

Introduction to OSPF

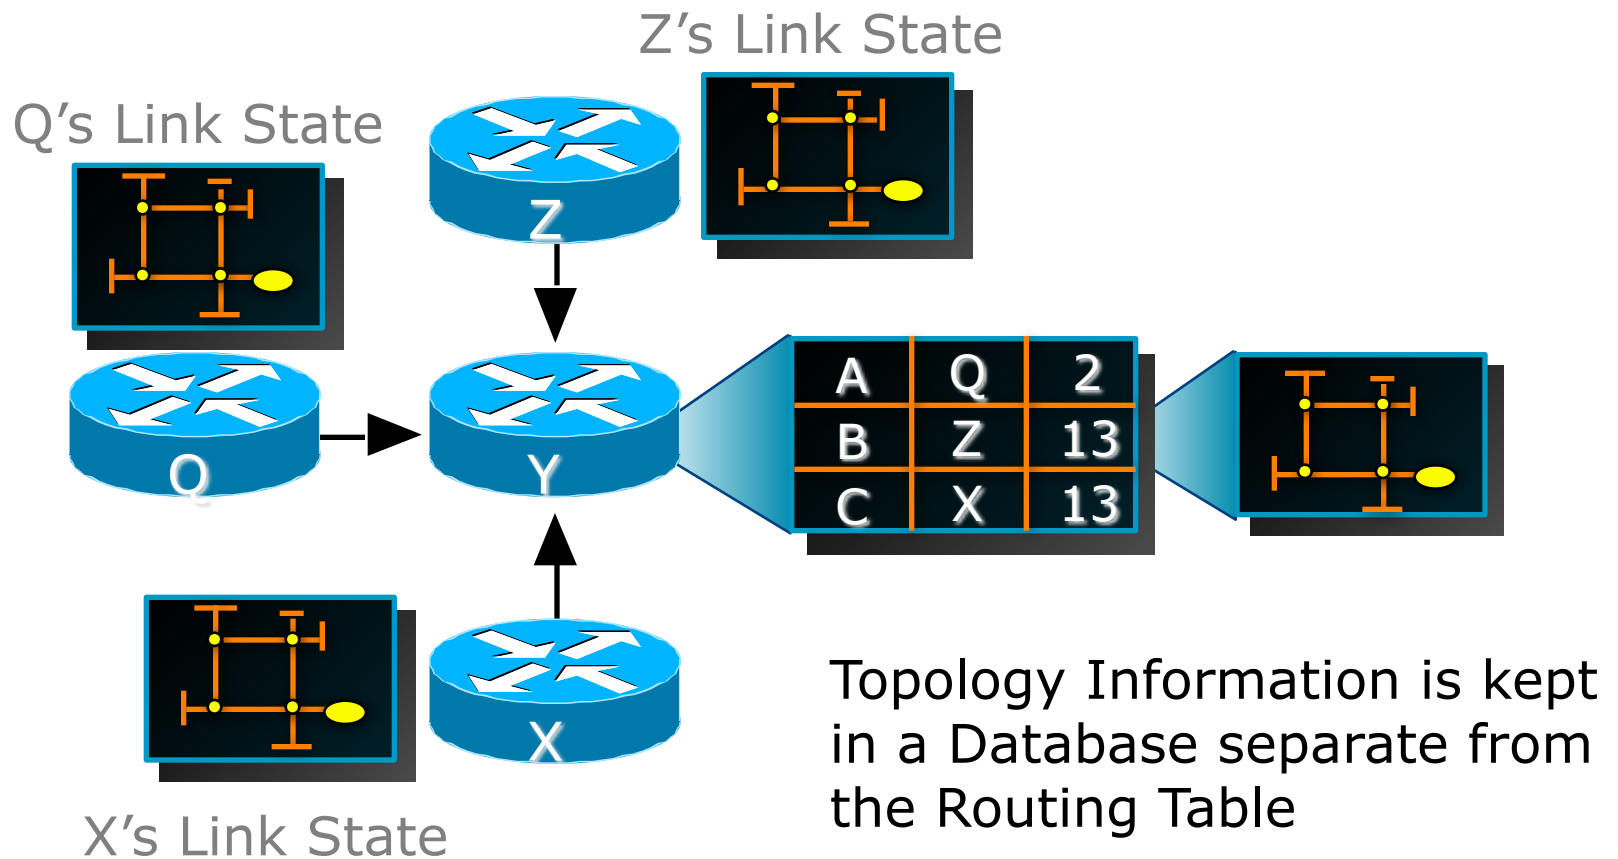


ISP Training Workshops

OSPF

- p Open Shortest Path First
- p Link state or SPF technology
- p Developed by OSPF working group of IETF (RFC 1247)
- p OSPFv2 standard described in RFC2328
- p Designed for:
 - TCP/IP environment
 - Fast convergence
 - Variable-length subnet masks
 - Discontiguous subnets
 - Incremental updates
 - Route authentication
- p Runs on IP, Protocol 89

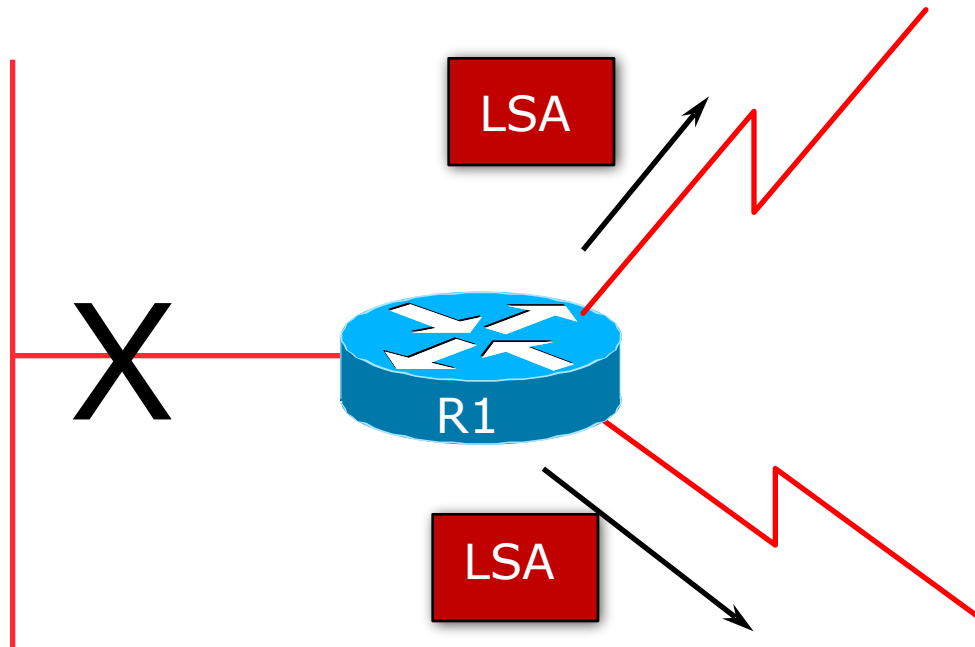
Link State



Link State Routing

- p Neighbour discovery
- p Constructing a Link State Packet (LSP)
- p Distribute the LSP
 - (Link State Announcement – LSA)
- p Compute routes
- p On network failure
 - New LSPs flooded
 - All routers recompute routing table

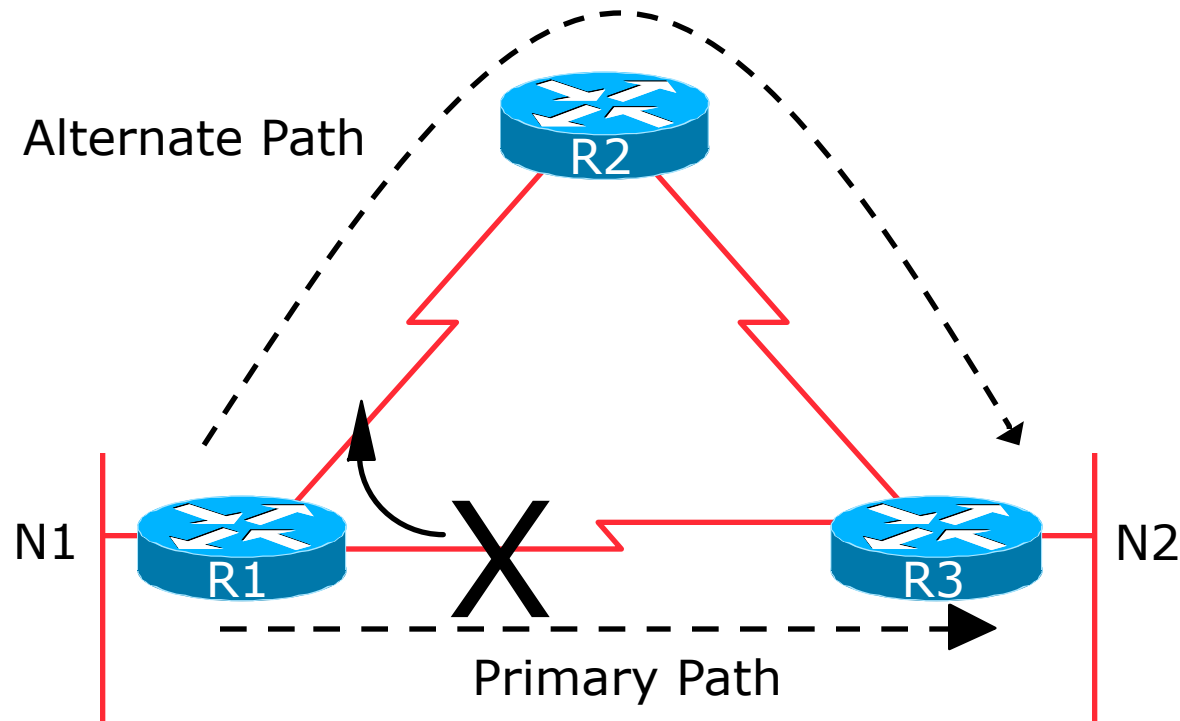
Low Bandwidth Utilisation



- ⌘ Only changes propagated
- ⌘ Uses multicast on multi-access broadcast networks

Fast Convergence

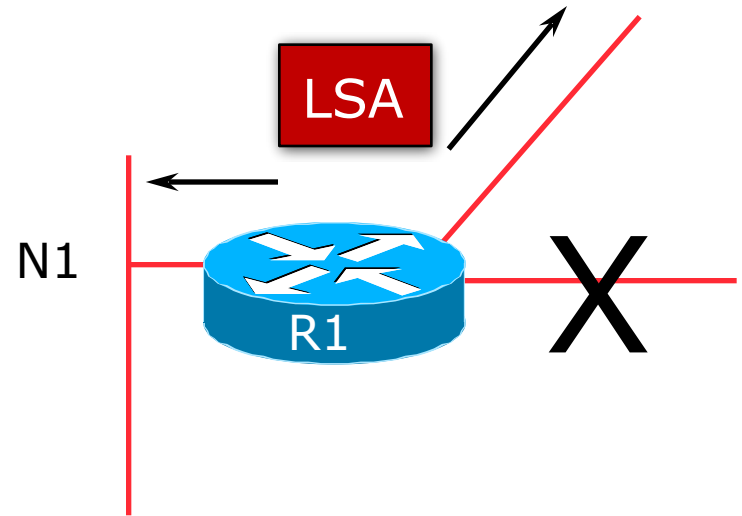
- Detection Plus LSA/SPF
 - Known as the Dijkstra Algorithm



Fast Convergence

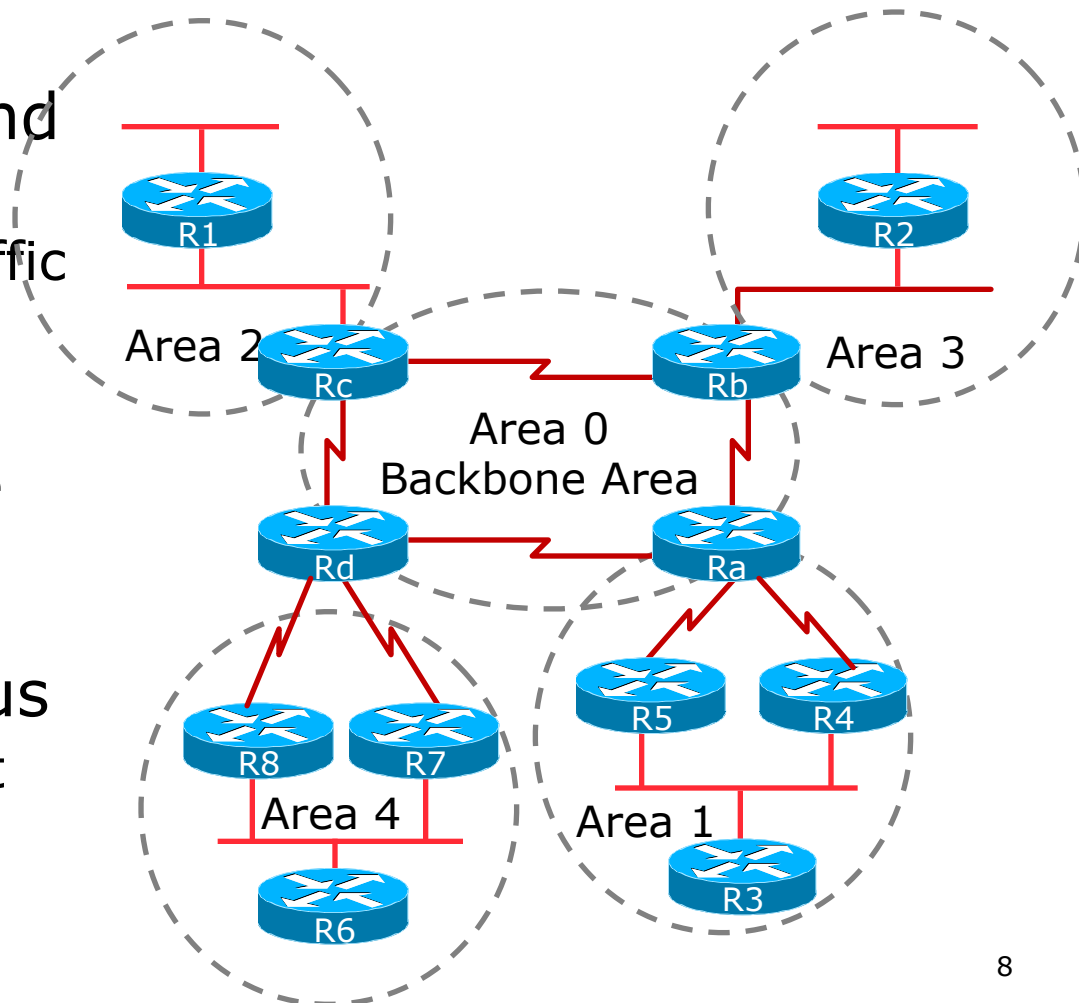
p Finding a new route

- LSA flooded throughout area
- Acknowledgement based
- Topology database synchronised
- Each router derives routing table to destination network



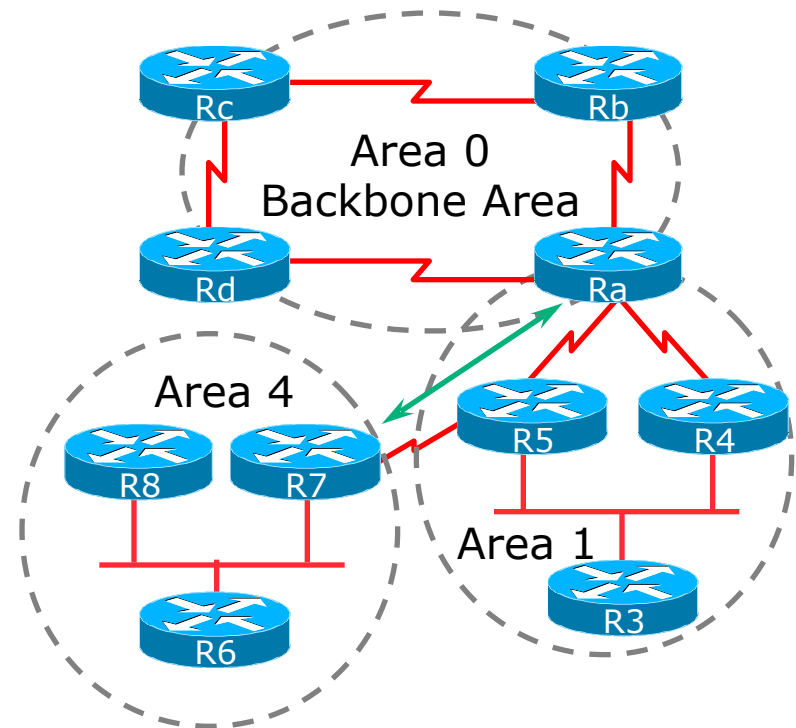
OSPF Areas

- Area is a group of contiguous hosts and networks
 - Reduces routing traffic
- Per area topology database
 - Invisible outside the area
- Backbone area **MUST** be contiguous
 - All other areas must be connected to the backbone

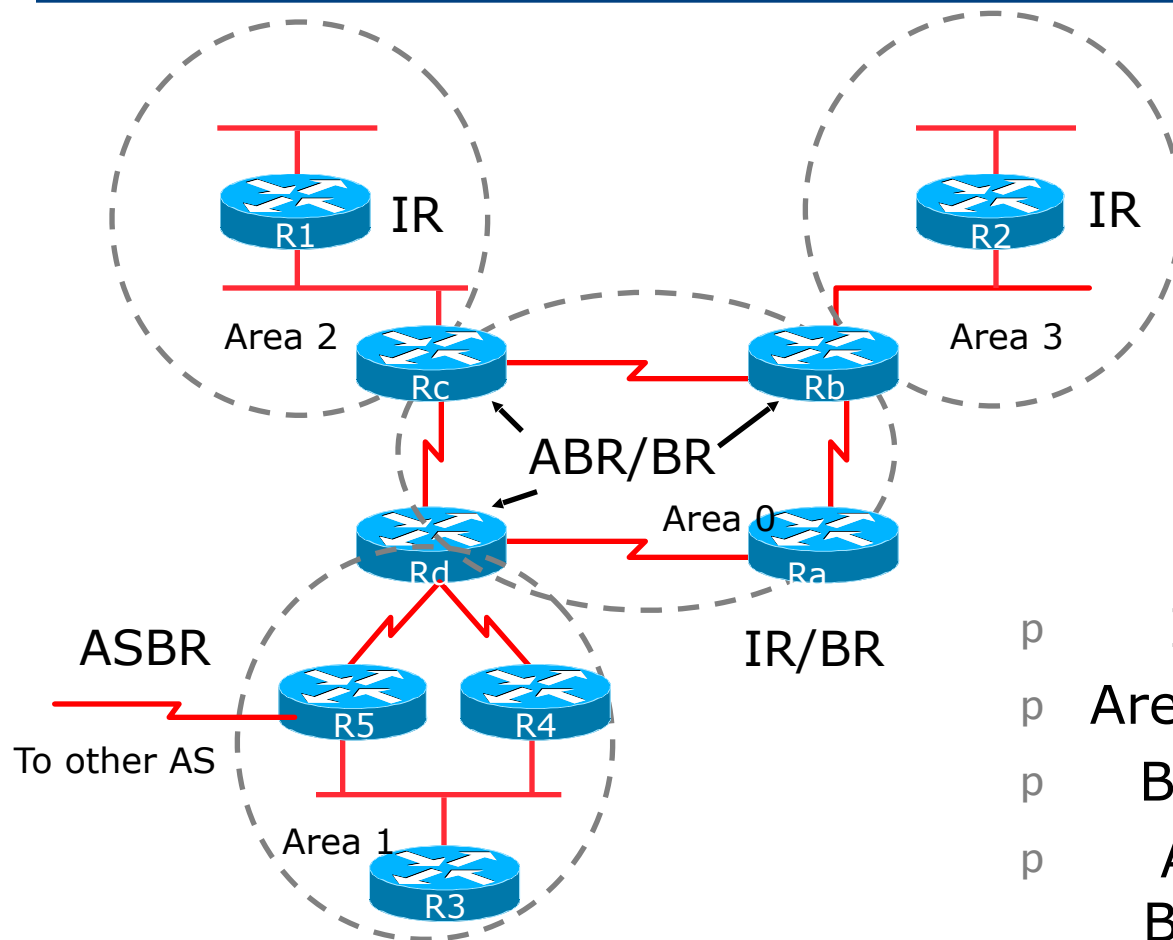


Virtual Links between OSPF Areas

- Virtual Link is used when it is not possible to physically connect the area to the backbone
- ISPs avoid designs which require virtual links**
 - Increases complexity
 - Decreases reliability and scalability

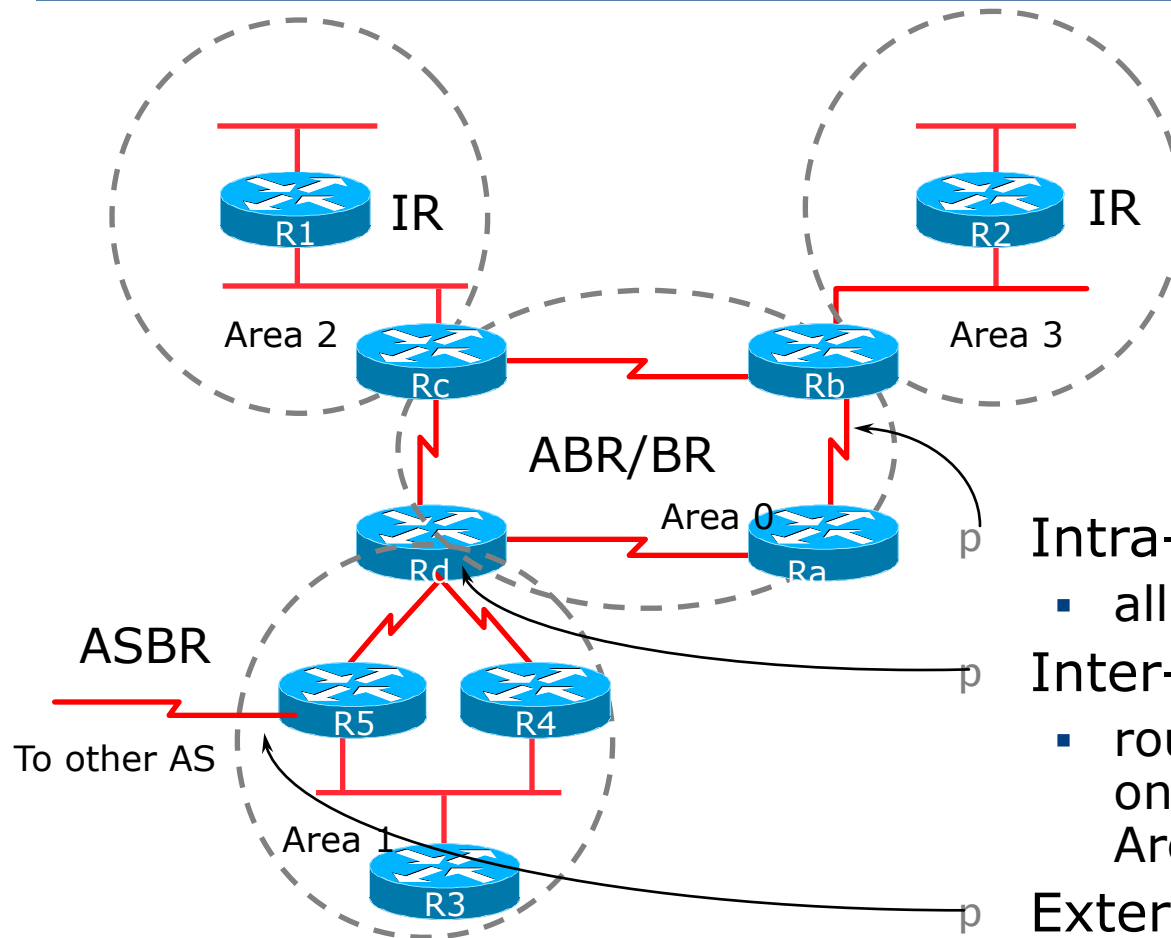


Classification of Routers



- p Internal Router (IR)
- p Area Border Router (ABR)
- p Backbone Router (BR)
- p Autonomous System Border Router (ASBR)

OSPF Route Types



Intra-area Route

- all routes inside an area

Inter-area Route

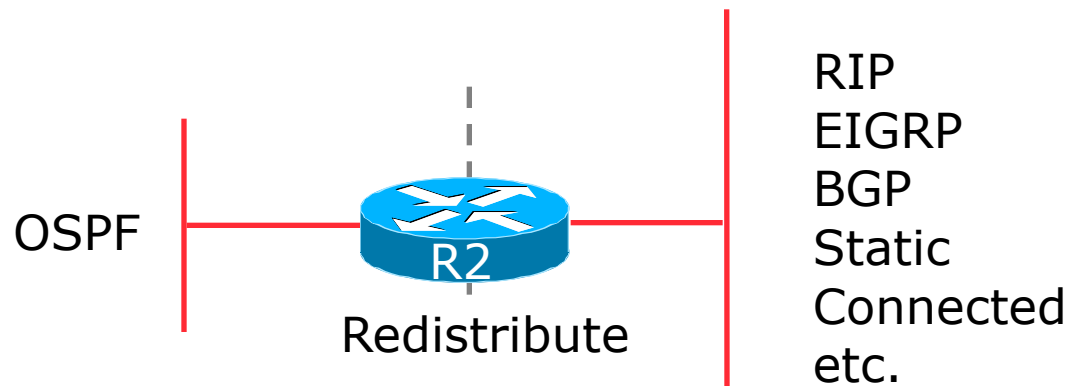
- routes advertised from one area to another by an Area Border Router

External Route

- routes imported into OSPF from other protocol or static routes

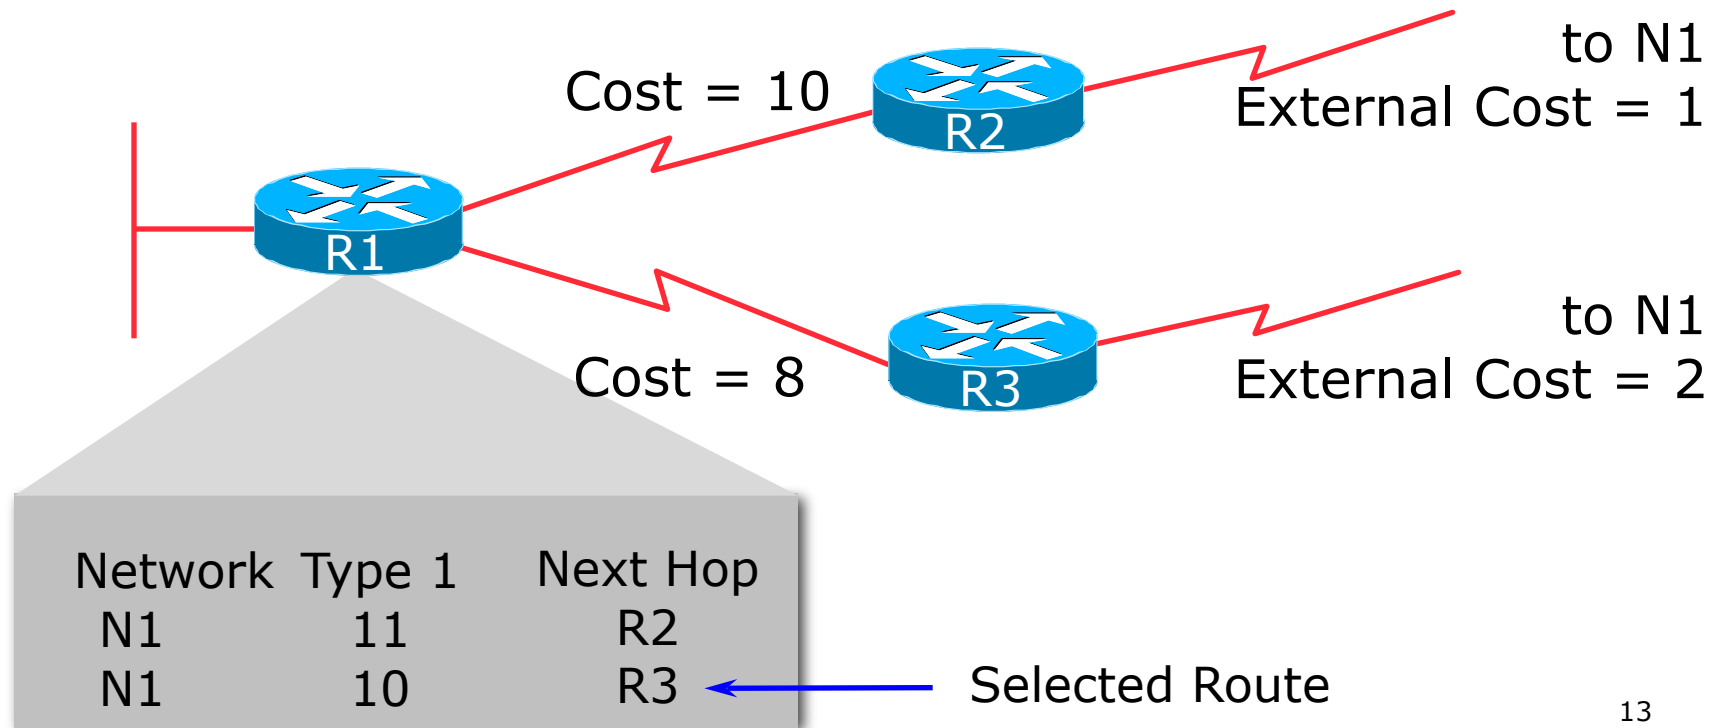
External Routes

- Prefixes which are redistributed into OSPF from other protocols
- Flooded unaltered throughout the AS
 - Recommendation: Avoid redistribution!!**
- OSPF supports two types of external metrics
 - Type 1 external metrics
 - Type 2 external metrics (Cisco IOS default)



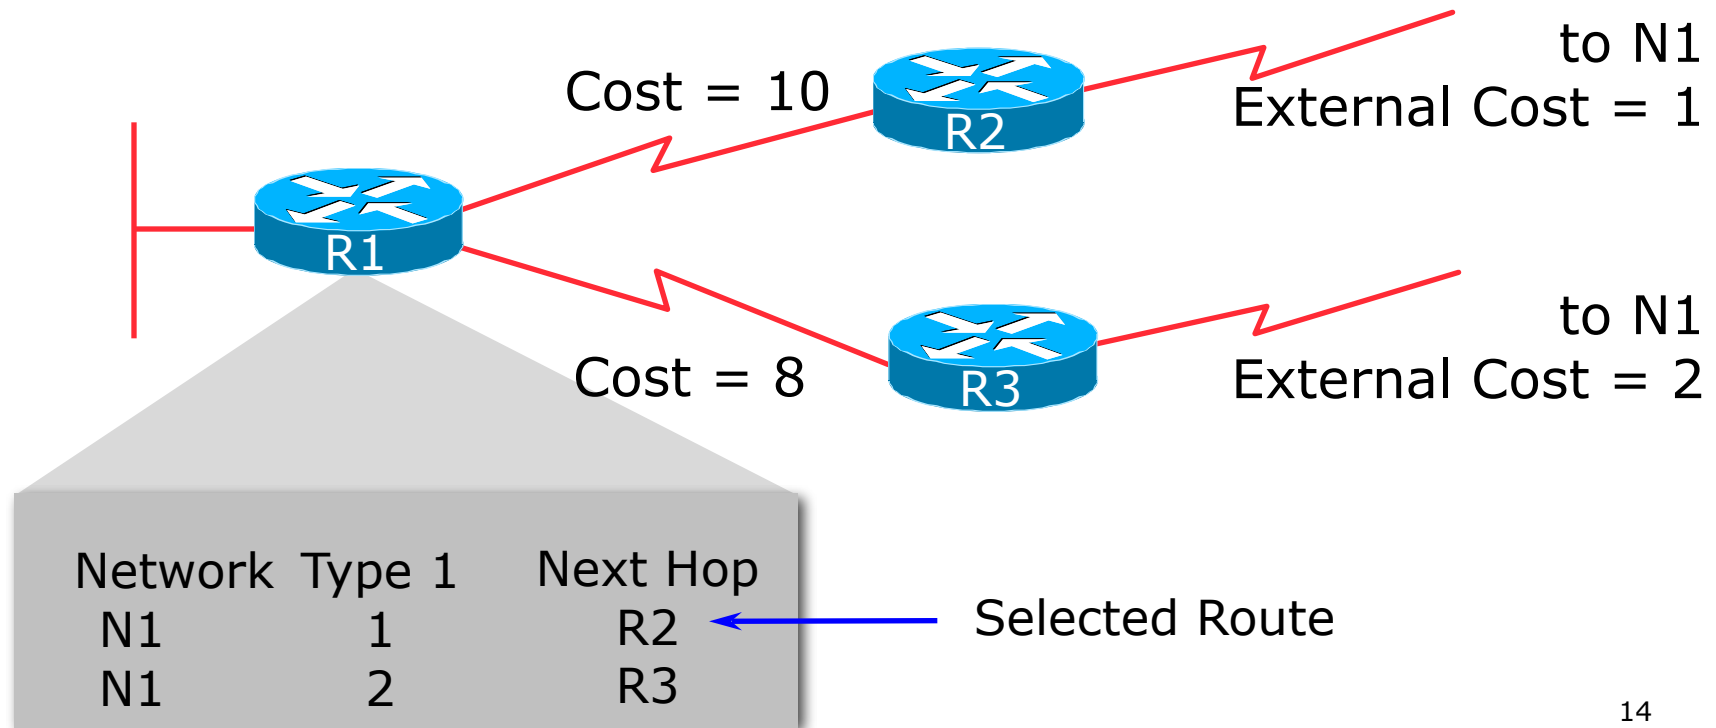
External Routes

- Type 1 external metric: metrics are added to the summarised internal link cost



External Routes

- Type 2 external metric: metrics are compared without adding to the internal link cost

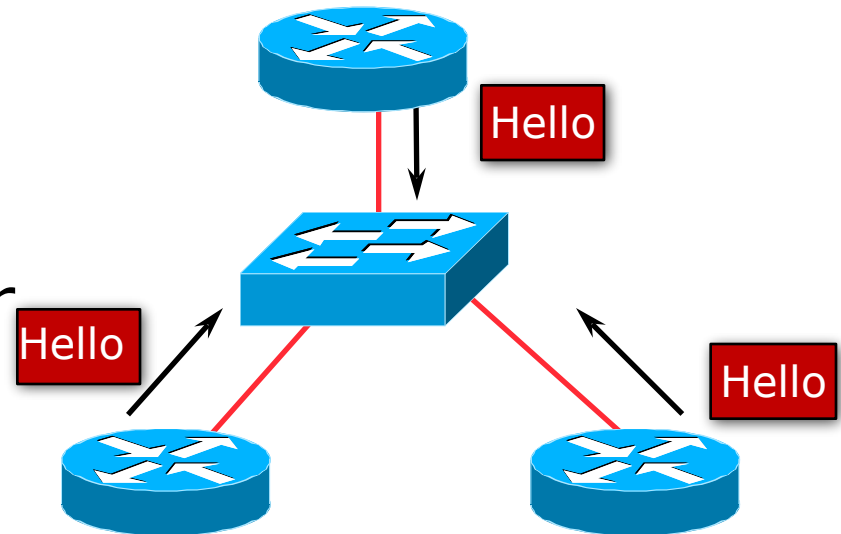


Topology/Link State Database

- ⌘ A router has a separate LS database for each area to which it belongs
- ⌘ All routers belonging to the same area have identical database
- ⌘ SPF calculation is performed separately for each area
- ⌘ LSA flooding is bounded by area
- ⌘ Recommendation:
 - Limit the number of areas a router participates in!!
 - 1 to 3 is fine (typical ISP design)
 - >3 can overload the CPU depending on the area topology complexity

The Hello Protocol

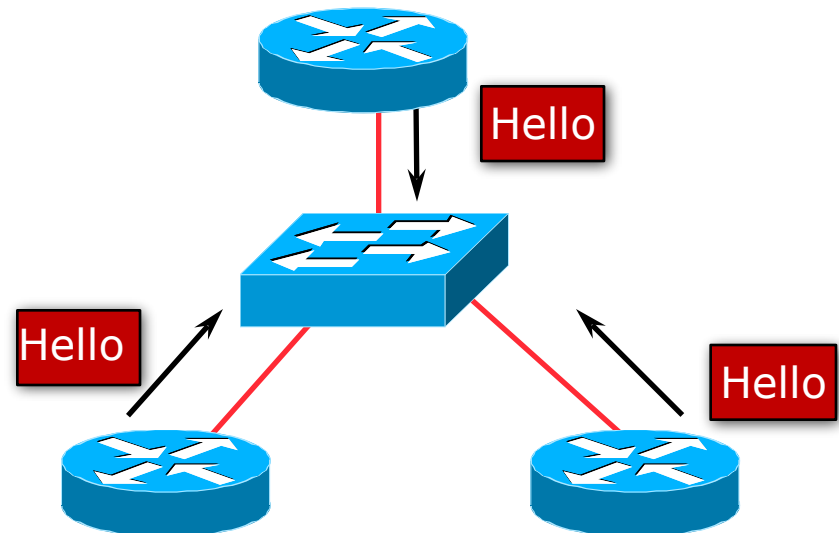
- ⌘ Responsible for establishing and maintaining neighbour relationships
- ⌘ Elects designated router on multi-access networks



The Hello Packet

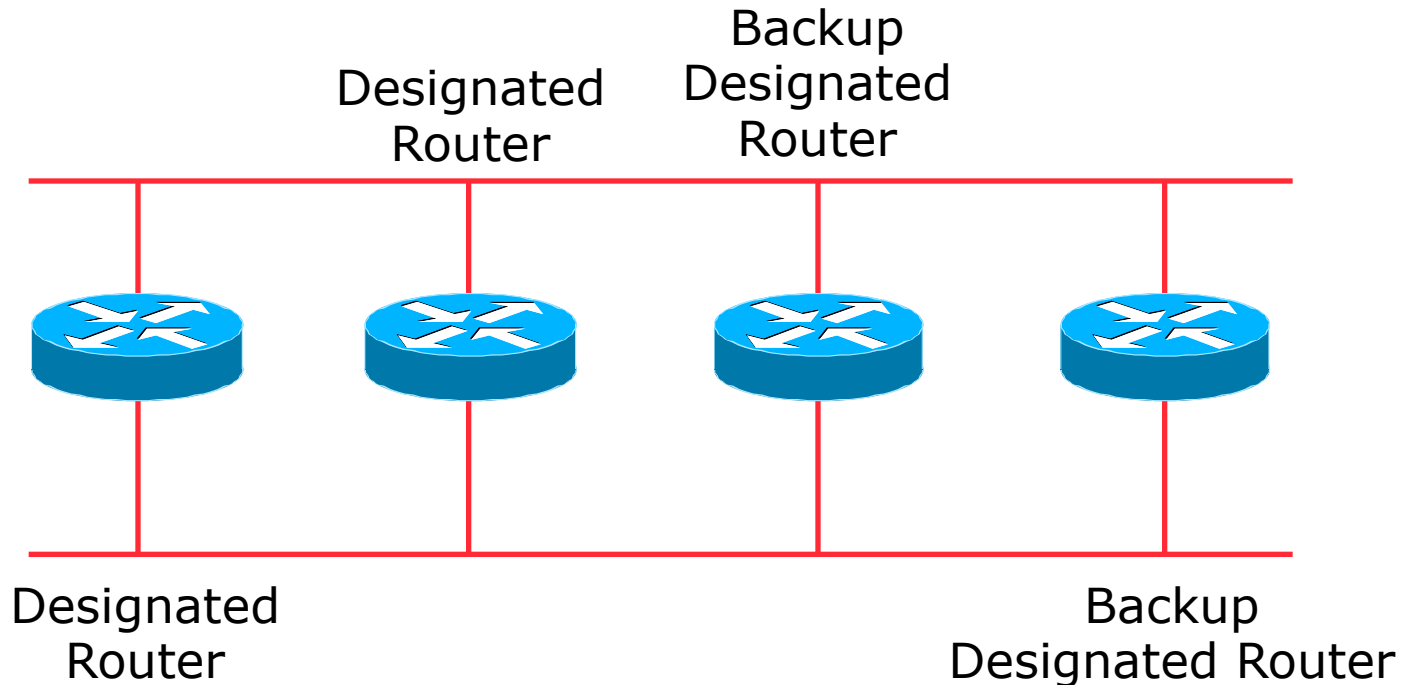
p Contains:

- Router priority
- Hello interval
- Router dead interval
- Network mask
- List of neighbours
- DR and BDR
- Options: E-bit, MC-bit,... (see A.2 of RFC2328)



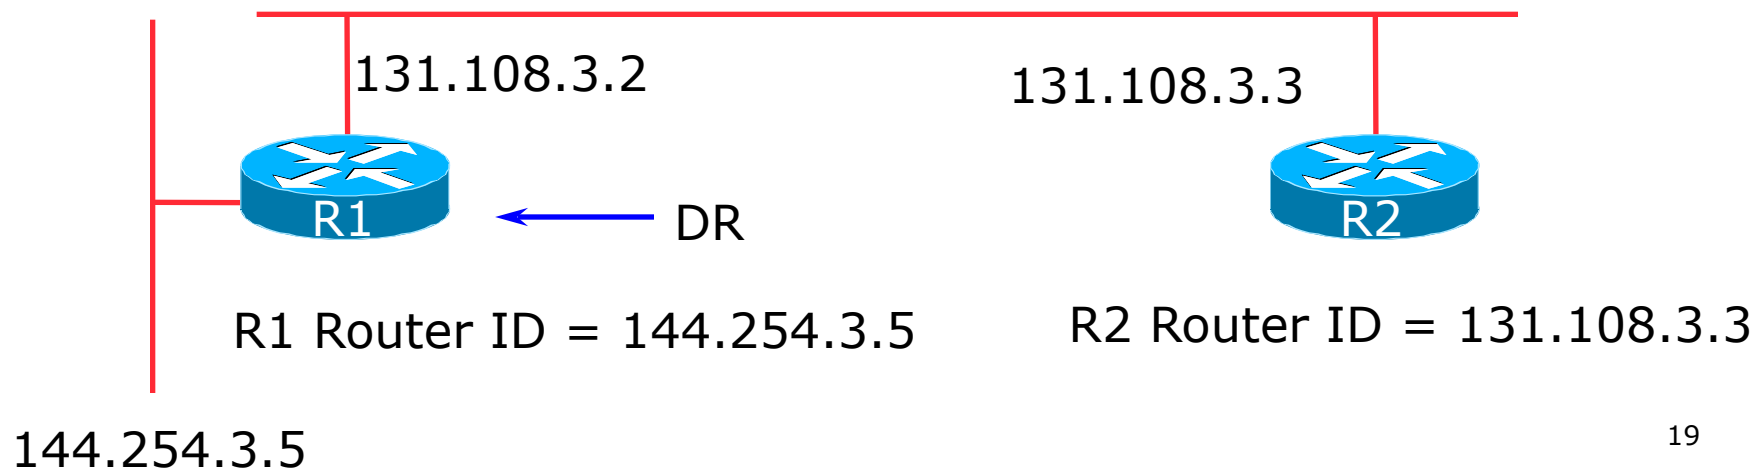
Designated Router

- There is ONE designated router per multi-access network
 - Generates network link advertisements
 - Assists in database synchronization



Designated Router by Priority

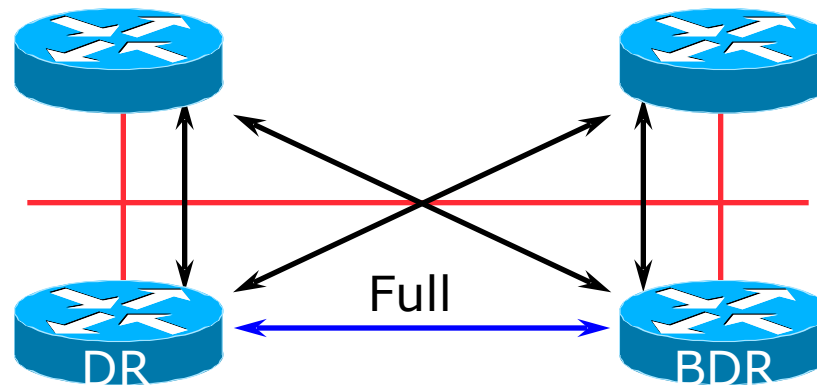
- Configured priority (per interface)
 - ISPs configure high priority on the routers they want as DR/BDR
- Else determined by highest router ID
 - Router ID is 32 bit integer
 - Derived from the loopback interface address, if configured, otherwise the highest IP address



Neighbouring States

p Full

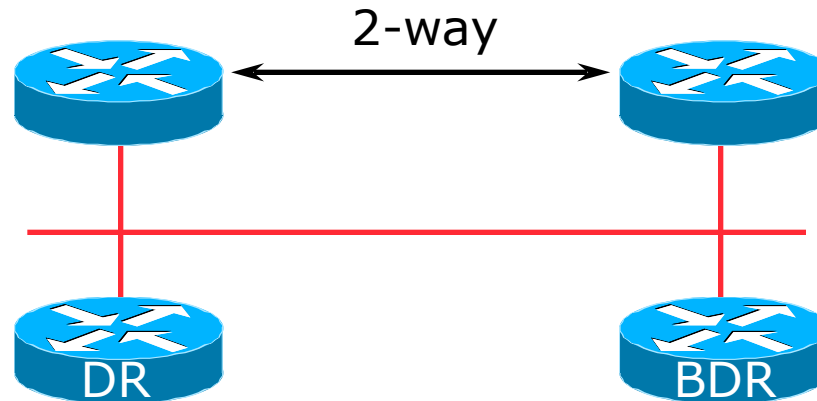
- Routers are fully adjacent
- Databases synchronised
- Relationship to DR and BDR



Neighbouring States

p 2-way

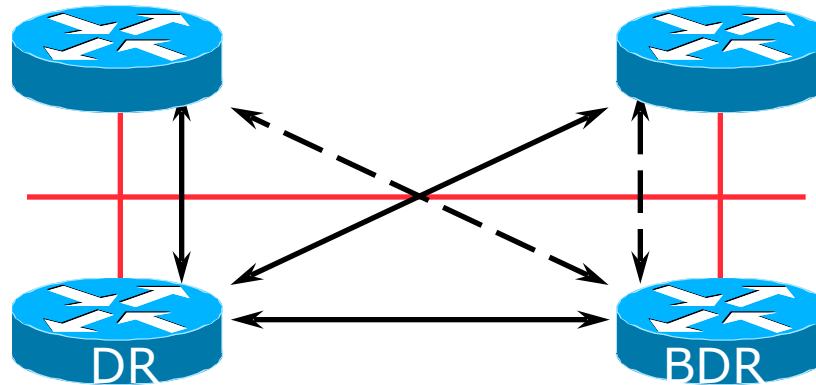
- Router sees itself in other Hello packets
- DR selected from neighbours in state 2-way or greater



When to Become Adjacent

- p Underlying network is point to point
- p Underlying network type is virtual link
- p The router itself is the designated router or the backup designated router
- p The neighbouring router is the designated router or the backup designated router

LSAs Propagate Along Adjacencies



- LSAs acknowledged along adjacencies

Broadcast Networks

- p IP Multicast used for Sending and Receiving Updates
 - All routers must accept packets sent to AllSPFRouters (224.0.0.5)
 - All DR and BDR routers must accept packets sent to AllDRouters (224.0.0.6)
- p Hello packets sent to AllSPFRouters (Unicast on point-to-point and virtual links)

Routing Protocol Packets

- p Share a common protocol header
- p Routing protocol packets are sent with type of service (TOS) of 0
- p Five types of OSPF routing protocol packets
 - Hello – packet type 1
 - Database description – packet type 2
 - Link-state request – packet type 3
 - Link-state update – packet type 4
 - Link-state acknowledgement – packet type 5

Different Types of LSAs

p Six distinct type of LSAs

- Type 1 : Router LSA
- Type 2 : Network LSA
- Type 3 & 4: Summary LSA
- Type 5 & 7: External LSA (Type 7 is for NSSA)
- Type 6: Group membership LSA
- Type 9, 10 & 11: Opaque LSA (9: Link-Local, 10: Area)

Router LSA (Type 1)

- ⌘ Describes the state and cost of the router's links to the area
- ⌘ All of the router's links in an area must be described in a single LSA
- ⌘ Flooded throughout the particular area and no more
- ⌘ Router indicates whether it is an ASBR, ABR, or end point of virtual link

Network LSA (Type 2)

- ⌘ Generated for every transit broadcast and NBMA network
- ⌘ Describes all the routers attached to the network
- ⌘ Only the designated router originates this LSA
- ⌘ Flooded throughout the area and no more

Summary LSA (Type 3 and 4)

- ⌘ Describes the destination outside the area but still in the AS
- ⌘ Flooded throughout a single area
- ⌘ Originated by an ABR
- ⌘ Only inter-area routes are advertised into the backbone
- ⌘ Type 4 is the information about the ASBR

External LSA (Type 5 and 7)

- p Defines routes to destination external to the AS
- p Default route is also sent as external
- p Two types of external LSA:
 - E1: Consider the total cost up to the external destination
 - E2: Considers only the cost of the outgoing interface to the external destination
- p (Type 7 LSAs used to describe external LSA for one specific OSPF area type)

Inter-Area Route Summarisation

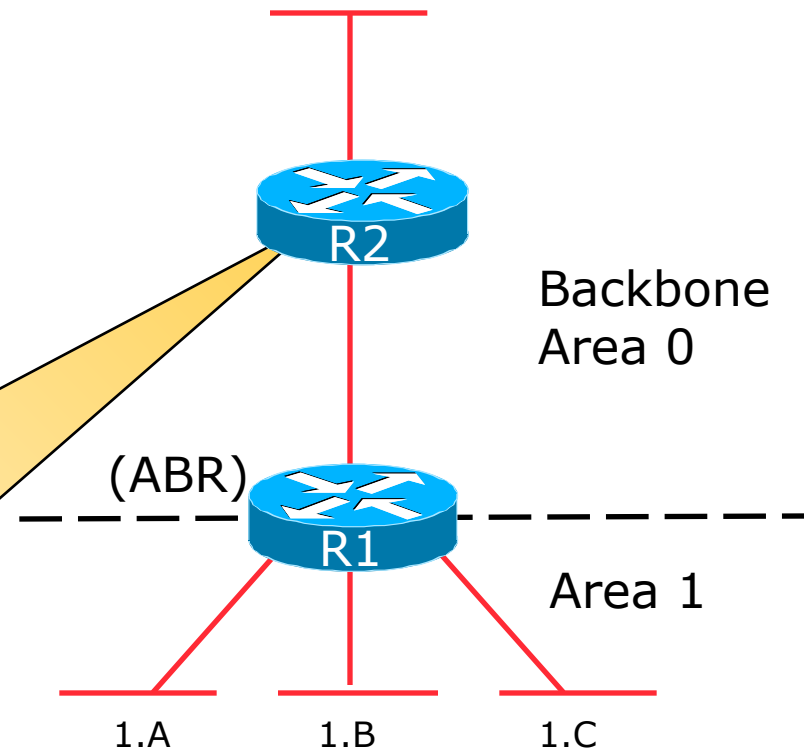
- Prefix or all subnets
- Prefix or all networks
- 'Area range' command

With
summarisation

Network	Next Hop
1	R1

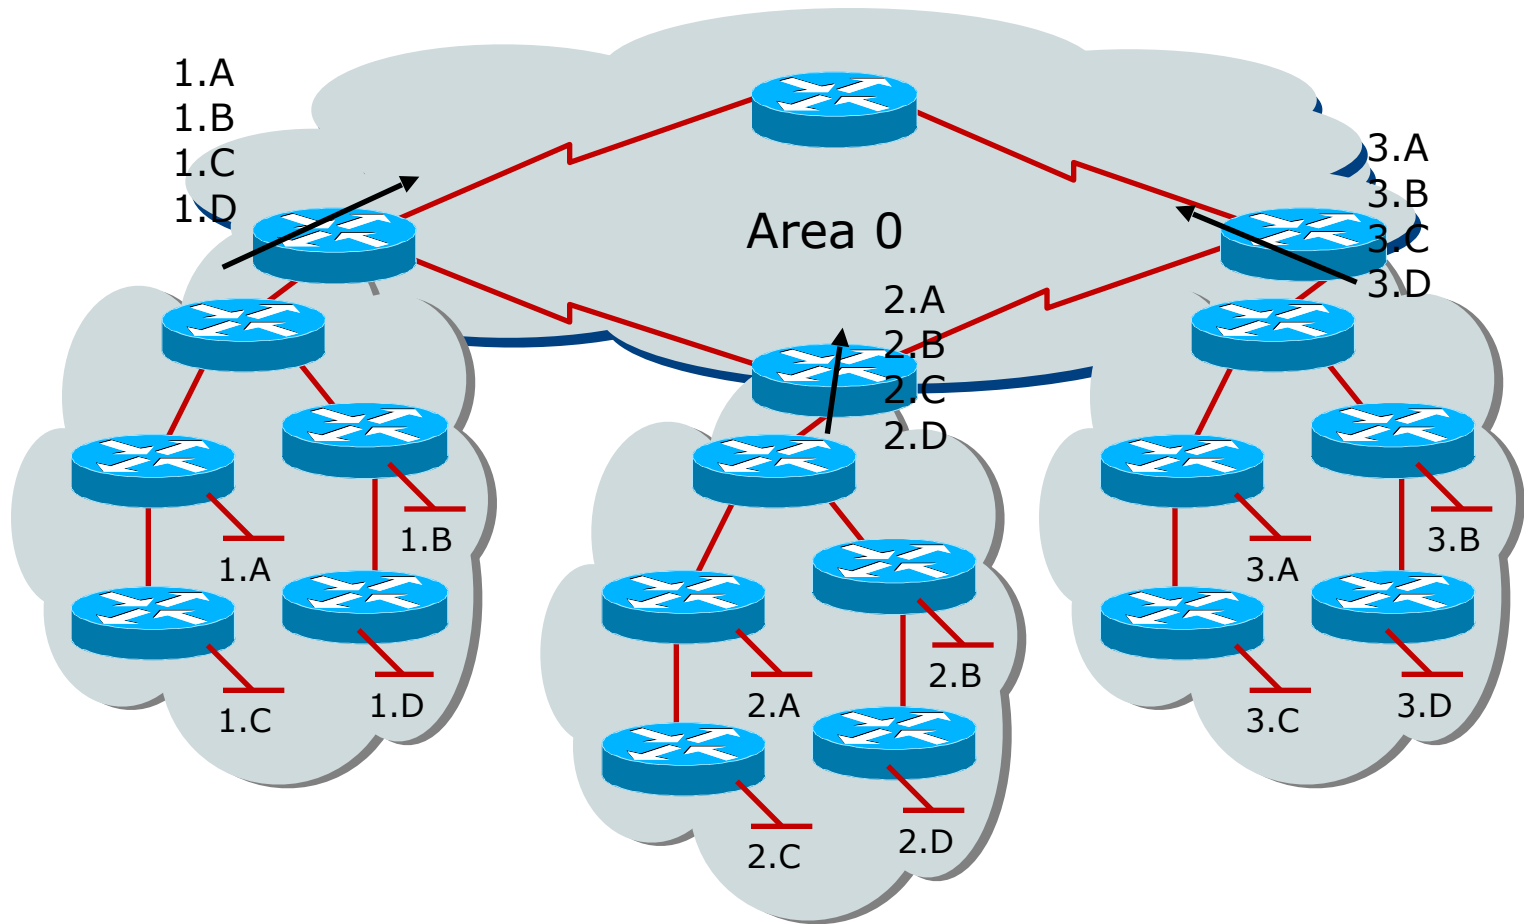
Without
summarisation

Network	Next Hop
1.A	R1
1.B	R1
1.C	R1



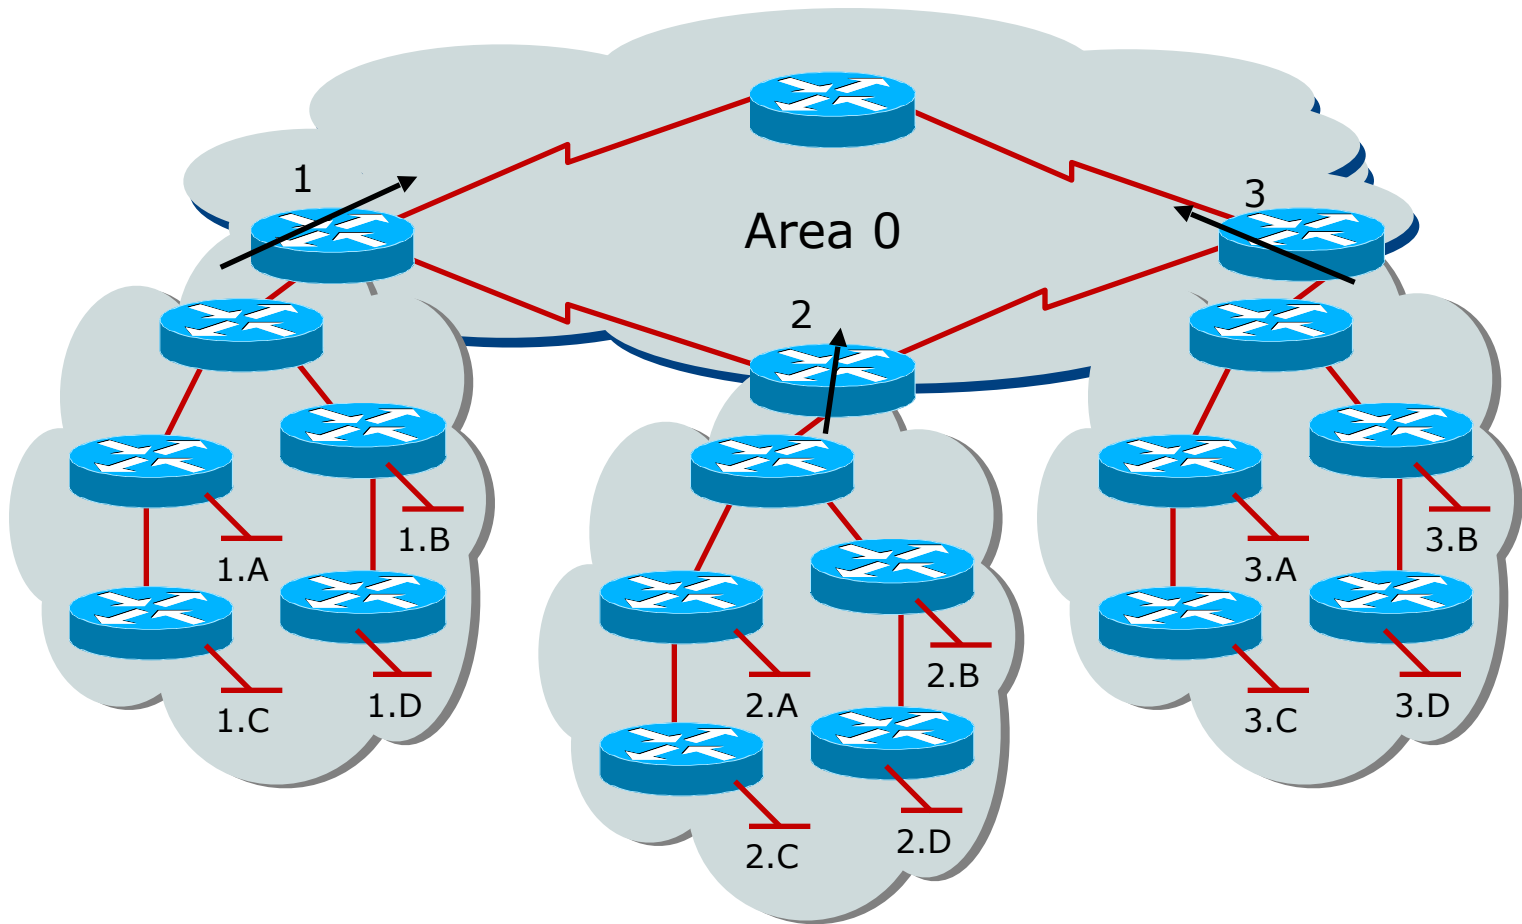
No Summarisation

- Specific Link LSA advertised out of each area
- Link state changes propagated out of each area



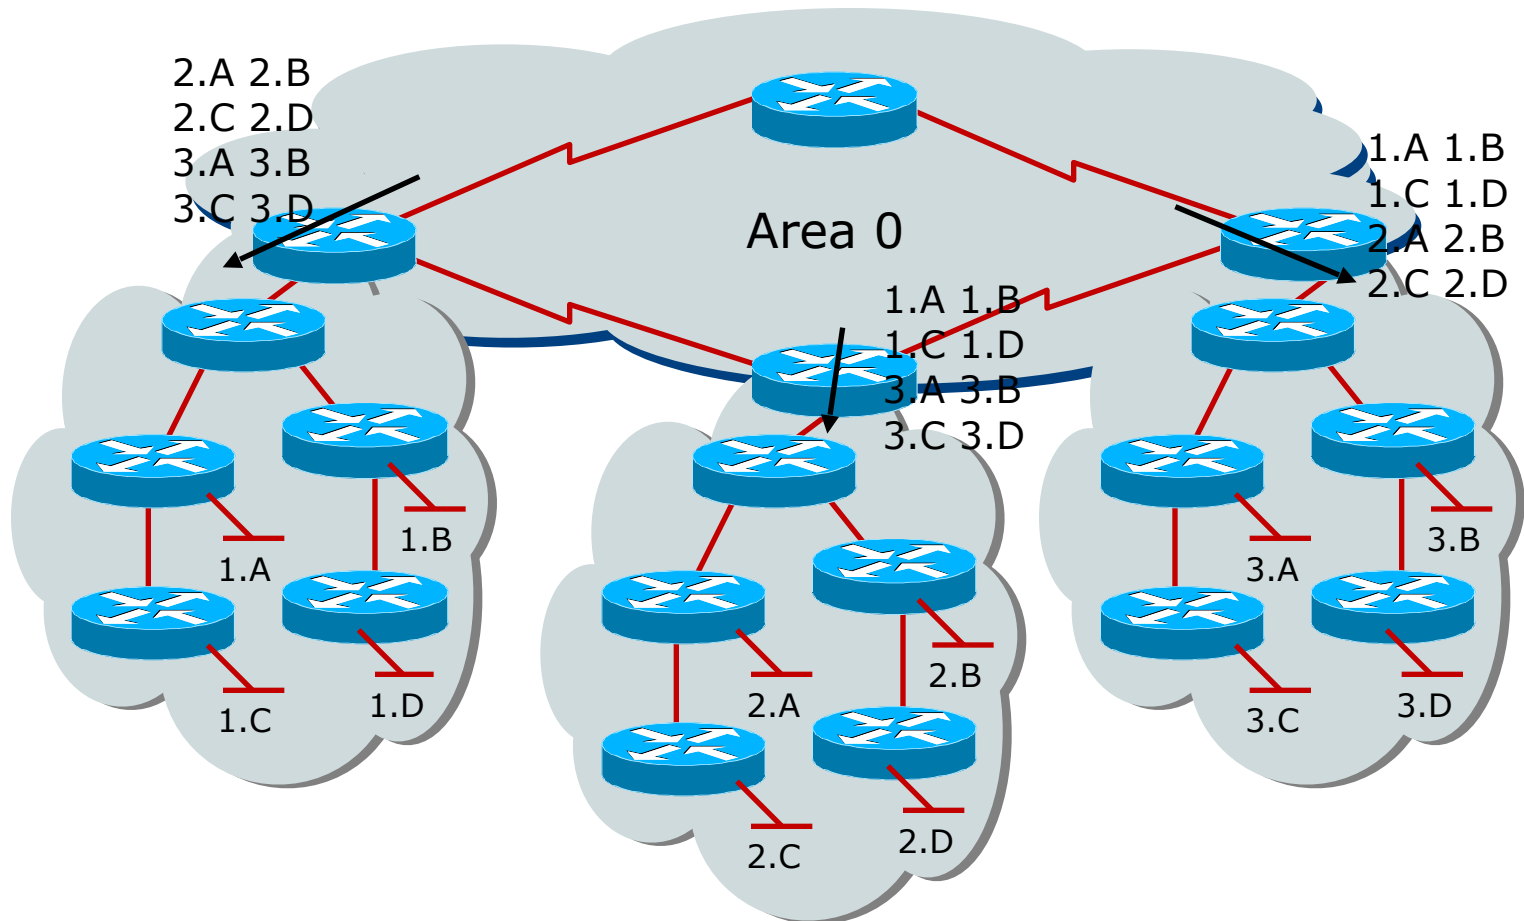
With Summarisation

- Only summary LSA advertised out of each area
- Link state changes do not propagate out of the area



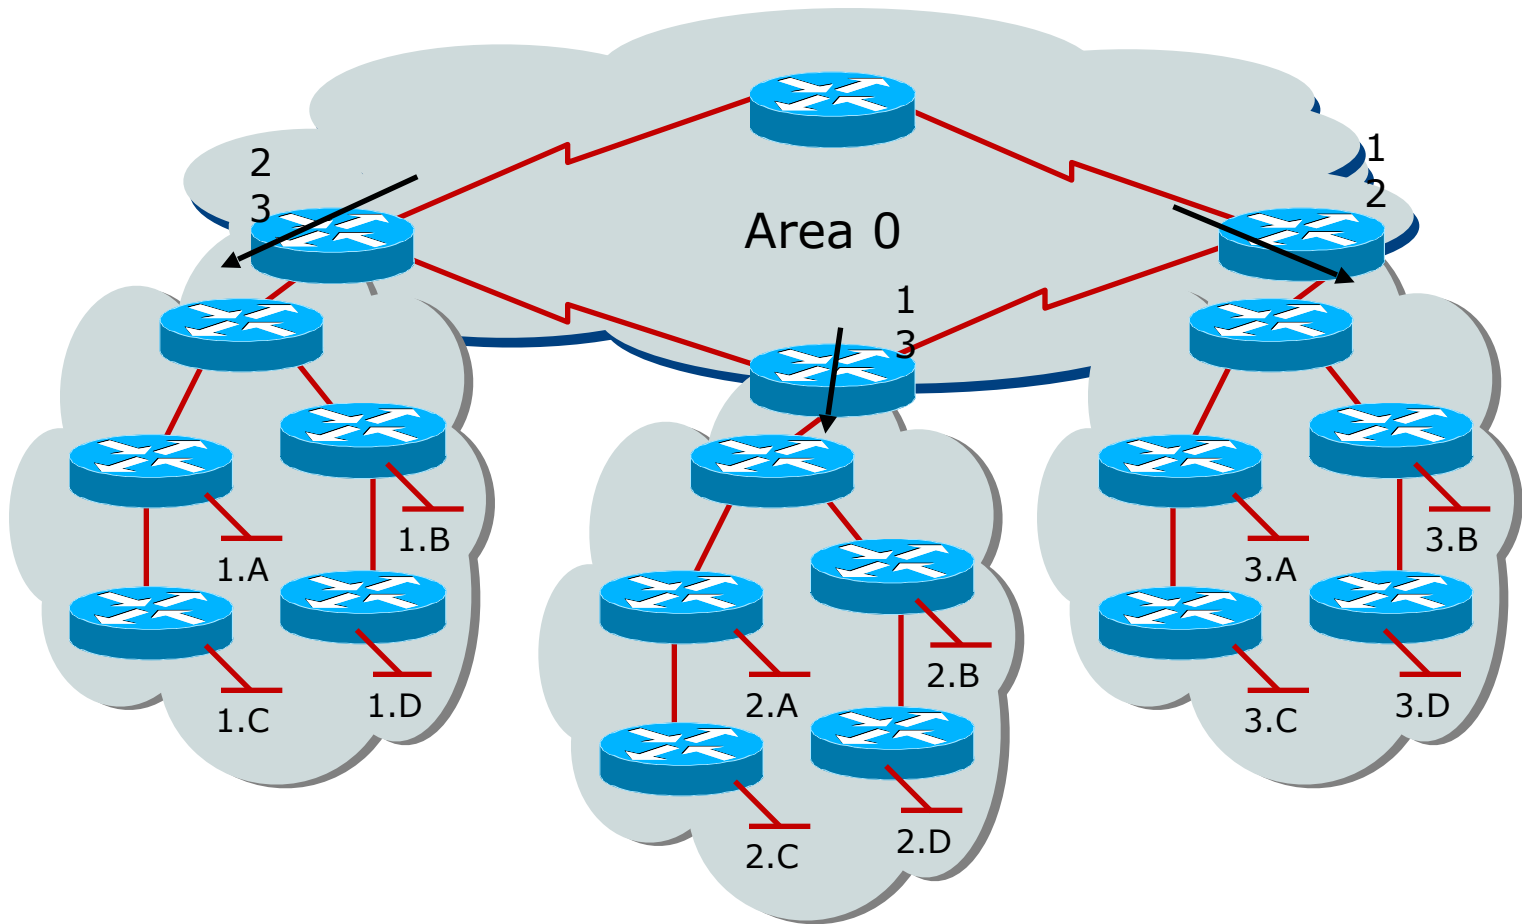
No Summarisation

- Specific Link LSA advertised in to each area
- Link state changes propagated in to each area



With Summarisation

- Only summary link LSA advertised in to each area
- Link state changes do not propagate in to each area

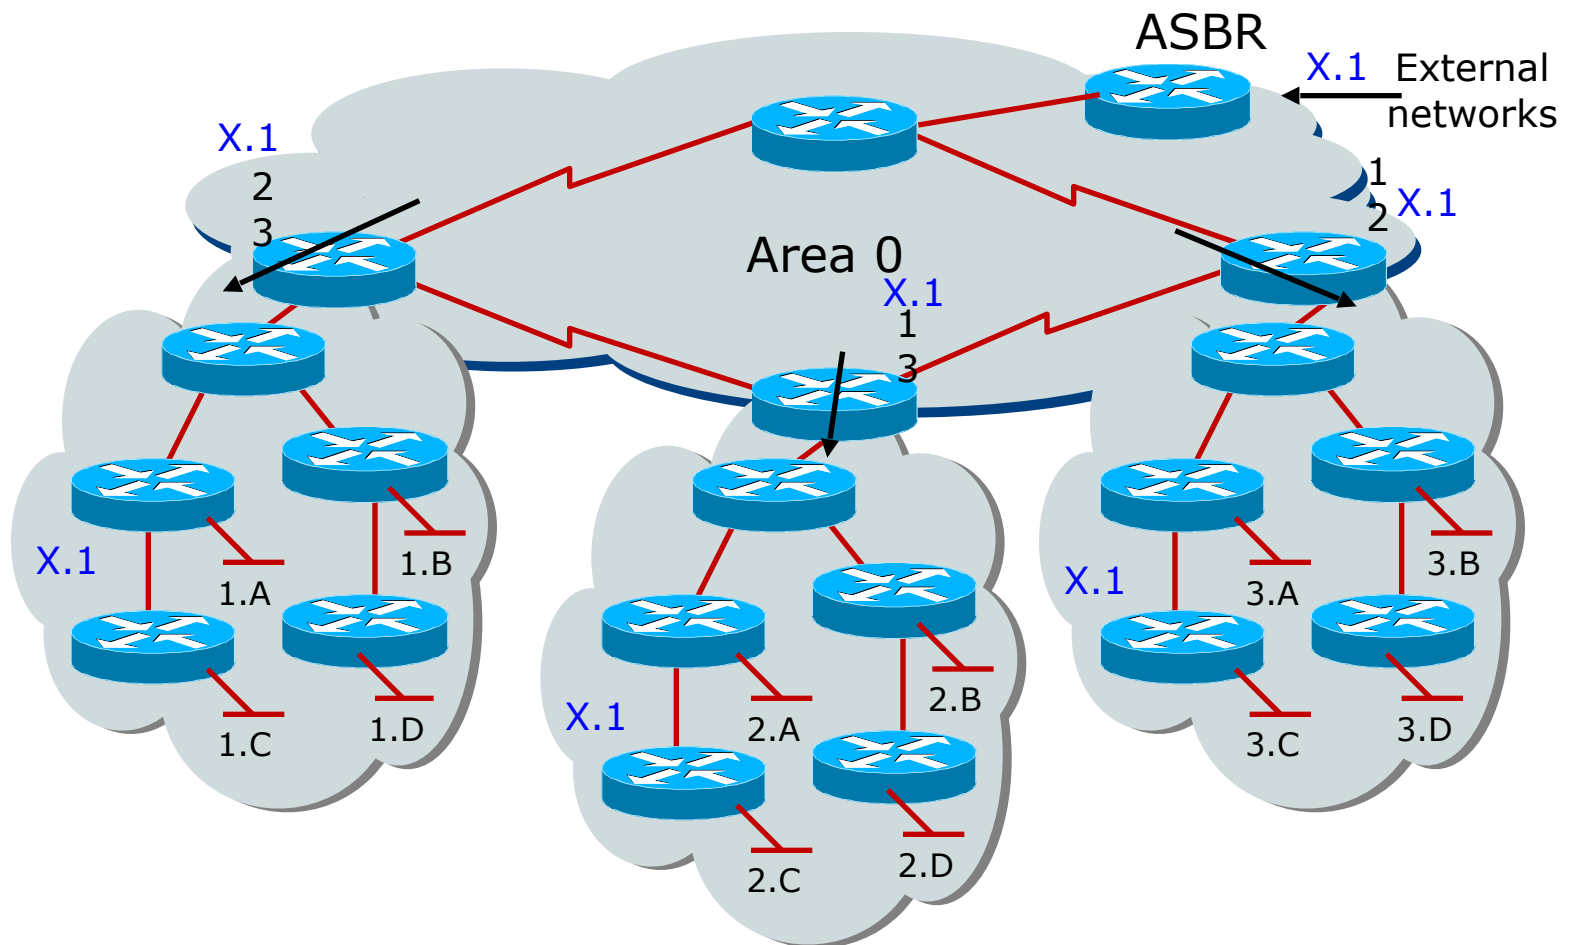


Types of Areas

- p Regular
- p Stub
- p Totally Stubby
- p Not-So-Stubby
- p **Only “regular” areas are useful for ISPs**
 - Other area types handle redistribution of other routing protocols into OSPF – ISPs don’t redistribute anything into OSPF
- p The next slides describing the different area types are provided for information only

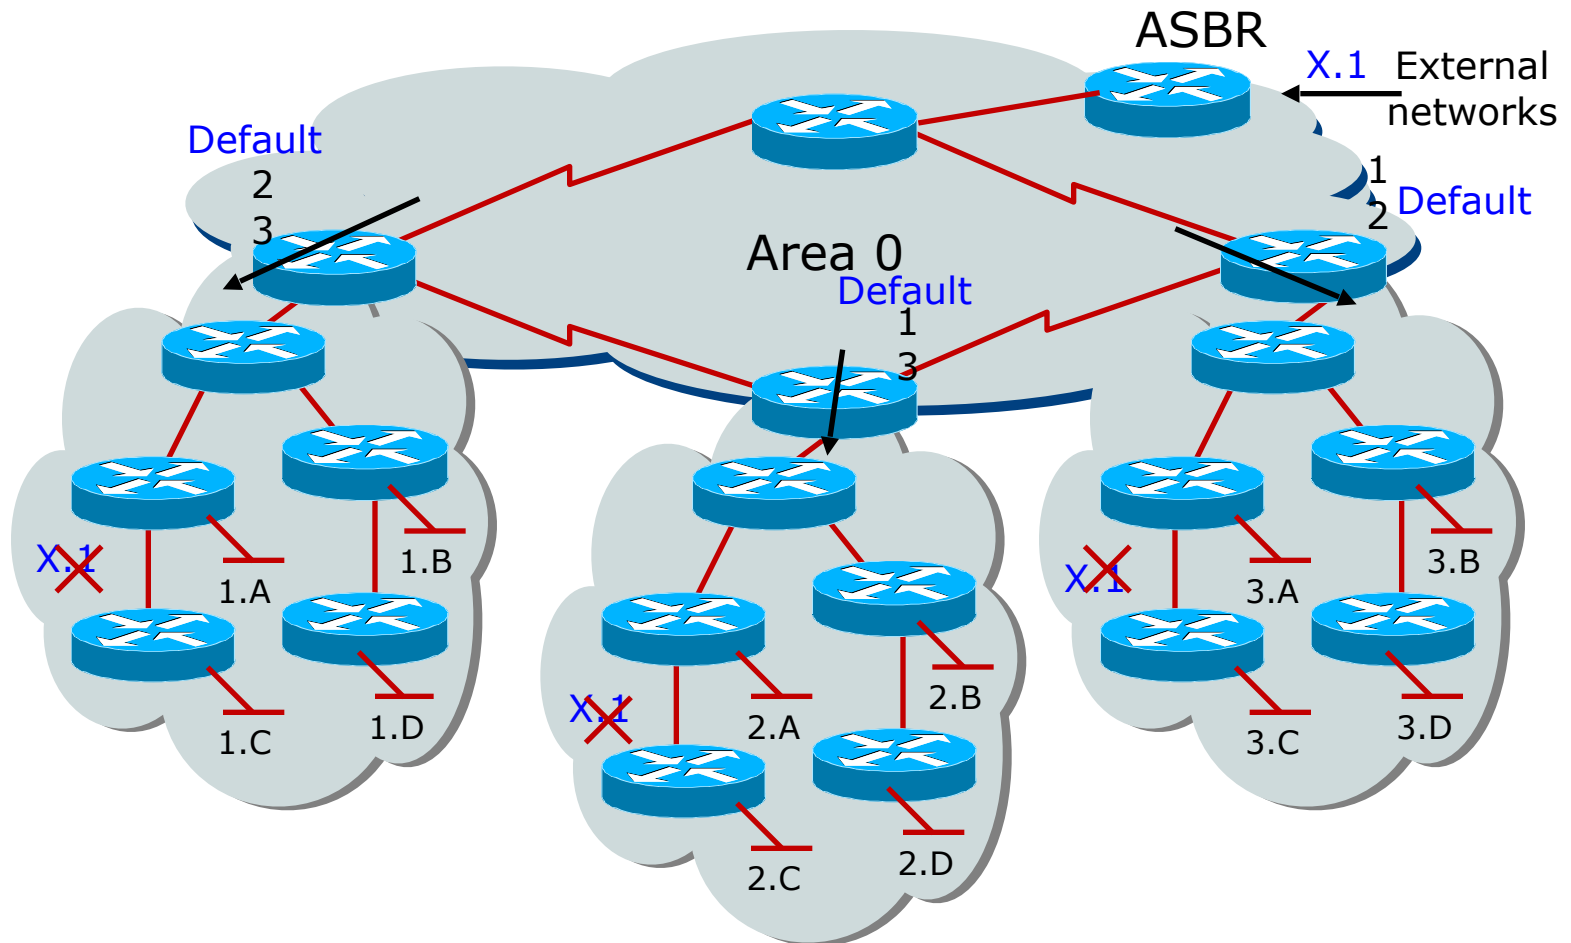
Regular Area (Not a Stub)

- From Area 1's point of view, summary networks from other areas are injected, as are external networks such as X.1



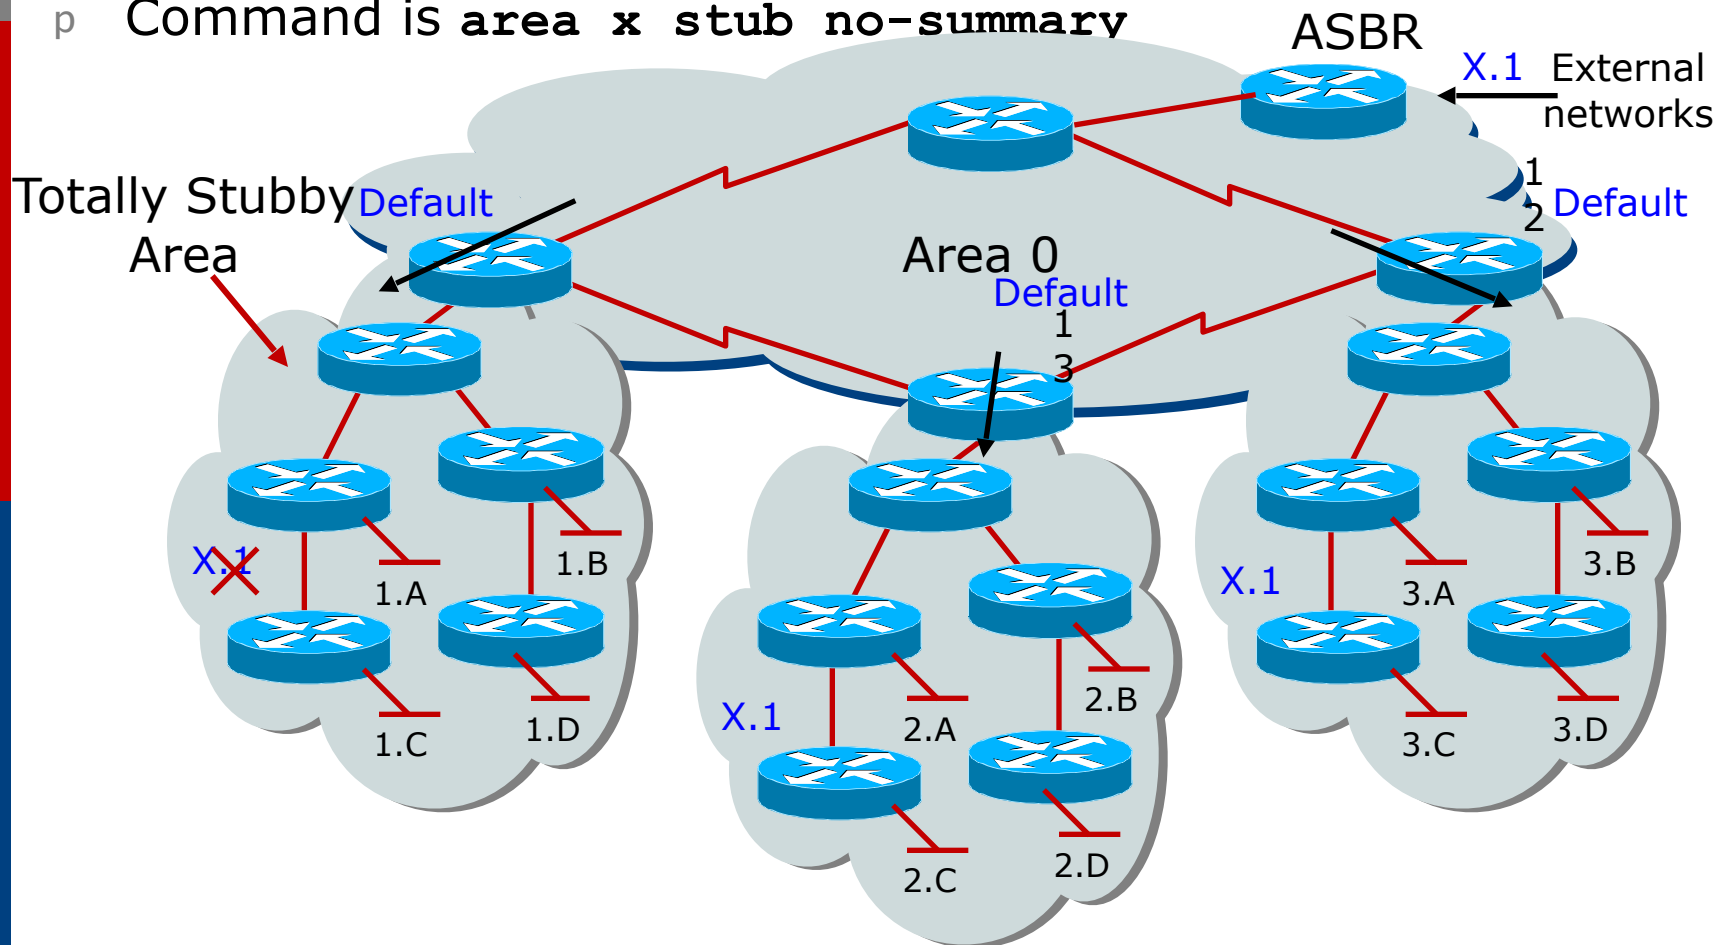
Normal Stub Area

- Summary networks, default route injected
- Command is **area x stub**



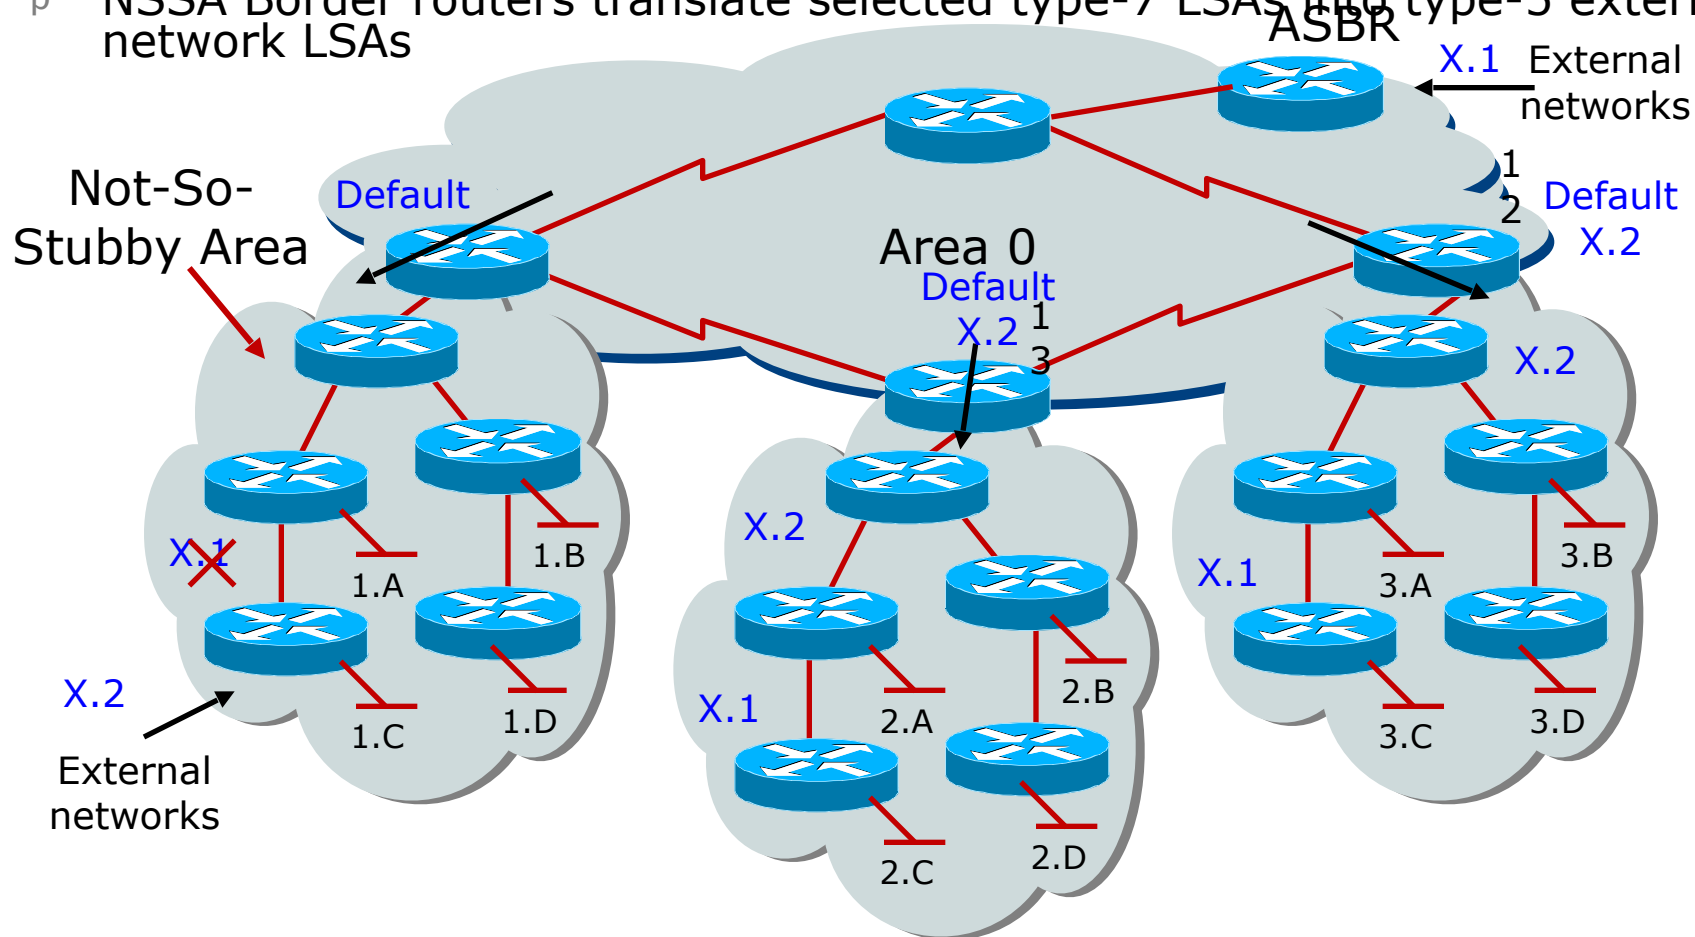
Totally Stubby Area

- Only a default route injected
 - Default path to closest area border router
- Command is **area x stub no-summary**



Not-So-Stubby Area

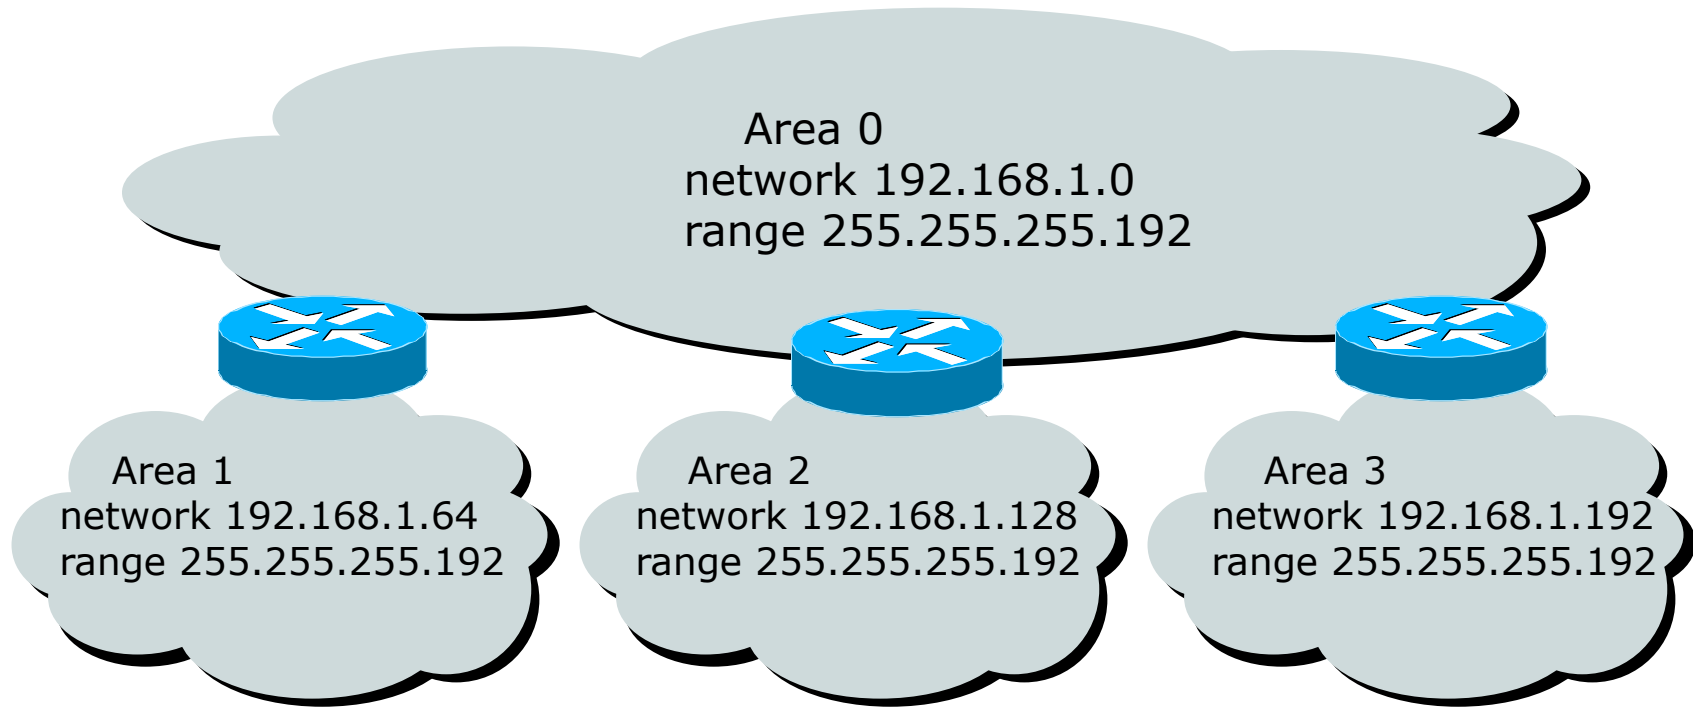
- Capable of importing routes in a limited fashion
- Type-7 LSA's carry external information within an NSSA
- NSSA Border routers translate selected type-7 LSAs into type-5 external network LSAs



ISP Use of Areas

- p ISP networks use:
 - Backbone area
 - Regular area
- p Backbone area
 - No partitioning
- p Regular area
 - Summarisation of point to point link addresses used within areas
 - Loopback addresses allowed out of regular areas without summarisation (otherwise iBGP won't work)

Addressing for Areas



- Assign contiguous ranges of subnets per area to facilitate summarisation

Summary

p Fundamentals of Scalable OSPF Network Design

- Area hierarchy
- DR/BDR selection
- Contiguous intra-area addressing
- Route summarisation
- Infrastructure prefixes only

Introduction to OSPF



ISP Training Workshops