E6998-02: Internet Routing

Lecture 11
OSPF continued (and finished).

John Ioannidis
AT&T Labs – Research

ji+ir@cs.columbia.edu

Announcements

Lectures 1-11 are available.

Start working on your project proposal (due 10/22).

- Literature search.
- Summary of approach.
- Division of labor.
- Validation of results.

Reading assignment: Read RFC 2328 twice.

Still looking for a TA.

Adjacencies Revisited

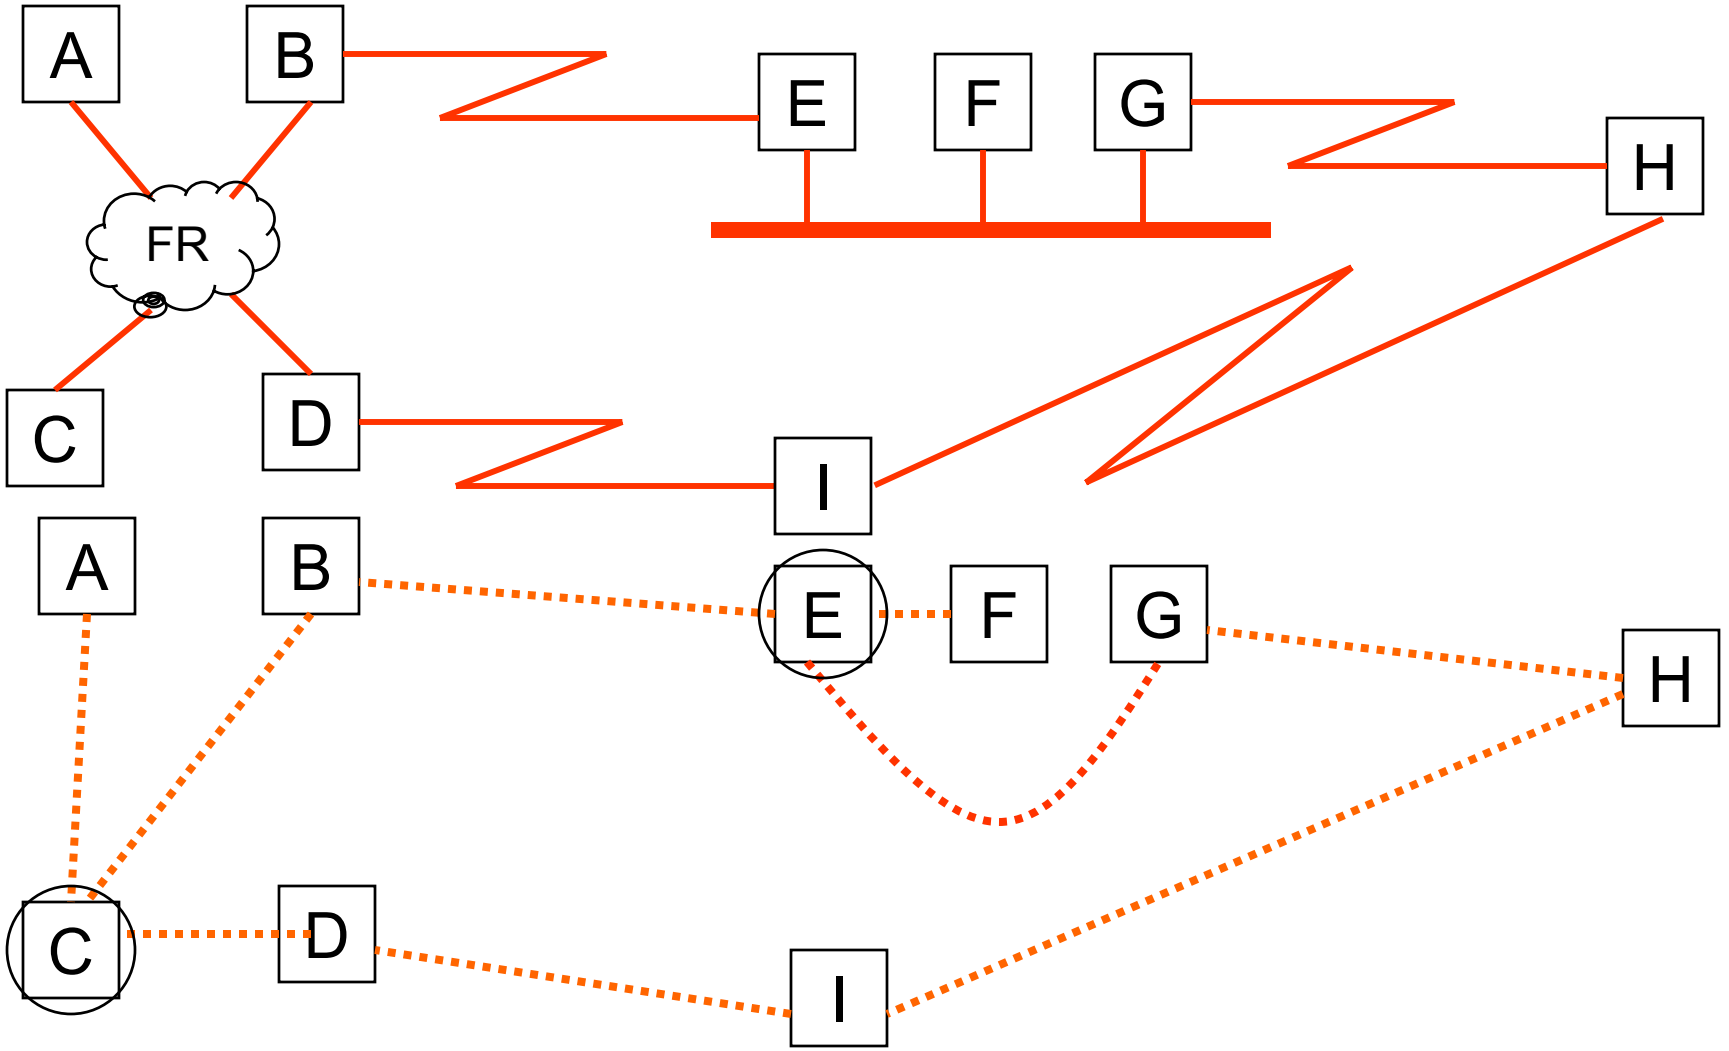

Routers connected by data links ⇔

nodes connected by adjacencies.

Flooding

- Link State Database: list of all LSAs the router has heard (and sent).
- Change in topology results in new or changed LSAs.
- Changed LSAs are flooded throughout the network:
  - Link State Updates.
  - Link State Acknowledgements.
- Each LSA reaches every router.
- Updates/Acks only flow between adjacent routers
  - I.E., it's not the update packets that get flooded, it's their contents.

Updates

- On point-to-point networks, multicast to `AllSPFRouters`.
- On broadcast networks:
  - DRothers multicast updates to `AllDRouters`.
  - The DR then multicasts an update to `AllSPFRouters`.
  - If the DR fails to do that, BDR takes over, otherwise BDR stays silent.
- On NBMA networks:
  - DRothers unicast updates to DR and BDR.
  - DR unicasts updates to all adjacent routers.

Reliable flooding

- Transmitted LSAs must be acked.
- Implicit acks: send the same LSA back.
  - Used when you would have sent it anyway.
- Explicit acks: OSPF packet type 5.
  - Carry only LSA header.
- When sending an LSA, put it in a retransmission queue in the neighbor data structure.
  - Retransmitted every RxmtInterval (or until adj. is broken).
- Delayed acks: more LSAs acked in a single update packet.
- Direct acks: sent immediately and are unicast.
  - When duplicate LSA received from neighbor.
  - Rxed LSA has MaxAge and router has no copy of it.

Sequence numbers

- Linear sequence number space.
  - Signed 32bit integers.
  - Start at `InitialSequenceNumber` (0x80000001).
  - End at `MaxSequenceNumber` (0x7fffffff).
- First LSA goes out with `InitialSequenceNumber`.
- Each new LSA adds 1 to the previous sequence number.
- If is `MaxSequenceNumber` reached:
  - LSA must be flushed out of other routers' list.
  - LSA is sent out with MaxAge.
  - When all neighbors (adj.) have acked, flush LSA and create new one.

Age

- Age of LSA in seconds.
- Unsigned 16-bit integer.
  - From 0 to `MaxAge` (3600).
- Set to 0 by originating router.
- At each router transit, incremented by `InfTransDelay`.
- Also incremented as it resides in database.
- When LSA reaches `MaxAge`, it is reflooded so it can be eliminated from the network.
- When the originating router wants to flush an LSA, it sets the age to `MaxAge` and floods it.
- LSAs are refreshed every `LSRefreshTime` (1800s).
  - With Sequence Number incremented by 1.
  - LSA group pacing.

LSA Comparison

- Highest sequence number is newest.
- Else highest checksum is newest.
- Else if one of the ages is `MaxAge`, it is newest.
- Else if ages differ by more than 15 minutes (`MaxAgeDiff`), lowest age is newest.
- Else LSAs are the same.

- An LSA in a router is replaced when a “newer” one is received.

LSA Types

1. Router
2. Network
3. Network Summary
4. ASBR Summary
5. AS External
6. Group Membership
7. NSSA External
8. External Attributes
9. Opaque (link-local scope)
10. Opaque (area-local scope)
11. Opaque (AS scope)

Router LSA

- Produced by every router.
- Flooded within an area.

- List of all of router's links (interfaces).
  - And corresponding costs.

- Type (=1)
- RouterID
- Number of links
- Link Descriptions (i/f address, link type, metric).

Network LSA

- Produced by the DR on MA networks.
- Flooded within an area.

- Represent the multiaccess network.
  - (MA network acts as a pseudonode).

- Type (=2)
- Network address and netmask.
- Addresses of attached routers.

Network Summary LSA

- Produced by Area Border Routers.
- Sent into an area to advertise prefixes outside that area.
  - One per destination (prefix).
  - If multiple paths known, lowest-cost LSA is advertised.
- When a NS LSA is received, the cost of the route to the ABR is added to the cost advertised in the NS LSA.
  - Distance-vector behavior!

- Type (=3)
- Prefix
- Metric

AS Boundary Router Summary LSA

- Produced by ABRs.
- Identical to NS (type 3) LSAs.
  - Advertise (host) routes to ASBRs.
  - Destination is a host address, prefix length is 32.

- Type (=4)
- ASBR IP address and mask (all-ones).
- Metric.

AS External LSA

- Produced by ASBRs.
- Advertise a destination (or a default route) external to the AS.
- Flooded throughout the AS (but not stub areas).
  - Since they are not associated with a particular area!

- Type (=5)
- Advertised prefix.
- Forwarding address (of external router).
- Metric.

Other LSAs

- Group membership.
  - Used for MOSPF.
- NSSA External.
  - Like AS External, but only flooded within the NSSA.
- External attributes.
  - Proposed as an alternative to IBGP.
- Opaque.
  - Proposed so that OSPF can be used to carry app-specific data to all routers in an AS.

Stub Areas

- Areas with no ASBRs.
- To reach ASBRs, you have to go through the ABR anyway.
- No point in advertising type 5 (AS External) LSAs.
  - No point in advertising type 4 (ASBR Summary) LSAs either.
- Just advertise a default route into the area.
- No virtual links can be configured through a Stub Area.

- Totally-stubby areas: type 3 (Network summary) LSAs are not advertised, except for a default route.

Stub Areas, cont'd

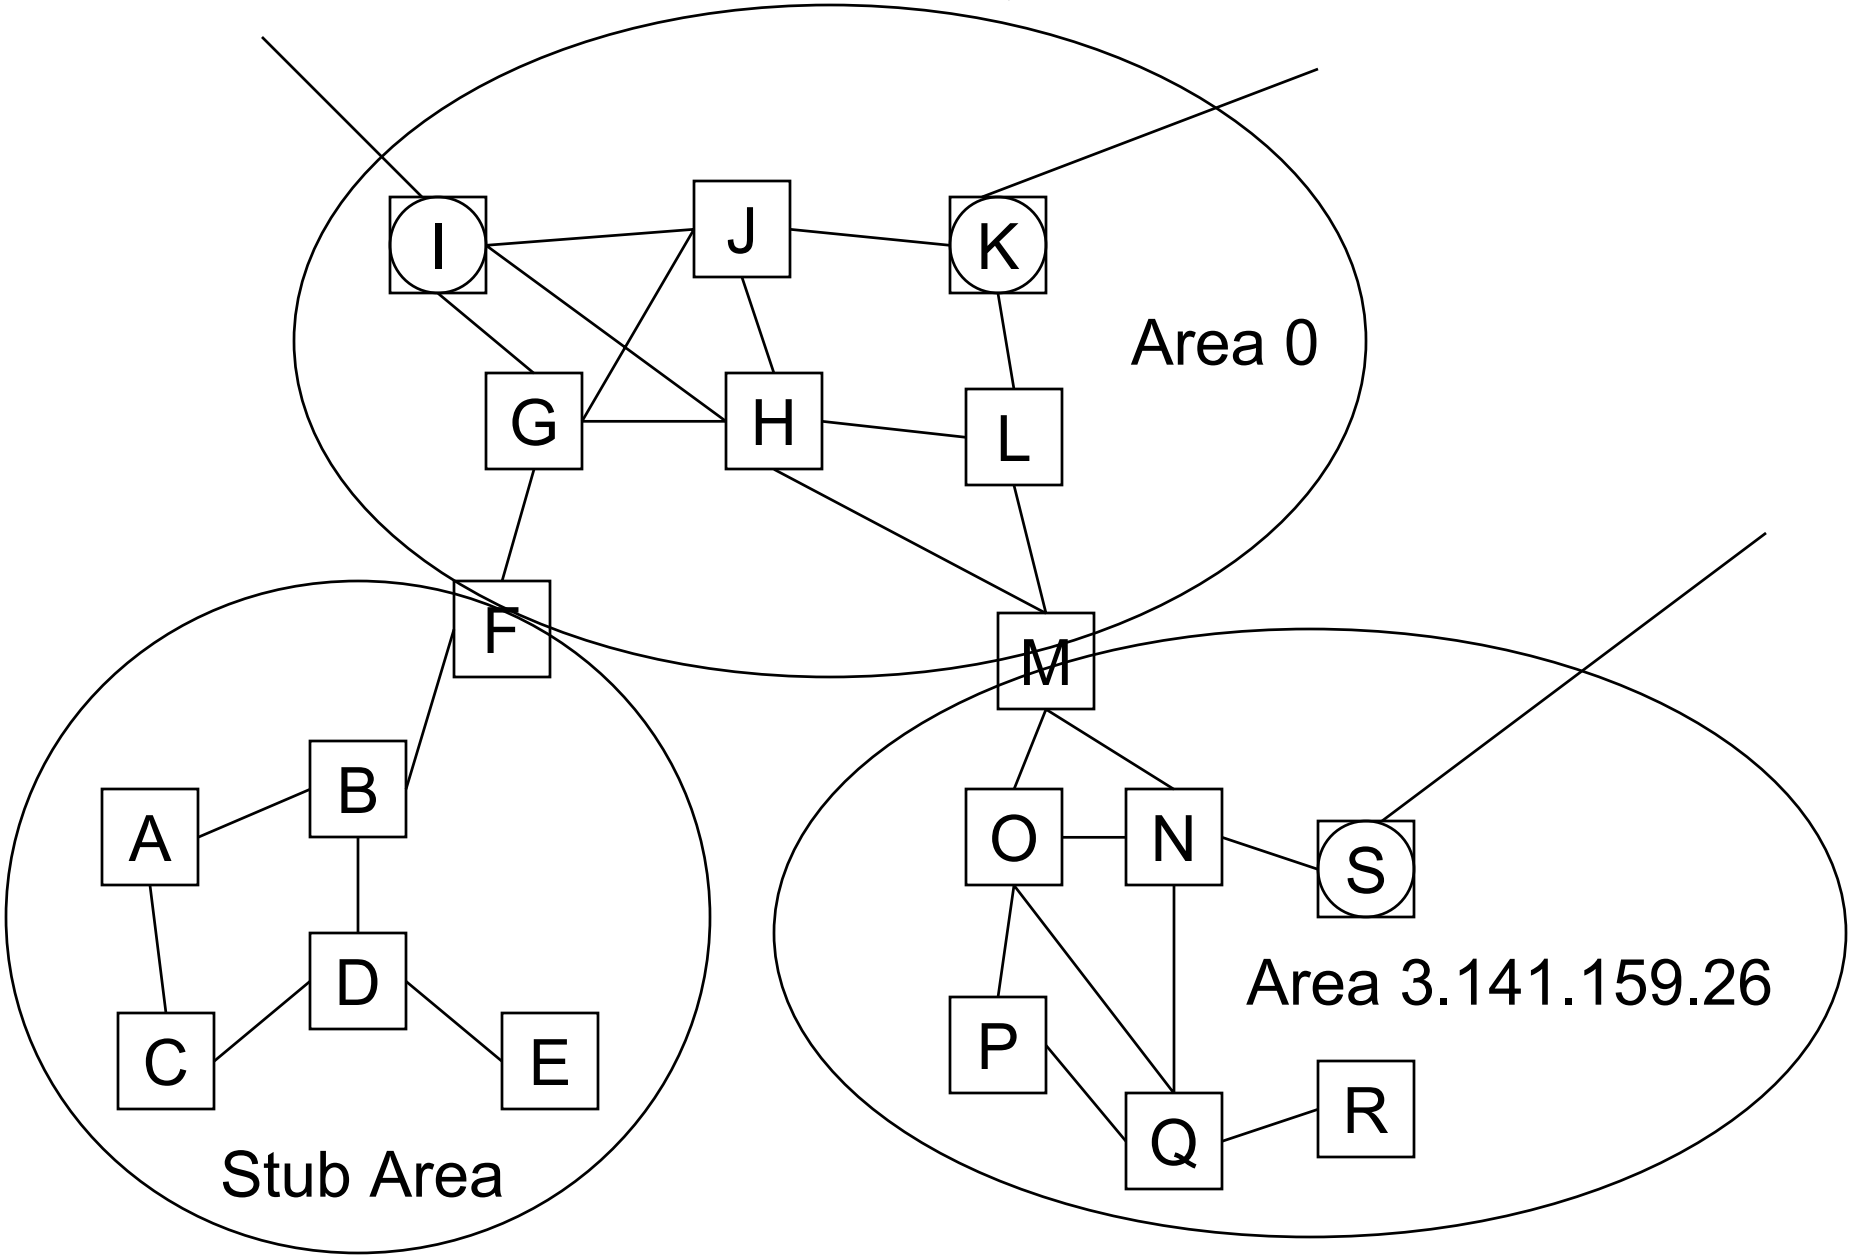

Not-So-Stubby Areas

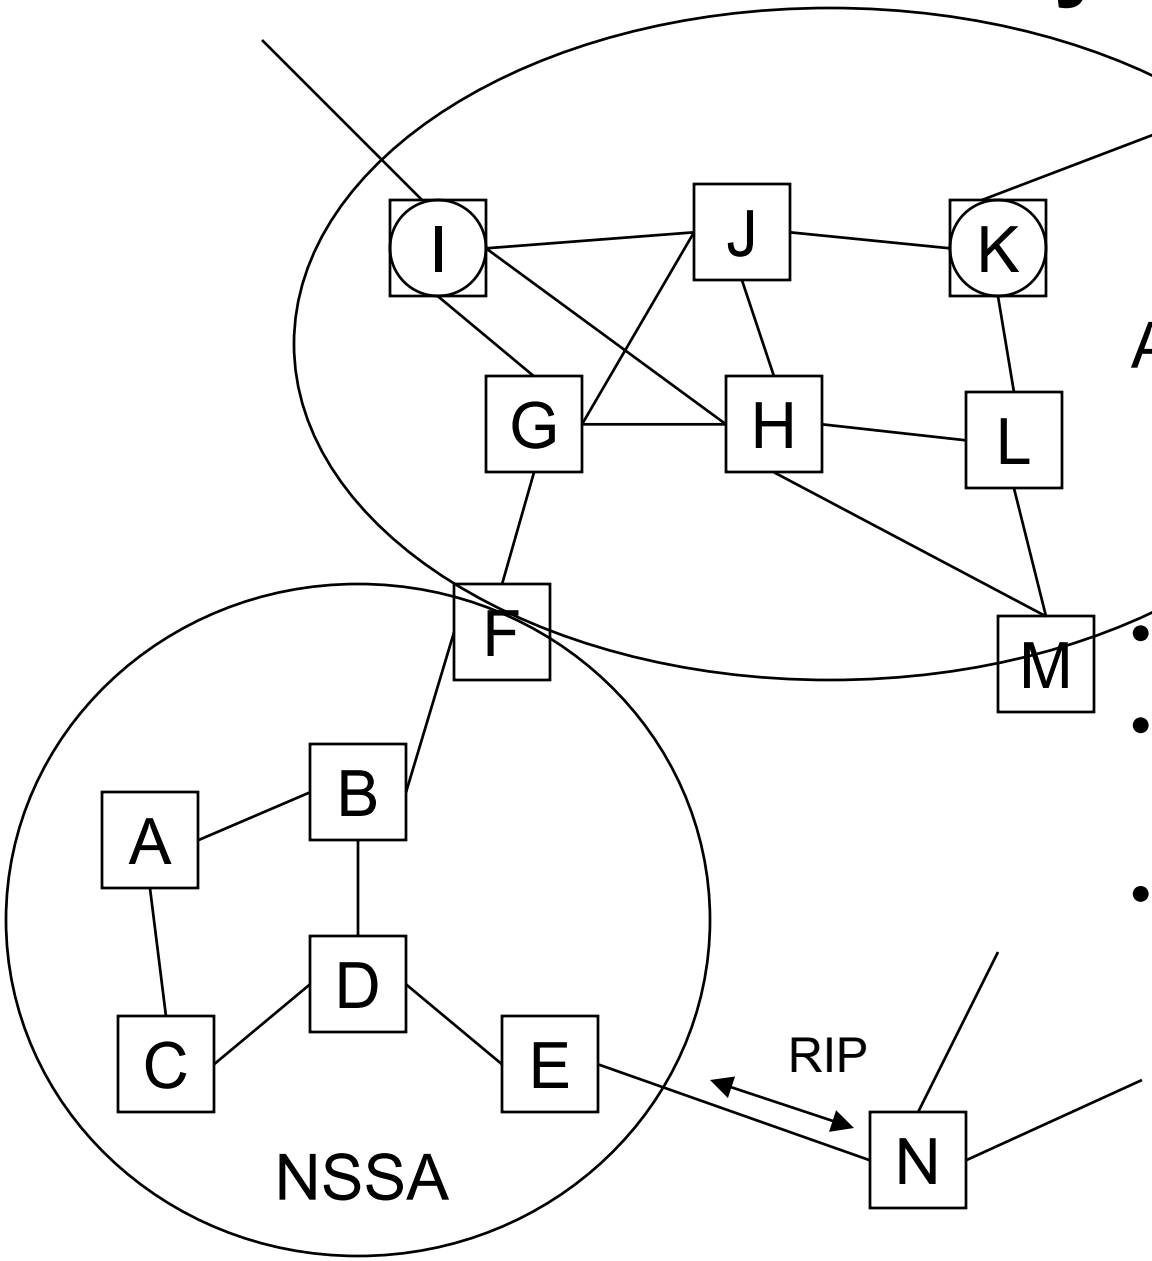

- N has a default route.
- E advertises type 7 LSAs to the area.
- F may block it (P=0) or translate into type 5 (P=1).

Options

- DC bit: Router is capable of supporting OSPF over demand circuits.
- EA bit: Router is capable of sending and receiving External Attributes (type 8) LSAs.
- N bit: Router can support NSSA LSAs. N=1 implies E=0.
- P bit: (Same position as N bit). ABR should translate a type 7 into a type 5 LSA.
- MC: Used by MOSPF.
- E: Router is capable of accepting AS External LSAs.
  - In hello packets, indicates ability to send/receive Type 5.
- T: capable of supporting TOS.