Picking a model for 802.11/802.1 bridging

Point-to-point links or emulated LAN?

Norman Finn

Version 1

Introduction

- This presentation is available at: http://www.ieee802.org/1/files/public/docs2012/new-nfinn-11-medium-choice-08-12-v01.pdf

Two 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”.

Set of point-to-point links

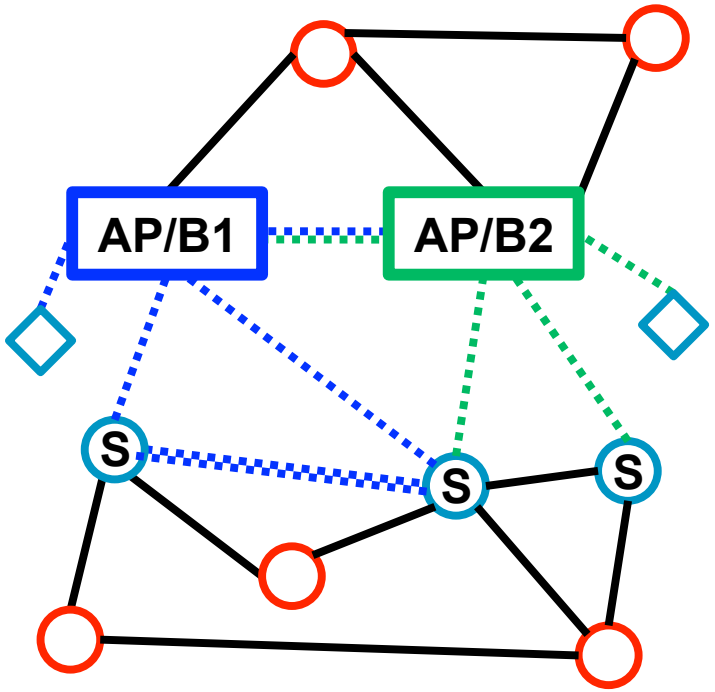

- The Access Points and their co-resident bridging functions become integrated AP bridges (AP/Bs).
- Devices with non-AP station capability(ies) and wired connections become “non-AP station bridges” (S).
- Of course, no all stations are bridges. (The diamonds are non-bridge non-AP stations.)

802.11 LAN emulation

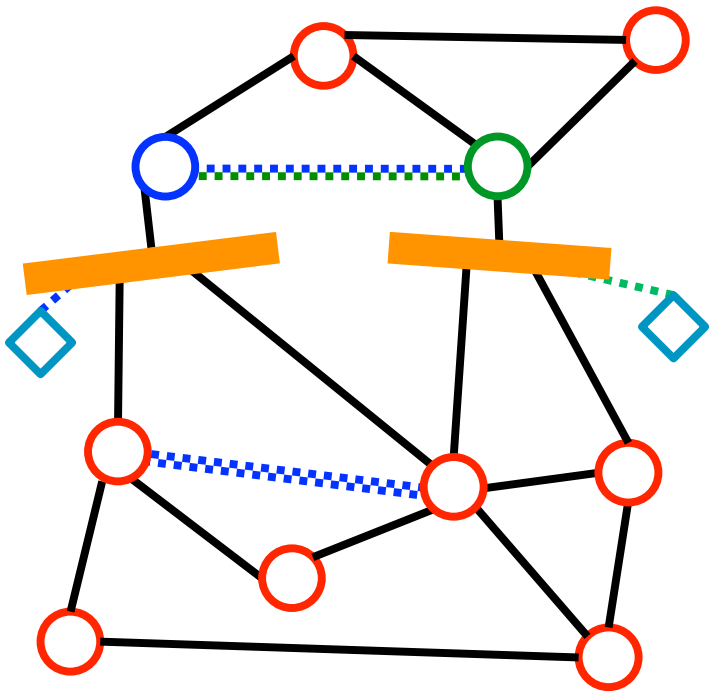

- Each AP and its stations emulate a shared medium LAN (fat yellow coax), as seen by the wired bridges.
- Each AP uses its bridge knowledge to optimize forwarding through the 802.11 medium, rather than broadcasting every frame.
- Direct AP-AP links have to be modeled separately from “coax”. Station-station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

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**?

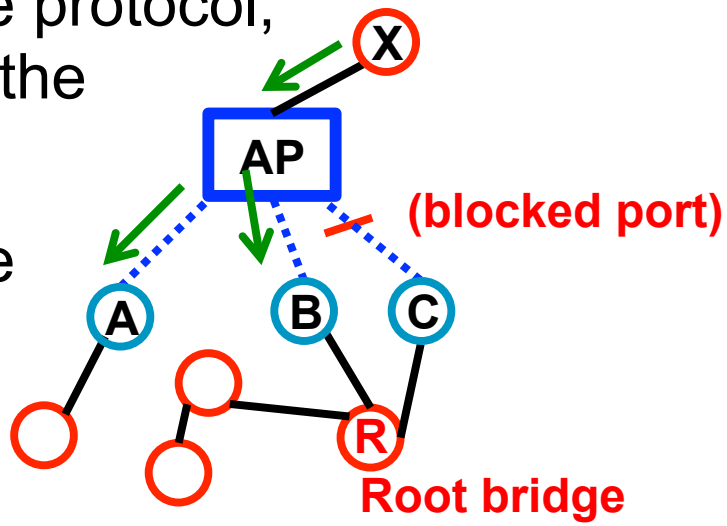

Multicast distribution

- 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.

Issues peculiar to LAN emulation model

Issue: Hidden costs

- There are issues with the actual vs. the apparent cost to cross the emulated LAN.

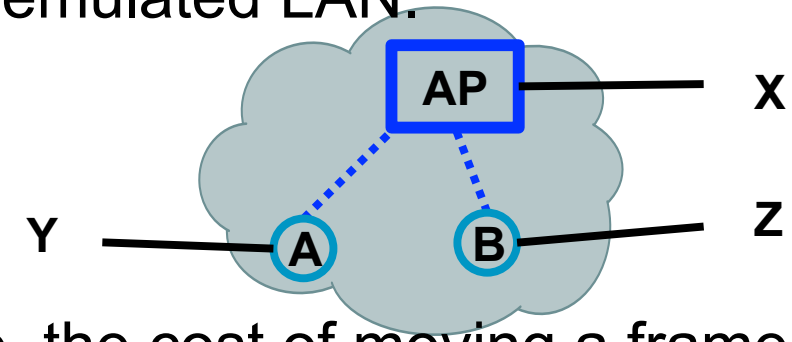

- For example, the cost of moving a frame between X and Y is typically half the cost of moving a frame between Y and Z.
- This cost difference can be important when dealing with video streams, and cannot be factored into forwarding decisions if the links are part of an emulated LAN, to which bridging protocols give a single cost.

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**.
- The Shortest Path Bridging protocol (SPB) **does not support 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.
- (Actually, that's not quite accurate – SRP does not support shared media.)

Issue: Protocols over emulated LANs 2

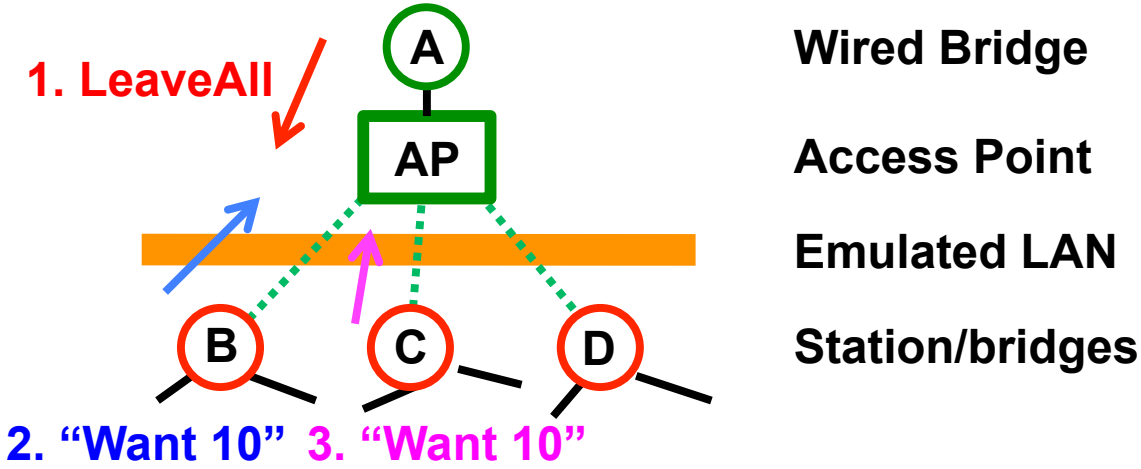

- Four bridges, A-D, attached to an 802.11 emulated LAN.
- All frames are broadcast to all bridges.
- Bridge A sends MVRP "**LeaveAll**" = "Reset & resend soon."
- All bridges (including A) start a short random timer.
- Bridge B times out first, & multicasts, "**I want VLAN 10**".
- Bridge C times out,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

Issue: Protocols over emulated LANs 2

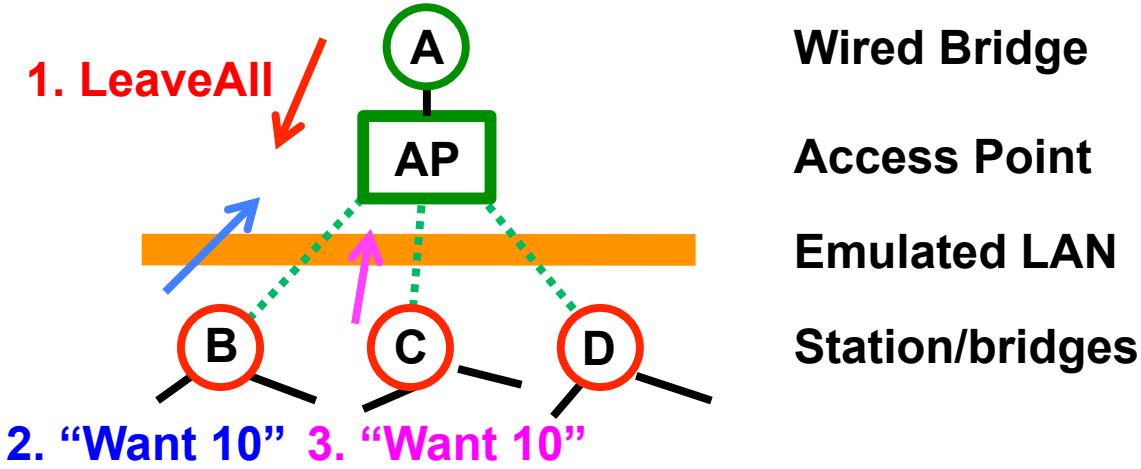

- Even if bridges A and D also need VLAN 10, **they say nothing**, because they see the others' transmissions, and thus know that everyone on the shared medium has seen the announcement (the clever part), so A and D will receive any VLAN 10 frames on the medium.
- The AP does not know whether A and D 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.**

Possible solutions for LAN emulation issues

Hidden costs

- Hidden costs are a largely inescapable consequence of layered abstractions. Where layered abstractions are used, the inability to make optimum decisions is accepted in the interests of simplifying the interactions between the abstracted entities.
- This tradeoff is made, for example, in the case of separate MSTP Regions that can interact only via the Common Spanning Tree, or in the case of 802.11s mesh networks and their interaction with wired networks.
- This author knows of no solution. That does not mean none exist, of course.

Protocols over emulated LANs

- In order to emulate a shared medium, while still providing the optimizations expected of a bridge (which are particularly to 802.11, because of limited bandwidth), the 802.11 stations must treat all of these protocols specially: MSTP BPDUs, MVRP, MMRP, MIRP, SRP, and IGMP.
- Also, SPB must be improved to handle shared media.
- If the AP and station/bridges are treating all of the 802.1Q protocols (and IGMP) as if they were bridges connected via point-to-point links, and if the data movement takes place over these same point-to-point links, then in what sense is are the APs and station bridges emulating a LAN, and not being a bridge?

Protocols over emulated LANs

- SRP is based on IS-IS.
- IS-IS handles a shared medium by replacing the shared medium in its calculations with a full mesh of point-to-point links among the members of the shared medium.
- Thus, IS-IS using an 802.11 emulated LAN has **all of the complexity of the point-to-point model** (**more**, actually, because not all stations have the direct links simulated by IS-IS), with **none of the benefits** of cost visibility.
- This is why, pending an understanding of the difficulties that 802.11 might have with the point-to-point solution, the LAN emulation model seems broken to 802.1.