



Network Layer Routing: RIP & OSPF

Prof. Anja Feldmann, Ph.D.

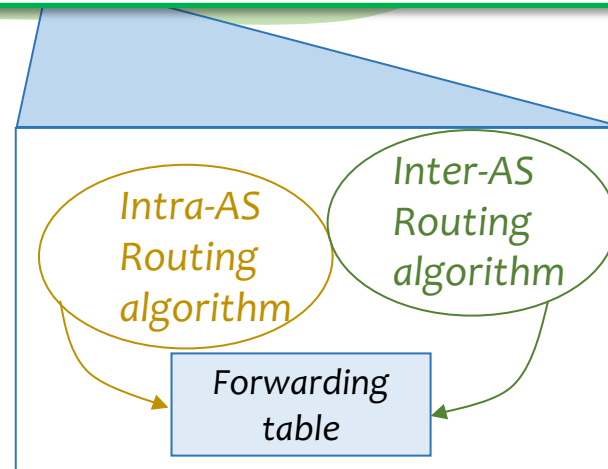


Interconnected ASes



Forwarding table is configured by both intra-AS and inter-AS routing algorithm.

- Intra-AS sets entries for *internal* destinations
- Inter-AS & Intra-AS set entries for *external* destinations




Intra-AS routing



Also known as *Interior Gateway Protocols (IGP)*

Most common Intra-AS routing protocols:

- *Routing Information Protocol (RIP)*
- ***Open Shortest Path First (OSPF)***




Intra-AS routing: OSPF



Open Shortest Path First (OSPF)

- *Link state* protocol (based on *Dijkstra*)
- Routers periodically *flood immediate reachability* info to all other routers
- Distance metric: *administrative weight*



Open Shortest Path First (OSPF)



State

- *Per router information about itself and attached networks*

OSPF advertisements


- *Propagates state*

Link state database

- *State of all routers*

Topology map

- *Derived from link state database*



OSPFv2: Components



Who is my neighbor?

- **Hello** Protocol

With whom I want to talk? (LAN!!!)

- **Designated router** or Backup designated router concept

What information am I missing?


- Database synchronization

How do I distribute info?

- Advertisements disseminated to **entire** Autonomous System (via **reliable flooding**)
- OSPF messages **directly over IP** (rather than TCP or UDP)

Route computation

- From link state database with Dijkstra's algorithm
- Supports equal-cost path routing



OSPF: Neighbor discovery & maintenance




Hello Protocol

- Ensures that neighbors can send packets to and receive packets from the other side: *bi-directional* communication
- Ensures that neighbors agree on parameters (*HelloInterval* and *RouterDeadInterval*)

How?

- **Hello packet** to fixed well-known multicast address
- **Periodic Hellos**
- Broadcast network: Electing designated router




Some multicast addresses



- 224.0.0.5 AllSPFRouters OSPF-ALL.MCAST.NET
- 224.0.0.6 AllDRouters OSPF-DSIG.MCAST.NET

- FF02::5 and FF02::6, respectively for OSPFv3.

- While we are at it:
 - 224.0.0.1 ALL- SYSTEMS. MCAST. NET
 - 224.0.0.2 ALL- ROUTERS. MCAST. NET
 - 224.0.0.9 RIP2- ROUTERS. MCAST. NET
 - 224.0.0.10 IGRP- ROUTERS. MCAST. NET
 - Look up some more (with dig -x address)



Hello Protocol: Phases



Down


- Neighbor is supposed to be “*dead*”
- No communication at all

Init


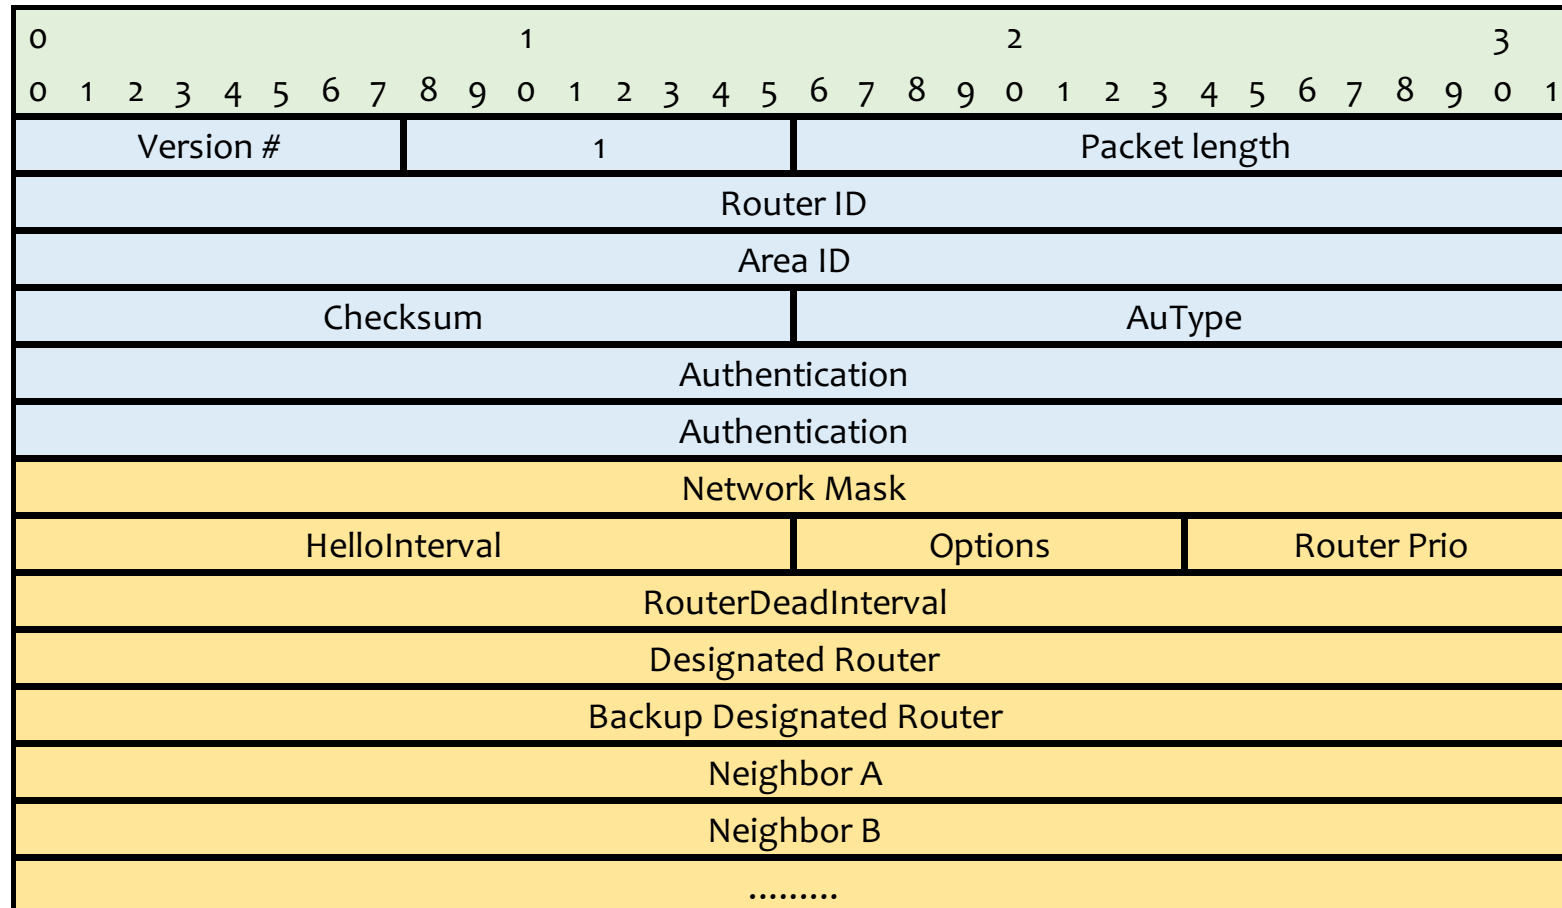
- “*I have heard of a Neighbor*”
- Uni-directional communication

ExStart or TwoWay


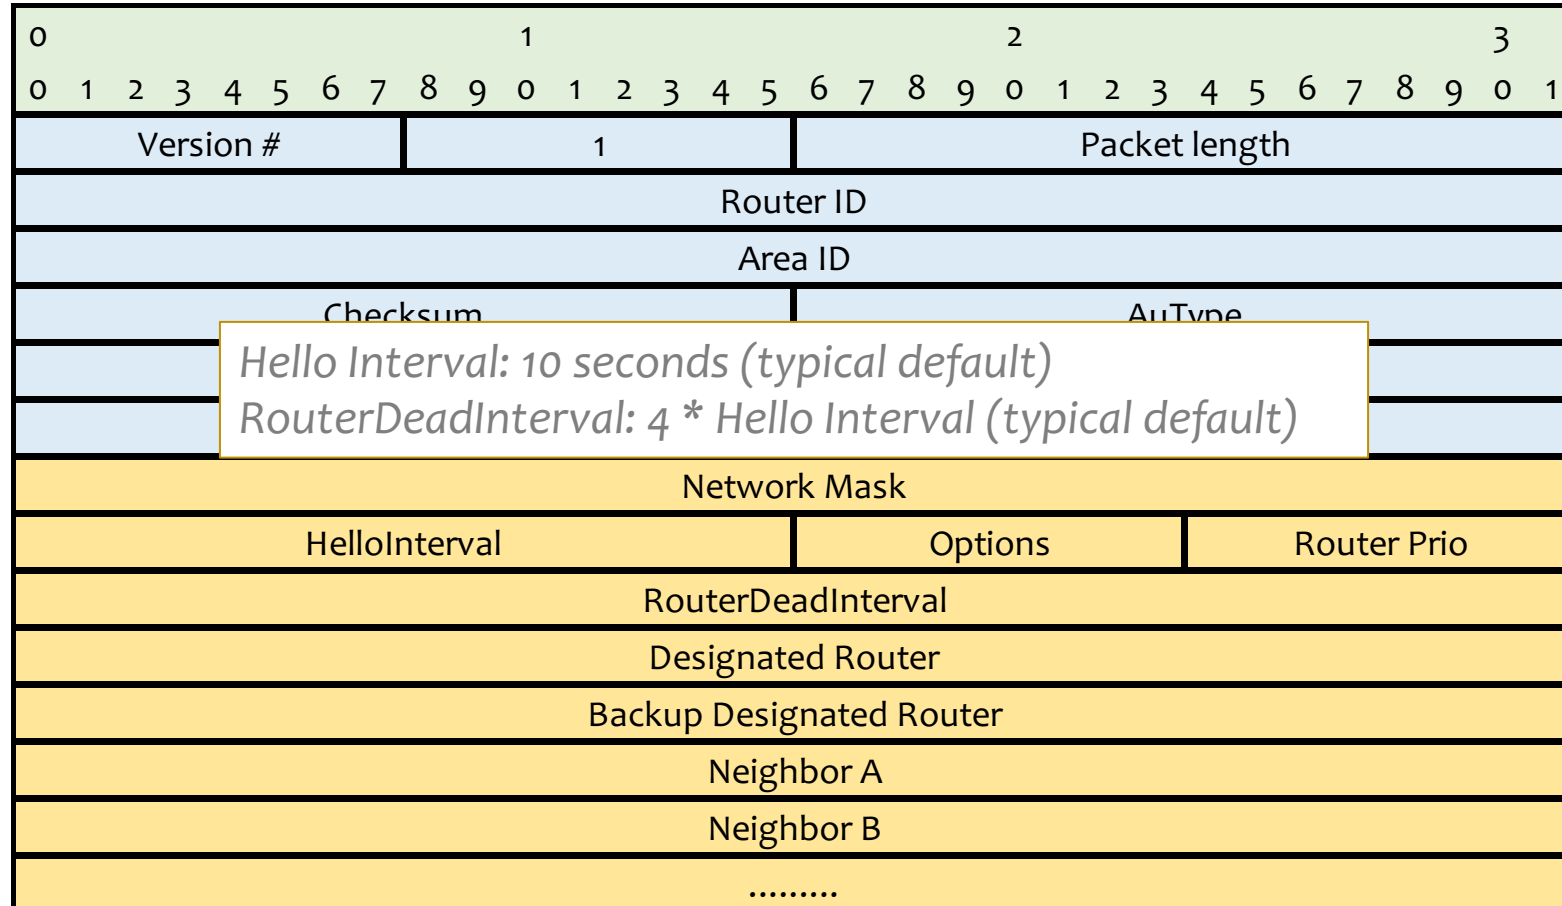
- Communication is bi-directional



Hello Protocol: Packet



Hello Protocol: Packet




OSPF: Packet



- IP Protocol **#89**
- Directly to neighbors using **multicast** address; **TTL 1**

Five packet types

- **Hello**
- **Database Description**
- **Link State Request**
- **Link State Update**
- **Link State Acknowledgement**



OSPF: Link state database




Based on link-state technology



- Local view of topology in a database

Database

- Consists of **Link State Advertisements (LSA)**
- **LSA**: Data unit describing local state of a network/router
- **Must be kept synchronized** to react to routing failures




OSPF: Example network



OSPF: Example link state database




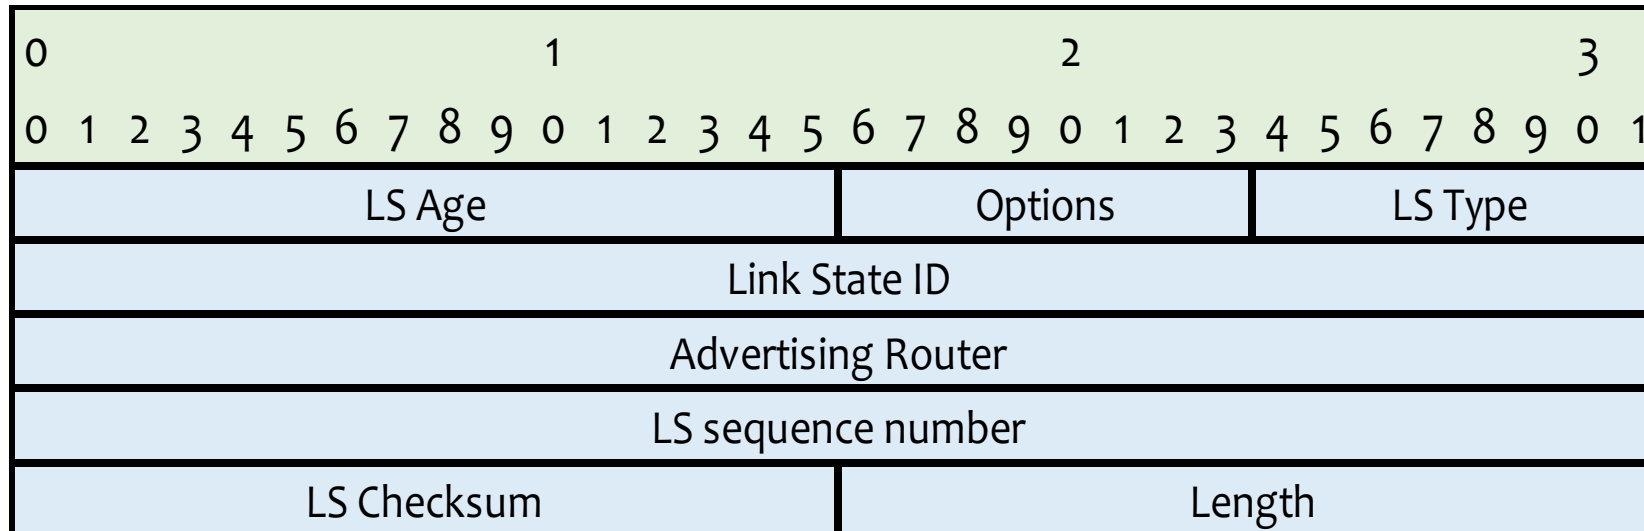
<i>LS-Type</i>	<i>Link State ID</i>	<i>Adv. Router</i>	<i>Checksum</i>	<i>Seq. No.</i>	<i>Age</i>
Router-LSA	10.1.1.1	10.1.1.1	0x9b47	0x80000006	0
Router-LSA	10.1.1.2	10.1.1.2	0x219e	0x80000007	1618
Router-LSA	10.1.1.3	10.1.1.3	0x6b53	0x80000003	1712
Router-LSA	10.1.1.4	10.1.1.4	0xe39a	0x8000003a	20
Router-LSA	10.1.1.5	10.1.1.5	0xd2a6	0x80000038	18
Router-LSA	10.1.1.6	10.1.1.6	0x05c3	0x80000005	1680



OSPF: LSAs



- LSAs consist of a **header** and a **body**
- Header size is **20 bytes**



OSPF: LSAs



Identifying LSAs


- LS Type Field
- Link State ID Field
- Advertising Router Field

Verifying LSA Contents

- LS Checksum Field

Identifying LSA Instances *(keeping in mind that the topology changes)*

- LS Sequence Number Field
- Linear sequence space
- Max. Seq.: New instance




OSPF: LSAs



LS Age Field (to ensure consistency)


- Goal: New sequence number every *30 minutes*
- Maximum value 1 hour
- Age > 1 hour \Rightarrow invalid \Rightarrow removal
- Enables premature aging
- Ensures removal of outdated information



OSPF: Example Router-LSA



0									1									2									3								
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1				
LS Age									Options									LS Type																	
Link State ID																																			
Advertising Router																																			
LS sequence number																																			
LS Checksum									Length																										
0			V	E	B	0									# Link																				
Link ID																																			
Link Data																																			
Type						# TOS						Metric																							
.....																																			




OSPF: Link state database



Is the database synchronized?

- Same number of LSAs?
- Sums of LSA LS Checksums are equal?




OSPF: Database synchronization



*Central aspect: all routers need to have **identical** databases!*

Two types of synchronization

- **Initial** synchronization
 - After hello
- **Continuous** synchronization
 - Flooding




OSPF: Initial synchronization



Explicit transfer of the database upon establishment of *neighbor-ship*

Once *bi-directional* communication exists

- Send *all LS headers* from database to neighbor
 - OSPF *database description (DD)* packets
 - Flood all future LSAs



OSPF: Initial synchronization



Database description (DD) exchange


- Only one DD at a time
- Wait for Ack

Control of DD exchange


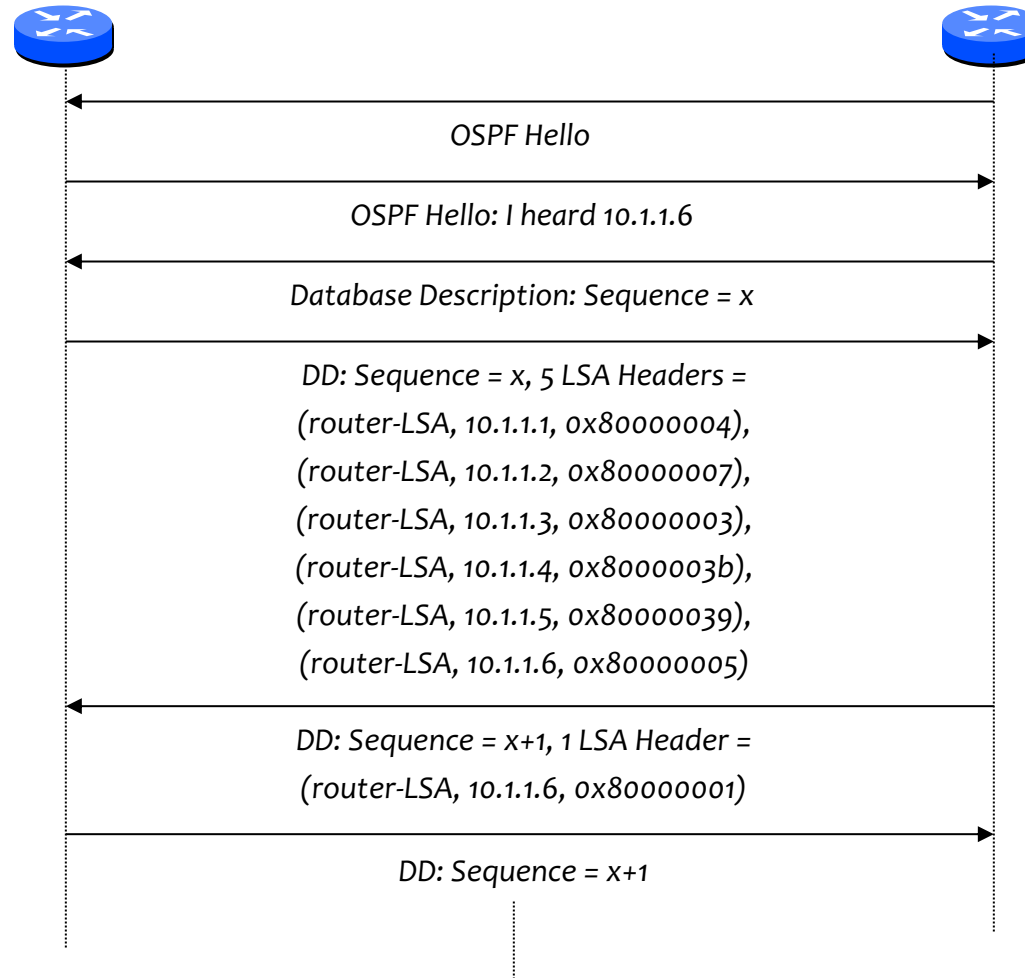
- Determine Master/Slave for DD exchange
- Determine which LSA's are missing in own DB
- Request those via link state request packets
- Neighbor sends these in link state update packets

Result:


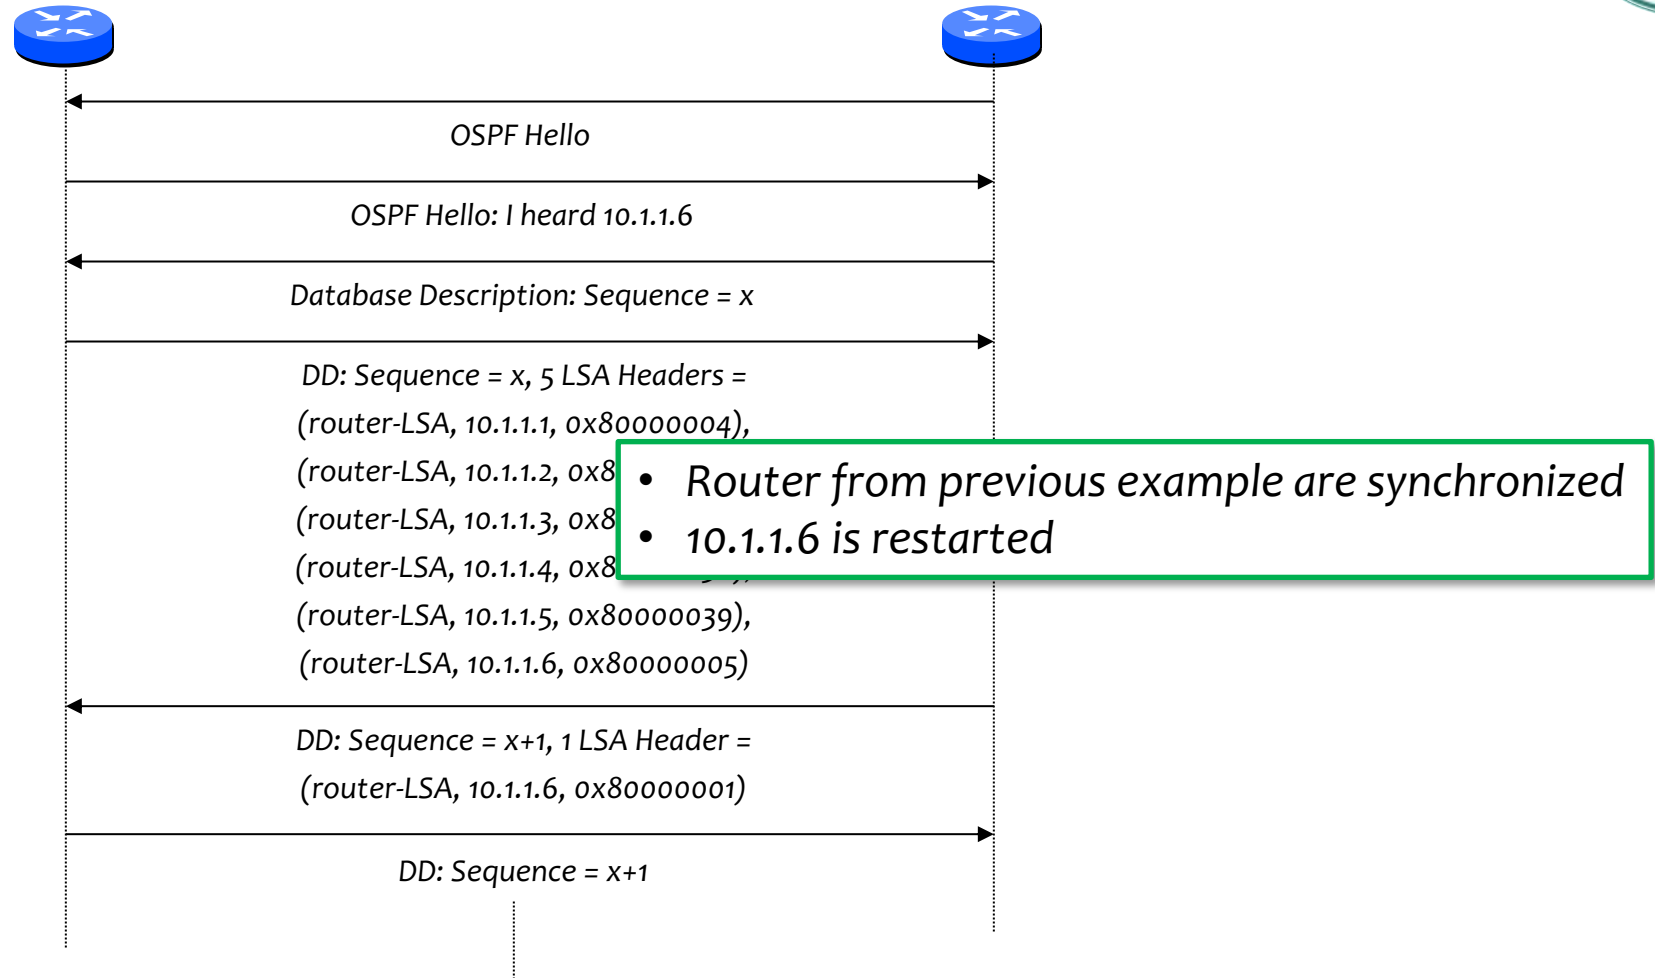
- Fully adjacent OSPF neighbors




OSPF: Example synchronization




OSPF: Example synchronization




OSPF: Reliable flooding




- 10.1.1.3 sends *LS Update*
- Same copy of an LSA is an *implicit ACK*
- Use *delayed* ACKs
- All LSAs *must* be acknowledged either *implicitly* or *explicitly*




OSPF: Reliable flooding




- 10.1.1.3 sends *LS Update*
- Same copy of an LSA is an *implicit ACK*
- Use *delayed* ACKs
- All LSAs *must* be acknowledged either *implicitly* or *explicitly*




OSPF: Reliable flooding




- 10.1.1.3 sends *LS Update*
- Same copy of an LSA is an *implicit ACK*
- Use *delayed* ACKs
- All LSAs *must* be acknowledged either *implicitly* or *explicitly*




OSPF: Reliable flooding




- 10.1.1.3 sends *LS Update*
- Same copy of an LSA is an *implicit ACK*
- Use *delayed* ACKs
- All LSAs *must* be acknowledged either *implicitly* or *explicitly*



OSPF: Reliable flooding




- 10.1.1.3 sends *LS Update*
- Same copy of an LSA is an *implicit ACK*
- Use *delayed* ACKs
- All LSAs *must* be acknowledged either *implicitly* or *explicitly*



OSPF: Robustness of flooding




- More robust than a *spanning tree*
- LSA refreshes every *30 minutes*
- LSAs ...
 - Have checksums
 - Are aged
 - Cannot be send at arbitrary rate: There are *timers!*



OSPF: LSA timers



MinLSArrival	1 second
MinLSInterval	5 seconds
CheckAge	5 minutes
MaxAgeDiff	15 minutes
LSRefreshTime	30 minutes
MaxAge	1 hour




OSPF: Routing table calculation





Link state database is a *directed graph with costs* for each link

Dijkstra's SPF algorithm


- Add all routers to shortest-path-tree
- Add all neighbors to candidate list
- Add routers with the smallest cost to tree
- Add neighbors of this router to candidate list
 - If not yet on it
 - If cost smaller
- Continue until candidate list empty




OSPF: Example routing table calc.




OSPF: Example routing table calc.





10.1.1.5 (1, 10.1.1.5)
10.1.1.2 (3, 10.1.1.2)
10.1.1.1 (5, 10.1.1.1)



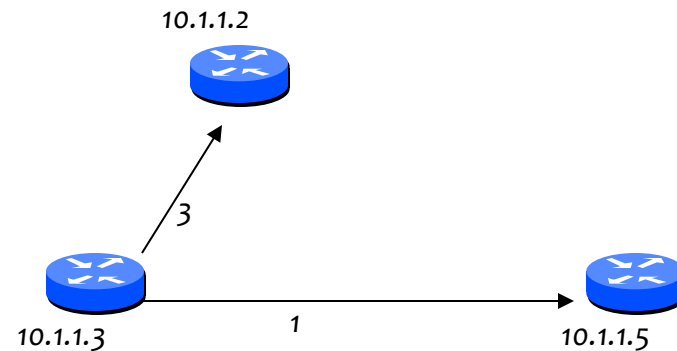

OSPF: Example routing table calc.




10.1.1.2 (3, 10.1.1.2)
10.1.1.4 (4, 10.1.1.5)
10.1.1.1 (5, 10.1.1.1)
10.1.1.6 (11, 10.1.1.5)





OSPF: Example routing table calc.




10.1.1.4 (4, 10.1.1.5/2)
10.1.1.1 (5, 10.1.1.1)
10.1.1.6 (11, 10.1.1.5)





OSPF: Example routing table calc.



10.1.1.1 (5, 10.1.1.1)
10.1.1.6 (10, 10.1.1.5/2)



OSPF: Example routing table calc.



OSPF: Example routing table calc.

