Picking a model for 802.11/802.1 bridging

Point-to-point links, emulated LAN, or emulated bridge?

Norman Finn

Version 4

Introduction

- This presentation is available at: http://www.ieee802.org/1/files/public/docs2012/new-nfinn-11-medium-choice-0812-v04.pdf

Three different models

1. An 802.11 AP and its non-AP stations export to the rest of the network, and utilize themselves, a view of the 802.11 medium as a **set of point-to-point links** such that every non-AP station has a link to the AP. There may also be links between APs or between some pairs of non-AP stations.

2. The 802.11 AP and its non-AP stations appear to the logical bridge functions that may reside in some or all of the AP and its associated non-AP stations to be a single **emulated LAN**, rather similar to the original 802.3 “fat yellow coax”.

3. The 802.11 AP and its non-AP stations appear to be a single **emulated bridge** to the rest of the network.

Set of point-to-point links

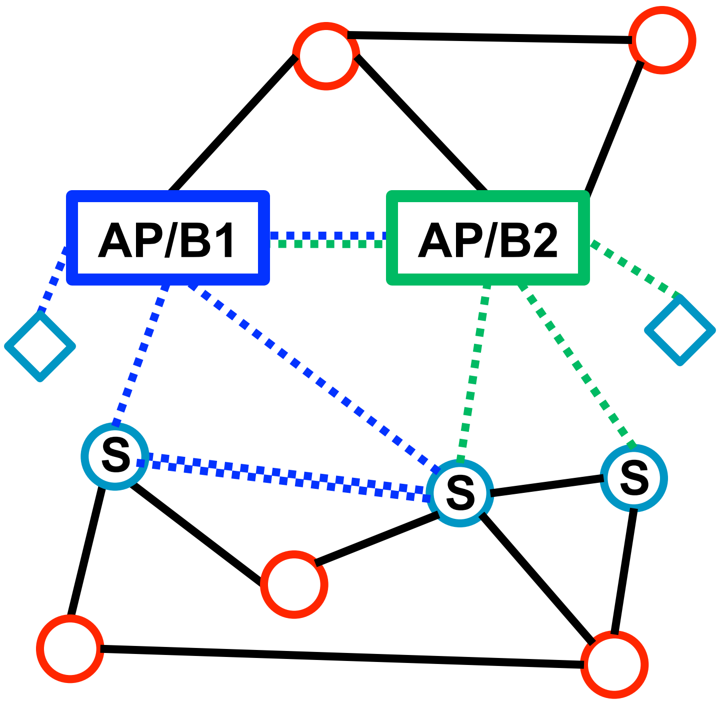

- The Access Points and their co-resident bridging functions become integrated AP bridges (AP/Bs).
- Devices with non-AP station capability(ies) and wired connections become "non-AP station bridges" (S).
- Of course, not all stations are bridges. (The diamonds are non-bridge non-AP stations.)

802.11 LAN emulation

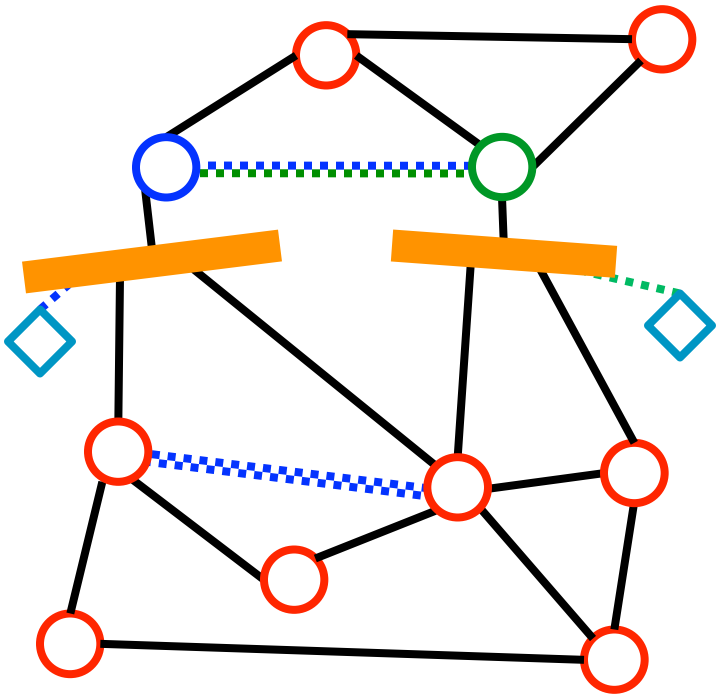

- Each AP and its stations emulate a shared medium LAN (fat yellow coax), as seen by the wired bridges.
- Each AP uses its bridge knowledge to optimize forwarding through the 802.11 medium, rather than broadcasting every frame.
- Direct AP-AP links have to be modeled separately from “coax”. Station-station links can be separate (shown) or part of emulated LAN.

Emulated Bridge

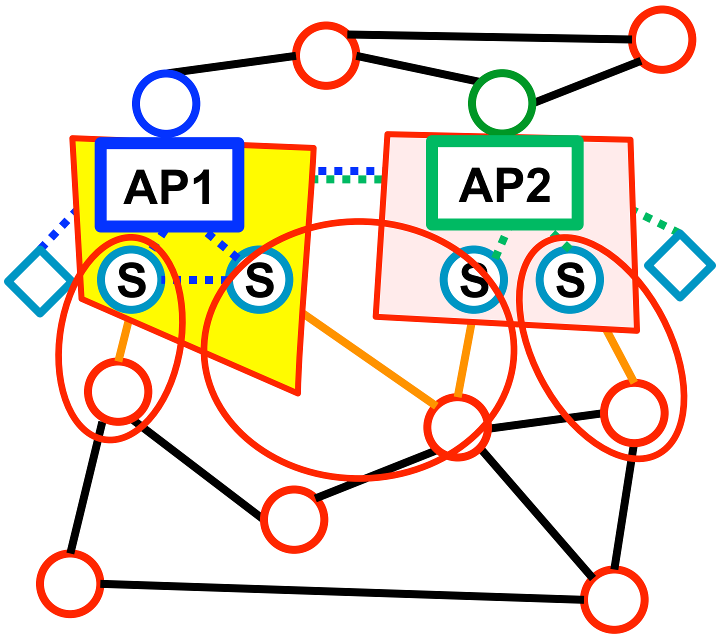

- Each Access Point and its non-AP stations emulate a single, separate bridge.
- An AP with multiple wired connections is logically separated into an AP and a wired bridge.
- Each station/bridge must be broken up with separate stations and a virtual wired bridge, with virtual wires to each component.

Comparing models

Simple LAN Emulation model rejected

- Suppose we fix SPB (and any other protocols that need fixing) to work properly over shared media.
- Several protocols – in particular MVRP, MMRP, and IGMP snooping – already work on shared media.
- Each of these three protocols avoids excessive PDU transmissions by observing that other devices have advertised their needs, and not re-advertising information that is known to the shared medium.
- The result of this anti-chatty behavior is that, **if these protocols are broadcast to all**, some potential transmitters will not speak up, so **the Access Point will be unable to prune VLANs or multicasts**. The AP must distribute all VLANs/multicasts to all stations. This is not acceptable.

More advanced LAN Emulation model?

- The APs non-AP station bridges could handle MVRP/SRP/ IGMP specially. They could always handle these protocols as if the wireless medium is a set of point-to-point links, not a shared medium.
- But, it is precisely in the handling of protocols that the difference between the LAN Emulation model and the point-to-point-link model become apparent. We want the AP and non-AP stations to handle data the same, in any case.
- I believe that **emulating a fat yellow coax** in all ways except 1) move the data more optimally than copying every frame to every receiver, and 2) use the protocols in point-to-point fashion, instead of shared media fashion, **is largely equivalent to, but more complex than, simply using the point-to-point model**.

Emulated bridge or point-to-point

- **How many bridges must be implemented** in these models?

10 (But see later section) **12**

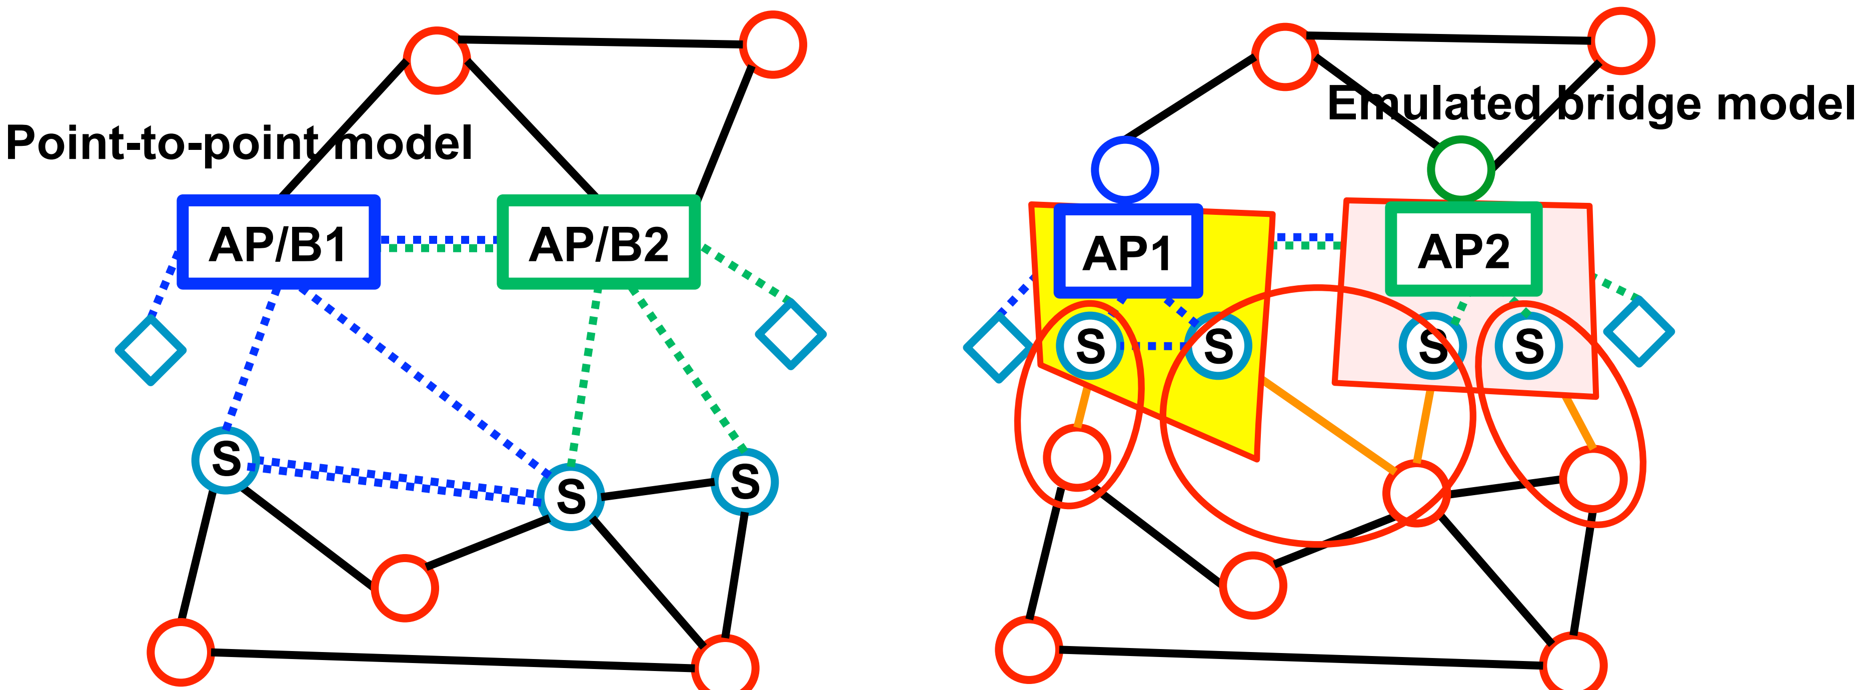

Emulated bridge or point-to-point

- **What must be implemented in a station/bridge** in these models?

A normal bridge with four ports.

A wired bridge with three ports, plus an emulated bridge component with one external and two internal ports.

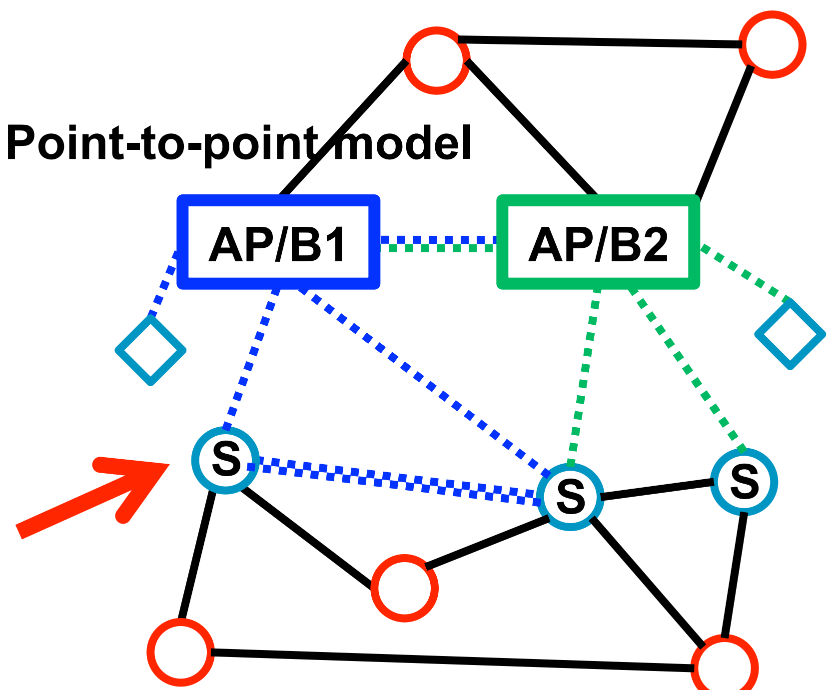

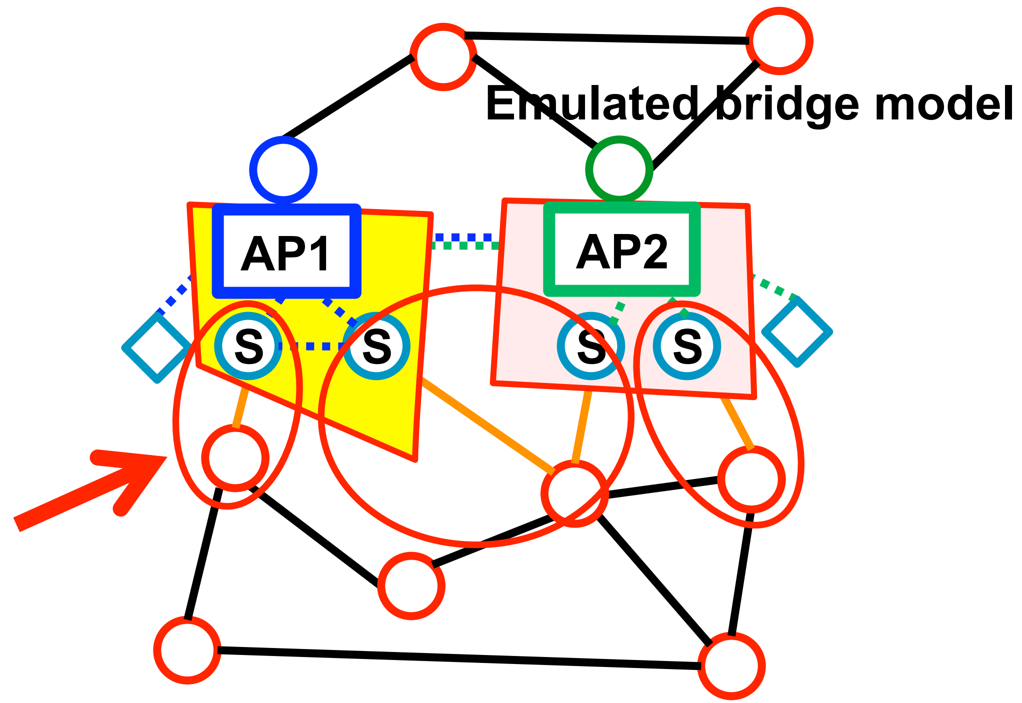

Emulated bridge or point-to-point

- **What must be implemented in a station/bridge** in these models?

A normal bridge with six ports.

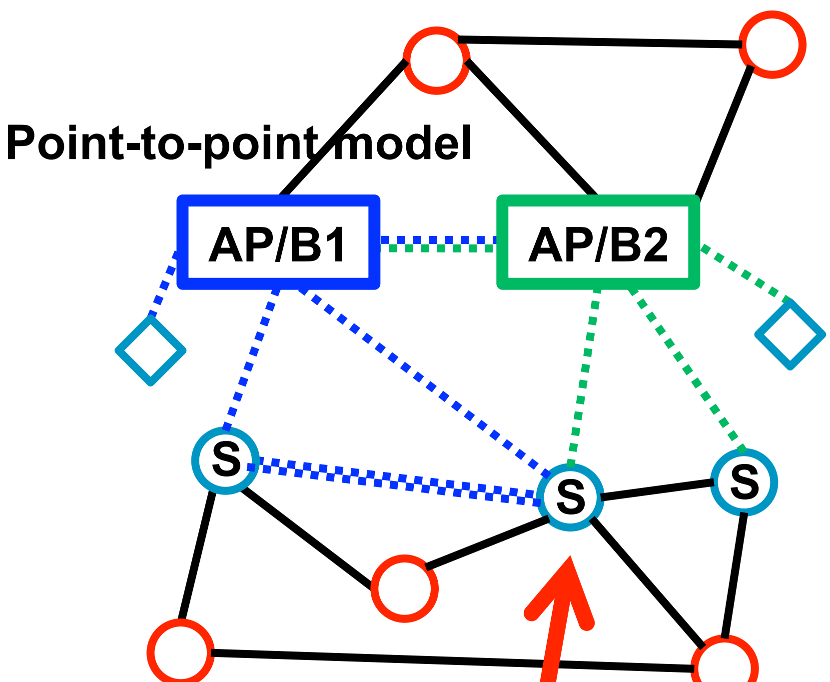

A wired bridge with five ports, plus two emulated bridge components serving different masters.

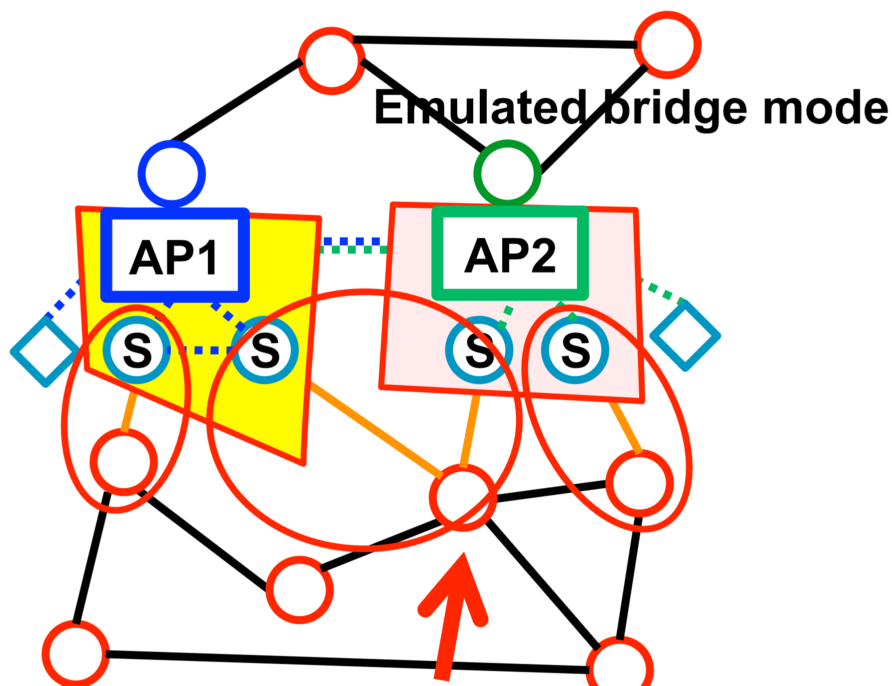

Emulated bridge or point-to-point

- **What must be implemented in a station/bridge** in these (simplified) models?

A Two-Port MAC Relay. **An emulated bridge component.**

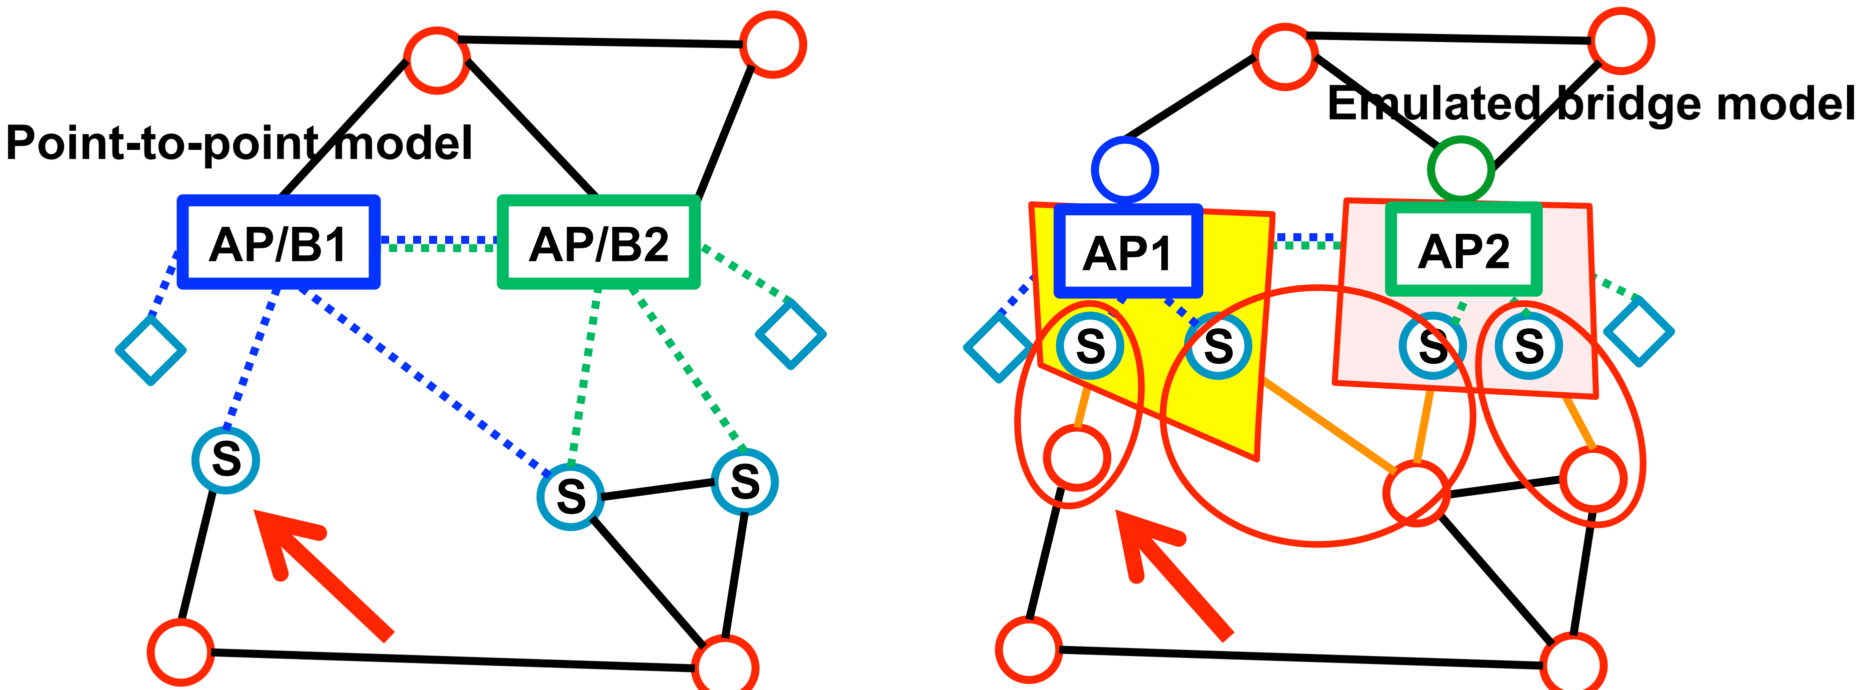

Emulated bridge or point-to-point

- I see no simplification to be obtained from inventing an emulated bridge and making the non-AP station/bridges elements of that emulated bridge.
- On the contrary, I see added complexity in the implementation, and added complexity in the specification.
- I recommend we use a model (the point-to-point model) in which the network map generated by SPB looks like the boxes and connections that, in fact, exist.

The central 802.1 / 802.11 issue

The single vs. dual function question

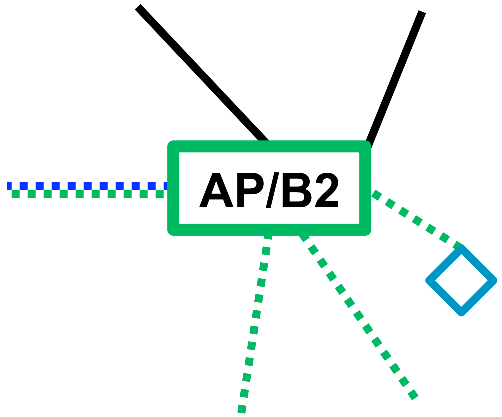

- Is the AP/bridge box a combined AP/bridge with both wired and wireless ports that are, for the purposes of routing frames, equivalent?

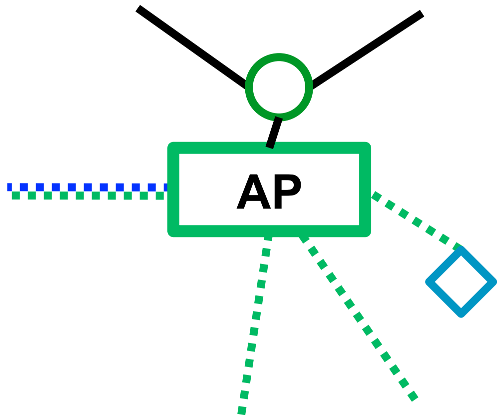

- Or, is the AP/bridge box an AP with a single wired “Portal” and a wired bridge?

The single vs. dual function question

- In my opinion at this writing, **the single function model is the preferable path**. Why?
- There are two parts to either a bridge or an AP:
  a. The forwarding part that uses MAC addresses, VLANs, etc., to decide on what port to transmit the frame; and
  b. The “QoS” (quality of service) part that decides, once the output port is selected, exactly when to transmit it.
- Let us consider these two parts separately.

The forwarding part of an AP/Bridge

- Both bridges and APs have definitions of forwarding tables.
- In bridges, these tables include the spanning tree port states, the filtering database, and the port/VLAN table.
- In APs, these tables are simpler. The handles required to interact with the bridge protocols are not present.
- The behavior of an AP separated from a bridge must, in respect to forwarding behavior, be identical to that of a bridge; otherwise the network doesn't work.
- So, **the AP's behavior with respect to forwarding has to change**. I would offer that the bridge forwarding model is the only sensible choice.

The QoS part of an AP/Bridge

- APs have very different QoS/queuing/delivery time requirements from Bridges. For example:
  - APs sometimes hold outgoing frames until the sleeping target station wakes up.
  - The actual physical 802.11 medium is, in effect, a single port (to the radio transmitter), with a single set of queues, not a set of individual point-to-point connections, each with its own set of queues.
  - In order to achieve acceptable reliability, APs and stations employ acknowledgements and retransmissions.
- So, **the Bridge's behavior with respect to QoS has to change**. I would offer that an amalgam of the Bridge's and the AP's QoS models is the only sensible choice. This will take some effort on the part of both groups, to see which .1Q QoS features work in 802.11.

The single vs. dual function question

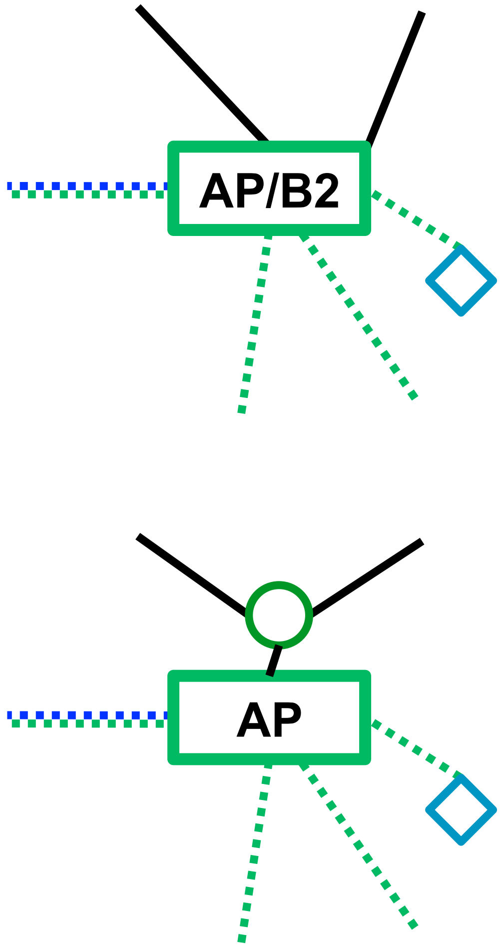

- Given the above arguments, I think that the single-function model for a combined AP/Bridge is the right answer, not a dual wired Bridge + modified AP model.

The art of standards writing

- I do not think it is practical to rewrite 802.1Q to accommodate the AP QoS structure.
- I do not think it is practical to rewrite 802.11 to accommodate the Bridge's forwarding structure.
- I think it may be practical for each standard to say, "802.1Q and 802.11 can be combined into a single AP/Bridge or station/bridge. For this part of the specification (QoS in .1Q, forwarding in .11), please look at the other document for guidance.

Issues peculiar to point-to-point model

Issue: Multicast distribution

- Each device below is a bridge, wireless connections are treated as point-to-point links, and a broadcast frame is sent by bridge **X**.
- Suppose bridge **R** is the spanning tree root, so that one of the AP's "ports" is blocked.
- In the standard spanning tree protocol, bridge **C** does not know that the AP's link to it is blocked.
- How does the AP forward the broadcast to **A** and **B** but not to **C**?

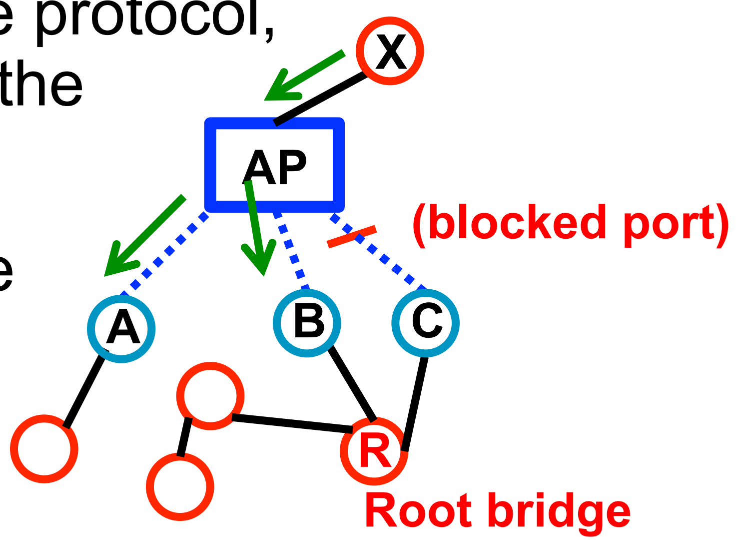

Multicast distribution

- One solution would be to extend/modify MSTP and/or Shortest Path Bridging to provide a handshake to tell bridge **C** that the AP end of the link is blocked.
- Another solution is to send multiple unicasts to the bridges, at least until the handshake (if any) is done.
- Another solution would be to provision a set multicast Receive Addresses, in frames sent by the AP, to specify sets of bridge / stations. (In this case, "**A** and **B** but not **C**".)
  - This latter idea has its own problems – either we must limit an AP to at most 24 bridge/stations (the number of bits available following the OUI in a MAC address), or define a protocol for distributing a mapping of vectors of stations to 24-bit IDs.

Multicast distribution

- Any of the above solutions have another consequence to bridges: In order to be efficient, 802.1Q needs to have the concept of a special port that provides multicast services to some number of individual ports.
- This same concept is required to support 802.1 EPON and MoCA media.
- This has been investigated by a number of 802.1 people over the years, and is believed to not be difficult for 802.11.

Issues peculiar to LAN emulation model

Issue: Protocols over emulated LANs 1

- On point-to-point links, MSTP performs a handshake with its neighbor, blocking a port briefly, in order to ensure against temporary forwarding loops. With this handshake, MSTP can converge in **milliseconds** after a topology change.
- There is no reliable handshake defined for a shared medium; instead, MSTP blocks a default timeout of **6 seconds**.
- This is one reason Shortest Path Bridging protocol (SPB) does not now support shared media. SPB uses a handshake to prevent forwarding loops that works very much like the MSTP handshake – it would have to time out on shared media.

Issue: Protocols over emulated LANs 2

- The **good news**: 802.1Q MRP (Multiple Registration Protocol, on which **MVRP**, **MMRP**, **MIRP**, and **SRP** are based) and IETF **IGMP** are very clever, in that they work on shared media with a minimum of transmitted frames.
- The **bad news**: **All such protocols require active assistance** by the LAN emulators in order to work on an emulated LAN unless the devices emulating the LAN take special steps.

Issue: Protocols over emulated LANs 2

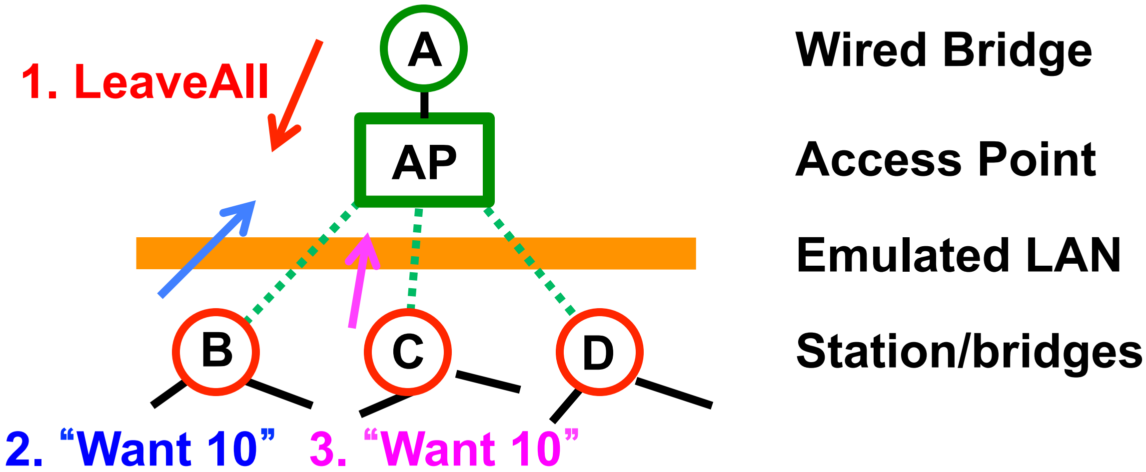

- Four bridges, A-D, attached to an 802.11 emulated LAN.
- All frames are broadcast to all bridges.
- Bridge A sends MVRP "**LeaveAll**" = "Reset & resend soon."
- All bridges (including A) start a short random timer.
- Bridge B times out first, & multicasts, "**I want VLAN 10**".
- Bridge C times out, & multicasts, "**I want VLAN 10**".

Issue: Protocols over emulated LANs 2

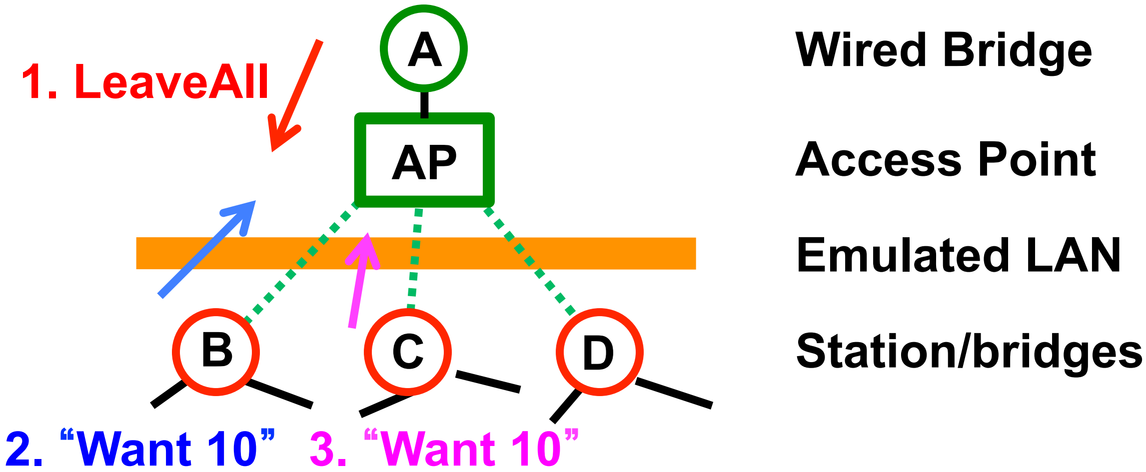

- Even if bridges A and D also need VLAN 10, **they say nothing**, because they see the others' transmissions, and thus know that everyone on the shared medium has seen the announcement (this stops unneeded chatter), so A and D will receive any VLAN 10 frames on the medium.
- The AP does not know whether A and D need VLAN 10, so it must send VLAN 10 to all four bridges, whether they need it or not! **Pruning doesn't work; bandwidth is wasted**.