




IP Forwarding Plane

Mark Seery

V1 August 22, 2020




Tutorial Purpose:

- High-level discussion of **FORWARDING** in IP networks
- This deck uses the term “forwarding plane”. “Data plane” also commonly used
- Addressing is illustrative, not actual
- Focus: principles and takeaways


Forwarding plane / data plane:

- **FORWARDING** IP packets from one host/end user to another
- User plane: A term sometimes used to refer to the user interface to a network, which sometimes has specialized signaling
- Also participates in lower-level link protocols as needed, for example Ethernet and PPP. Not discussed in this deck.



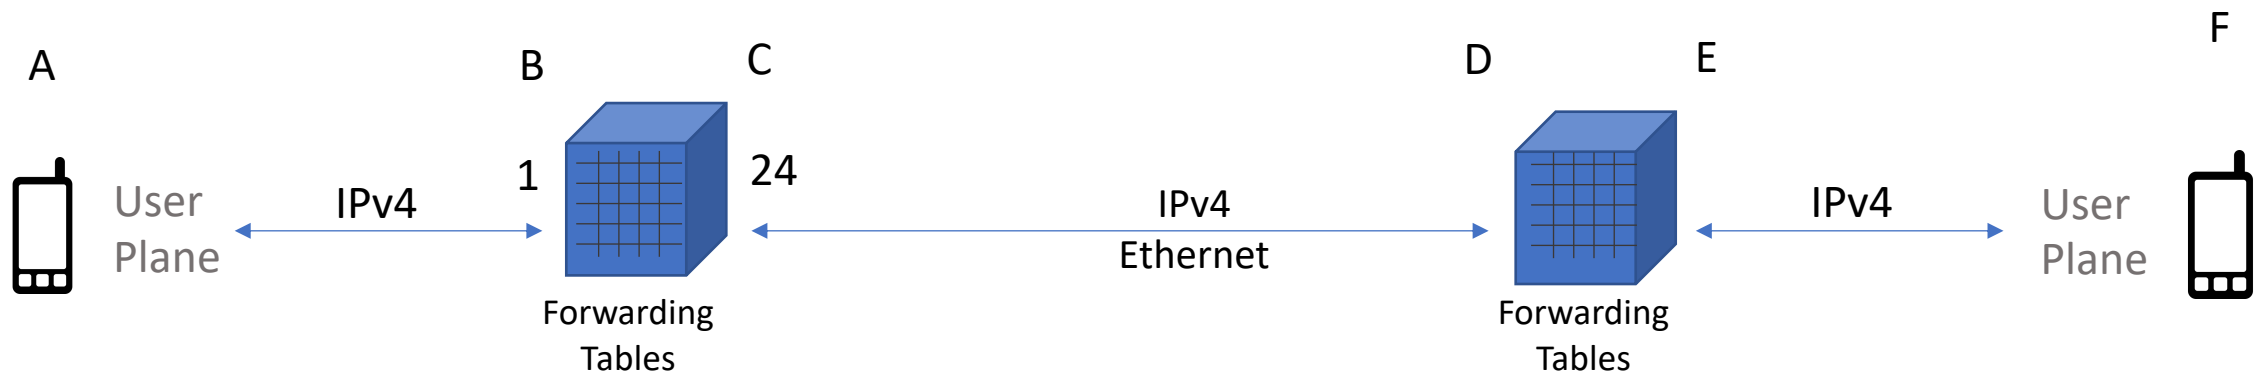
Making forwarding decisions:

- Routers have multiple ports attached to them
- Routers could forward IP packets randomly to other ports
- But that would not lead to a high success rate
- Packets may endlessly loop creating congestion
- Packets may never be delivered to the intended destination
- Worse, packets may be delivered to the wrong destination



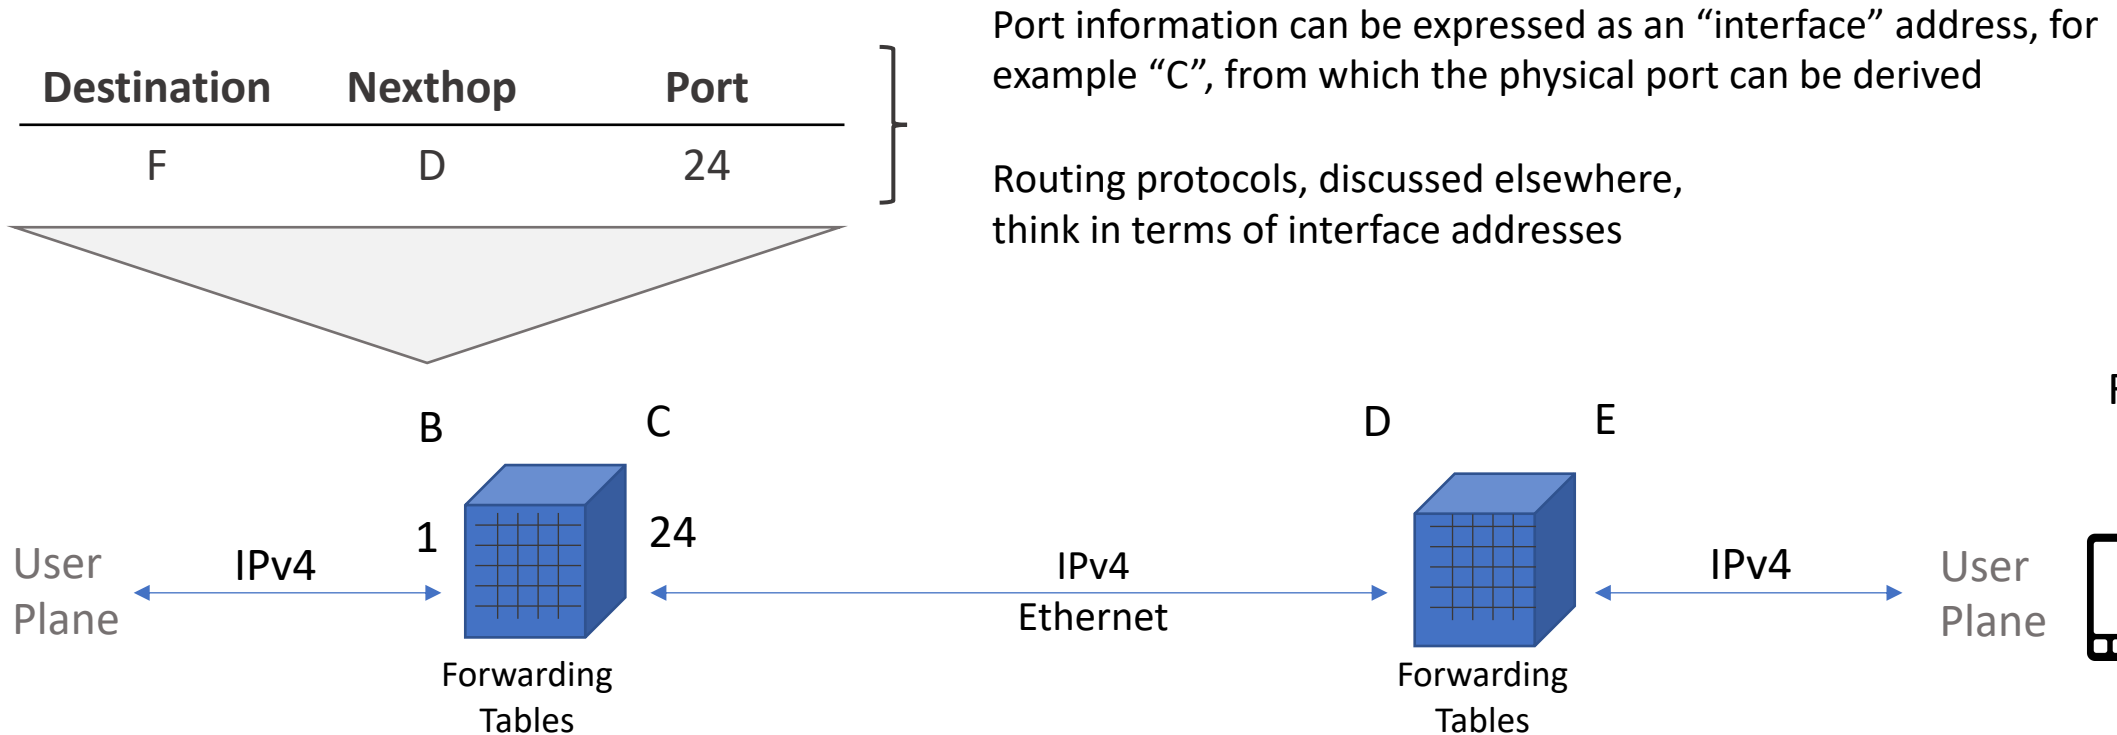
Required ***INFORMATION*** added to routers:

- The next router that the packet has to be forwarded to (nexthop)
- The port on the router that the nexthop is connected to
- For the destination of the IP packet
- This ADDED information is stored in forwarding tables, often referred to as a forwarding ***INFORMATION*** base (FIB)



How forwarding decisions are made:

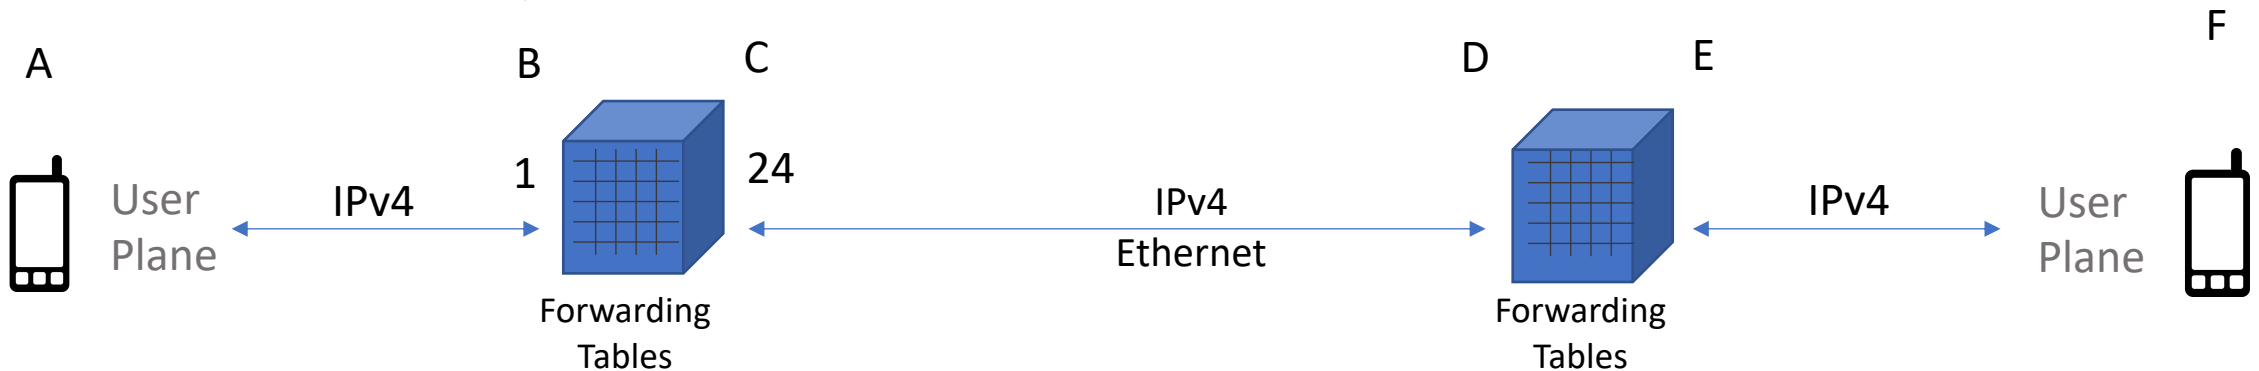
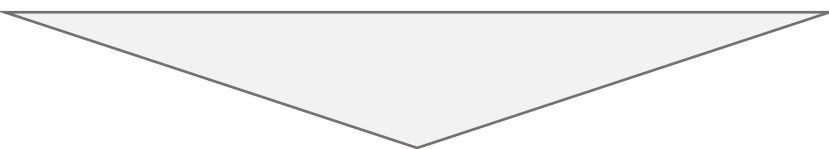
- Traditionally, only the IP destination address is used
- The destination address is then mapped to a nexthop and port



How forwarding tables are created:

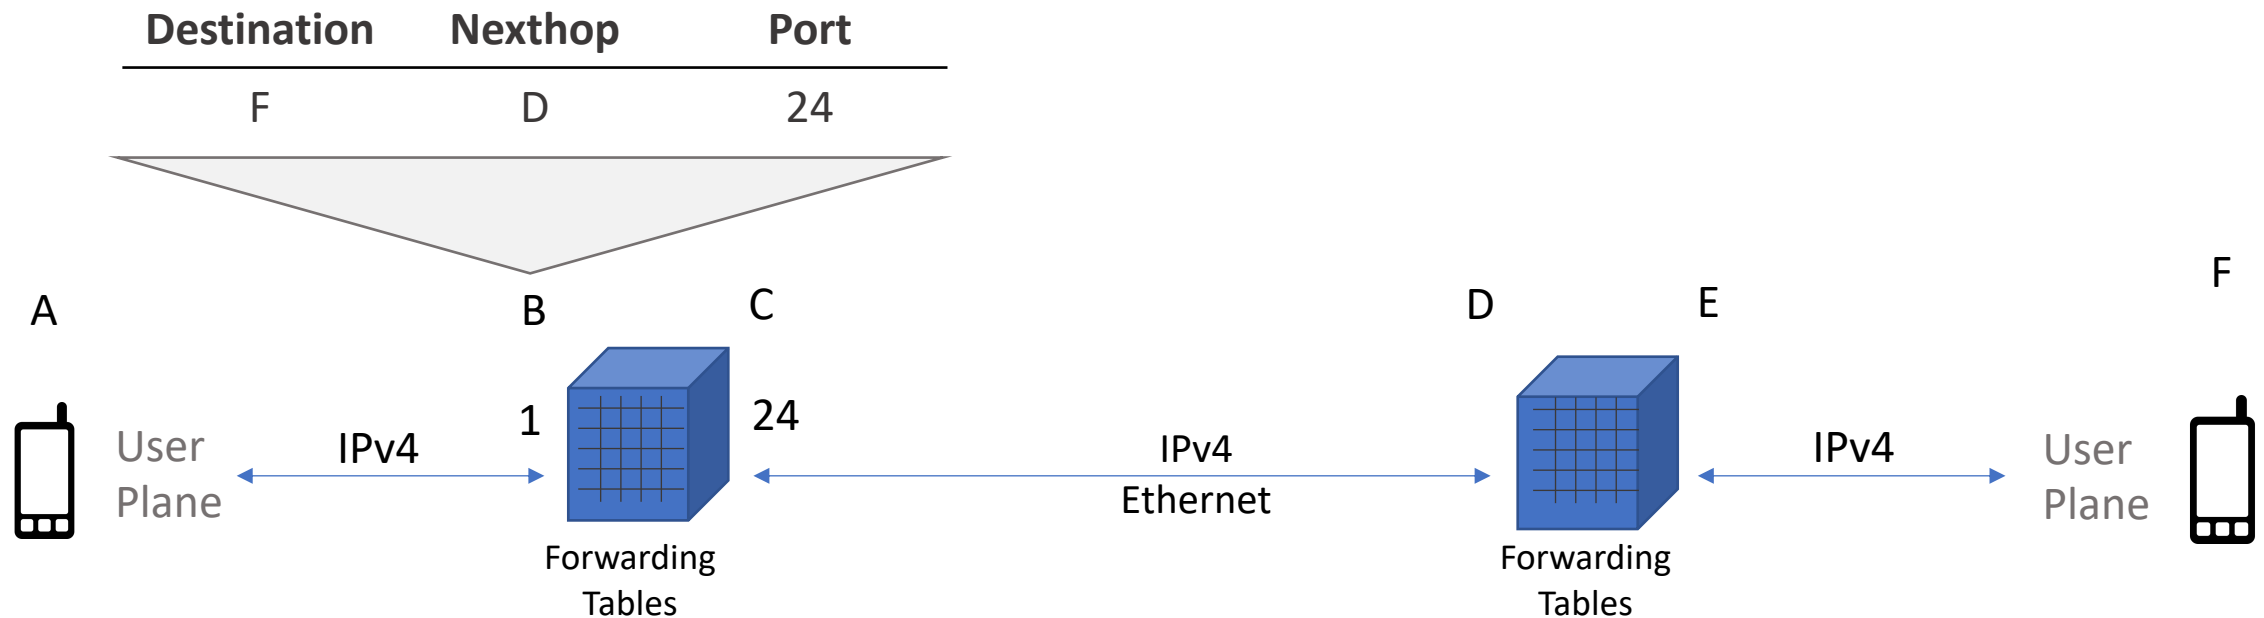
- Forwarding tables can be manually configured
- Traditionally, the heavy lifting has been done by routing protocols

Destination	Nexthop	Port
F	D	24



Manual configuration:


- As the network grows, more people have to be added for config
- Configuration has to be coherent across the network
- People are prone to typing / logic errors
- The network can change faster than people can respond



Automated forwarding table configuration:

- Routing protocols are the most common approach
- Centralized controllers / management approaches also used

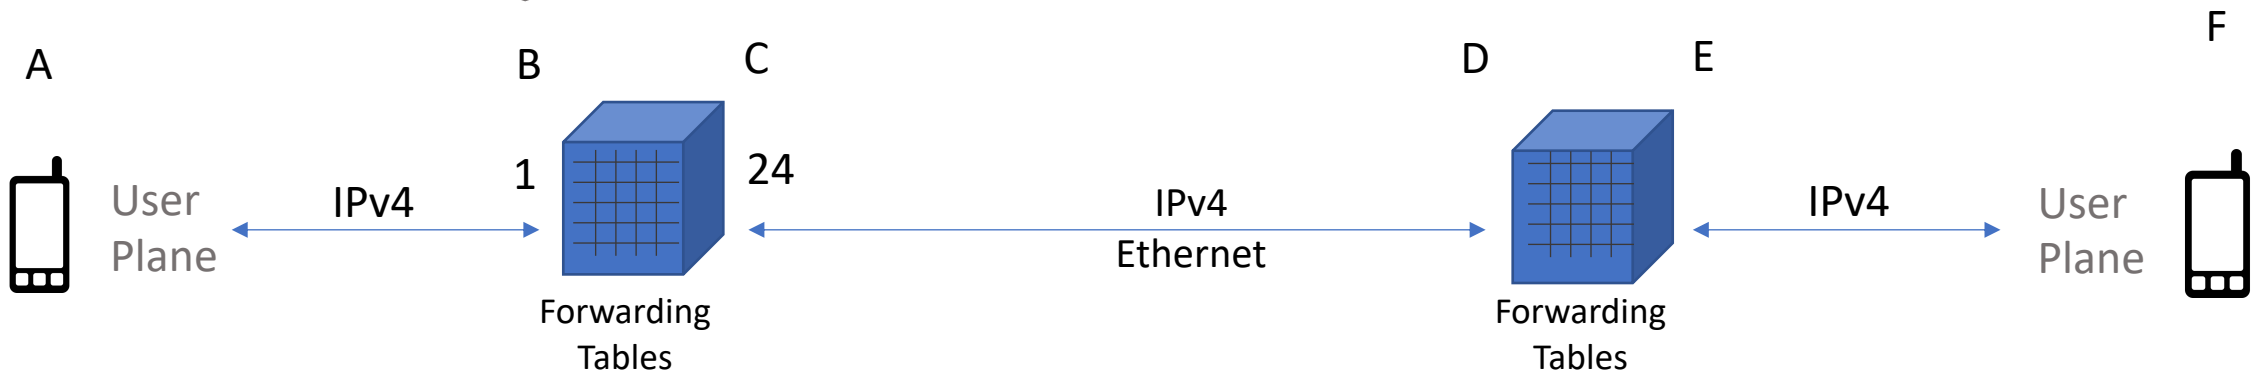

Destination	Nexthop	Port
F	D	24



Information summarization:


- Forwarding tables use expensive technology for performance
- Often, destinations are summarized as (sub)networks, reducing forwarding table size requirements / expense

Destination	Nexthop	Port
F	D	24



Network Addressing:

- If multiple hosts / end users exist on the same link, which Ethernet supports, then the forwarding table can have one entry for the (sub)network




Challenges with summarization:

- If 2.G is later moved to a new link, but the same summarization is maintained, then 2.G will not receive packets destined for it

Destination	Nexthop	Port
2.ALL	D	24

Normally the address would be changed to something like 3.G. This is illustrative of issues that occur in more complex IP networks, in different aspects of routing




Fixing summarization problems:

- Add forwarding information specific to 2.G
- IP forwarding uses the most specific information it has

Destination	Nexthop	Port
2.ALL	D	24
2.G	DD	23


Normally the address would be changed to something like 3.G. This is illustrative of issues that occur in more complex IP networks, in different aspects of routing



Longest prefix match:

- Longest prefix match (LPM) is the term applied to making forwarding decisions based on a tree of less/more specific information

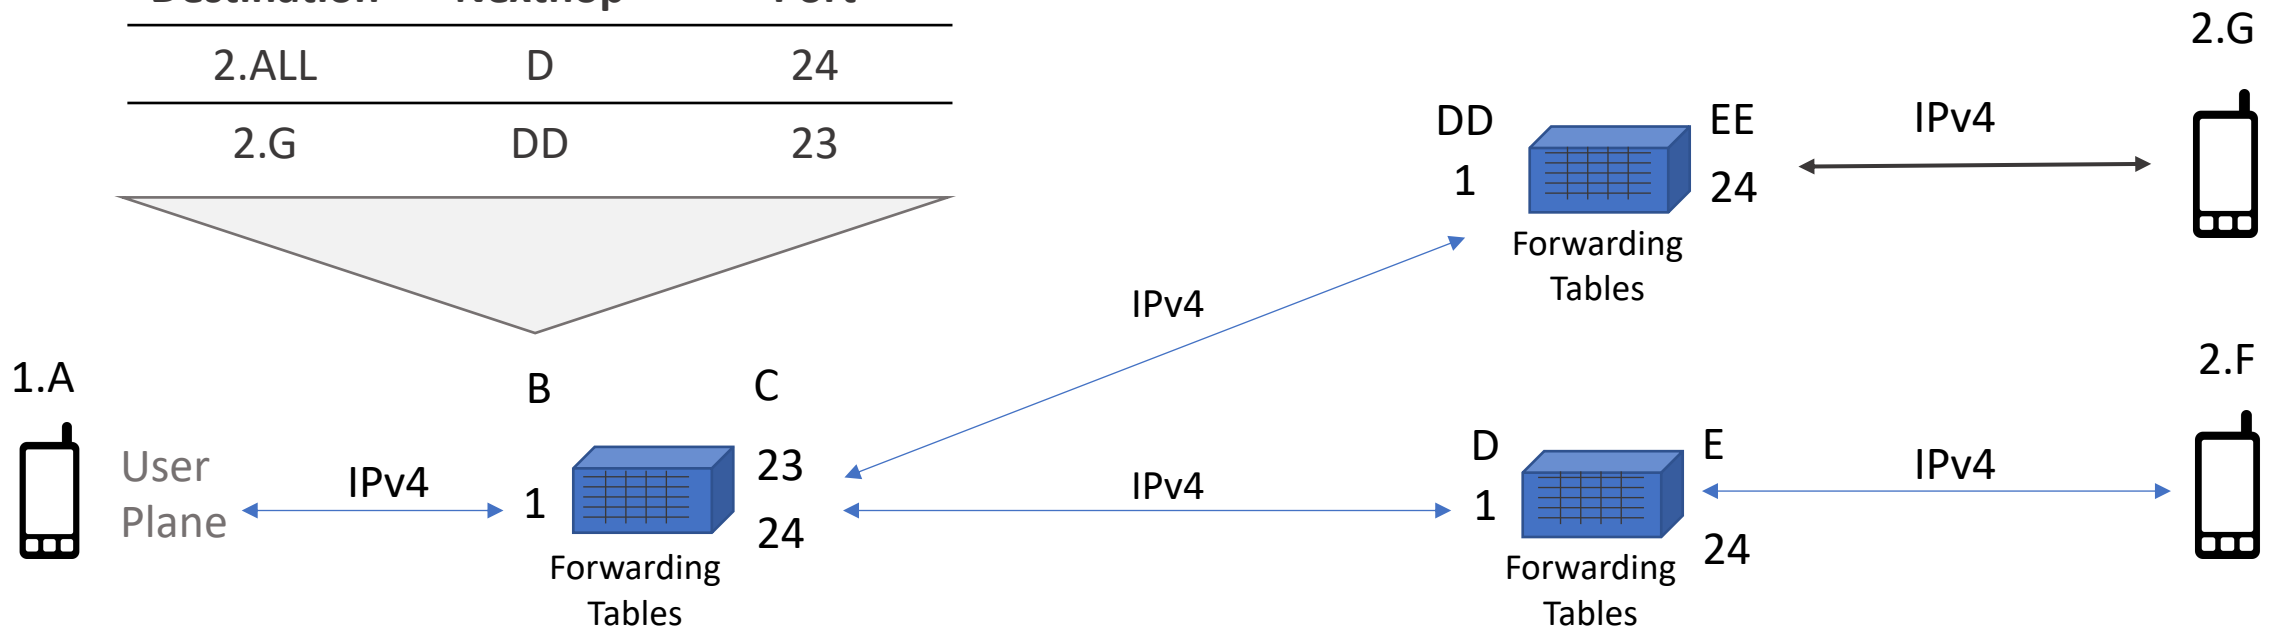
Destination	Nexthop	Port
2.ALL	D	24
2.G	DD	23



Longest prefix match (LPM) tradeoffs:

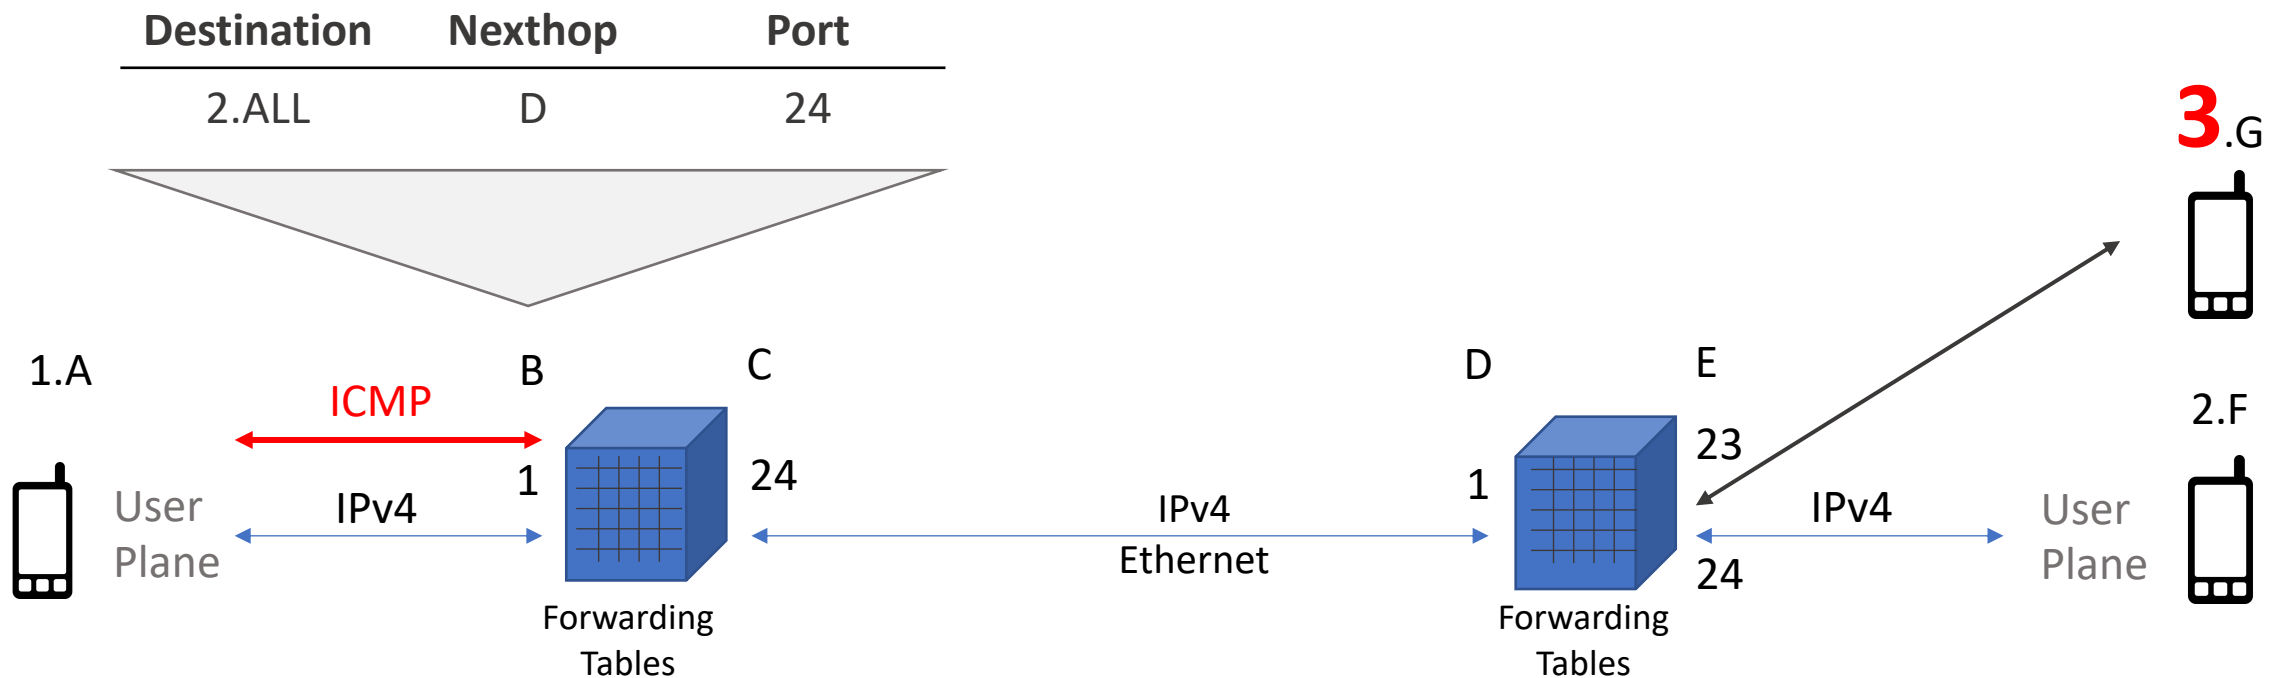
- LPM adds forwarding plane computational complexity / cost
- Provides flexibility to reduce overall forwarding information
- Increases the need for address planning, which is outside the scope of routing protocols (potentially unautomated burden)

Destination	Nexthop	Port
2.ALL	D	24
2.G	DD	23



Forwarding information not found:

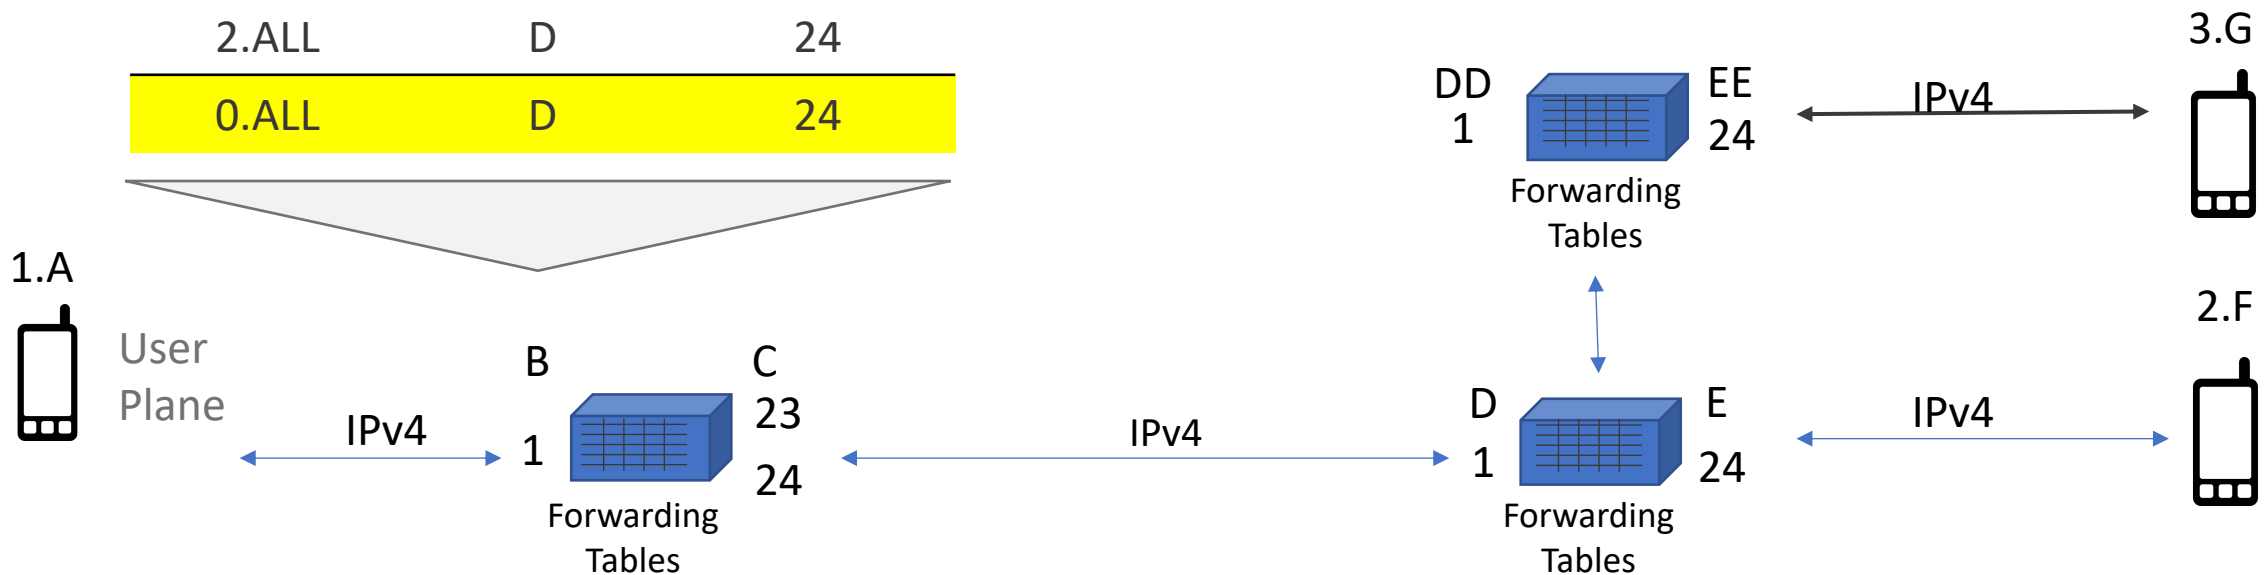
- If network 3 is created, but there is no information in forwarding tables for network 3, then an error is returned to the sender via a protocol called ICMP (Internet Control Message Protocol)



Default routing:


- Sometimes, by design, full routing information only exists in select routers. A packet can be sent to that router to resolve.
- This may tradeoff the most optimal routing for less expensive routers, but there may only be one upstream route anyway

Destination	Nexthop	Port
2.ALL	D	24
0.ALL	D	24



Summary

- Forwarding: info in IP packet PLUS info added to forwarding tables
- Summarization: reduces info required, while reducing clarity/certainty
- Summarization errors:
 - add more specific information where needed
 - increasing forwarding complexity
- Longest prefix match: increases flexibility, increases computation requirements
- Info creation & distribution: tradeoffs - people / scale and automation
- Manual configuration: what is not automated, increases burden on people, and compromises network efficacy



Thank You