




An IX-Anchored Internet

Richard T. B. Ma
Department of Computer Science
National University of Singapore

Motivation: Lindbergh operation

- Successfully conducted between New York and France in 2001
- Why tele-surgery is still not common today?
 - Bandwidth reservation via asynchronous transfer mode (ATM) network
 - Operated by a single admin domain: France Telecom
- End-to-end QoS is needed, but difficult for TCP/IP network



Motivation: End-To-End Quality of Service (QoS)

- Network layer service model is “best-effort”
 - QoS wasn’t a goal at the design
 - There is no native support from the Internet architecture
- End-to-end routing go through multiple Autonomous Systems (Ases)
 - Early architectural solutions, e.g., IntServ and DiffServ, didn’t get adopted
 - Economic incentives of ASes need to be addressed

Motivation: Do We Have End-To-End QoS?

- In theory, no; but in practice, we are using Netflix, Zoom, and etc.
- Large cloud or content providers (CPs) achieve QoS by
 - Bypass Tier 1 ISPs, leading to the flattening of the Internet
 - Building private wide-area networks and content delivery networks
 - Engage private peering with last mile access providers to reach users
- Lessons learned
 - No changes in the Internet architecture, but using (private) infrastructure
 - Currently can only be done if you are BIG, and still take a long time-scale

Goal and Big Ideas

- Goal: enable end-to-end QoS over the public Internet
- Idea 1: An infrastructural rather than an architectural approach
- Idea 2: Design around the Internet Exchanges (IXs)
- Idea 3: An economic plane of the Internet


Architecture vs. Infrastructure

- **Internet Architecture:**

- The vertical layered protocol stack
- IP layer and above

- **Internet Infrastructure:**


- The horizontal below the IP layer
- now consists of over a gigameter (a million of kilometers) of submarine cables



Why not an architecture approach?

- The core of the architecture is the network layer
 - Relies on the only routing protocol: BGP
 - Designs are centered around Ases, e.g., economic incentive via peering
 - An architectural approach must also be an AS-centric approach via BGP
- Limitations of BGP:
 - Routing is designed for universal reachability, not QoS
 - BGP enables bilateral peering, but not global value distribution along a path

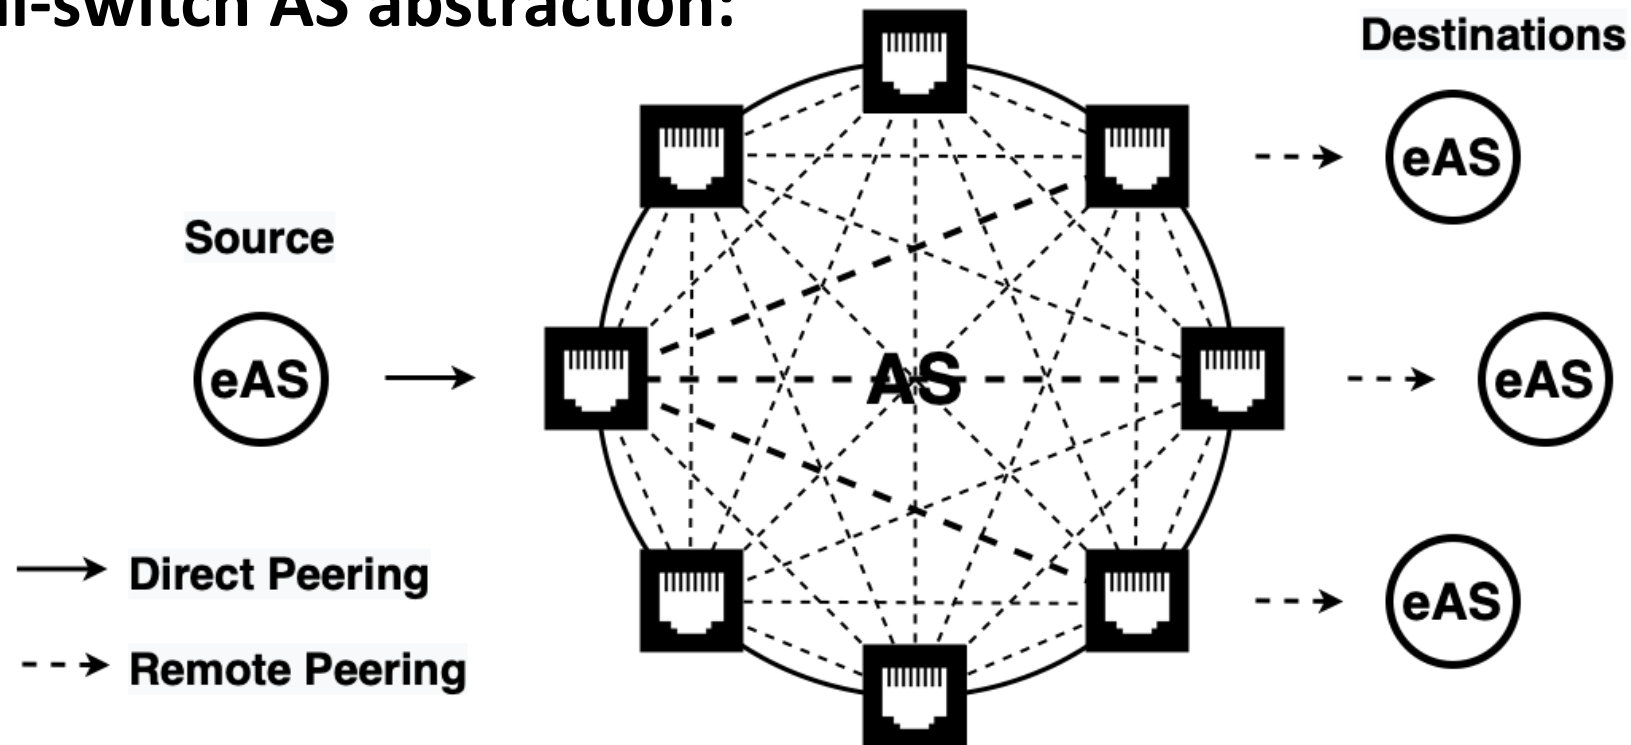
Design Principles



- Principle 1: Separation of Core and Edge
- Principle 2: Separation of Infrastructure and Architecture
- Principle 3: Building Economic Plane on Infrastructure

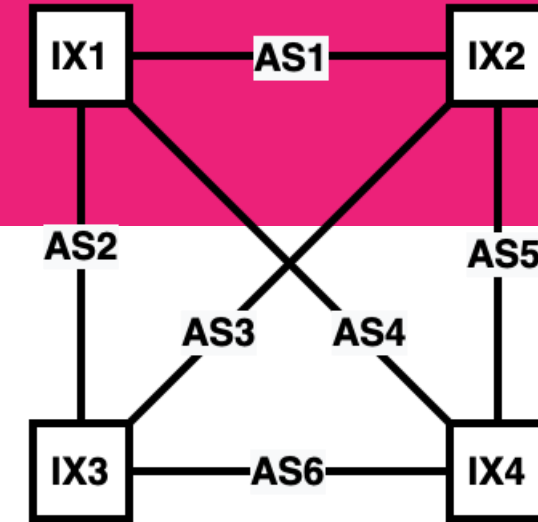
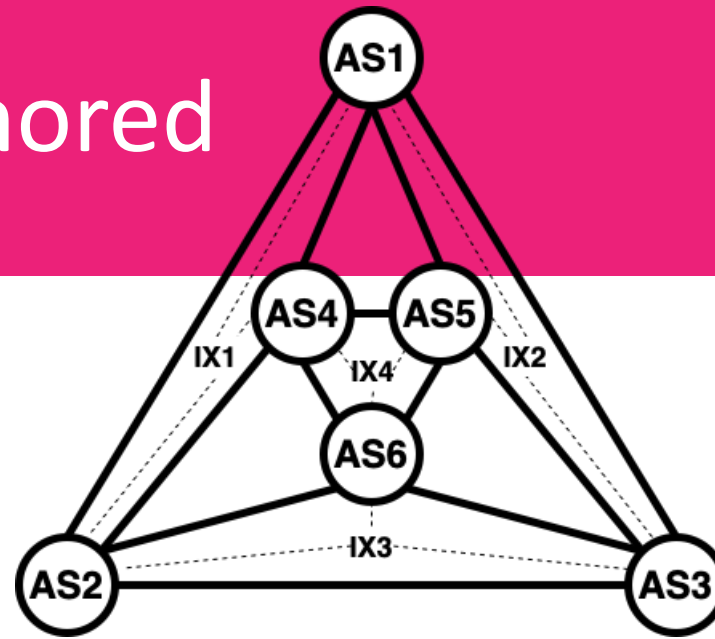
Abstraction of the Transport Core

- A global-switch AS abstraction:



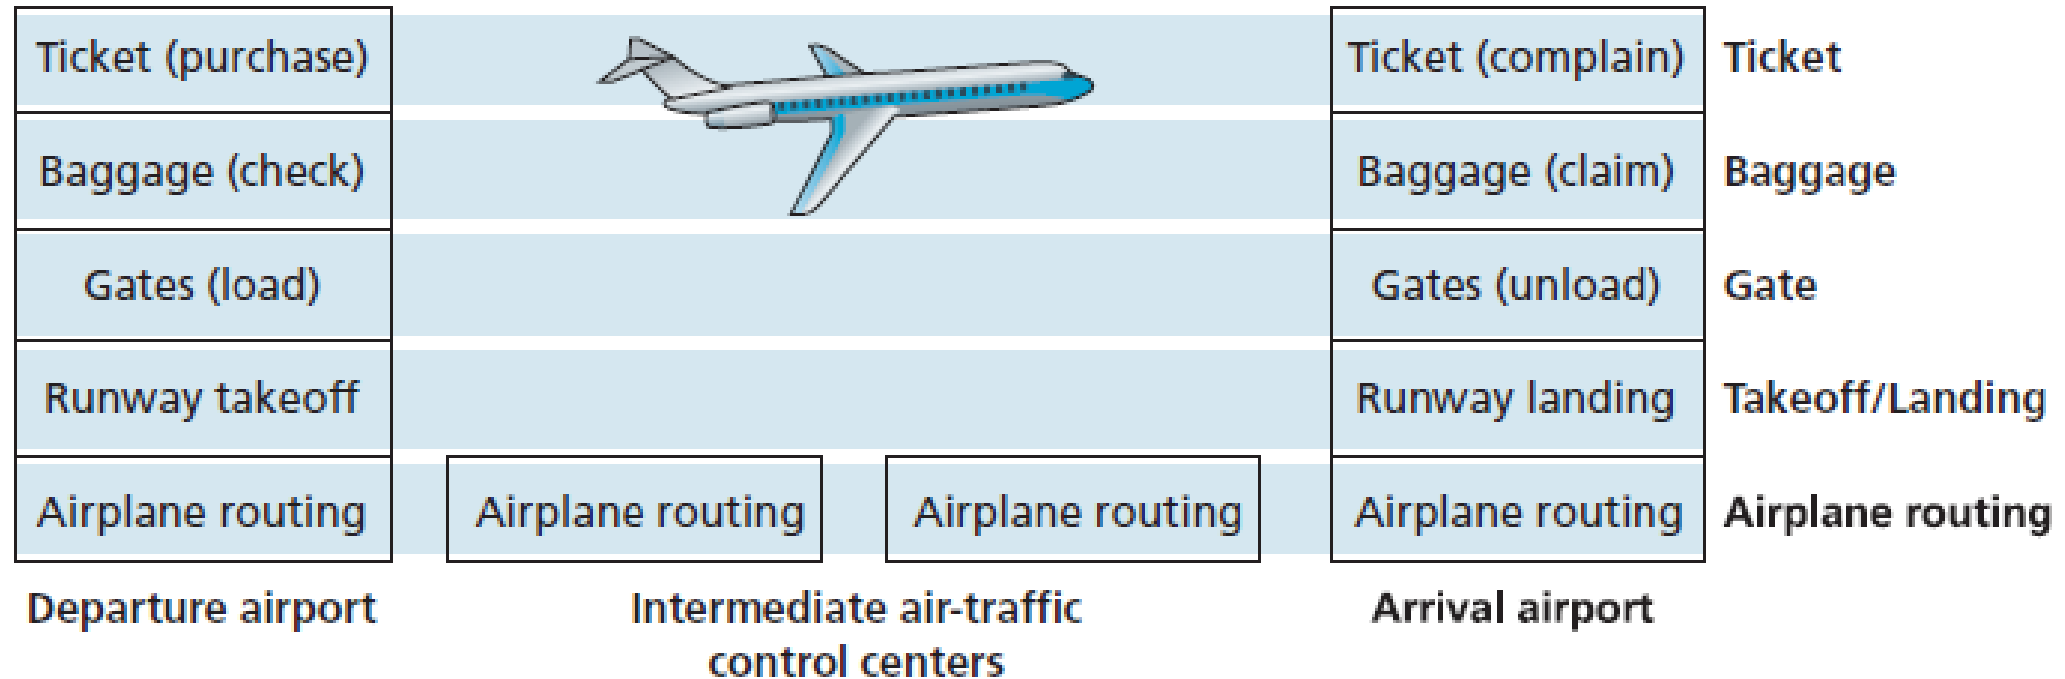
AS-Centric vs. IX-Anchored

- More natural graph mapping
 - IX is a L2 net with geolocation
 - AS is a virtual link between IXs



- An IX-path specifies **where** traffic will pass through
 - rather than by whom traffic will be carried
- An infrastructural approach, as IX does not belong to architecture
- Analogy to the public transportation systems: IX = airports

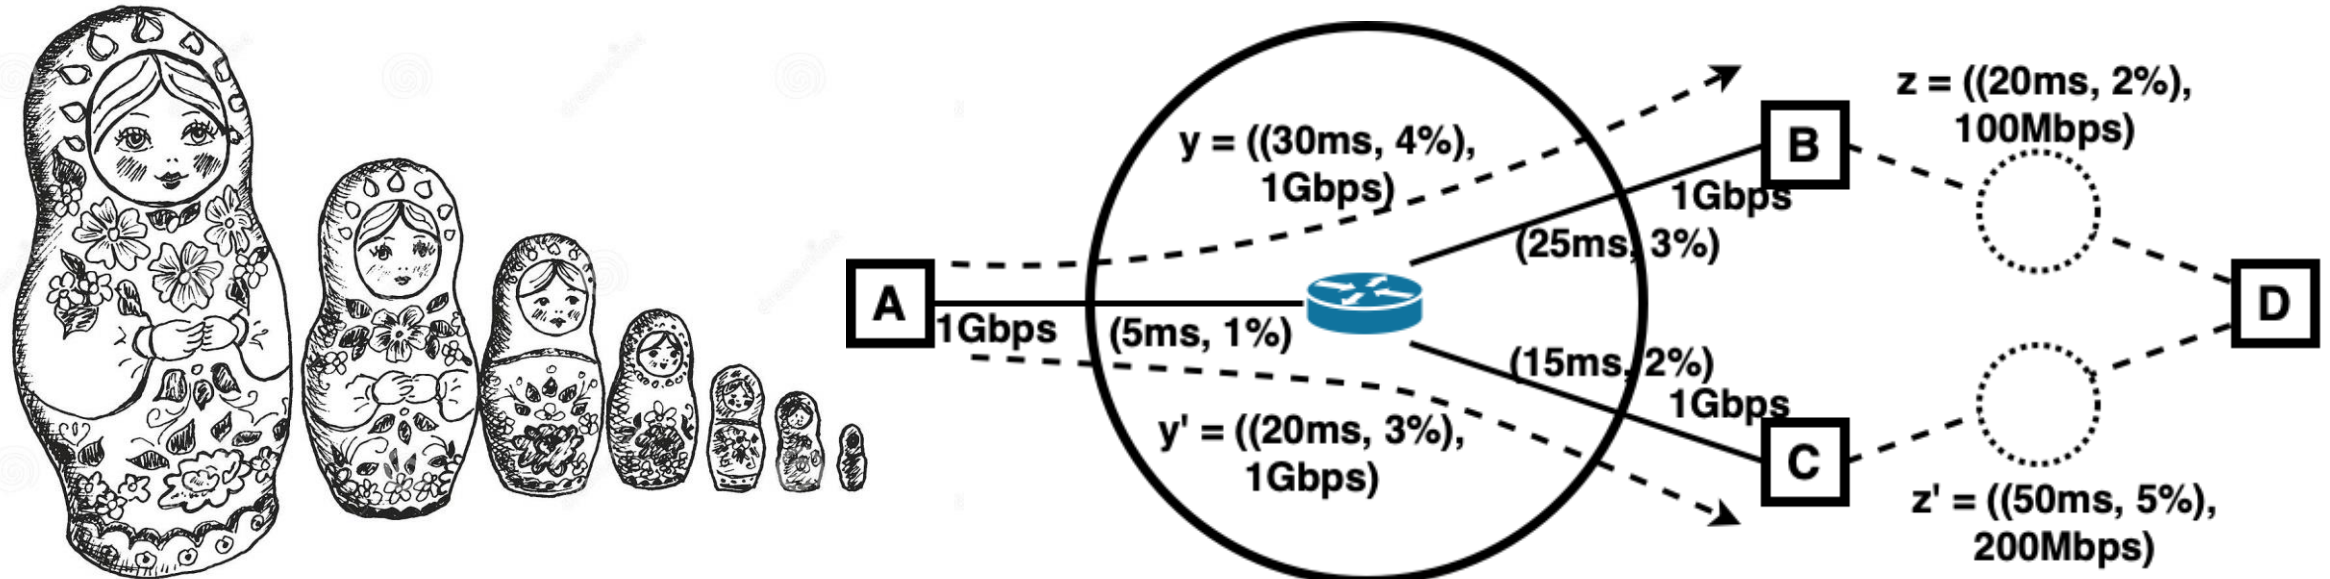
AS-Centric vs. IX-Anchored



- Analogy to the public transportation systems: IX = airports

Service Abstraction and Composition

- Matryoshka Model for service composition
- IX anchors enable distributed market structure
 - Transport services are sold at their source IXs only



Market Design, Implementation & Deployment

- Market Design
 - Long-term reservation market and short-term spot market design
- Implementation
 - Inter-domain: Blockchain/Smart Contracts; Intra-domain QoS: MPLS/SDN
- Evolutionary pathway
 - Backward compatible and Incremental deployment
- Conclusion
 - IXs should play a much more important role
 - Exchange beyond traffic, but services, more like the role of financial exchange