802.11 AP Bridges and MSTP

Two MSTP problems are encountered when non-AP 802.11 stations can be bridges

Rev. 1

Norman Finn

nfinn@cisco.com

References

- This presentation is available at: http://www.ieee802.org/1/files/public/docs2008/new-nfinn-ap-bridge-mstp-1108-v1
- For a simple description of the station bridge problem and the **4-address solution** on which this presentation is based, see: http://www.ieee802.org/1/files/public/docs2008/avb-nfinn-802-11-bridging-0308-v3.pdf
- For a more complete description of the station bridge problem and other possible solutions, see: http://www.ieee802.org/1/files/public/docs2008/avb-nfinn-802-11-bridging-0308-v2.pdf

Executive summary

- When an 802.11 Access Point is also a bridge, and connects to 802.11 non-AP stations that are also bridges, there is a problem when the virtual point-to-point AP-station links are blocked at the AP's end.
- In this situation it is tricky for the AP to optimize the distribution of multicasts and unknown unicasts to the station bridges over the wireless medium.
- The best solution seems to be to add a "**Proxy blocking**" trick to RSTP/MSTP.
- There is also a problem with **sub-optimal paths** and Access Points.
- **MSTP-based solutions are given here to both problems, but alternatives are solicited.**

Proxy Blocking

Reference diagram

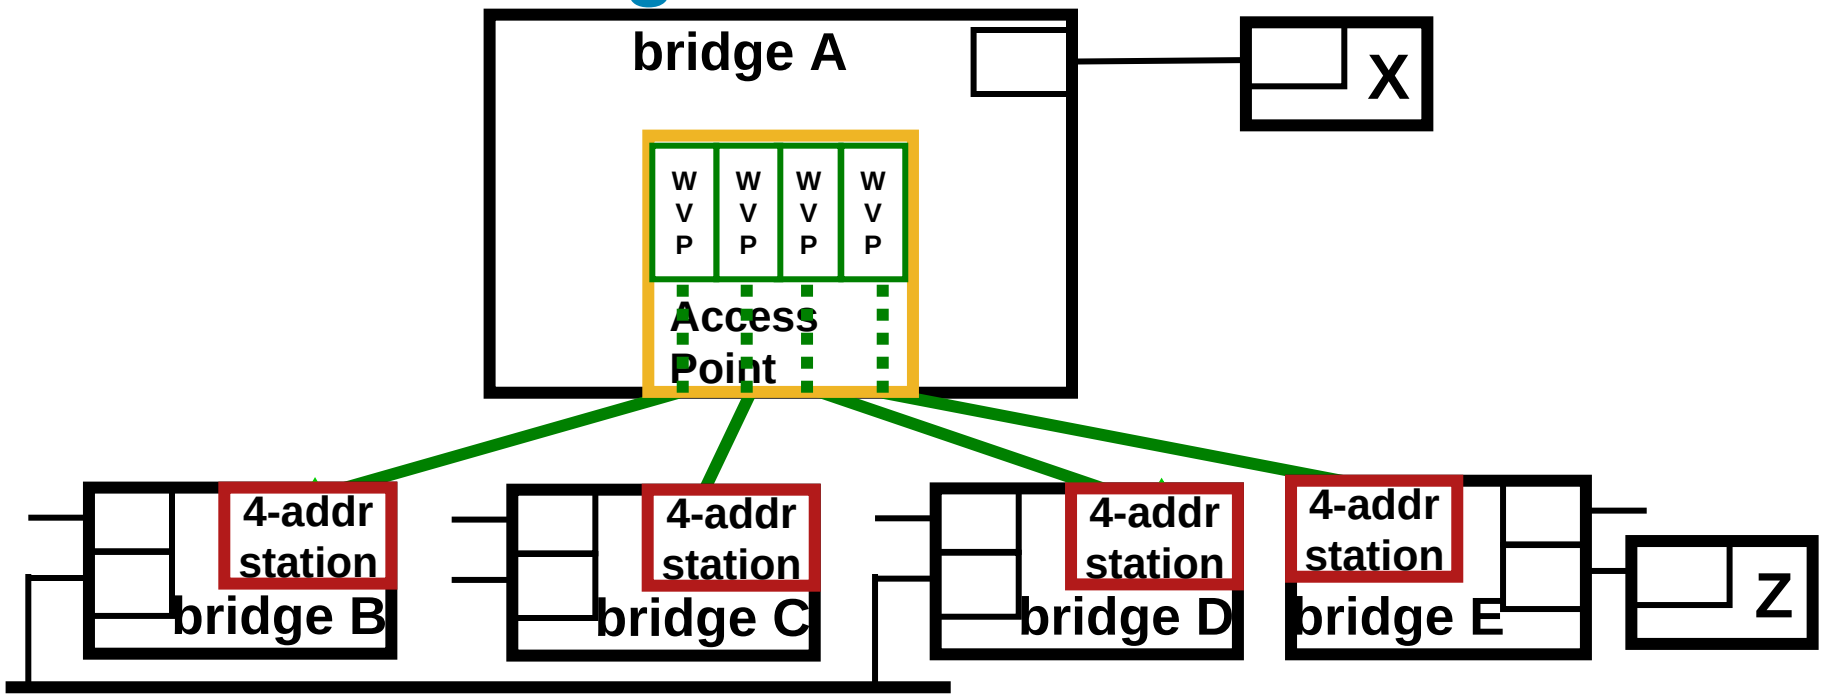

- Access Point is integrated into the Bridge A.
- “Wireless Virtual Ports” are presented to the bridge’s relay function.

What do we mean by, "An 802.11 station bridge is not supported"?

- IEEE Std. 802 leaves it up to each medium whether frames transmitted on that medium are reflected back and received at the source.
- IEEE Std. 802.1D and 802.1Q make it clear that a bridge does not work on any medium that reflects frames back to the source.
- An IEEE 802.11 wireless Access Point reflects frames (with a time delay!) back to the source non-AP station.
- On this reflecting medium, a station bridge cannot distinguish between frames it should discard as reflections, and frames from which it should learn.

ASSUMPTIONS

- To solve that problem, and to handle unknown unicasts, MVRP, GVRP, and new protocols now in progress, the Access Point cannot be at “arm’s length” from the bridge.

  The Access Point is a bridge.

  The bridge and the access point are integrated more closely than having a wire (whether virtual or real) connecting them.

- Each non-AP station associated with the Access Point, whether a bridge or not, looks like it’s connected to a separate port on the bridge with a **point-to-point LAN**.
- The Access Point continues to manage the peculiar aspects of the wireless media.
- The AP can optimize sending so that multiple copies of one frame on multiple “ports” can be sent only once.

Proxy Blocking

RSTP blocks a port

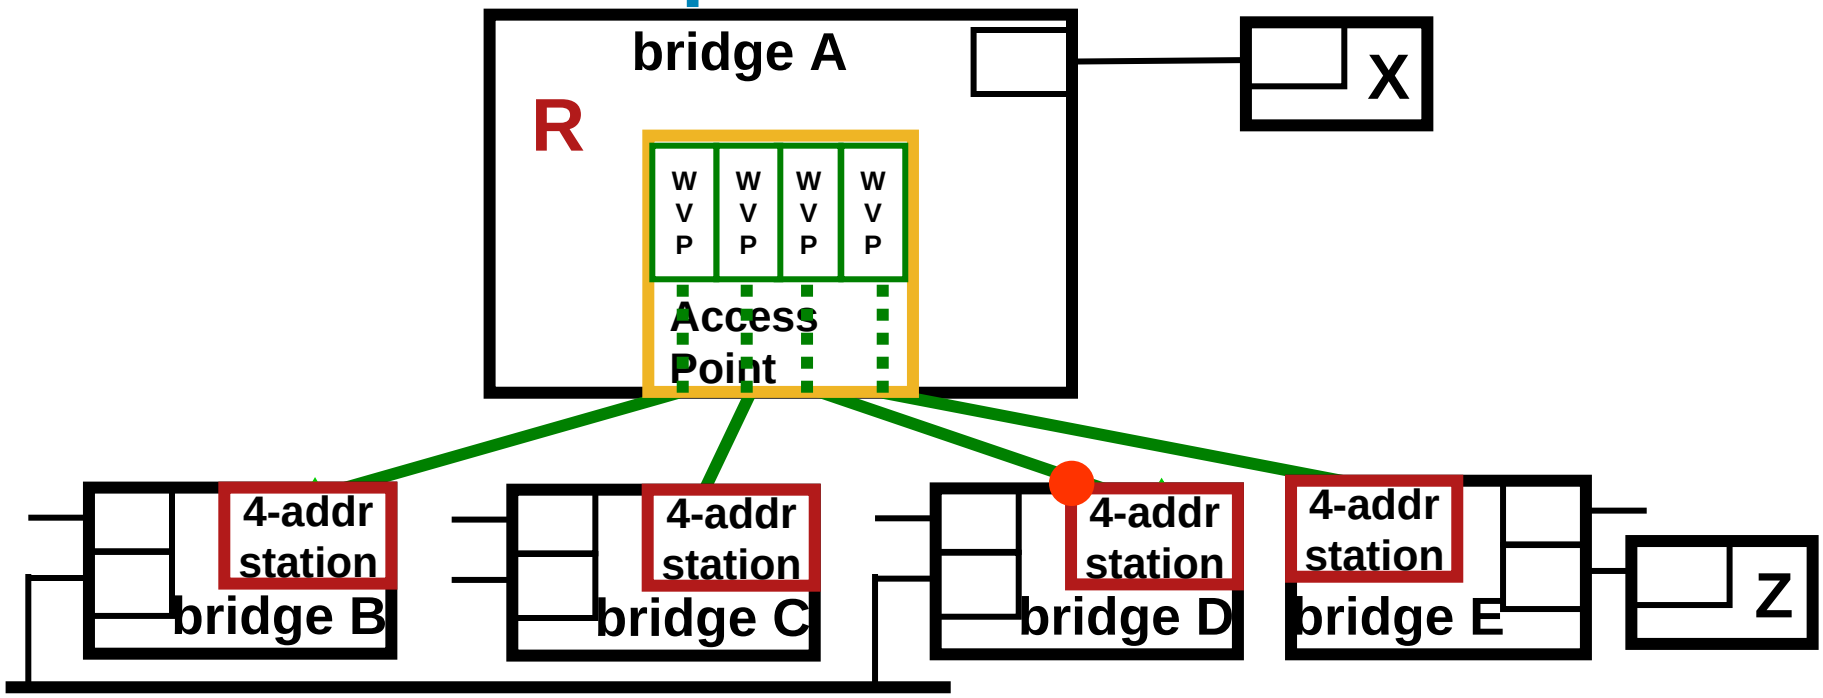

- Suppose we're running **RSTP**. Bridge A is the **R**oot.
- Suppose that **Bridge D's** MSTP is configured to **block** its wireless port to prevent loops.

Everything is easy, so far

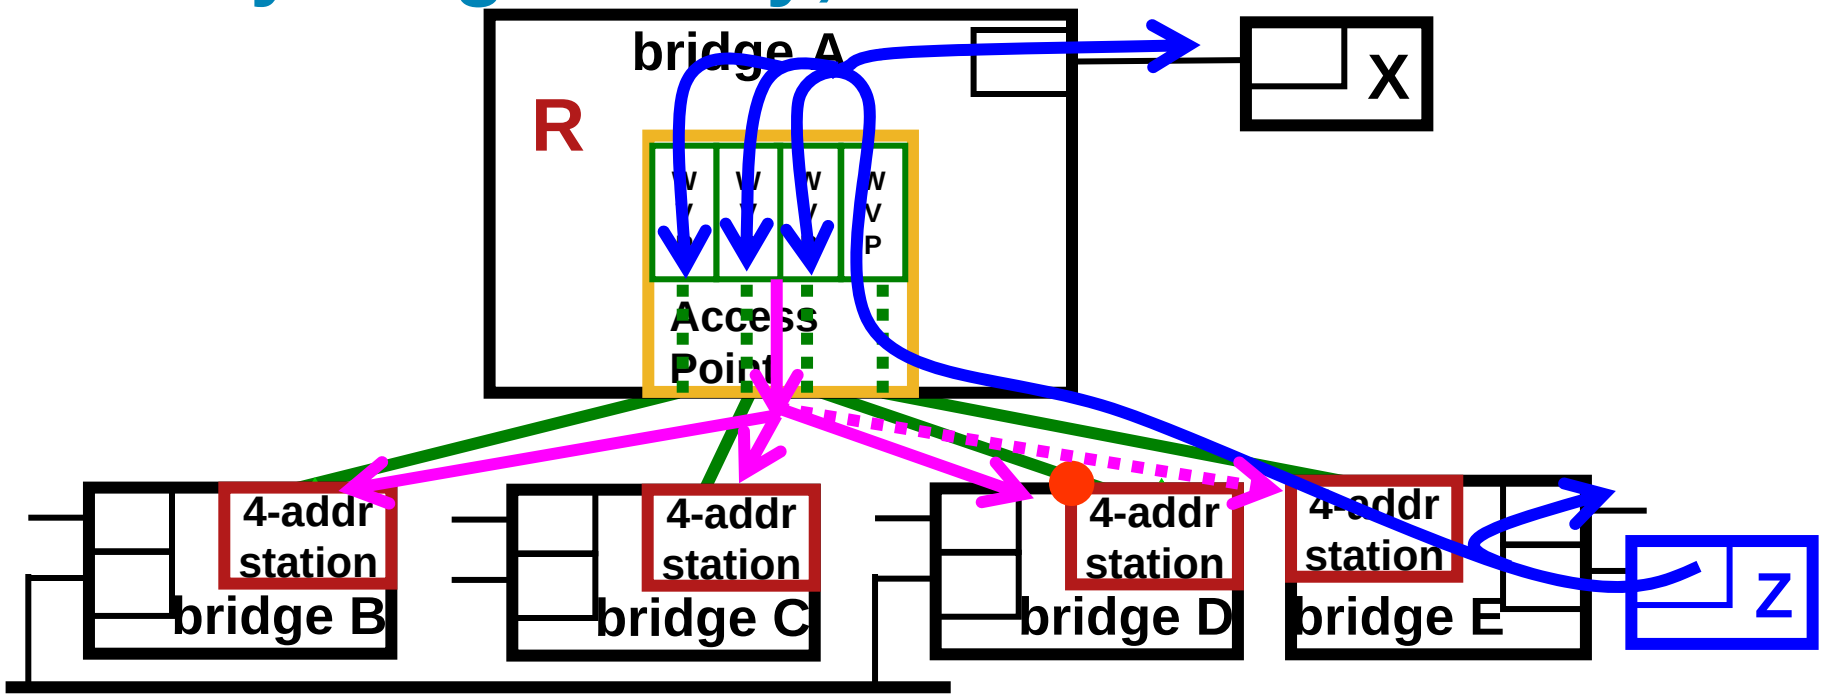

- **End station Z** sends a broadcast. Bridge A returns it on the other three virtual ports. The **Access Point** wants to send **just one copy**, to economize on wireless air time.
- AP uses the "**Not E**" Receiver Address, so Bridge E ignores it. Bridge D ignores it because port is blocked.

A little harder

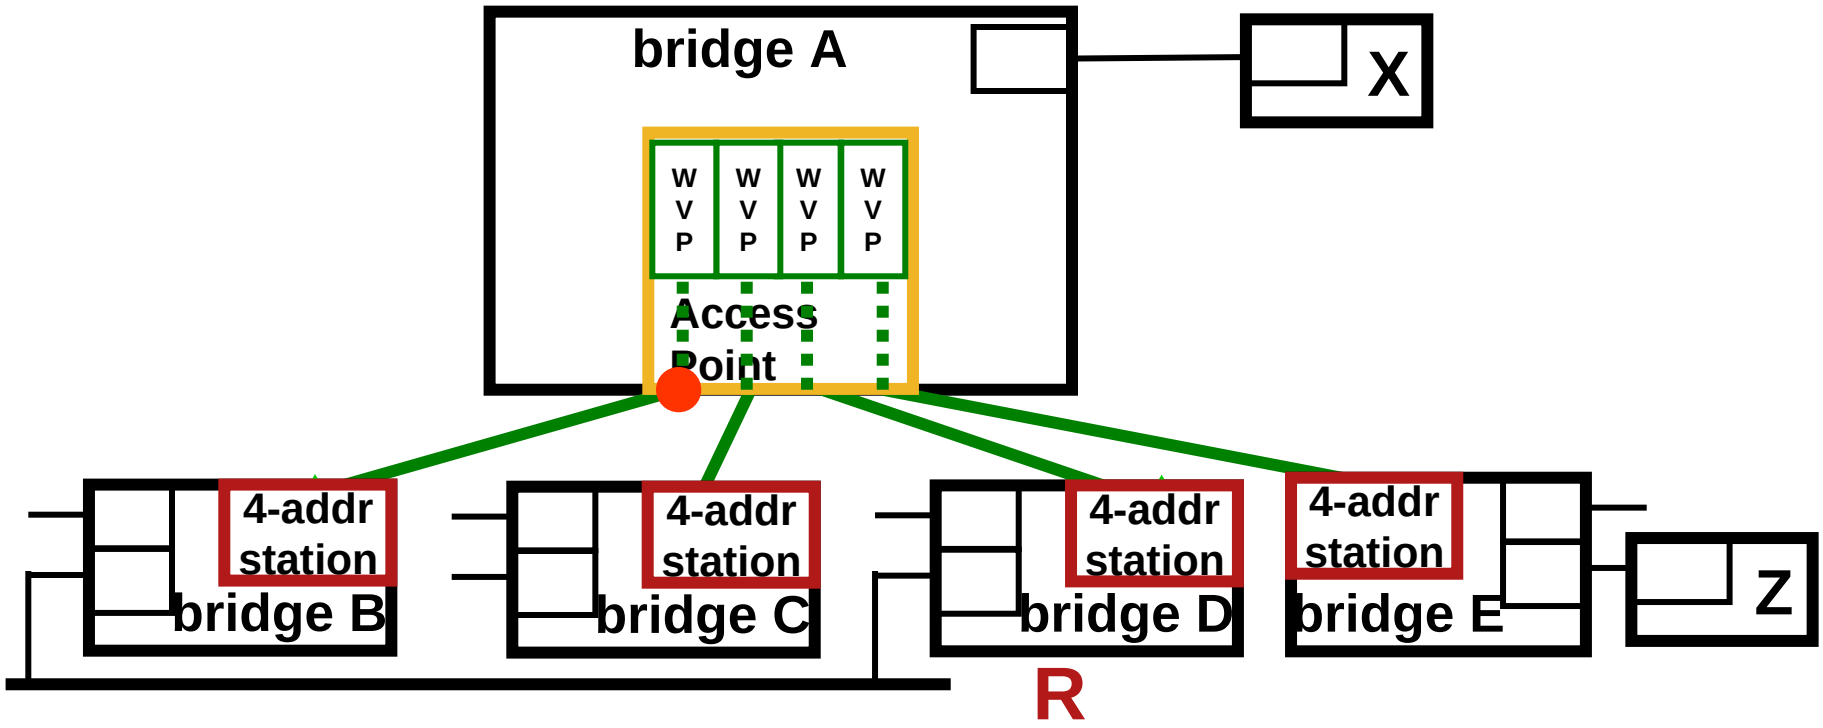

- But what if Bridge D is the **R**oot, and Bridge A must **block** one the **virtual port** to Bridge **B**?

A little harder

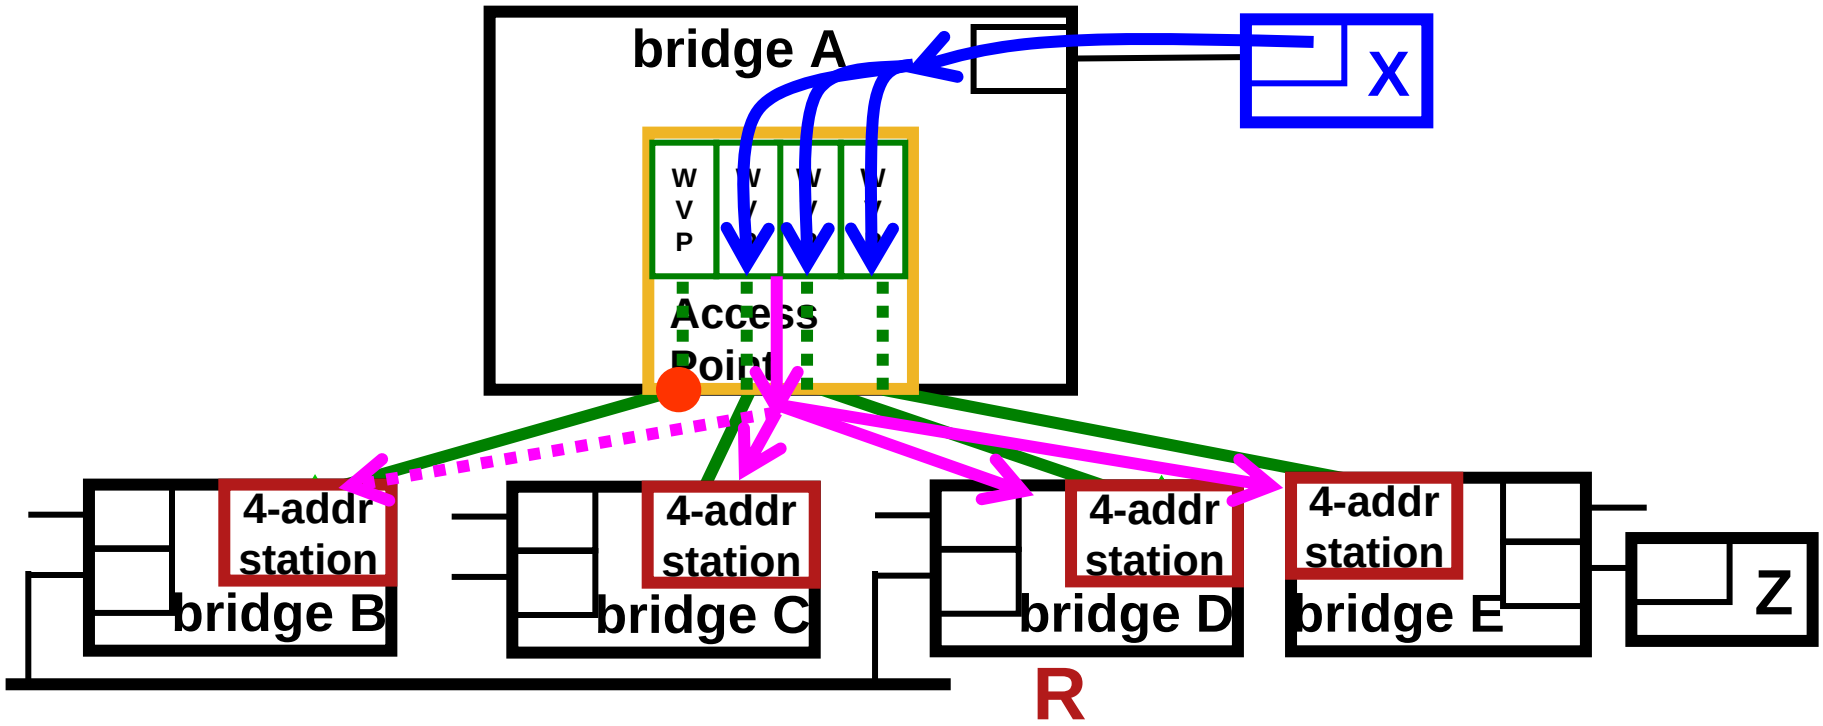

- But what if Bridge D is the **R**oot, and Bridge A must **block** one the **virtual port** to Bridge **B**?
- If **end station X** sends a broadcast, there is no problem; the Access Point uses the "**Not B**" Receiver Address, so Bridge B ignores the frame.

Ouch!

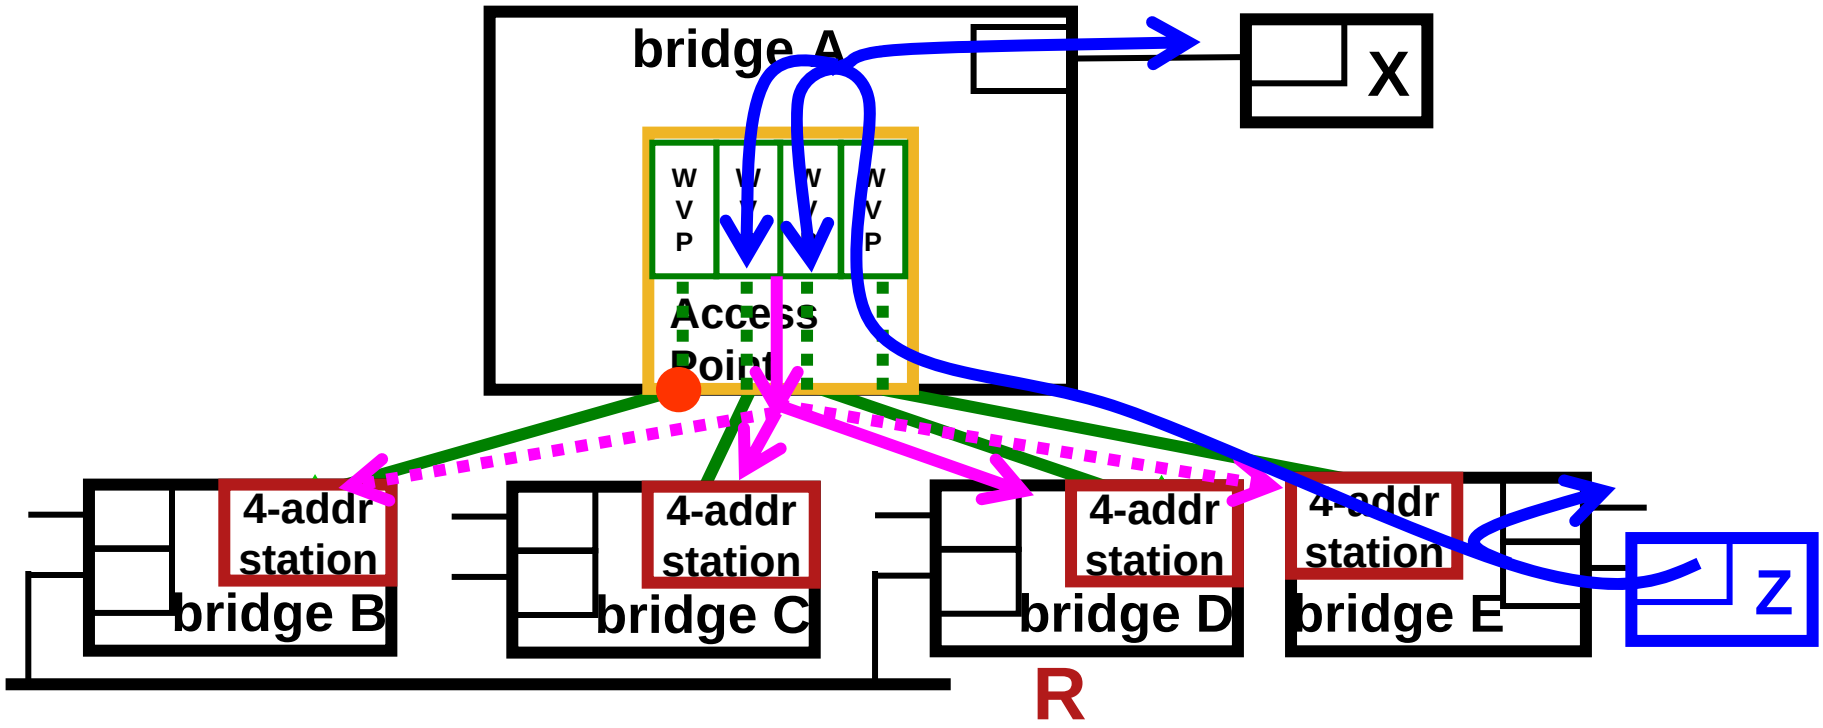

- But, if **end station Z** sends the broadcast, what does the Access Point use for a Receiver Address?
- Over time, the Access Point may need a repertoire of almost $2^{\text{number of station bridges}}$ Receiver Addresses for all possible combinations of blocked virtual AP ports!

Solution 1: Unicast Receiver Address

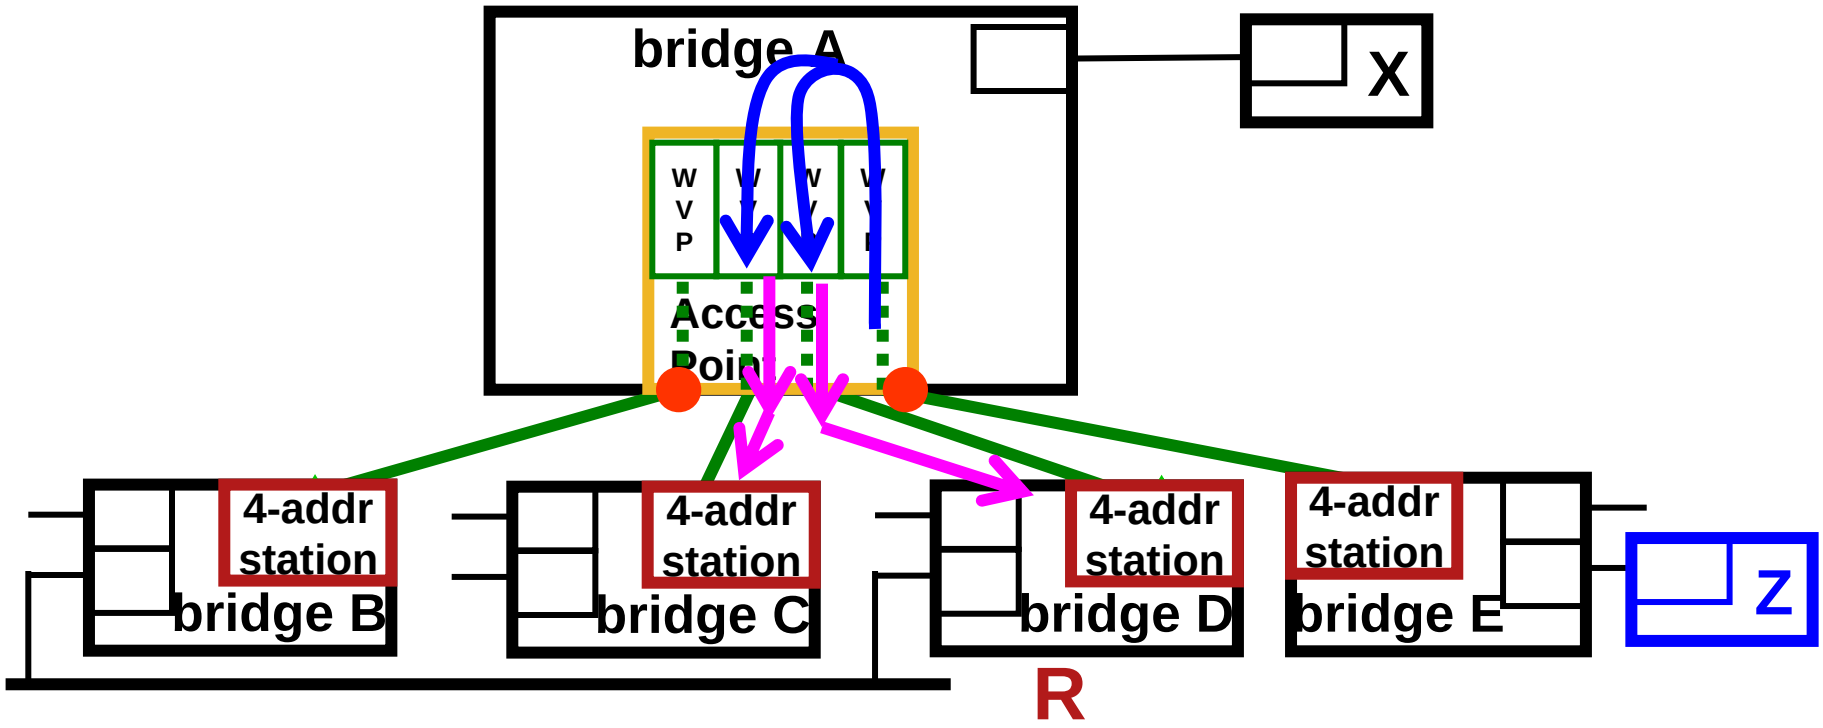

- The Access point could send **two separate frames**, one with **C** as the Receiver Address, and one with **D**.
- But, this takes up **more air time**, especially when there are more than a very few station bridges.

Solution 2: Bit vector in Receiver Address

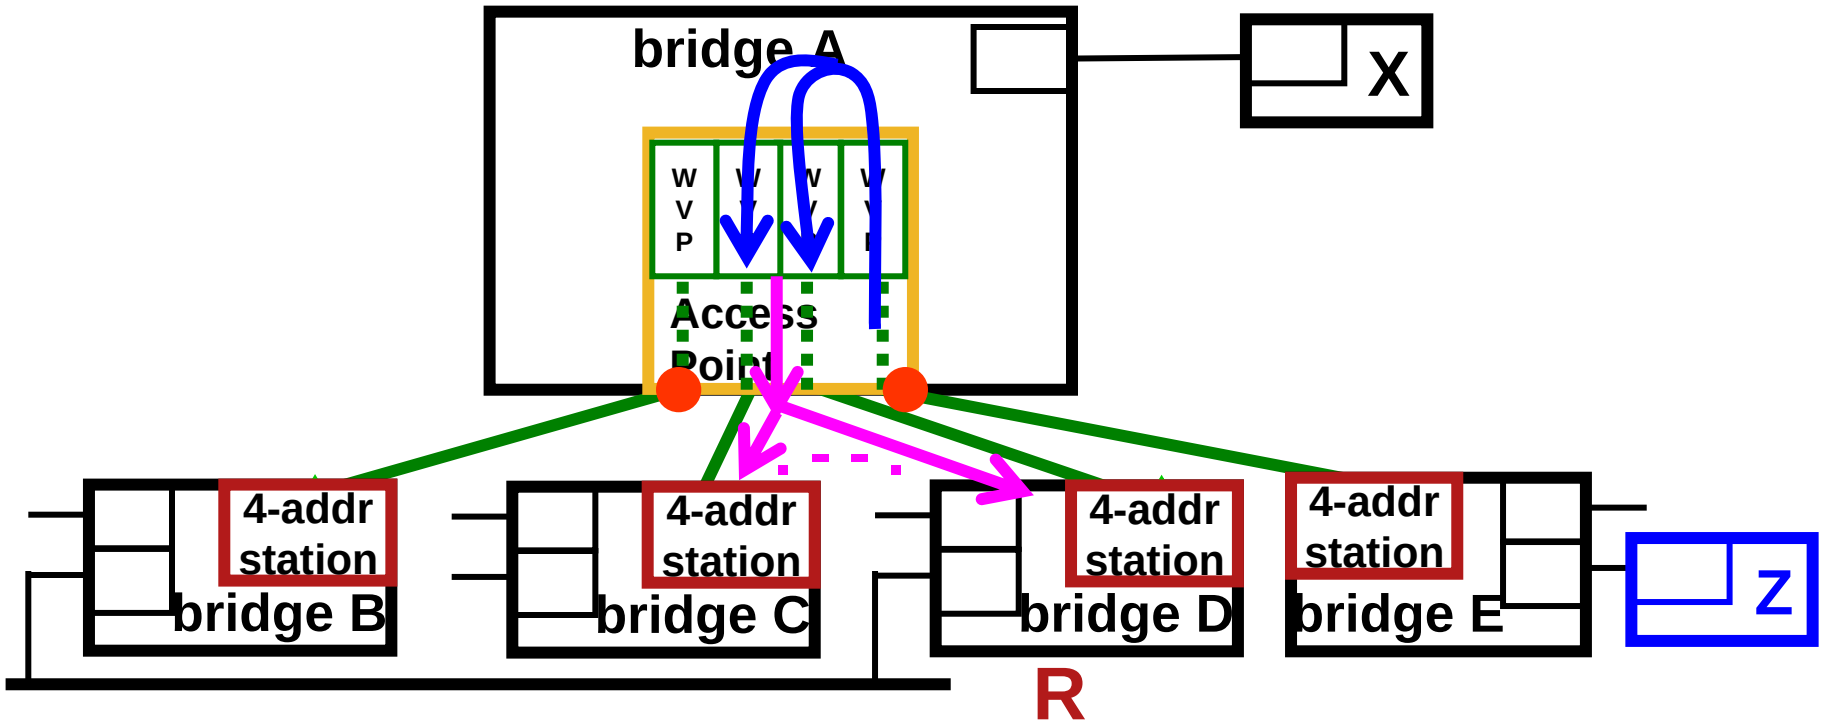

- If there are 24 or fewer station bridges, and each is assigned one bit of the lower half of a reserved Receiver Address, then the AP has a **bit vector** to control to which stations a frame is addressed.
- That is probably sufficient for 802.11 (the bit assignments could be carried in BPDUs), but **not general**.

Solution 3: Bit vector distribution protocol

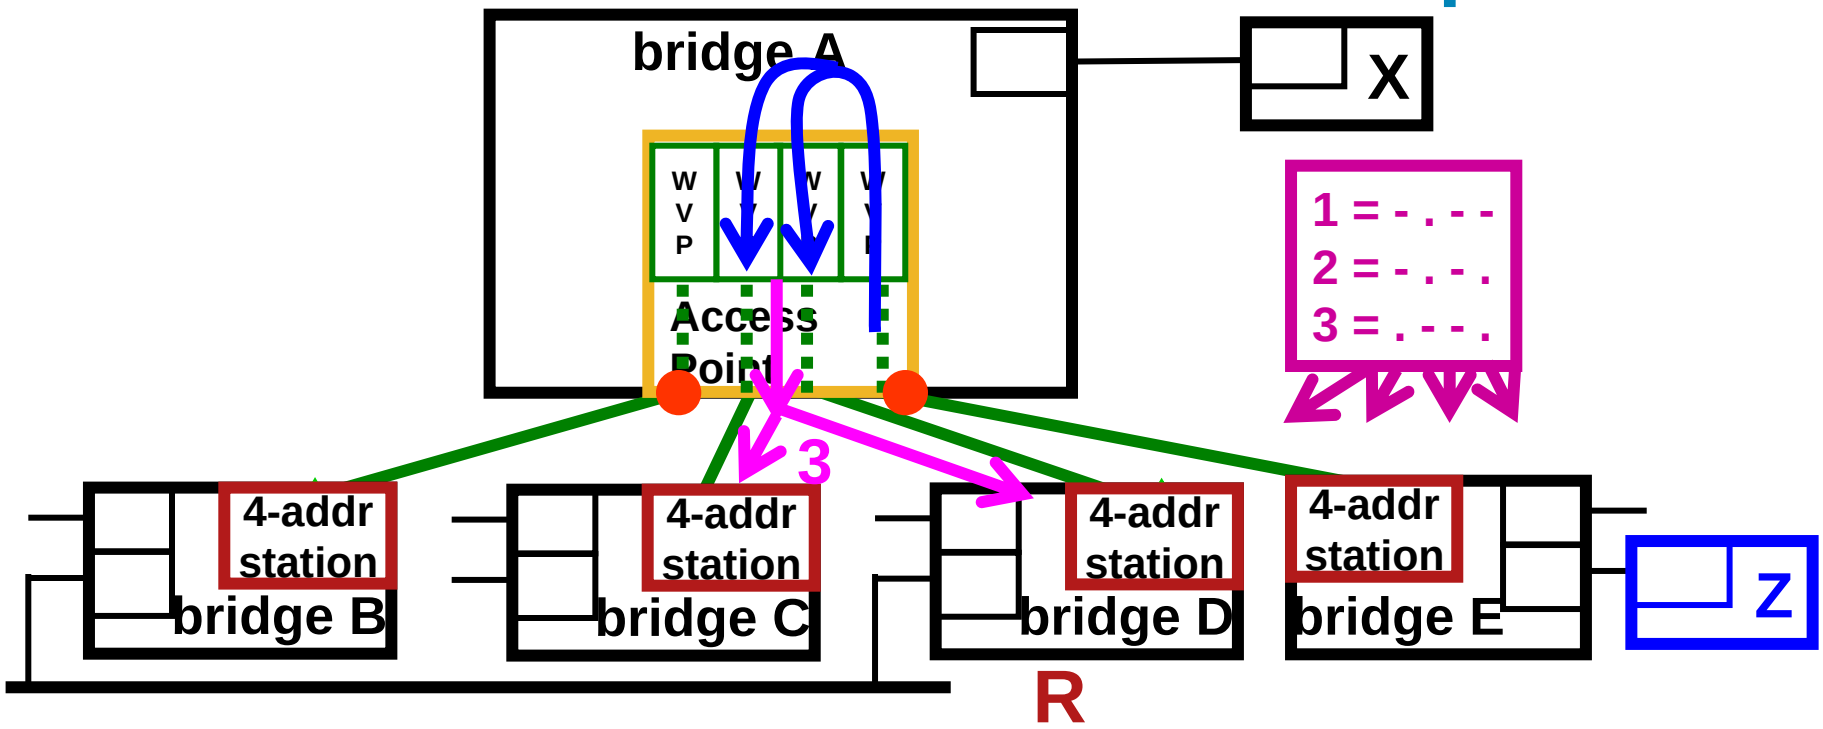

- The AP can distribute a **list of bit vectors** via a **protocol**; the bit vector index goes in the **Receiver Address**.
- This is a one-source multiple-destination protocol with (potentially) lots of data to move. Much more general, but **rather complex**.

Solution 4: MSTP handshake

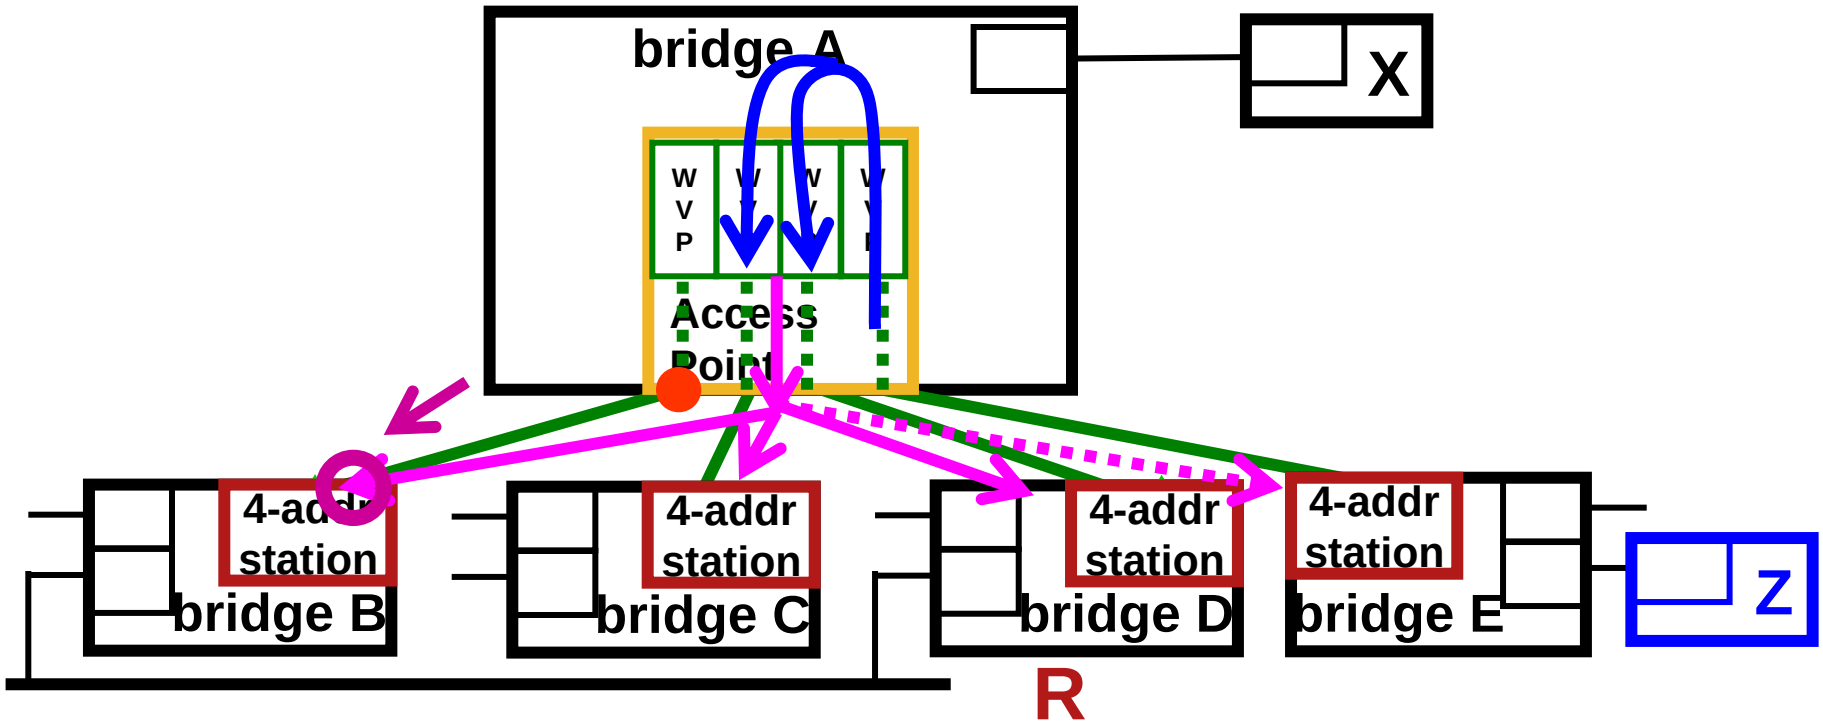

- The Bridge (not the AP) could **distribute knowledge** of its **blocked port** to the Bridge B, and let **B** do the **blocking by proxy**, while the AP continues to use the **broadcast** or the "**Not *P***" Receiver Address, as usual.

Signaling proxy blocking

- Only a station bridge that is a Designated Bridge might have to block, but it might not, also.
- The AP bridge can send unsolicited BPDUs that indicate that its end of the virtual point-to-point wireless link is blocked (Role == Alternate/Backup).
- **But, the AP's bridge needs to know that the station bridge knows that the AP's end of the port is blocked (!)** before the AP can send a broadcast or *Not P*-cast.

  Otherwise, the station bridge might not yet be proxy blocking.

  Then, you'd have a loop.

- **How can the station bridge signal, "OK, I'm proxy blocking"?**

One method to signal proxy blocking (Other suggestions are welcome, of course!)

- The Forwarding Flag in the CSTI and MSTI flags byte is, at present, ignored on receipt.
- A non-AP station bridge, on its wireless port(s), could:

  When the port role becomes Designated, not turn on forwarding (and not forward data) until it hears from the AP bridge whether it is in the Root role or Alternate/Backup role.

  If the last BPDU received by the non-AP station bridge from the AP bridge indicates Alternate/Backup, the non-AP station bridge does not forward data, and does not set the forwarding bit.

- An AP bridge could:

  If the Root port facing a station bridge becomes an Alternate port, the AP bridge sends BPDUs, and does not send a Receiver Address that the station bridge might accept, until the station bridge's Forwarding Flag goes false.

  The AP uses unicast Receiver Addresses until the Forwarding Flag is 0.

Alternatives to proxy blocking signaling?

- Whatever method is used to signal proxy blocking, it should fit in an MSTI configuration message in the MSTP BPDU.

  There are 8 free bits in the priority bytes, but one hopes that the Flags byte is sufficient.

  Of course, other clever people may think up other clever uses for these bits to facilitate other clever features.

- And, it must fit in the RSTP information, as well.

  This is easier, because the Master Flag bit (8) is unused in RSTP.

- But, the Forwarding Flag seems a natural.

- We could also invent an additional protocol, but if the MSTP BPDU can hold it ...

Sub-optimal paths

Wrong link is blocked

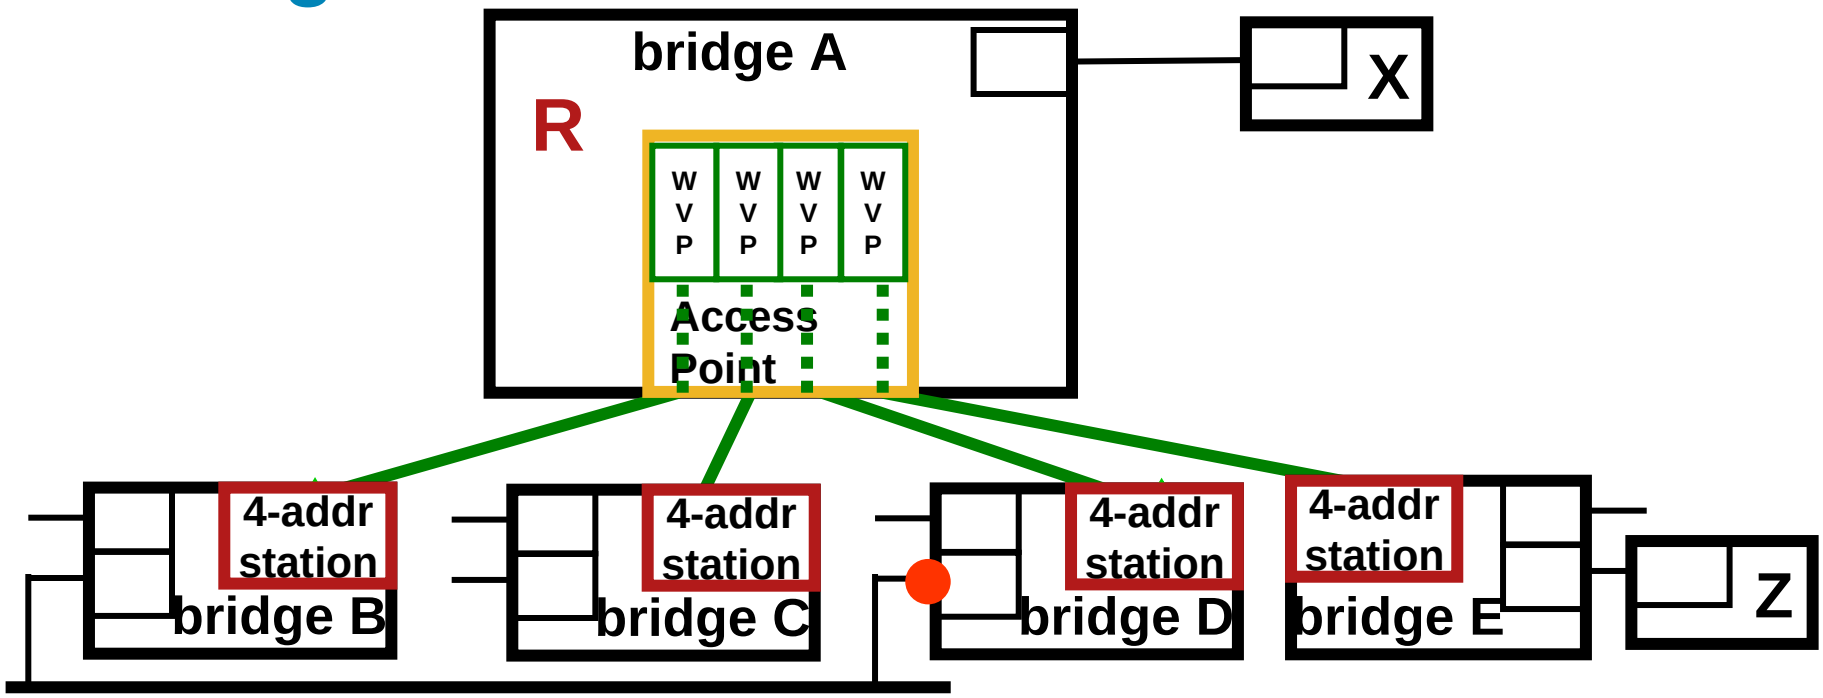

- Although Bridge D could be configured to block its wireless port, in the original example, for out-of-the-box bridges with a default configuration, **Bridge D would block its wired port**, even if that were a 10 Gbps wired link! This is almost always a sub-optimal choice.

Regions could be the answer

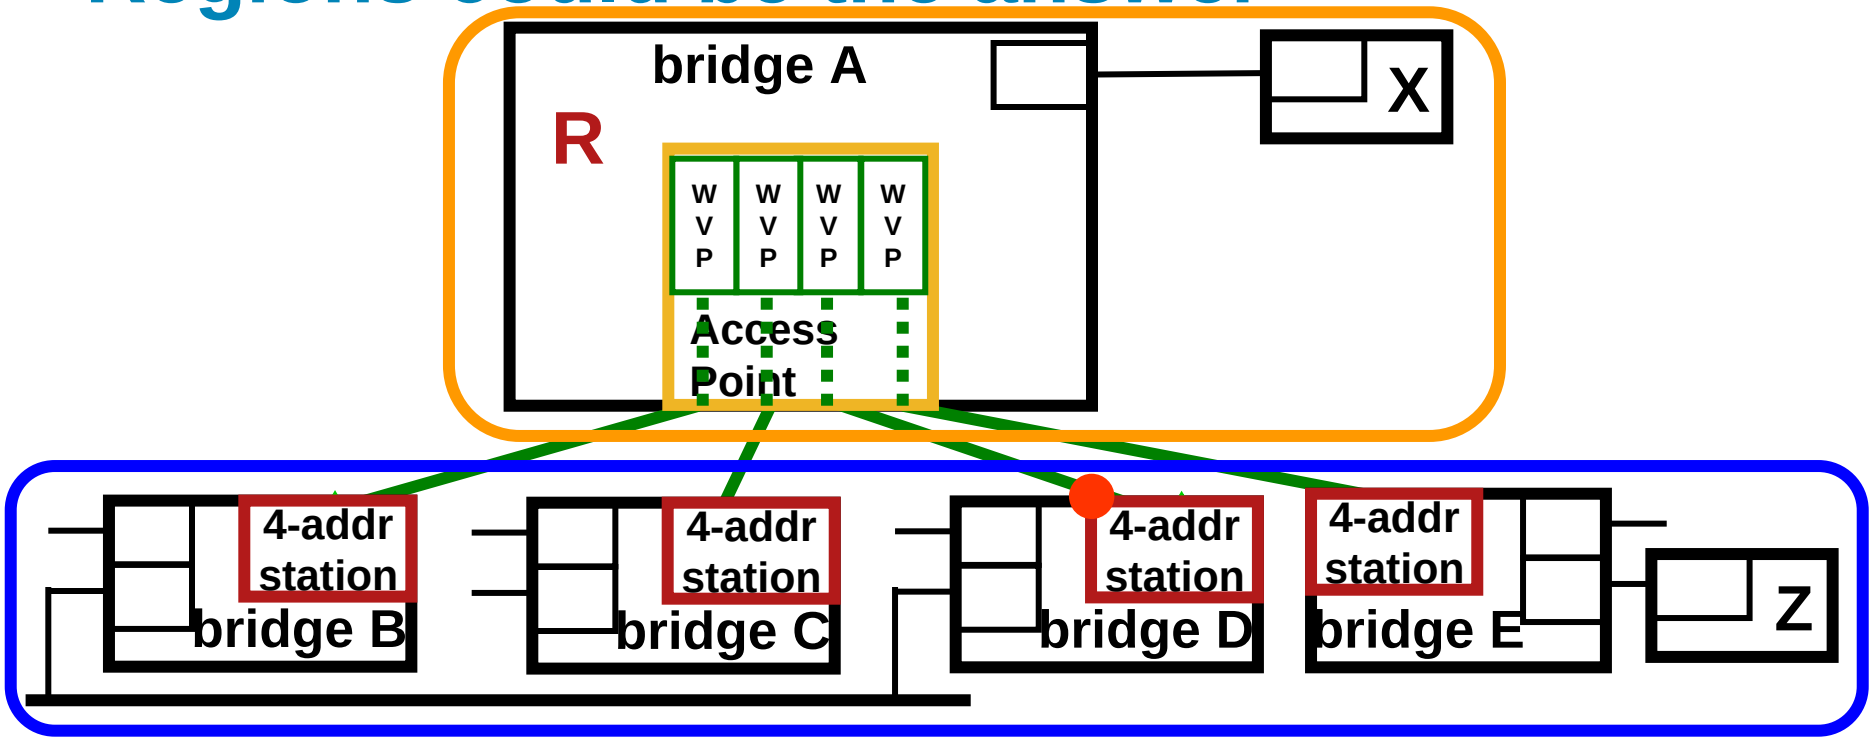

- On the other hand, if all of the station bridges were in one **MSTP Region**, and the AP bridge were in a **different MSTP Region** ...
- Then the wireless link would block. (This is François Tallet's idea.)

Why is this important?

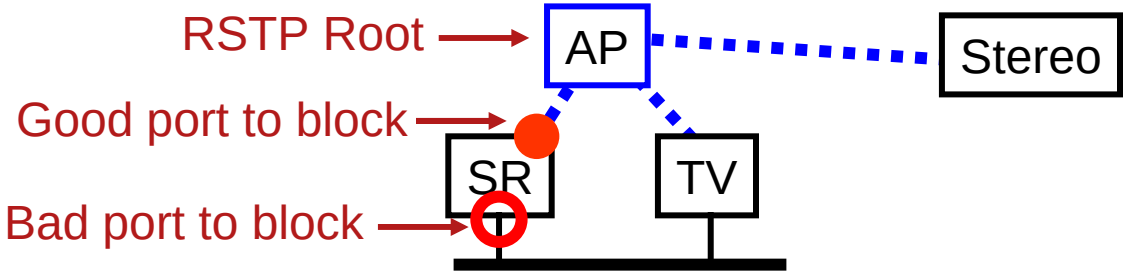

- We expect **exactly** this situation to occur in the home networks envisioned by the AVB Task Group.
- We don't want the Satellite Receiver that is sitting on top of a TV, and connected to it via a 10 Gbps wired link, to use the 2-hop 54 Mbit 802.11 wireless link to send a video program!

  Even if that means that the Satellite Receiver talks to the stereo by taking an extra hop through the TV before getting to the AP.

  That's a small price to pay for the reduction in wireless traffic.

A simple hack to make the Regions different

- If I and the other bridge are both exhibiting the default MSTP Configuration information, and we're connected over a wireless medium slower than 100 Mbps, then we are in different Regions, not the same Region.

  Explicit configuration easily overrides this hack.

  This trick works even in the default case, where there are no MSTIs configured! (I prefer bridges in my own Region even for the CIST.)

Other answers in abundance!

- One of the general-purpose spanning tree cut-through algorithms proposed over the years might solve this problem.
- IS-IS routing, either through P802.1aq or through TRILL, certainly solve this problem.

  (IS-IS also solves the proxy blocking problem, since every device knows the whole spanning tree.)
- **Other suggestions are certainly welcome!!**
- **But, an answer is needed soon.**

One last thought ...

- RSTP/MSTP aren't sent by Access Points towards their non-AP stations today, because:
  - The BPDUs take up precious bandwidth.
  - The non-AP stations can't be bridges, anyway.
- Non-AP station bridges don't send RSTP/MSTP towards the Access Points, because:
  - The BPDUs take up precious bandwidth.
  - Non-AP stations bridges aren't legal.
- So, we have an opportunity to alter (hack?) MSTP **before** these devices are deployed.