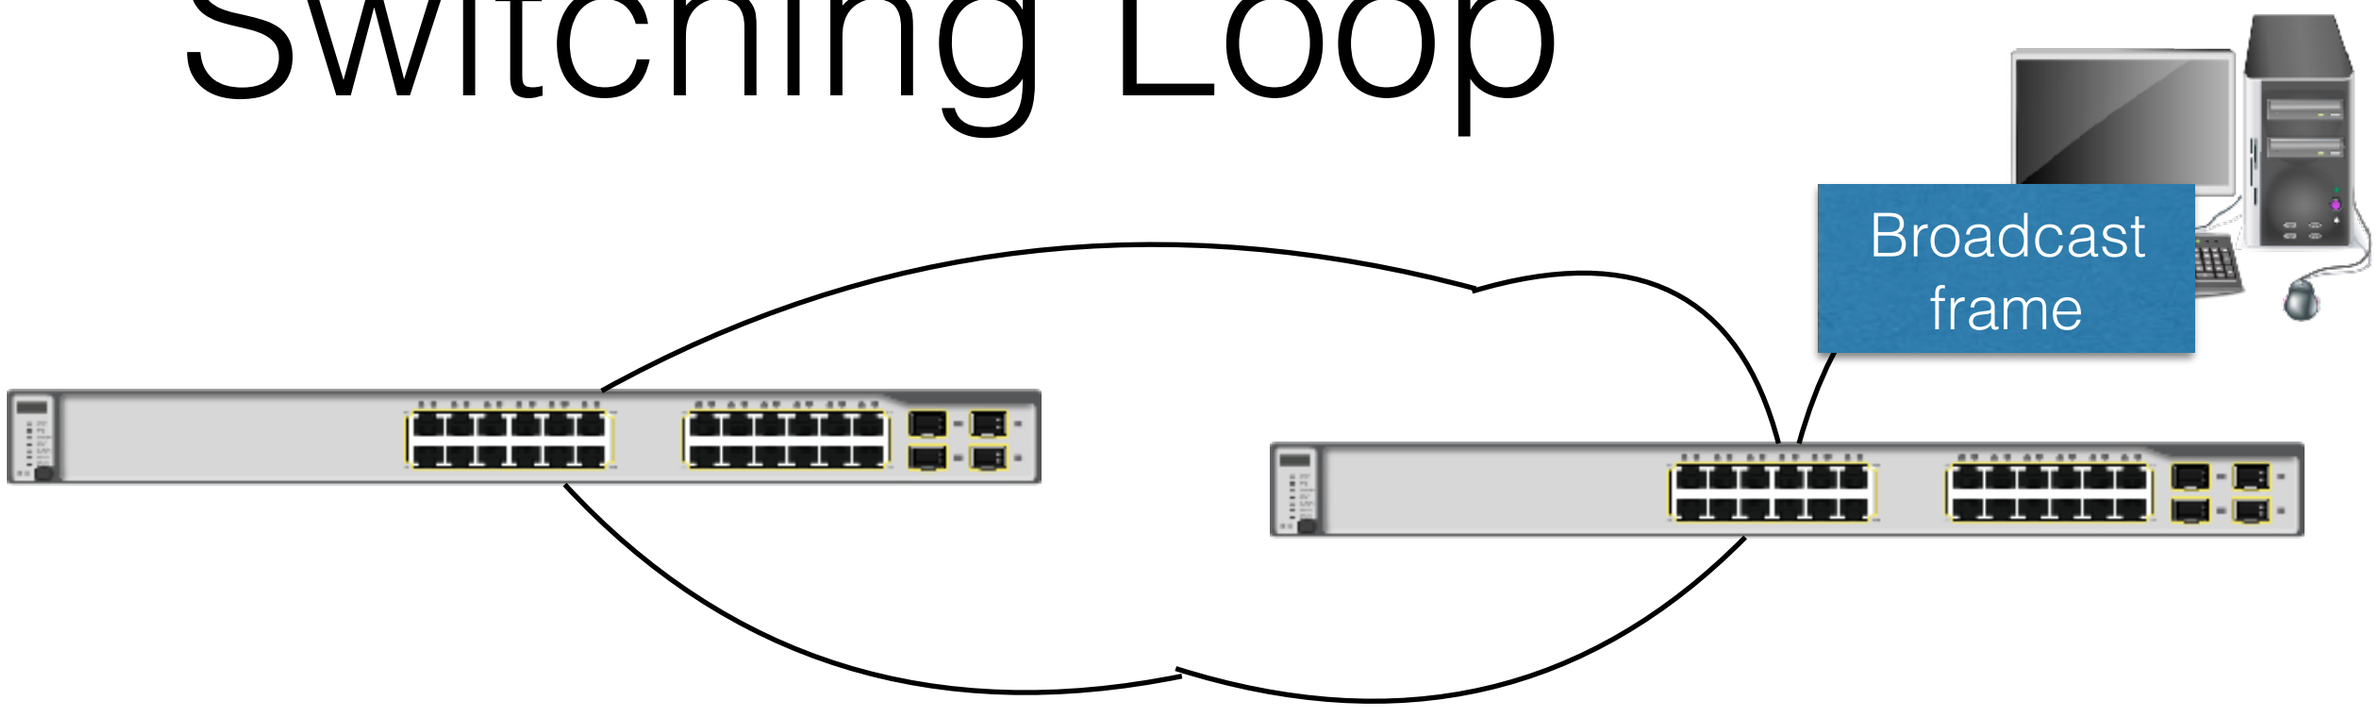


- Switching loop = a cycle in the “graph” created by the switches and the links (can involve multiple switches)

- What can happen?

  - Broadcast storm: broadcast frames repeatedly copied and sent to all ports, until all bandwidth is consumed.

  - ARP table is no longer reliable (MAC flipping or seeing MAC on more than one port) —> unicast frames sent to wrong port (frame loss)

# Q: RD & Special Project Division Want Separate Networks

RD Switch

RD Server 1

RD Server 2

RD Server 3

Special Project Switch

But this is too expensive!
1. Each subnet has limited # of hosts
2. More subnets in the future

Special Project 1

Special Project 2

# VLAN

## Ethernet Frame Format

| Preamble | Start of frame | MAC destinatio | MAC source | Length (IEEE 80 | **802.1Q ta g** | Payload | Frame check |
|---|---|---|---|---|---|---|---|
| 7 octets | 1 octet | 6 octets | 6 octets | 2 octets | (4 octets) | 42–1500 octets | 4 octets |

802.1Q tag = subnet number

RD Server 1

VLAN 100

Switch with no VLAN support

RD Server 2

VLAN 100

Special Project 1

VLAN 200

Host network interface (driver) would drop packets with other VLAN number (tag).

RD Server 3

VLAN 100

Special Project 2

VLAN 200

“But RD geeks can still see our packets!”

# VLAN

RD Server 1

RD Server 2

RD Server 3

Switch with VLAN support

Port 1: VLAN 100
Port 2: VLAN 100
Port 3: VLAN 100
Port 4: VLAN 200
Port 5: VLAN 200

Special Project 1

Special Project 2

Network Administrator

- Switch would automatically create multiple logical subnets (VLANs).
- Packets from different VLANs would be separately forwarded.
- No 802.11Q in packets! (supported by ordinary hosts too!)

"But NA wants to see all traffic!"

# Trunk

RD Server 1

RD Server 2

RD Server 3

Switch with VLAN support

Port 1: VLAN 100
Port 2: VLAN 100
Port 3: VLAN 100
Port 4: VLAN 200
Port 5: VLAN 200

Special Project 1

Special Project 2

**Trunk**

Network Administrator

VLAN 200

VLAN 100

Port 6: VLAN 100, VLAN 200
One physical network interface receives packets from multiple VLANs with multiple logical network interfaces. **(tagged traffic)**

# Trunk

Switch with VLAN support

Switch with VLAN support


Port 1: VLAN 100
Port 2: VLAN 100
Port 3: VLAN 100
Port 4: VLAN 200
Port 5: VLAN 200
Port 7: VLAN 400

1 physical cable per VLAN?

No. **Trunk** to the rescue!
Only one **physical cable/port** is needed!

Port 1: VLAN 400
Port 2: VLAN 400
Port 3: VLAN 300
Port 4: VLAN 100

Port 8: VLAN 100, 400 <——> Port 5: VLAN 100, 400

# Link Aggregation

- Addresses two problems:

    1. Bandwidth limitation (when both links work)
    2. Lack of resilience (when one of the links fails)

- Both open standards (LACP in 802.1AX and 802.1aq) and proprietary protocols (e.g., Cisco)

- Useful for switches and hosts

- Limitations: ports on same switch & with same link speed

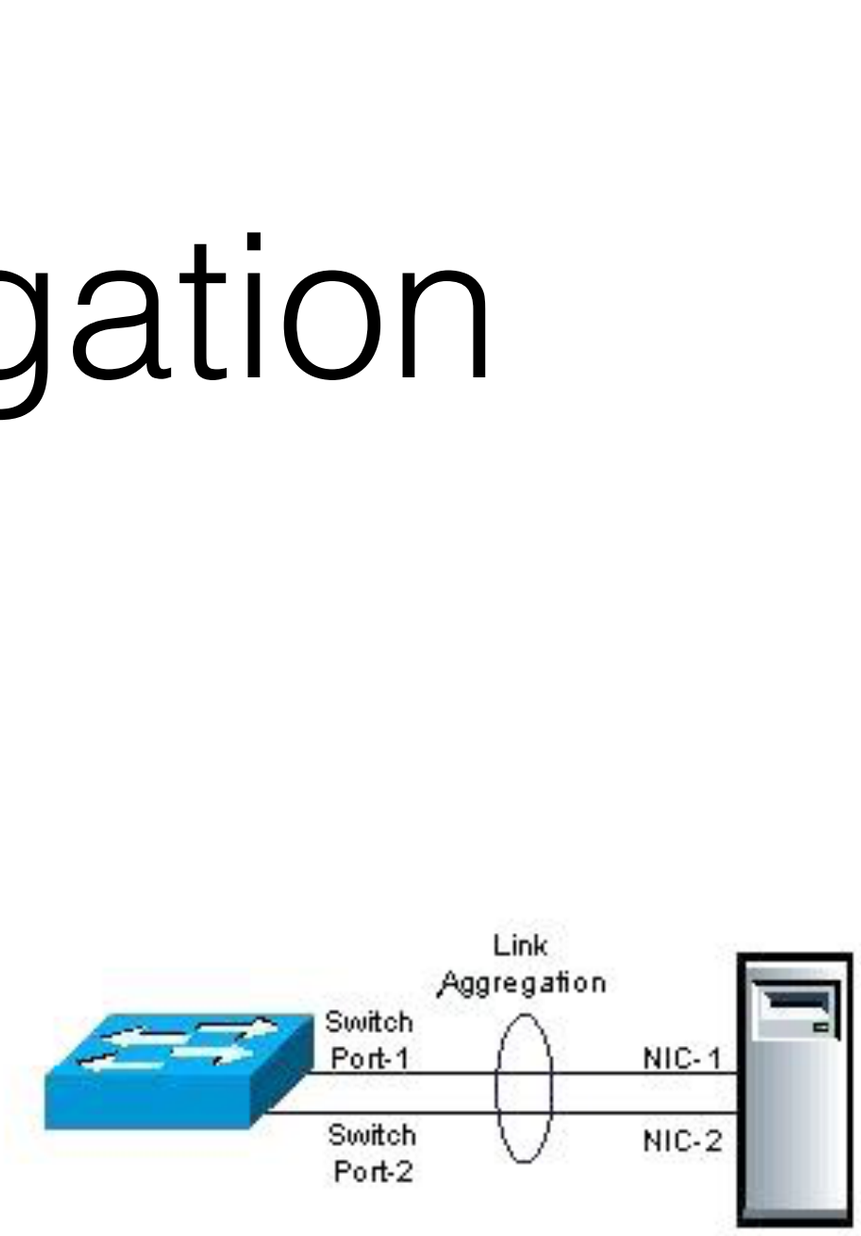

# VLAN & Trunk & LA

- Why are these important?
- Keep the cost low / the **physical infrastructure** simple, while providing
  - better security
  - better performance
- LA allows us to add additional cabling to increase backbone bandwidth

# Misc Information

- Power over Ethernet (802.3af): Power + Ethernet on the same physical cable
  - Important: power budget & availability of PoE switch
  - Not to confuse with Powerline Ethernet (what is this?)
- Jumbo frames (Gigabit ethernet): **header** is always an overhead
  Idea: 1518 byte frames —> 9018 bytes frame (or even larger)
  Measured gain from the experience: 10% gain
  (Limitation: all equipment on the subnet must support it)
- Autonegotiation: determine the speed of the other side
  - Must-use for all interfaces capable of 1Gb/s or above
  - < 100 Mb/s, auto or manual on **both** sides