Lab Course „RouterLab“

Open Shortest Path First (OSPF)

Miscellaneous

- Don't set enable passwords on Cisco router
- ...

OSPF (Open Shortest Path First)

- "Open": specification publicly available
  - RFC 1247, RFC 2328
  - Working group formed in 1988
  - Goals:
    - Large, heterogeneous internetworks
- Uses the Link State algorithm
  - Topology map at each node
  - Route computation using Dijkstra's algorithm
- Hierarchy

OSPFv2: Tasks (to be filled in)

- Neighbor discovery and maintenance
  - „Hello“ Protocol
- Link state database
  - kept at each router
  - stores topological information
  - needs to be synchronized with neighbors
- Calculation of routing table
  - Dijkstra

OSPFv2: Components

- Hello Protocol: “Who is my neighbor?”
- Database Synch: “What info am I missing?”
  - after establishing OSPF neighborhood
- Reliable flooding algo: “How do I distribute info?”
  - after something has changed (link failure, cost changes)
- Route computation
  - Based on link state database
  - Using Dijkstra’s algorithm

OSPF Packets

- 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

Neighbor Discovery and 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

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 -> **different LSA types!**
  - Must kept synchronized to react to routing failures

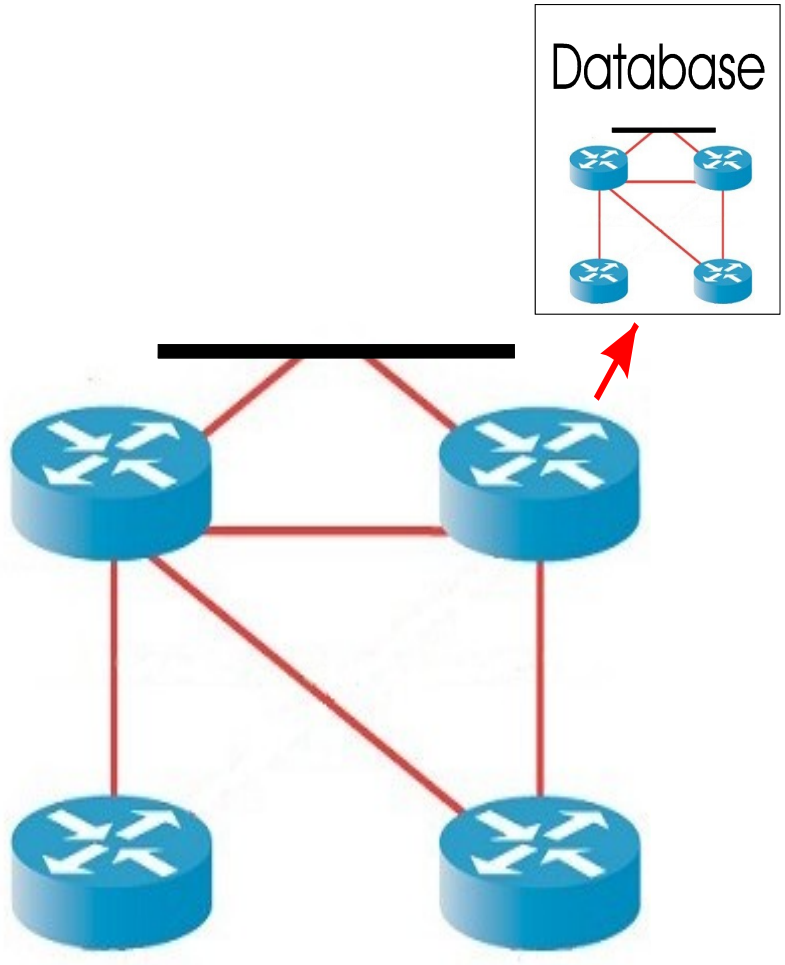

Database synchronization

- Central aspect:
  all routers need to have identical databases!
- 2 types of synchronization
  - Initial synchronization
    - After hello
  - Continuous synchronization
    - Flooding

Initial Synchronization

- Explicit transfer of the database upon establishment of neighborship
- Once bi-directional communication exists
- Send all LS **header** from database to neighbor
  - OSPF database description packets (DD pkt)
  - Flood all future LSA's

Initial Synchronization (2.)

- Database description (DD) exchange
  - Only one DD at a time
  - Wait for Ack
- Control of 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

Database Synchronization - Example

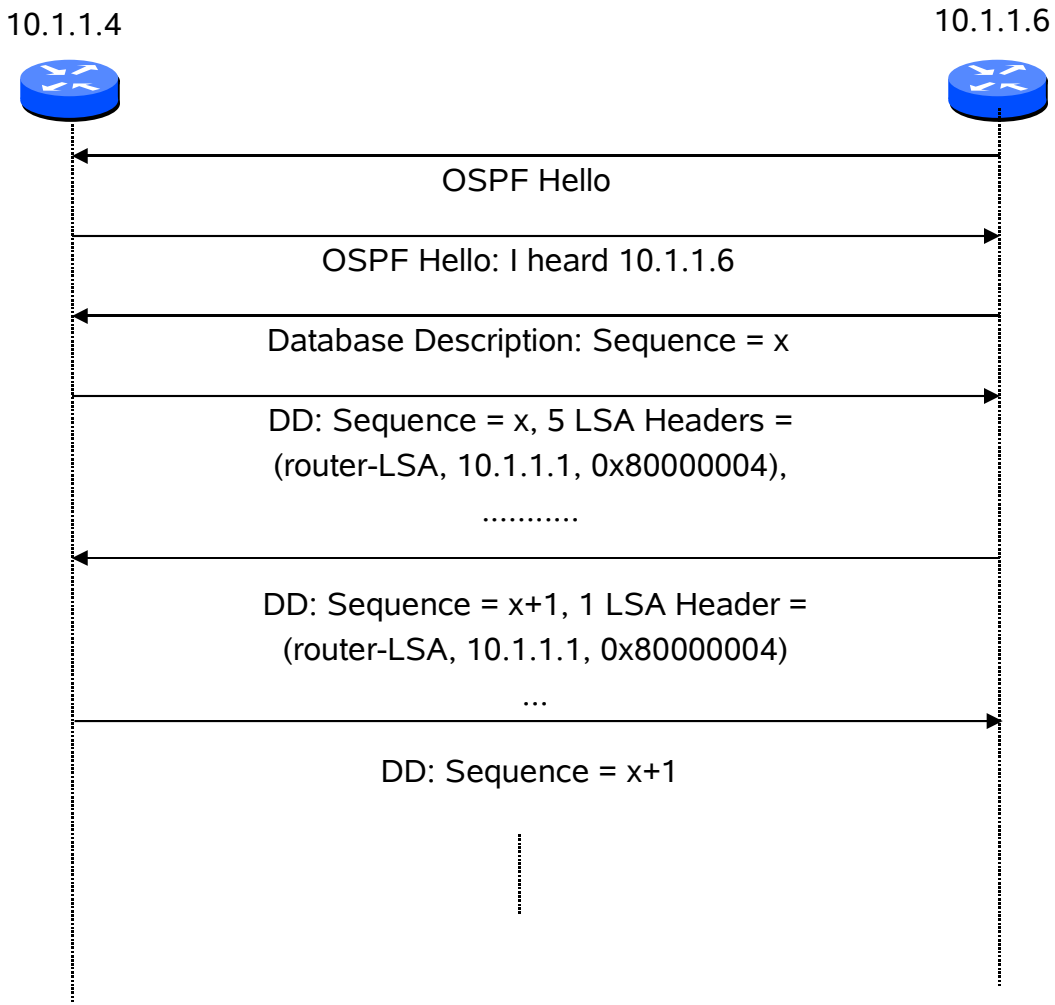

Reliable Flooding

- E.g., after something changes
  - link failure
  - OSPF cost change for a link
- Robustness
  - LSA refreshes every 30 minutes
  - LSAs have checksums
  - LSAs are aged
  - LSAs cannot be send at arbitrary rate: there are timers

Calculation of routing table

- Link state database is a directed graph with costs for each link
- Use Dijkstra to compute paths from source to all destinations
- More info on Dijkstra: Check the web ...

Network Types

- So far only point-to-point
- Many other technologies
  - Point-to-point
  - Broadcast
  - ...

Hierarchical OSPF

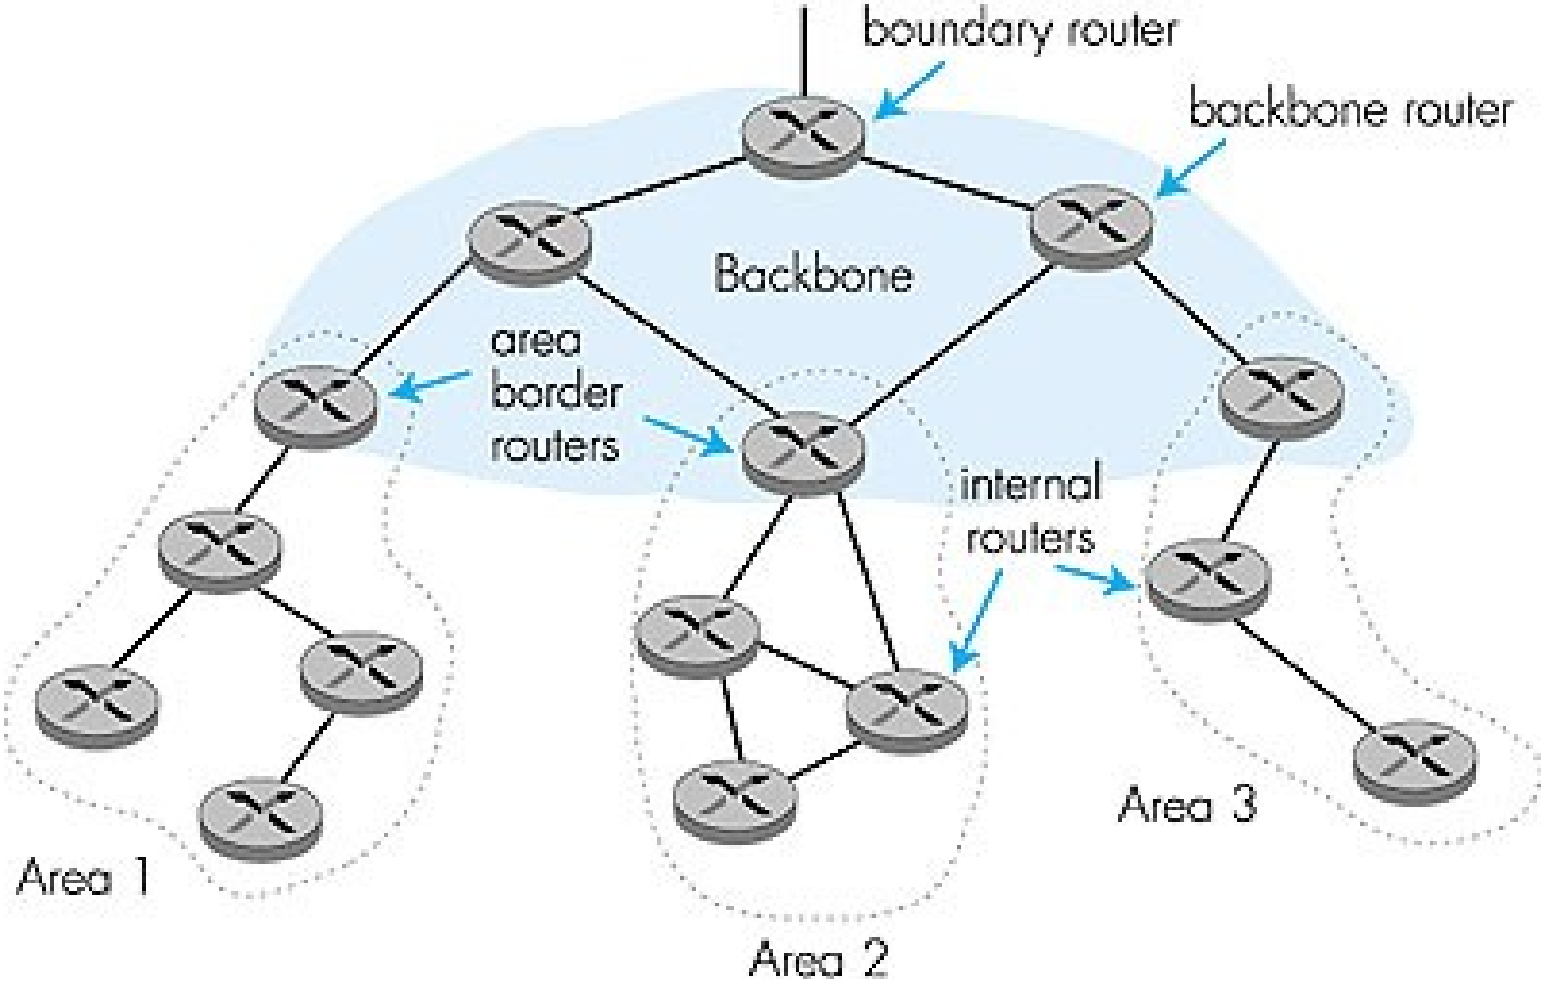

Hierarchical OSPF

- Two-level hierarchy: local area and backbone.
  - Link-state advertisements do not leave respective areas.
  - Nodes in each area have detailed area topology; they only know direction (shortest path) to networks in other areas.
- Area Border routers: "summarize" distances to networks in the area and advertise them to other Area Border routers.
- Backbone routers: run an OSPF routing algorithm limited to the backbone.
- Boundary routers: connect to other ASs.

Areas

- An AS (or Routing Domain) is divided into areas.
- Group of routers
- “Close” to each other.
- Reduce the extend of LSA flooding
- Intra-area traffic
- Inter-area traffic
- External traffic: injected from a different AS
- OSPF requires a backbone area (Area 0)
  - Routing between areas only via backbone area
  - Strict area hierarchy (no loops allowed)

OSPF: Summary

- Neighbors
  - Discovery — Multicast group
  - Maintenance — Hello protocol
- Database
  - Granularity — Link state advertisements (LSA)
  - Synchronization — Initial synchronization
Reliable flooding
- Routing table
  - Calculation — Local shortest path calculation

Lab Course „RouterLab“

Work sheet 3: Questions

Quagga (Question 1)

- „Software router“ for various platforms
- Supports many routing protocols
- Collection of several daemons
  - *zebra*: „communication“ with kernel and integration of all daemons
  - *ripd*: RIP support
  - *ospfd*: OSPF support
  - ...
- User interface shell: VTY
- Apart from that: Similiar configuration as on previous work sheets
  - if you wish you can reuse existing dumps

Work sheet 3

- Question 2
  - basic OSPF setup
  - reachability over multiple hops
- Question 3
  - Link-State Database
  - Database Synchronization

Question 4

- Link state database is a directed graph with costs for each link
- Dijkstra‘s SPF algorithms
  - 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
- “Run the algorithm by hand“!

Question 5

- OSPF costs
- OSPF hierarchy, areas
- intra- and interarea paths