



Wired – Wireless Bridging

IEEE 802 standards needed

Norman Finn

Version 2

Introduction


- This presentation is available at:
<http://www.ieee802.org/1/files/public/docs2012/new-nfinn-wired-wireless-bridges-0612-v02.pdf>
- We will describe the case for integrating **802.11** into 802.1Q VLAN bridging as “**just another medium**”, touching upon:
 - **Use cases**;
 - Selecting a **model** for wired-wireless bridging;
 - **Issues** and **potential solutions**.
 - A short-term **work plan**.
- This work addresses only **VLAN MAC bridging**; **IP routing is also important**, but not addressed.

Executive summary

- Contrary to the current 802.11 architecture, there are use cases for (and examples of) non-AP stations with both wired and wireless connectivity. These stations, as well as Access Points, need to be 802.1 bridges.
- The largest identified markets are in the **home** and in **industrial control** networks.
- Allowing APs and non-AP stations to be bridges opens up new markets and heads off the growing number of non-standard deployed solutions.
- There are several possible models for making 802.11 work just like any other 802-style medium.
- There are a number of issues to be addressed, but all perfectly solvable with 802.1/802.11 cooperation.
- **Time is of the essence – the market is moving:**
 - **802.1 PAR for wired-wireless bridging in November.**
 - **802.11 PAR to support 802.11 stations that are bridge ports. November?**
 - **Schedule weekly joint .1/.11 virtual meetings to progress this work.**


Use Cases

802.1 Audio Video Bridging world view



- In a home or small studio, Ethernet-like links: **802.3**, **802.11**, **MoCA**, **1901 power line**, etc. will replace other modes for exchanging data.
- You expect wired stacks connected via wireless.
- For **plug-and-play connectivity**, **every device with multiple ports** can usefully be an **802.1 bridge**.

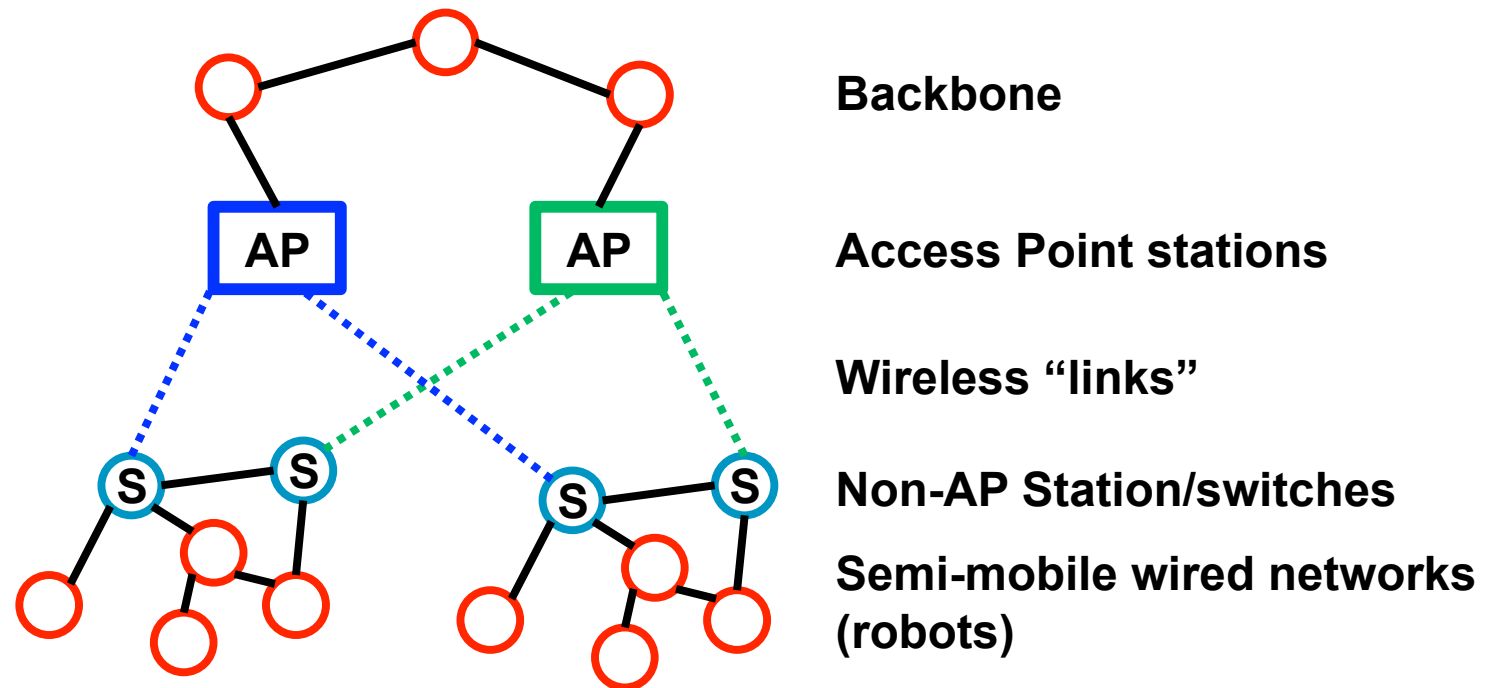
Simplified common case



- Most connected via Wi-Fi.
- Most connected via IEEE 1901 broadband (Ethernet) over power lines.
- Bridging is necessary for full connectivity.

Example: Wi-Fi as Just Another Medium


- Consider a possible **all-L2** industrial control network:



- IEEE 802.1 bridging technologies (e.g., 802.1aq SPBV) can make this a **single MAC Layer network**.

Wi-Fi is not at the edge of the network!

- 802.11ac, gigabit Wi-Fi, makes this even more imperative. **A Gb/s link is not always at the edge of the network!**



.11ac Access Point

1 Gb/s Wireless “links”

.11g Access Points

54 Mb/s Wireless “links”

Wired bits of a single company
in a multi-tenant building.

- Simply make these devices ordinary switches, and it works just fine.

Model: What kind of medium is 802.11?

Selecting a model

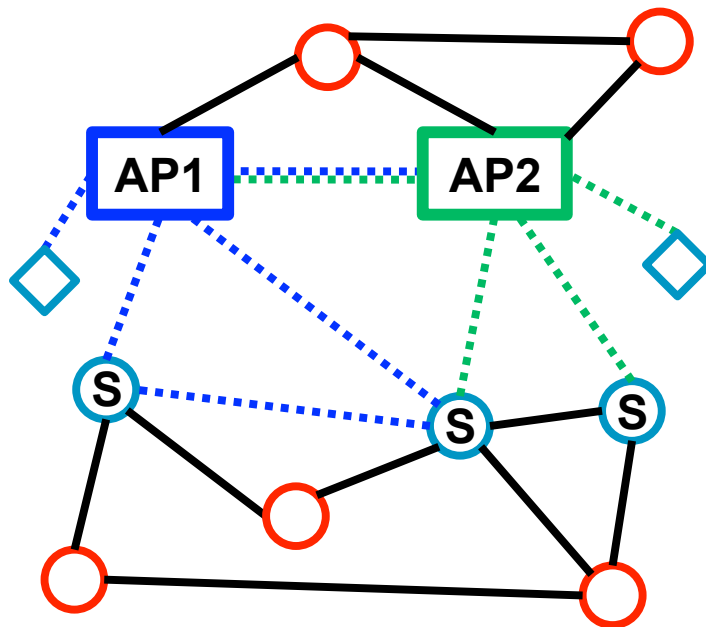
- There are many ways to make 802.11 “just another medium”.
- Several will be presented in this section.
- This author cannot claim to have identified all possible models for this idea; all of those presented here have been suggested in IEEE 802.1 Working Group meetings.
- There are many variants possible with each model. Only one variant for each model is presented, here. The choice is arbitrary.

Four different models

1. Form an **802.11s mesh** out of a subnetwork of interconnected APs and their station/bridges. This subnetwork peers with other (wired) subnetworks.
2. The collection of a single AP/bridge and all of its associated non-AP station/bridges (but excluding its non-bridge non-AP stations) cooperate to appear, to the rest of the network, to be a single **virtual bridge** with only wired external ports.
3. 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 **shared medium**, rather similar to the original 802.3 “fat yellow coax”.
4. 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 among some pairs of non-AP stations.

Models: “802.11 is just another medium”

- In order to compare these ideas, we will use a reference network:



Wired bridge


Access Points (loosely)

Non-bridge non-AP stations

non-AP station bridges (loosely)

wired bridges

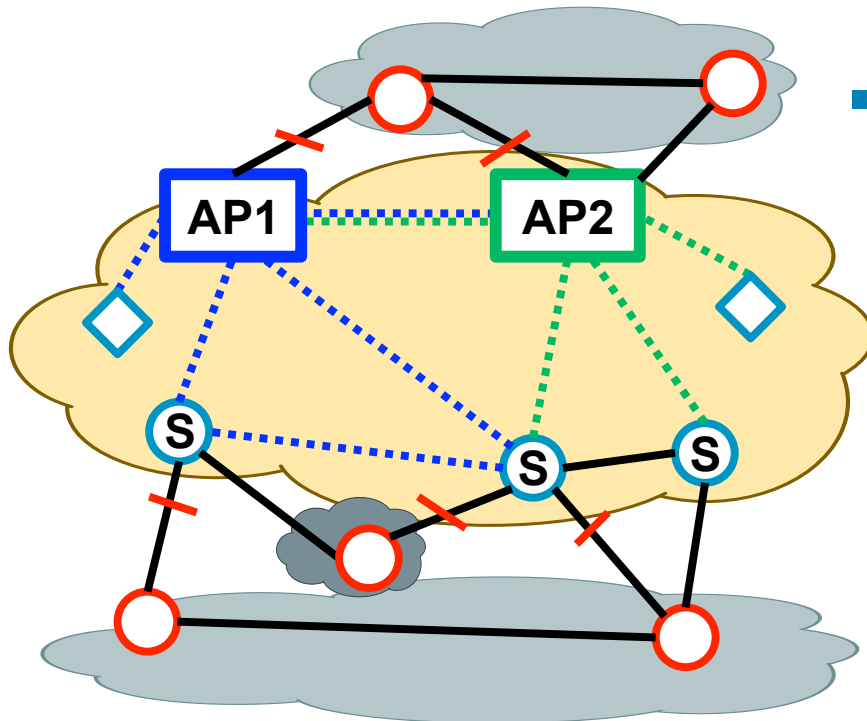
Current 802.11 architecture



- In the current 802.11 architecture, an Access Point with multiple wired connections should normally be modeled as a wired bridge with a single connection to the AP.
- Non-AP station bridges cannot exist.

Mesh subnetwork + wired subnetwork

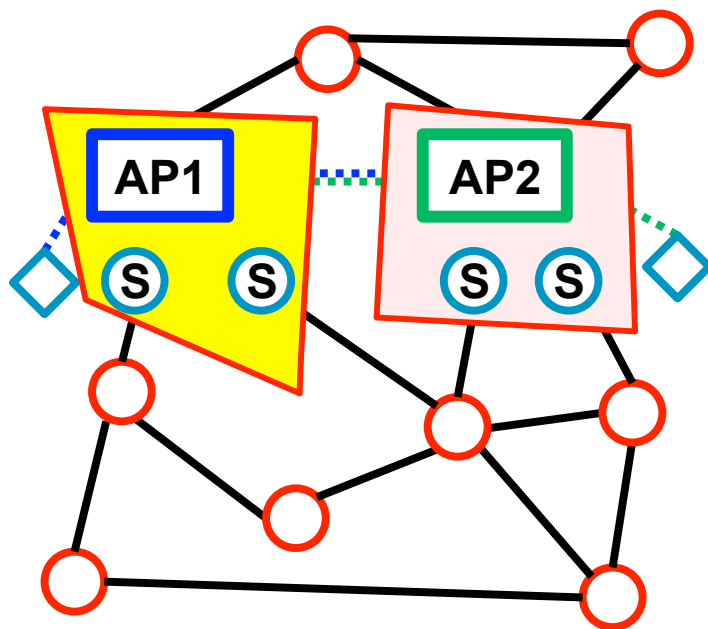
- There can be only one active connection between each pair of clouds.
- The interior of each cloud is opaque to the other cloud(s).



Mesh subnetwork + wired subnetwork

- The wireless world is integrated into an 802.11s mesh subnetwork. Each wired subnet forms its own independent subnetwork.
- Direct links between stations can be utilized within the mesh network.
- Problems with this approach are listed in the issues section, following.

Virtual Bridge




- Each AP and its station bridges become a single bridge.
- Note that each station/bridge is split into two parts, the wired and wireless part, with a simulated wire connecting them.
- One device (the station bridge that is a station to both AP1 and AP2) required two separate wireless parts.

Virtual Bridge

- This offers a general solution to the problem.
- We will assume that the boundaries of the virtual bridges do not contain the whole of each non-AP station / bridge element and their wired connections.
 - This would not be possible for the case of a non-AP station bridge that is connected to two different APs.
 - This would greatly amplify the “hidden costs” problem described in the next section; some pairs of wires would have 0 cost across the virtual bridge, and some very high cost.
- Problems with this approach are listed in the issues section, following.

802.11 shared medium

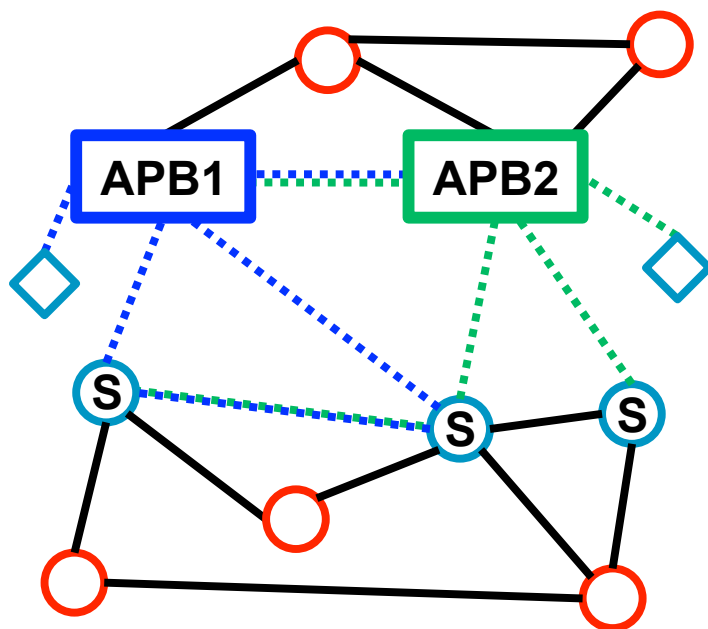


- Each AP and all of its stations become a single shared medium (fat yellow coax), as seen by the wired bridges.
- Each AP uses its bridge knowledge to optimize forwarding through the 802.11 medium, rather than broadcasting every frame.
- Direct station-station or AP-AP links have to be modeled independently.

802.11 shared medium

- This solution may be the least-difficult to standardize for bridges employing the Multiple Spanning Tree Protocol. The techniques are well-known, and require relatively little coordination among the non-AP stations.
- Problems with this approach are listed in the issues section, following.

Set of point-to-point links



- The Access Points and their co-resident bridging functions become integrated AP bridges (APBs).
- Devices with non-AP station capability(ies) and wired connections become “non-AP station bridges”.

Set of point-to-point links

- This is the most general solution to the problem.
- Problems with this approach are listed in the issues section, following.

Comparison

- **In this author's opinion, the “set of point-to-point links” model is the clear winner, and should be pursued:**
 - It is perfectly compatible with any bridging or routing topology control protocol; it has the least impact on any existing protocol.
 - Full visibility by all bridges of all parts of the network yields the optimum least-cost forwarding decisions.
 - Full visibility allows protocols (e.g. ISIS) to make use of other criteria than least-cost (e.g. reliability) in making forwarding decisions.
 - Point-to-point connectivity converges faster than shared medium connectivity for any topology control protocol.
 - In the absence of manual intervention, full visibility results in much lower overall convergence times than any network involving mutually opaque clouds.
 - The “virtual bridge” model cannot make the internal costs of an 802.11 medium visible to other bridges in the network.
 - Only the point-to-point model supports the full range of existing and future 802.1 protocols and features with little or no need for extra standards work (VLAN pruning, multicast address pruning, traffic-shaped queuing, time-scheduled queuing, congestion notification, ...)

Issues and potential solutions

Issues and solutions*

- There are a number of issues that must be addressed by the standard(s) that define how bridges work with both wired and wireless media.
- Some issues are peculiar to the model employed (see previous section), and some arise for all models.
- This section will give a brief summary of each issue, and list one or more viable solutions.
- On the next slide is a matrix, showing which models must deal with which issues. For some models, the list of issues depends upon the control protocol used (Multiple Spanning Tree Protocol vs. Shortest Path Bridging).

*** You may wish to skip down to the last two slides; this gets pretty deep!**

Model vs. Issues

Issue:	Model:	Mesh + Wired	Virtual Bridge	Shared medium	Pt-to-pt media
1. Network interactions		issue	OK	OK	OK
4. VLAN tag		issue	issue	issue	issue
5. Multicast distribution		issue	issue	OK	issue
6. Hidden costs		issue	issue	issue	OK
7. Shared medium convergence		OK	OK	issue	OK
8. Unreliable connections		?	issue	issue	issue
9. Mixed broadcasts		?	issue	issue	issue

Model/protocol: Issue:	Mesh + wired	VB, SM, or PPM protocol = MSTP	VB, SM, or PPM protocol = SPB
2. Spanning tree	issue	issue	OK
3. Reflected frames	OK	issue	issue

Issue 1: Network Interactions

- If the mesh + wired model is used, the wireless network is an opaque entity to the wired networks, and vice-versa. This causes several problems:
 - There must be some means of preventing loops in the whole network. If there are two kinds of clouds, then four clouds can create a forwarding loop.
 - If there are multiple connections between two clouds, some means of identifying them and restricting their use is necessary.
 - Making a cloud opaque means that multi-cloud path selection will often be sub-optimal, due simply to a lack of detailed path cost information.

Solution 1: Network Interactions

- There are known solutions for **loop prevention** over connected subclouds. All have problems with convergence time and/or stability, and are inferior to a single integrated cloud solution unless scaling and/or ownership issues force the use of subclouds:
 - A global cloud-based control protocol;
 - Manual configuration and the consequent loss of automatic recovery;
 - Making one cloud type transparent to another's control messages.

Solution 1: Network Interactions


- There are known solutions for **load sharing** multiple connections between connected subclouds. There is an open PAR for this task (P802.1AX-REV). However:
 - Plug-and-play identification of situations requires either configuration, passing one network's control packets over the other network, or an intimate knowledge of the other subcloud's control protocol. The last alternative, where possible, is known to work best.
 - If one or the other network uses MAC address learning, then the connections can be load-shared only by VLAN. A plug-and-play solution to this issue would be required.

Solution 1: Network Interactions

- There are no known solutions for the problem of **suboptimal routing** caused by lack of knowledge of the details of an adjacent cloud.
- To phrase this another way, one would like to load-share connections among clouds, or to use another cloud to interconnect two nodes of the same cloud in preference to a torturous in-cloud path. But, without knowledge of the details of the other cloud and protocols to support them, this is not possible.
- Any solution to this problem is necessarily equivalent to erasing the boundaries among subclouds.

Issue 2: Spanning Tree


- As has been known for a long time, spanning tree has issues in simple networks with links of widely disparate data rates. For example:



- This diagram illustrates the problem in the home.

Solution 2: Shortest Path Bridging

- **IEEE Std 802.1aq-2012** solves this problem by using the ISIS (Intermediate System – Intermediate System) protocol to make it possible to **utilize all links**.



- Note that this is a **software solution**; the existing 802.1Q data plane (used by STP) remains unchanged.


Issue 3: Reflected frames

- A non-AP station uses its own MAC address as both the Ethernet source address and the transmitter address; hence three addresses (that plus receiver address and Ethernet destination address) are sufficient. If the non-AP station is a bridge, the bridge's own MAC address cannot be included in the three addresses.




Issue 3: Reflected frames

- **CASE 1:** Suppose a non-AP station/bridge **B** is forwarding data for attached wired device **X**.
- Suppose **X** sends a frame (a broadcast, for example) up through bridge **B**.




CASE 1

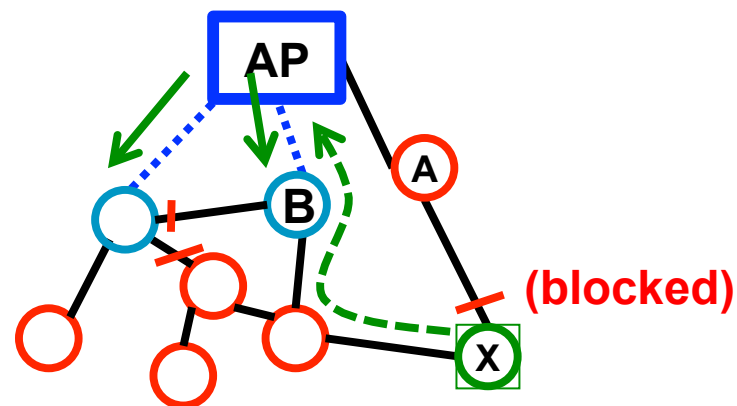


Issue 3: Reflected frames

- The Access Point reflects the frame back down to all of the AP's stations, including **X**.
- Bridge **B** needs to **discard** the frame. (Its portion of the network has already seen it.)




CASE 1

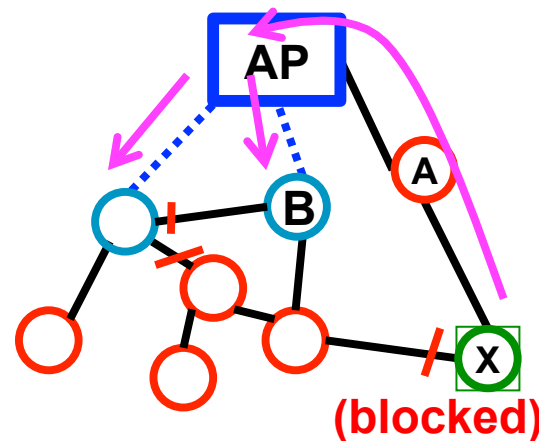


Issue 3: Reflected frames

- **CASE 2:** Suppose instead, that **X** has moved, and transmits that same broadcast frame.
- The Access Point transmits the broadcast to all of its stations, including bridge **B**.





CASE 2



Issue 3: Reflected frames

- The problem is that Case 1 and Case 2 result in **exactly the same frame** from the AP to bridge **B**.
- Bridge **B** doesn't know whether to **discard** the frame (Case 1) or to **forward it and learn X's** new location (Case 2).




Solution 3: Four address format

- If all four addresses are used, there is no problem:

Both UP




CASE 1




CASE 2



CASE 1



CASE 2



Solution 3:

Four address format

- The “~B” address is the “not B” address. This could be, for example, a fixed multicast address with the low-order bits set to station/bridge **B**’s association identifier. A station/bridge accepts all frames in this set of addresses **except** the one with its own association identifier.
- This is not the only solution; there are other viable solutions to the generalized problem of recognizing whether a given frame is a reflection of one’s own previous transmission. (Note, however, that neither “remember who is behind you,” or “remember what frames you sent, recently” are viable solutions.)
- **Such broadcasts are used widely** to signal the new location of a migratory 802.11 station to bridges.

Issue 4: VLAN tag


- Native 802.11 frames use the IEEE 802.2 LLC format.
- Therefore, adding a VLAN tag to a frame requires adding 10 bytes to the frame (8-byte SNAP encoding + 2 byte payload).
- Between bridges, a large fraction (often, all) of the traffic carries VLAN tags. When both the AP and the non-AP station are bridges, the 6 extra bytes are a problem.

Solution 4: VLAN tag (and learning!)

- The IEEE Std 802.11n-2009 A-MSDU format carries the extra addresses needed to fix the learning/discarding Issue 3 (more than necessary, actually), and uses a 2-byte EtherType to signal the presence of the VLAN tag to fix Issue 4.
- So, **802.11n could be the fix for both Issues 3 and 4.**

Issue 5: 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**?




Solution 5: Multicast distribution

- One solution would be to extend/modify MSTP and/or Shortest Path Bridging to provide a handshake informing bridge **C** that the link is blocked on the AP side.
- Another solution would be to extend the set of receiver MAC addresses mentioned in Issue 3, in frames sent by the AP, to specify sets of bridge / stations beyond “not **B**”. (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 a protocol for distributing maps of vectors of stations to 24-bit IDs is necessary.

Issue 6: Hidden costs

- If a wireless medium is abstracted into either an opaque network cloud (mesh + wired model), a single virtual bridge (virtual bridge model), or a single shared medium (shared medium model) there are issues with the actual vs. the apparent cost to cross the abstracted object.



- In the trivial case, in the absence of direct connections between non-AP stations, the cost of moving a frame between the X and Y is typically half the cost of moving a frame between Y and Z. This cost difference cannot be factored into forwarding decisions if the links are part of an abstracted object.

Solution 6: 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.
- There is no solution to this issue; either the inefficiencies must be accepted, or the model rejected.

Issue 7: Shared medium convergence

- The Multiple Spanning Tree Protocol (MSTP) converges much more slowly (several seconds) when connected to a shared medium, versus the sub-second convergence possible with point-to-point media.
- The Shortest Path Bridging protocol (SPB) does not, at present, support shared media.

Solution 7: Shared medium convergence

- It would be possible to add protocol elements to MSTP to fix this issue.
 - In this case, access to the shared medium is controlled by a single entity (the AP). This opens the possibility of a bridge connected to this medium to obtain a list of other bridges, so that a multi-point agreement among the connected bridges could replace the point-to-point agreement now used by MSTP.
 - This solution, however, brings one perilously close to the point-to-point model, which does not have this problem.
- It would (presumably) be possible to add the ability to use shared media to SPB.
 - Details To Be Determined.

Issue 8: Unreliable connections

- Wireless media are inherently less dependable than wired media.
- The effective speed and availability of a given connection can vary over a timescale so short that it is very undesirable to export that speed/availability information to the rest of the network; the result could easily be that the whole network becomes unstable.

Solution 8: Unreliable connections

- It is conceivable that labels like, “unreliable” or “intermittent” could be applied to certain connections and be handled by a general purpose protocol. However, several such attempts have failed over history.
- It is probably a better approach for the AP to filter what links are used and/or exported. Links between bridges that vary on a timescale that would damage the stability of the network are not used, or are given a very high cost.

Issue 9: Mixed broadcasts

- We must assume that there will always be legacy stations (that is, stations that do not understand non-AP station bridges) will be attached to Access Points.
- If a solution to any of the above issues (such as Issue 3 reflected frames) involves a novel encapsulation, then a broadcast or multicast from the AP to stations that include both bridges and legacy stations are a problem; the bridge requires the encapsulation, and the legacy stations cannot support it.

Solution 9: Mixed broadcasts

- If an encoding scheme can be found that legacy stations support, the problem goes away.
- At the worst, the AP must send broadcasts twice, once with the encapsulation once with the legacy 3-address format. Legacy stations ignore what they do not understand, non-AP station bridges ignore the legacy format.
- At best, we create a means for the AP and all stations, whether bridges or not, exchange capabilities, so that the AP knows whether a given broadcast or multicast needs to be transmitted twice.

Additional media

- There are at least three other media types that are of interest in this space:
 - IEEE 802.3 Ethernet over Passive Optical Networks (EPON),
 - Multimedia over Coax Alliance (MoCA),
 - IEEE 1901 Broadband (Ethernet) over power line.
- Some of these media share with 802.11 the following characteristics:
 - A master node is in charge of admitting non-master stations to the medium.
 - The master node has a set of point-to-point connections to every non-master.
 - The master has the capability of sending a broadcast/multicast message to multiple non-master nodes.

Plan of action

Plan of action. Given that:

- The use cases justify integrating 802.11 with 802.3 media into 802.1 bridges.
- There are at least four viable models by which this can be accomplished.
- All models require additional standards work by 802.1 and 802.11.
- The amount of work to be done depends, at least in part, on what tradeoffs we wish to make regarding capabilities versus complexity.
- It may not be difficult to include MoCA, 1901, and 802.3 EPON into bridging, if the use cases exist.

This author proposes that:

- 802.1 initiate a PAR (target for approval at November plenary) for incorporating additional media into 802.1Q, specifically:
 - At least 802.11, in the sense of any AP or non-AP station can be part of a bridge port.
 - Multimedia over Coax Alliance (MoCA), IEEE 1901 Ethernet over power, and IEEE 802.1 Ethernet over Passive Optical Networks (EPON), if interest and capable volunteers are found.
- 802.11 initiate a PAR (target for approval at November or March plenary) to support the 802.1 PAR.
- We establish a series of regular weekly virtual meetings to refine these PARs, and to enable 802.1 and (at least) 802.11 to work together. Others would be welcome.