Reference: IS-IS vs OSPF

Shivkumar Kalyanaraman
Rensselaer Polytechnic Institute
shivkuma@ecse.rpi.edu

Abstracted from NANOG talks by Dave Katz (Juniper) and Abe Martey (Cisco)

IS-IS Overview

- The Intermediate Systems to Intermediate System Routing Protocol (IS-IS) was originally designed to route the ISO Connectionless Network Protocol (CLNP) . (ISO10589 or RFC 1142)
- Adapted for routing IP in addition to CLNP (RFC1195) as Integrated or Dual IS-IS (1990)
- IS-IS is a Link State Protocol similar to the Open Shortest Path First (OSPF). OSPF supports only IP
- IS-IS competed neck-to-neck with OSPF.
  - OSPF deployed in large enterprise networks
  - IS-IS deployed in several large ISPs

IS-IS Overview

- 3 network layer protocols play together to deliver the ISO defined Connectionless Network Service
  - CLNP
  - IS-IS
  - ES- IS - End System to Intermediate System Protocol
- All 3 protocols independently go over layer 2

CLNS Addressing

NSAP Format

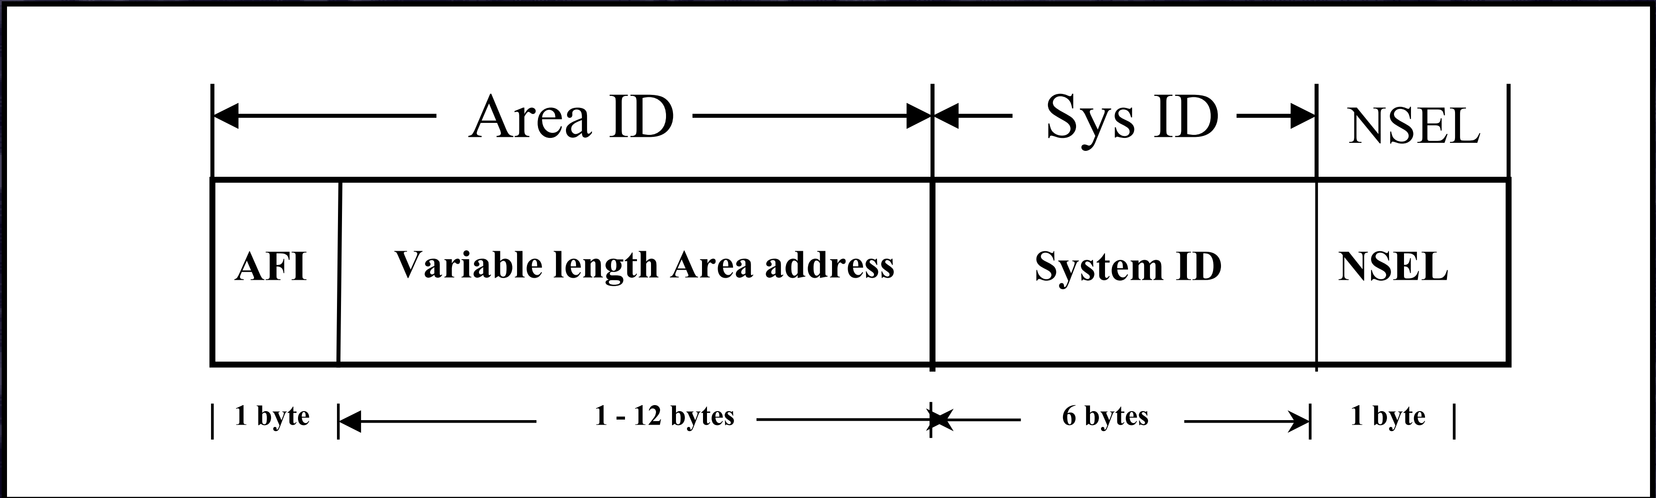

- NSAP format has 3 main components
  - Area ID
  - System ID
  - N-Selector (NSEL) - value is 0x00 on a router
- NSAP of a router is also called a NET

CLNS Addressing

Requirements and Caveats

- At least one NSAP is required per node
- All routers in the same area must have a common Area ID
- Each node in an area must have a unique System ID
- All level 2 routers in a domain must have unique System IDs relative to each other
- All systems belonging to a given domain must have System IDs of the same length in their NSAP addresses

IS-IS Terminology

Intermediate system (IS) - *Router*

Designated Intermediate System (DIS) - *Designated Router*

Pseudonode - *Broadcast link emulated as virtual node by DIS*

End System (ES) - *Network Host or workstation*

Network Service Access Point (NSAP) - *Network Layer Address*

Subnetwork Point of attachment (SNPA) - *Datalink interface*

Packet data Unit (PDU) - *Analogous to IP Packet*

Link State PDU (LSP) - *Routing information packet*

Level 1 and Level 2 – *Area 0 and lower areas*

IS-IS Protocol Concepts: Network Nodes

- Hosts
- Level-1 Routers
- Level-2 Routers
- Level-1 and Level-2 Pseudonodes on broadcast links only

IS-IS Protocol Concepts: Network Nodes

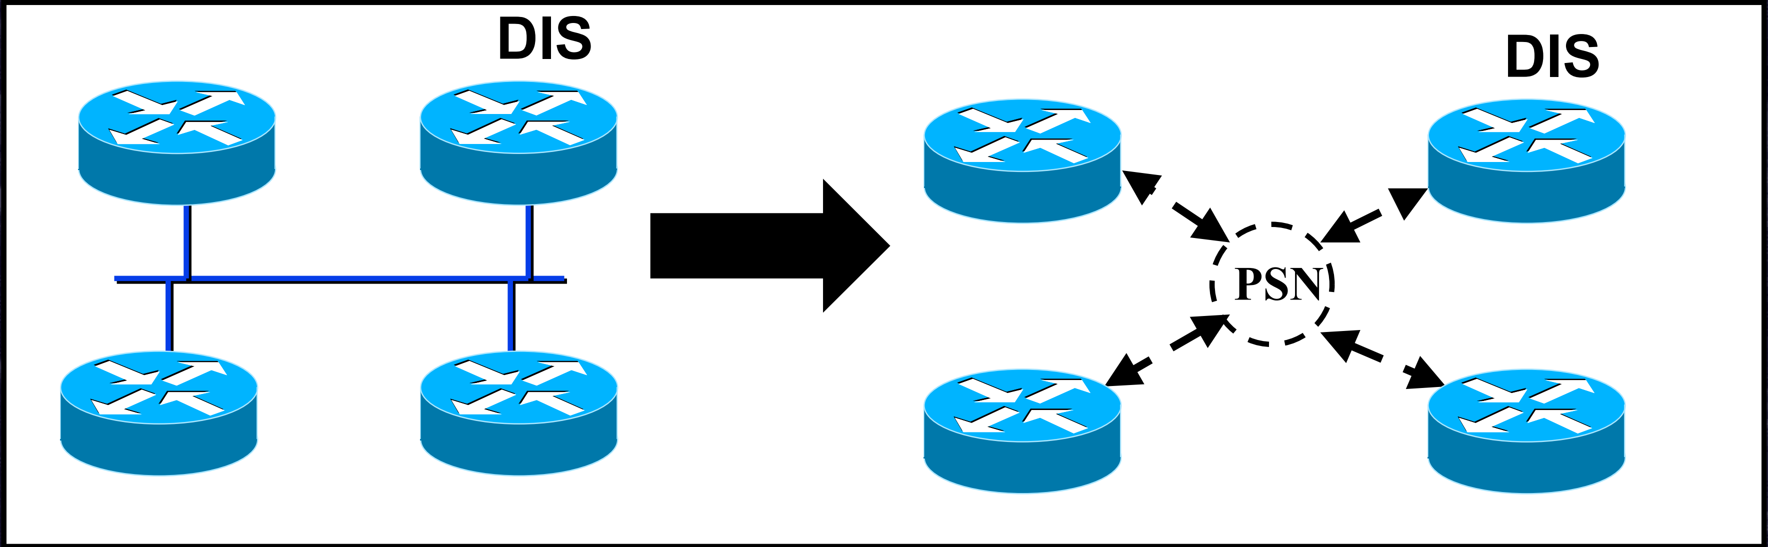

- Broadcast link represented as virtual node, referred to as Pseudonode (PSN)
- PSN role played by the Designated Router (DIS)
- DIS election is preemptive, based on interface priority with highest MAC address being tie breaker
- IS-IS has only one DIS. DIS/PSN functionality supports database synchronization between routers on a broadcast type link

Shivkumar Kalyanaraman

IS-IS Protocol Concepts: Areas

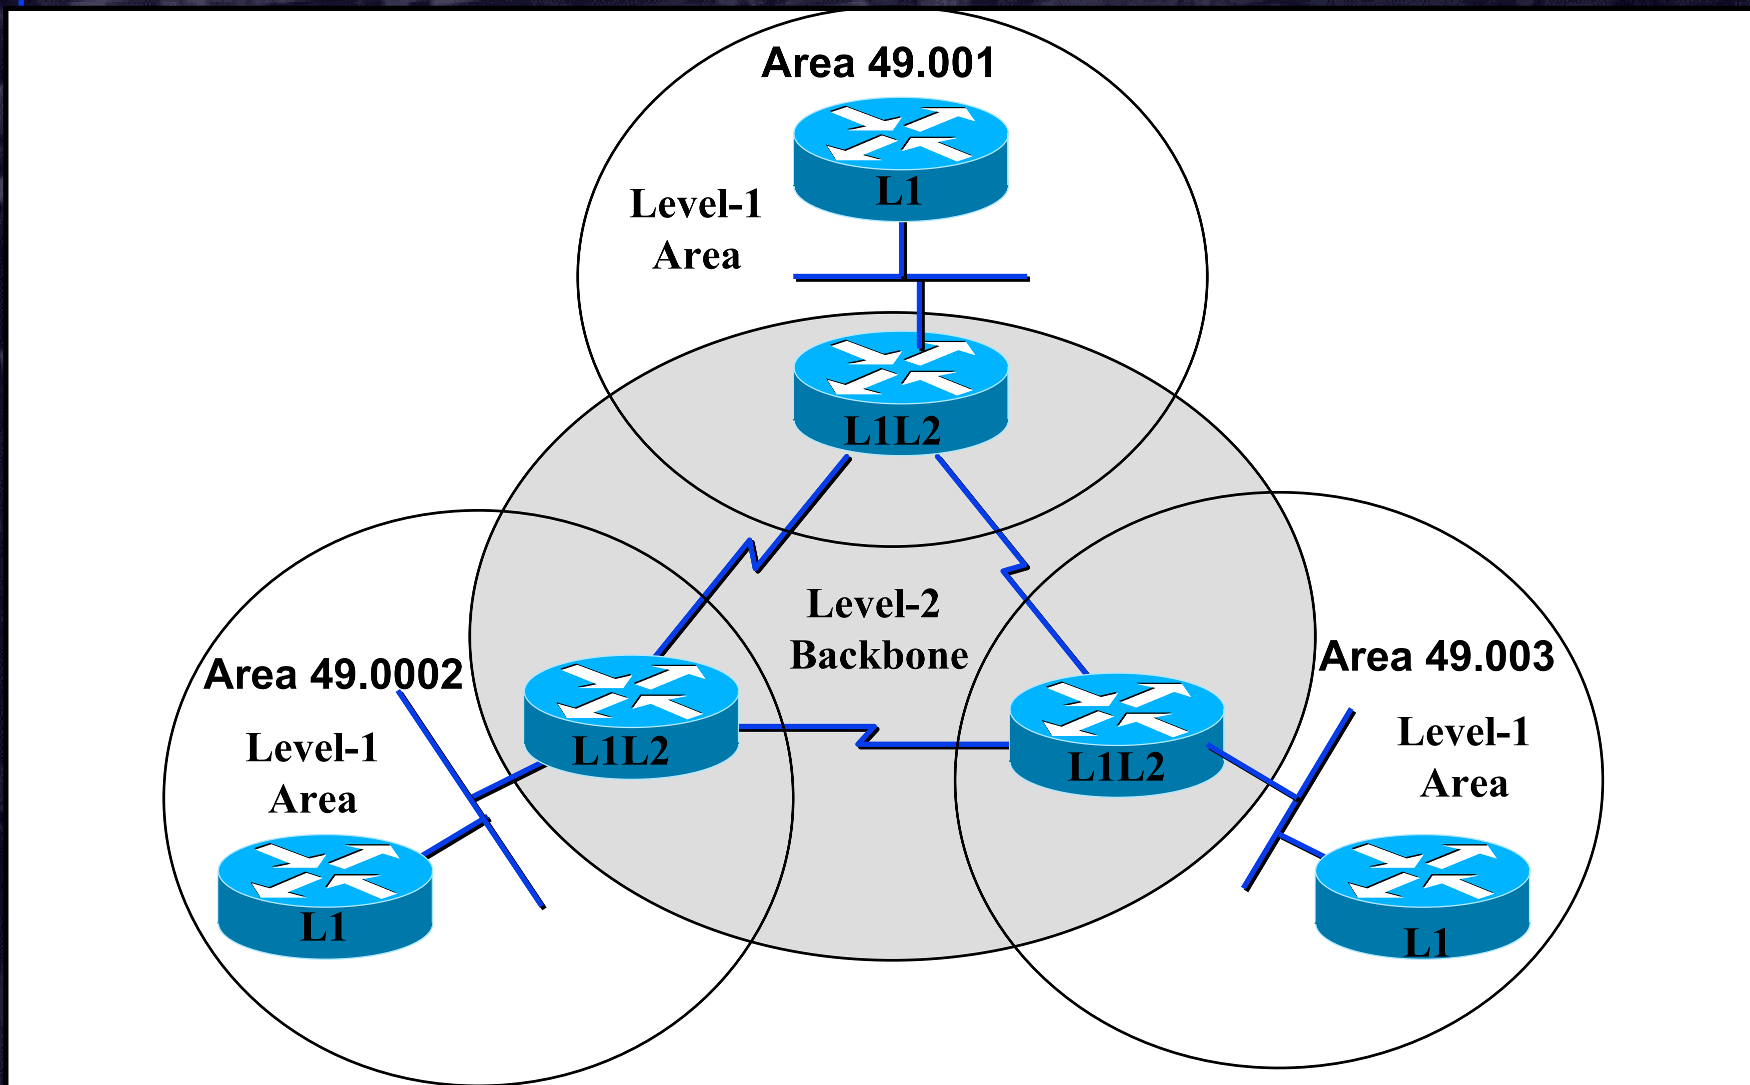

IS-IS Protocol Concepts: Hierarchical Routing

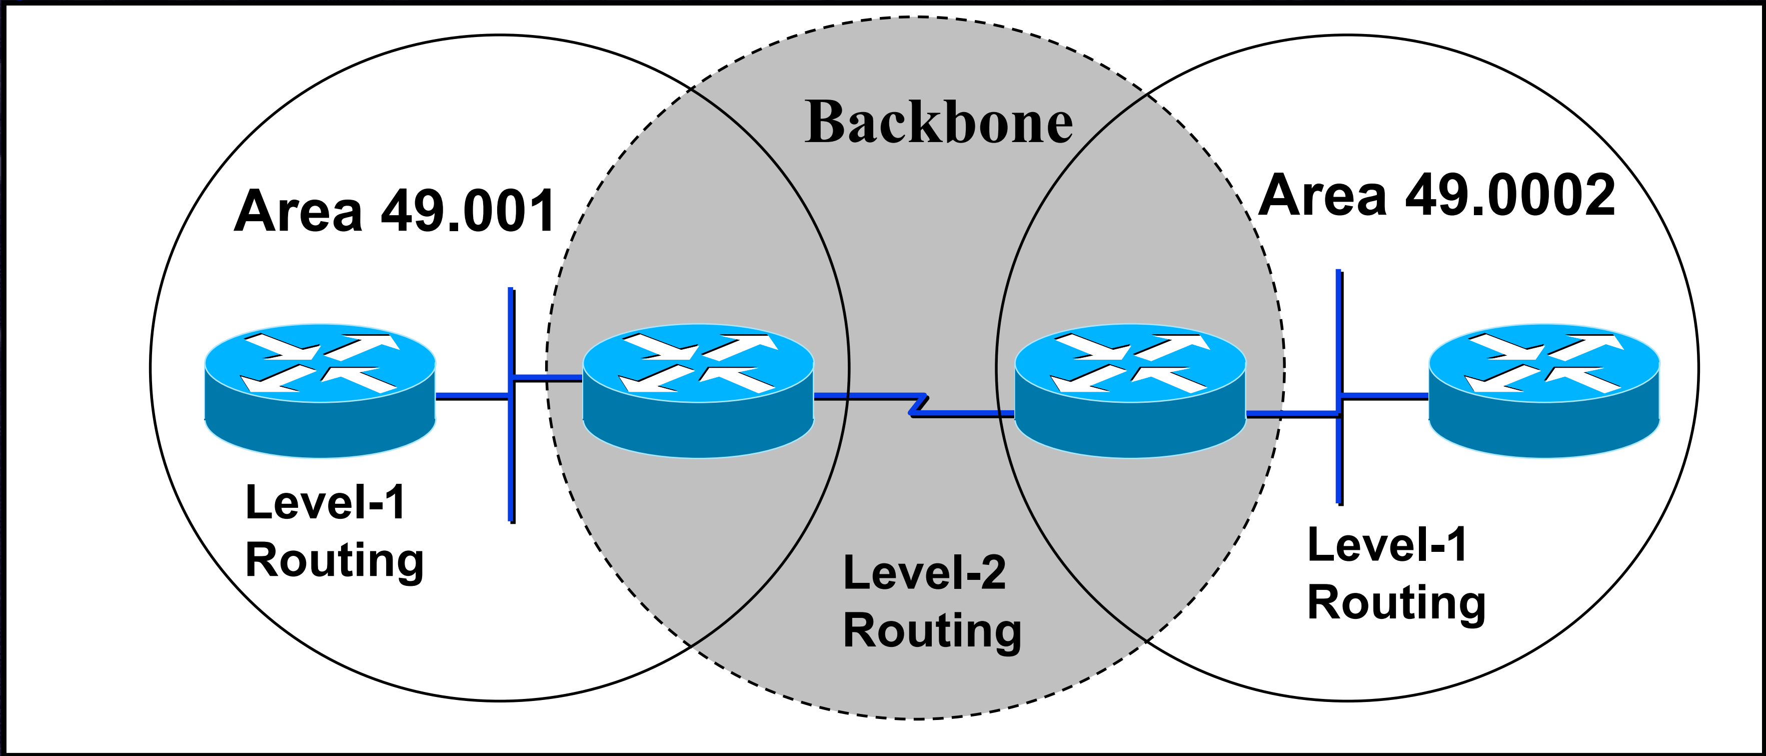

- IS-IS supports 2-level routing hierarchy
- Routing domain is carved into areas. Routing in an area is level-1. Routing between areas is level-2
- All ISO 10589/RFC1195 areas are stubs

IS-IS Protocol Concepts: IS-IS Packet Types

- **IS-IS Hello Packets (IIH)**
  - Level 1 LAN IS-IS Hello
  - Level 2 LAN IS-IS Hello
  - Point-to-point Hello
- **Link State Packets (LSP)**
  - Level 1 and Level 2
- **Complete Sequence Number packets (CSNP)**
  - Level 1 and Level 2
- **Partial Sequence Number Packets (PSNP)**
  - Level 1 and Level 2

IS-IS LS Database: IS-IS Packet Format

- A Fixed Header
  - Contains generic packet information and other specific information about the packet
- Type, Length, Value (TLV) Fields
  - TLVs are blocks of specific routing-related information in IS-IS packets

Shivkumar Kalyanaraman

IS-IS LS Database: Generic Packet Format

Field	No. of Octets
Intra-domain Routing Protocol Discriminator	1
Length Indicator	1
Version/Protocol ID Extension	1
ID Length	1
R | **R** | **R** | **PDU Type**	1
Version	1
Reserved	1
Maximum Area Addresses	1
Packet-Specific Header Fields	
TLV Fields	

IS-IS LS Database: LSP Format

Field	Octets
Intradomain Routing Proto Descriminator	1
Lenth Indicator	1
Version/Protocol ID Extension	1
ID Length	1
R | R | R | PDU Type	1
Version	1
Reserved	1
Maximum Area Addresses	1
PDU Length	2
Remaining Lifetime	2
LSP ID	ID Length + 2
Sequence Number	4
Checksum	2
P | ATT | LSPDBOL | IS Type	1
TYPE LENGTH VALUE FIELDS	Variable

Level-1 TLVs

TLV Name	Type	Origin
Area Address	1	ISO 10589
Intermediate System Neighbors	2	ISO 10589
End System Neighbors	3	ISO 10589
Authentication information	10	ISO 10589
IP Internal Reachability Information	128	RFC 1195
Protocols Supported	129	RFC 1195
IP Interface Address	132	RFC 1195

Level-2 TLVs

TLV Name	Type	Origin
Area Address	1	ISO 10589
Intermediate System Neighbors	2	ISO 10589
Partition Designated Level-2 IS	4	ISO 10589
Prefix Neighbors	5	ISO 10589
Authentication information	10	ISO 10589
IP Internal Reachability Information	128	RFC 1195
Protocols Supported	129	RFC 1195
IP External Reachability Information	130	RFC 1195
Inter-domain Routing Protocol Information	131	RFC 1195
IP Interface Address	132	RFC 1195

High-level Comparison w/ OSPF

- Protocols are recognizably similar in function and mechanism (common heritage)
- Link state algorithms
- Two level hierarchies
- Designated Router on LANs
- Widely deployed (ISPs vs enterprises)
- Multiple interoperable implementations
- OSPF more "optimized" by design (and therefore significantly more complex)
- IS-IS not designed from the start as an IP routing protocol (and is therefore a bit clunky in places)

Detailed comparison points

- Encapsulation
  - OSPF runs on top of IP=> Relies on IP fragmentation for large LSAs
  - IS-IS runs directly over L2 (next to IP) => fragmentation done by IS-IS
- Media support
  - Both protocols support LANs and point-to-point links in similar ways
  - IS-IS supports NBMA in a manner similar to OSPF pt-mpt model: as a set of point-to-point links
  - OSPF NBMA mode is configuration-heavy and risky (all routers must be able to reach DR; bad news if VC fails)

Comparison: Packet Encoding

- OSPF is "efficiently" encoded
  - Positional fields, 32-bit alignment
  - Only LSAs are extensible (not Hellos, etc.)
  - Unrecognized types not flooded. Opaque-LSAs recently introduced.
- IS-IS is mostly Type-Length-Value (TLV) encoded
  - No particular alignment
  - Extensible from the start (unknown types ignored but still flooded)
  - All packet types are extensible
  - Nested TLVs provide structure for more granular extension

Shivkumar Kalyanaraman

Comparison: Area Architecture

- Both protocols support two-level hierarchy of areas
- OSPF area boundaries fall within a router
  - Interfaces bound to areas
  - Router may be in many areas
  - Router must calculate SPF per area
- IS-IS area boundaries fall on links
  - Router is in only one area, plus perhaps the L2 backbone (area)
  - Biased toward large areas, area migration
  - Little or no multilevel deployment (large flat areas work so far)

Shivkumar Kalyanaraman

Comparison: Database Granularity

- OSPF database node is an LSAdvertisement
  - LSAs are mostly numerous and small (one external per LSA, one summary per LSA)
  - Network and Router LSAs can become large
  - LSAs grouped into LSUpdates during flooding
  - LSUpdates are built individually at each hop
  - Small changes can yield small packets (but Router, Network LSAs can be large)

Comparison: Database Granularity

- IS-IS database node is an LSPacket
  - LSPs are clumps of topology information organized by the originating router
  - Always flooded intact, unchanged across all flooding hops (so LSP MTU is an architectural constant--it must fit across all links)
  - Small topology changes always yield entire LSPs (though packet size turns out to be much less of an issue than packet count)
  - Implementations can attempt clever packing

Comparison: Neighbor Establishment

- Both protocols use periodic multicast Hello packets, "I heard you" mechanism to establish 2-way communication
- Both protocols have settable hello/holding timers to allow tradeoff between stability, overhead, and responsiveness
- OSPF requires hello and holding timers to match on all routers on the same subnet (side effect of DR election algorithm) making it difficult to change timers without disruption
- IS-IS requires padding of Hello packets to full MTU size under some conditions (deprecated in practice)
- OSPF requires routers to have matching MTUs in order to become adjacent (or LSA flooding may fail, since LSUpdates are built at each hop and may be MTU-sized)

Neighbor Adjacency Establishment

- OSPF uses complex, multistate process to synchronize databases between neighbors
  - Intended to minimize transient routing problems by ensuring that a newborn router has nearly complete routing information before it begins carrying traffic
  - Accounts for a significant portion of OSPF's implementation complexity
  - Partially a side effect of granular database (requires many DBD packets)
- IS-IS uses its regular flooding techniques to synchronize neighbors
  - Coarse DB granularity => easy (a few CSNPs)

Designated Routers and Adjacency

- Both protocols elect a DR on multi-access networks to remove O(N^2) link problem and to reduce flooding traffic
- OSPF elects both a DR and a Backup DR, each of which becomes adjacent with all other routers
  - BDR takes over if DR fails
  - DRship is sticky, not deterministic
- In IS-IS all routers are adjacent (adjacency less stateful)
  - If DR dies, new DR must be elected, with short connectivity loss (synchronization is fast)
  - DRship is deterministic (highest priority, highest MAC address always wins)
  - DRship can be made sticky by cool priority hack (DR increases its DR priority)

Comparison: LAN Flooding

- OSPF uses multicast send, unicast ack from DR
  - Reduces flood traffic by 50% (uninteresting)
  - Requires per-neighbor state (for retransmissions)
  - Interesting (but complex) acknowledgement suppression
  - Flood traffic grows as O(N)
- IS-IS uses multicast LSP from all routers, CSNP from DR
  - Periodic CSNPs ensure databases are synced (tractable because of coarse database granularity)
  - Flood traffic constant regardless of number of neighbors on LAN
  - But big LANs are uninteresting

Comparison: Routes and Metrics

- IS-IS base spec used 6-bit metrics on links
  - Allowed an uninteresting SPF optimization (CPUs are fast these days)
  - Proved difficult to assign meaningful metrics in large networks
  - Wide metric extension fixes this

- Dual IS-IS spec advertises only default into L1 areas
  - Inter-area traffic routed sub-optimally
  - Route leaking extension addresses this

Comparison: Pragmatic Considerations

- OSPF is much more widely understood
  - Broadly deployed in enterprise market
  - Many books of varying quality available
  - Preserves our investment in terminology
- IS-IS is well understood within a niche
  - Broadly deployed within the large ISP market
  - Folks who build very large, very visible networks are comfortable with it