

# CS 268: Lecture 9

## Intra-domain Routing Protocols

Ion Stoica  
Computer Science Division  
Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley  
Berkeley, CA 94720-1776


(\*Based in part on Aman Shaikh's slides)

## Internet Routing

---

- Internet organized as a two level hierarchy
- First level – autonomous systems (AS's)
  - AS – region of network under a single administrative domain
- AS's run an intra-domain routing protocols
  - Distance Vector, e.g., Routing Information Protocol (RIP)
  - Link State, e.g., Open Shortest Path First (OSPF)
- Between AS's runs inter-domain routing protocols, e.g., Border Gateway Routing (BGP)
  - De facto standard today, BGP-4

## Example



3


## Intra-domain Routing Protocols

- Based on unreliable datagram delivery
- Distance vector
  - Routing Information Protocol (RIP), based on Bellman-Ford
  - Each neighbor periodically exchange reachability information to its neighbors
- Link state
  - Open Shortest Path First (OSPF), based on Dijkstra
  - Each network periodically floods immediate reachability information to other routers

4

## Routing


- Goal: determine a “good” path through the network from source to destination
  - Good means usually the shortest path
- Network modeled as a graph
  - Routers → nodes
  - Link → edges
    - Edge cost: delay, congestion level,...



5

## Routing Problem

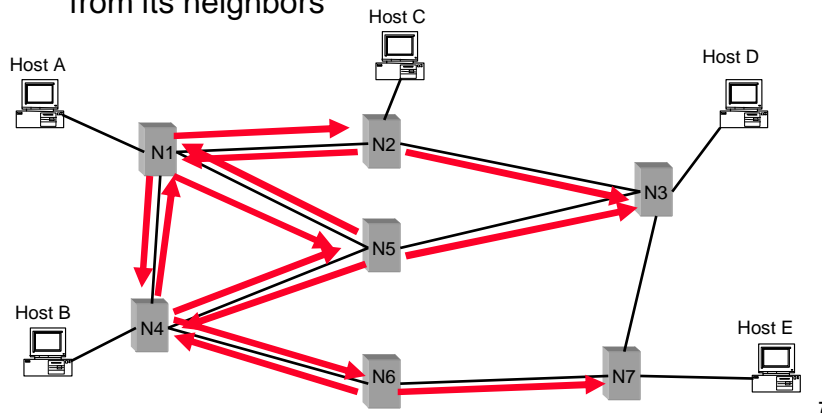
- Assume
  - A network with N nodes, where each edge is associated a cost
  - A node knows only its neighbors and the cost to reach them
- How does each node learn how to reach every other node along the shortest path?



6


## Distance Vector: Control Traffic

- When the routing table of a node changes, the node sends its table to its neighbors
- A node updates its table with information received from its neighbors



7

## Example: Distance Vector Algorithm



Node A

Dest.	Cost	NextHop
B	2	B
C	7	C
D	$\infty$	-

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	3	D

Node C

Dest.	Cost	NextHop
A	7	A
B	1	B
D	1	D

Node D


Dest.	Cost	NextHop
A	$\infty$	-
B	3	B
C	1	C

```

1 Initialization:
2 for all neighbors V do
3   if V adjacent to A
4     D(A, V) = c(A, V);
5 else
6   D(A, V) =  $\infty$ ;
...
    
```

8

## Example: 1<sup>st</sup> Iteration (C → A)



Node A

Dest.	Cost	NextHop
B	2	B
C	7	C
D	∞	-

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	3	D

...  
7 loop:

```

12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                    D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
    
```

Node C

Dest.	Cost	NextHop
A	7	A
B	1	B
D	1	D


Node D

Dest.	Cost	NextHop
A	∞	-
B	3	B
C	1	C

(D(C,A), D(C,B), D(C,D))

9

## Example: 1<sup>st</sup> Iteration (C → A)



Node A

Dest.	Cost	NextHop
B	2	B
C	7	C
D	8	C

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	3	D

...  
7 loop:

```

12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                    D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
    
```

Node C

Dest.	Cost	NextHop
A	7	A
B	1	B
D	1	D

Node D


Dest.	Cost	NextHop
A	∞	-
B	3	B
C	1	C

$$D(A, D) = \min(D(A, D), D(A, C) + D(C, D)) = \min(\infty, 7 + 1) = 8$$

(D(C,A), D(C,B), D(C,D))

10

## Example: 1<sup>st</sup> Iteration (C → A)



Node A

Dest.	Cost	NextHop
B	2	B
C	7	C
D	8	C

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	3	D

...  
7 loop:

```

12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                     D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
  
```

Node C


Dest.	Cost	NextHop
A	7	A
B	1	B
D	1	D

Node D

Dest.	Cost	NextHop
A	∞	-
B	3	B
C	1	C

11

## Example: 1<sup>st</sup> Iteration (B → A, C → A)



Node A

Dest.	Cost	NextHop
B	2	B
C	3	B
D	5	B

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	3	D

...  
7 loop:

$$D(A, D) = \min(D(A, D), D(A, B) + D(B, D)) = \min(8, 2 + 3) = 5$$

$$D(A, C) = \min(D(A, C), D(A, B) + D(B, C)) = \min(7, 2 + 1) = 3$$

```

12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                     D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
  
```

Node C


Dest.	Cost	NextHop
A	7	A
B	1	B
D	1	D

Node D

Dest.	Cost	NextHop
A	∞	-
B	3	B
C	1	C

12

## Example: End of 1<sup>st</sup> Iteration



...  
7 loop:

```

...
12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                    D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
  
```

Node A

Dest.	Cost	NextHop
B	2	B
C	3	B
D	5	B

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	2	C

Node C


Dest.	Cost	NextHop
A	3	B
B	1	B
D	1	D

Node D

Dest.	Cost	NextHop
A	2	B
B	3	B
C	1	C

13

## Example: End of 3<sup>rd</sup> Iteration



...  
7 loop:

```

...
12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y),
18                    D(A, V) + D(V, Y));
19   if (there is a new minimum for dest. Y)
20     send D(A, Y) to all neighbors
21 forever
  
```

Node A

Dest.	Cost	NextHop
B	2	B
C	3	B
D	4	B

Node B

Dest.	Cost	NextHop
A	2	A
C	1	C
D	2	C

Node C

Dest.	Cost	NextHop
A	3	B
B	1	B
D	1	D

Node D

Dest.	Cost	NextHop
A	4	C
B	2	C
C	1	C



Nothing changes → algorithm terminates

14

## Distance Vector: Link Cost Changes

```

7 loop:
8 wait (link cost update or update message)
9 if (c(A,V) changes by d)
10 for all destinations Y through V do
11   D(A, Y) = D(A, Y) + d
12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y), D(A, V) + D(V, Y));
18   if (there is a new minimum for destination Y)
19     send D(A, Y) to all neighbors
20 forever
    
```





15

## Distance Vector: Count to Infinity Problem

```

7 loop:
8 wait (link cost update or update message)
9 if (c(A,V) changes by d)
10 for all destinations Y through V do
11   D(A, Y) = D(A, Y) + d
12 else if (update D(V, Y) received from V)
13   for all destinations Y do
14     if (destination Y through V)
15       D(A, Y) = D(A, V) + D(V, Y);
16     else
17       D(A, Y) = min(D(A, Y), D(A, V) + D(V, Y));
18   if (there is a new minimum for destination Y)
19     send D(A, Y) to all neighbors
20 forever
    
```





16



## Distance Vector: Poisoned Reverse


- If C routes through B to get to A:
  - C tells B its (C's) distance to A is infinite (so B won't route to A via C)
  - Will this completely solve count to infinity problem?



17


## Link State: Control Traffic

- Each node floods its local information to every other node in the network
- Each node ends up knowing the entire network topology → use Dijkstra to compute the shortest path to every other node



18


## Link State: Node State



19

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
→ 0	A	2,A	5,A	1,A	$\infty$	$\infty$
1						
2						
3						
4						
5						




- 1 **Initialization:**
- 2  $S = \{A\};$
- 3 for all nodes  $v$
- 4 if  $v$  adjacent to  $A$
- 5 then  $D(v) = c(A,v);$
- 6 else  $D(v) = \infty;$
- ...

20

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
0	A	2,A	5,A	1,A	$\infty$	$\infty$
→ 1	AD		4,D		2,D	$\infty$
2						
3						
4						
5						



```


...
8 Loop
9   find w not in S s.t. D(w) is a minimum;
10  add w to S;
11  update D(v) for all v adjacent
    to w and not in S:
12     $D(v) = \min( D(v), D(w) + c(w,v) );$ 
13  until all nodes in S;

```

21

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
0	A	2,A	5,A	1,A	$\infty$	$\infty$
1	AD		4,D		2,D	$\infty$
→ 2	ADE		3,E			4,E
3						
4						
5						



```


...
8 Loop
9   find w not in S s.t. D(w) is a minimum;
10  add w to S;
11  update D(v) for all v adjacent
    to w and not in S:
12     $D(v) = \min( D(v), D(w) + c(w,v) );$ 
13  until all nodes in S;

```

22

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
0	A	2,A	5,A	1,A	$\infty$	$\infty$
1	AD		4,D		2,D	$\infty$
2	ADE		3,E			4,E
→ 3	ADEB					
4						
5						




```

...
8 Loop
9 find w not in S s.t. D(w) is a minimum;
10 add w to S;
11 update D(v) for all v adjacent
    to w and not in S:
12   D(v) = min( D(v), D(w) + c(w,v) );
13 until all nodes in S;
    
```

23

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
0	A	2,A	5,A	1,A	$\infty$	$\infty$
1	AD		4,D		2,D	$\infty$
2	ADE		3,E			4,E
3	ADEB					
→ 4	ADEBC					
5						




```

...
8 Loop
9 find w not in S s.t. D(w) is a minimum;
10 add w to S;
11 update D(v) for all v adjacent
    to w and not in S:
12   D(v) = min( D(v), D(w) + c(w,v) );
13 until all nodes in S;
    
```

24

## Example: Dijkstra's Algorithm

Step	start S	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
0	A	2,A	5,A	1,A	$\infty$	$\infty$
1	AD		4,D		2,D	$\infty$
2	ADE		3,E			4,E
3	ADEB					
4	ADEBC					
→ 5	ADEBCF					



```

...
8 Loop
9 find w not in S s.t. D(w) is a minimum;
10 add w to S;
11 update D(v) for all v adjacent
   to w and not in S:
12   D(v) = min( D(v), D(w) + c(w,v) );
13 until all nodes in S;
    
```

25

## Link State vs. Distance Vector

### Message complexity

- LS:  $O(n^2 \cdot e)$  messages
  - n: number of nodes
  - e: number of edges
- DV:  $O(d \cdot n \cdot k)$  messages
  - d: node's degree
  - k: number of rounds

### Time complexity

- LS:  $O(n \cdot \log n)$
- DV:  $O(n)$

### Convergence time

- LS:  $O(1)$
- DV:  $O(k)$

### Robustness: what happens if router malfunctions?

- LS:
  - node can advertise incorrect *link* cost
  - each node computes only its *own* table
- DV:
  - node can advertise incorrect *path* cost
  - each node's table used by others; error propagate through network

26

## Open Shortest Path First (OSPF)

- All routers in the domain come to a consistent view of the topology by exchange of Link State Advertisements (LSAs)
- Router describes its local connectivity (i.e., set of links) in an LSA
  - Set of LSAs (self-originated + received) at a router = topology
- Hierarchical routing
  - OSPF domain can be divided into areas
  - Hub-and-spoke topology with area 0 as hub and other non-zero areas as spokes

27

## OSPF Performance

- OSPF processing impacts convergence, (in)stability
  - Load is increasing as networks grow
- Bulk of OSPF processing is due to LSAs
  - Sending/receiving LSAs
  - LSAs can trigger Route calculation (Dijkstra's algorithm)
- Understanding dynamics of LSA traffic is key for a better understanding of OSPF

28

## Objectives for OSPF Monitor

- Real-time analysis of OSPF behavior
  - Trouble-shooting, alerting, validation of maintenance
  - Real-time snapshots of OSPF network topology
- Off-line analysis
  - Post-mortem analysis of recurring problems
  - Generate statistics and reports about network performance
  - Identify anomaly signatures
  - Facilitate tuning of configurable parameters
  - Analyze OSPF behavior in commercial networks

29

## Categorizing LSA Traffic

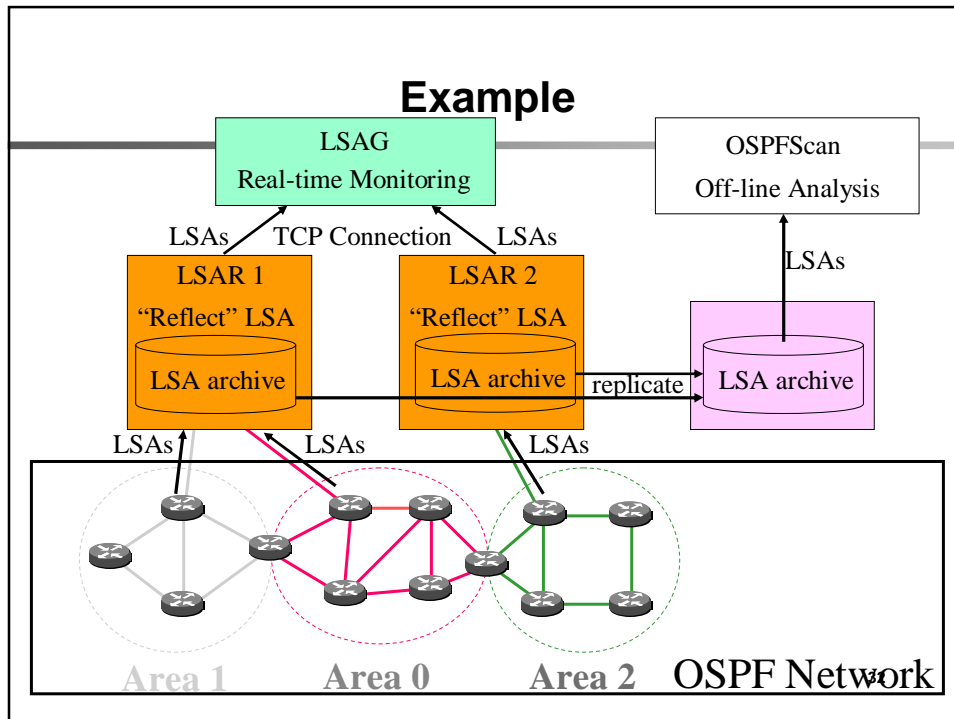
- A router originates an LSA due to...
  - Change in network topology **Change LSAs**
    - Example: link goes down or comes up
    - Detection of anomalies and problems
  - Periodic soft-state refresh **Refresh LSAs**
    - Recommended value of interval is 30 minutes
    - Forms baseline LSA traffic
- LSAs are disseminated using reliable flooding
  - Includes change and refresh LSAs
  - Flooding leads to duplicate copies of LSAs being received at a router **Duplicate LSAs**
  - Overhead: wastes resources

30

## Components

- Data collection: LSA Reflector (LSAR)
  - Passively collects OSPF LSAs from network
  - “Reflects” streams of LSAs to LSAG
  - Archives LSAs for analysis by OSPFScan
- Real-time analysis: LSA aGgregator (LSAG)
  - Monitors network for topology changes, LSA storms, node flaps and anomalies
- Off-line analysis: OSPFScan
  - Supports queries on LSA archives
  - Allows playback and modeling of topology changes
  - Allows emulation of OSPF routing

31





## How LSAR attaches to Network

- Host mode: Join multicast group
- Full adjacency mode: form full adjacency (= peering session) with a router
- Partial adjacency mode: keep adjacency in a state that allows LSAR to receive LSAs, but does not allow data forwarding over link


33

## How LSAR attaches to Network

- Host mode
  - Join multicast group
  - **Adv:** completely passive
  - **Disadv:** not reliable, delayed initialization of LSDB
- Full adjacency mode
  - Form full adjacency (= peering session) with a router
  - **Adv:** reliable, immediate initialization of LSDB
  - **Disadv:** LSAR's instability can impact entire network
- Partial adjacency mode
  - Keep adjacency in a state that allows LSAR to receive LSAs, but does not allow data forwarding over link
  - **Adv:** reliable, LSAR's instability does not impact entire network, immediate initialization of LSDB
  - **Disadv:** can raise alarms on the router

34

## Partial Adjacency for LSAR



- Router R does not advertise a link to LSAR
- LSAR does not originate any LSAs
- Routers (except R) not aware of LSAR's presence
  - Does not trigger routing calculations in network
  - LSAR's going up/down does not impact network
- LSAR↔R link is not used for data forwarding

35

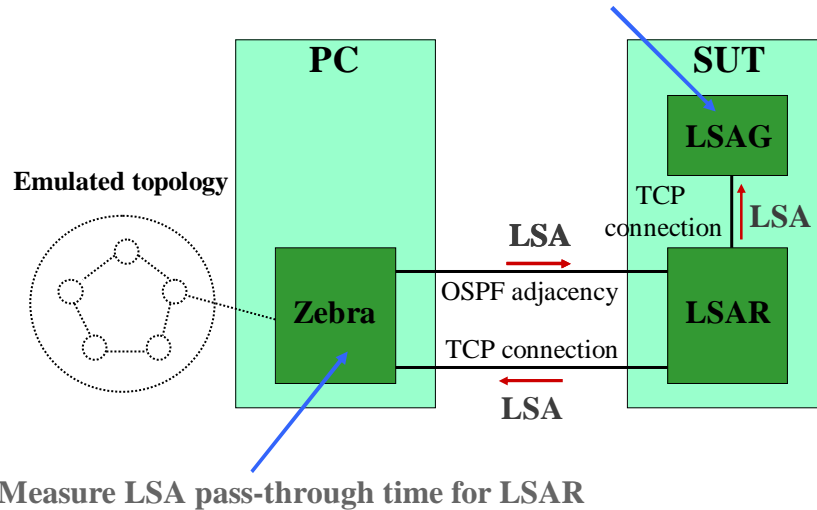
## Performance Evaluation

- Performance of LSAR and LSAG through lab experiments
  - LSAR and LSAG are key to real-time monitoring
- How performance scales with LSA-rate and network size

36

## Experimental Setup

Measure LSA processing time for LSAG



37

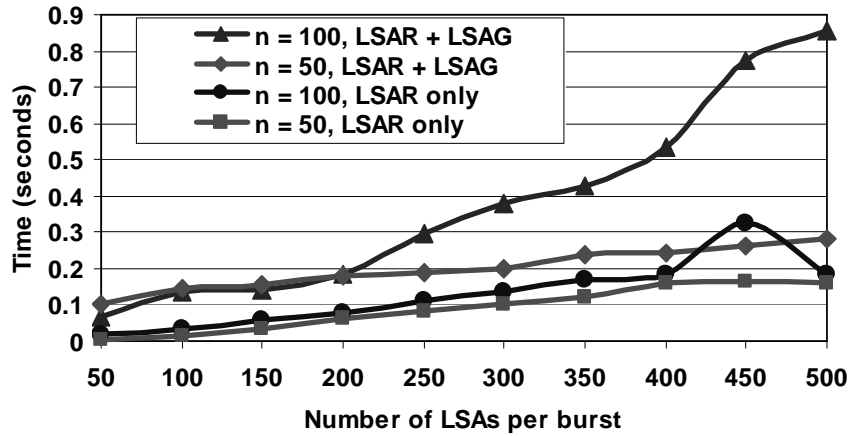
## Methodology

- Send a burst of LSAs from Zebra to LSAR
  - Vary number of LSAs ( $l$ ) in a burst of 1 sec duration
- Use of fully connected graph as the emulated topology
  - Vary number of nodes ( $n$ ) in the topology
- Performance measurements
  - LSAR performance: LSA “pass-through” time
    - Zebra measures time difference between sending and receiving an LSA from LSAR
  - LSAG performance: LSA processing time
    - Instrumentation of LSAG code

38

## LSAR Performance

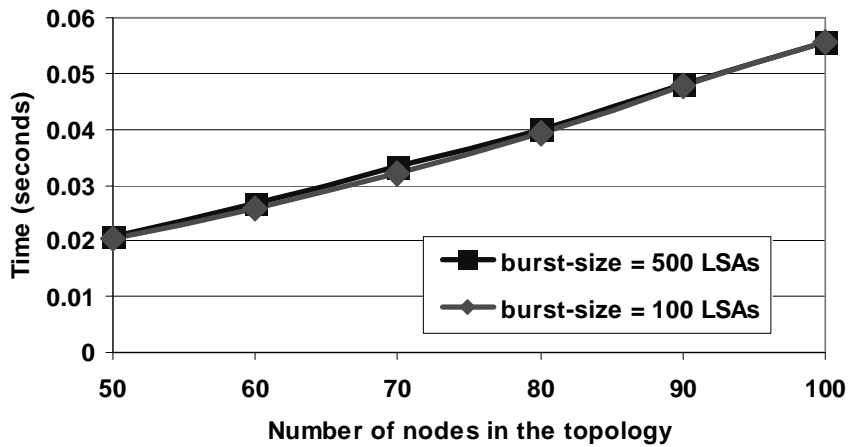
Mean LSA pass-through time (LSAR) v/s burst-size



39

## LSAG Performance

Mean LSA processing time (LSAG) v/s network size

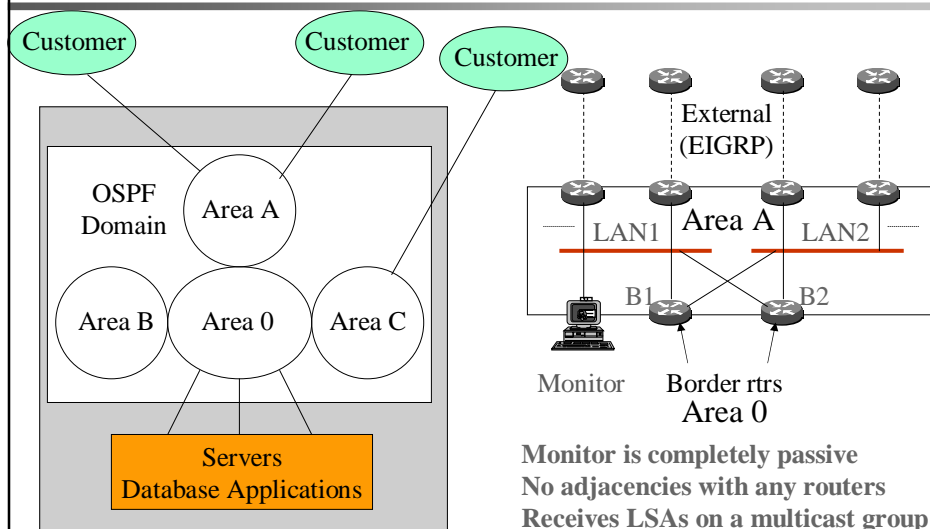


## Enterprise Network Case Study

- The network provides customers with connectivity to applications and databases residing in the data center
- OSPF network
  - 15 areas, 500 routers
    - This case study covers 8 areas, 250 routers
    - One month: April 2002
  - Link-layer = Ethernet-based LANs
- Customers are connected via leased lines
  - Customer routes are injected via EIGRP into OSPF
    - The routes are propagated via external LSAs
    - Quite reasonable for the enterprise network in question

41

## Enterprise Network Topology

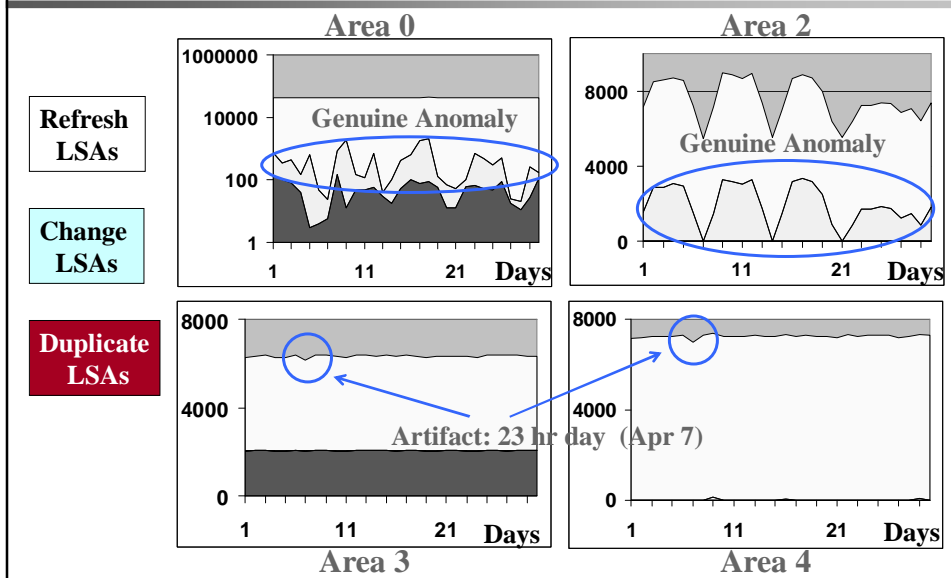


## Highlights of the Results

- Categorize, baseline and predict
  - Categories: Refresh, Change, Duplicate; External, Internal
  - Bulk of LSA traffic is due to refresh
  - Refresh LSA traffic is smooth: no evidence of refresh synchronization across network
  - Refresh LSA traffic is predictable from router configuration info
- Detect, diagnose and act
  - Almost all LSAs arise from persistent yet partial failure modes
  - Internal LSA spikes
    - Indicate router hardware degradation
    - Carry out preventive maintenance
  - External LSA spikes
    - Indicate degradation in customer connectivity
    - Call customer before customer calls you
- Propose Improvements
  - Simple configuration changes to reduce duplicate LSA traffic



43

## LSA Traffic in Different Areas



## Baseline LSA Traffic: Refresh LSAs


- Refresh LSA traffic can be reliably predicted using information available in router configuration files
  - Important for workload modeling
  - See paper for details



45

## Refresh process is not synchronized


Hour 12 of Apr 10, 2002 for area 3



- No evidence of synchronization
  - Contrary to simulation-based study in [Basu01]
- Reasons
  - Changes in the topology help break synchronization
  - LSA refresh at one router is not coupled with LSA refresh at other routers
  - Drift in the refresh interval of different routers

46

## Anomaly Detection: Change LSAs





- Internal to OSPF domain versus external
  - Change LSAs due to external events dominated
  - Not surprising due to large number of leased lines used to import customer routes into OSPF
    - Customer volatility → network volatility

47


## Root Causes of Change LSAs

- Persistent problem → flapping → numerous change LSAs
  - Internal LSA spikes → hardware router problems
    - OSPF monitor identified a problem early and led to preventive maintenance
  - External LSA spikes → customer route volatility
    - Overload of an external link to a customer between 8 pm – 4 am causes EIGRP session on that link to flap






# Overhead: Duplicate LSAs




- Why do some areas witness substantial duplicate LSA traffic, while other areas do not witness any?
  - OSPF flooding over LANs leads to control plane asymmetries and to imbalances in duplicate LSA traffic

# OSPF Operations over Broadcast Networks

- Each node sends an LSA to multicast group *DR-rtrs*
  - Both designated router (DR) and backup designated router BDR subscribe to this group



- DR floods the LSA back to all routers on the network
  - Send to *all-rtrs* multicast group to which all nodes subscribe

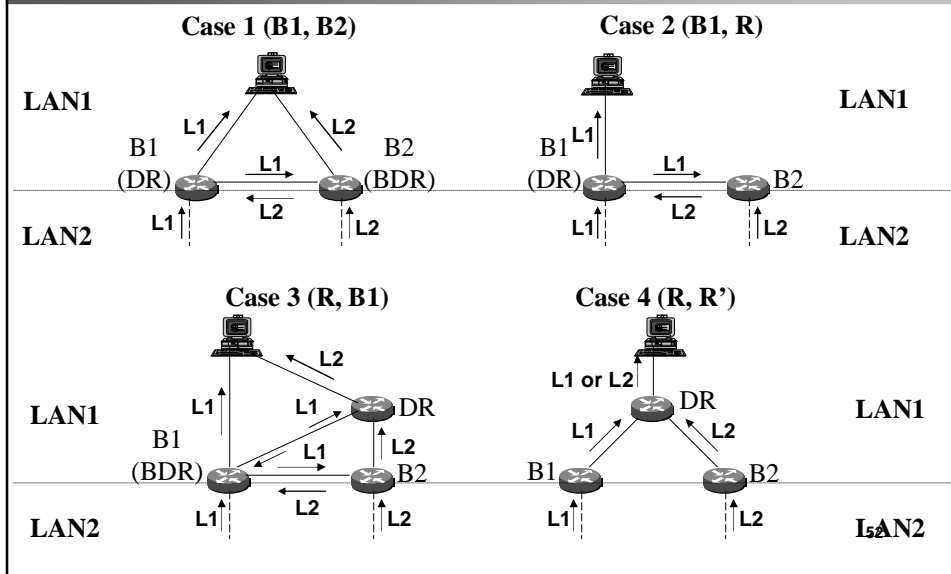


## Control Plane Asymmetry

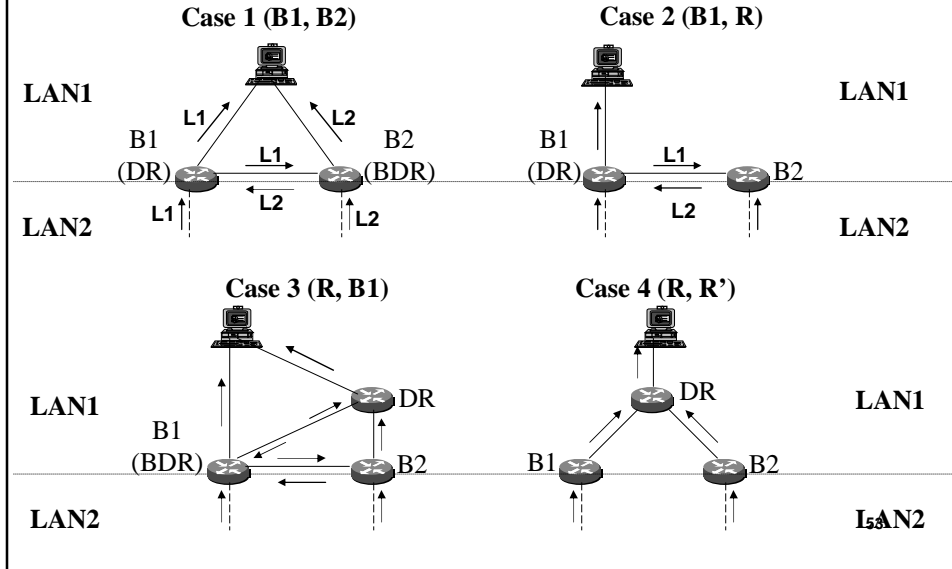
- Two LANs (LAN1 and LAN2) in each area
- Monitor is on LAN1
- Routers B1 and B2 are connected to LAN1 and LAN2
- LSAs originated on LAN2 can get duplicated depending on which routers have become DR and BDR on LAN1
  - Leads to control plane asymmetry
  - Four cases
- Note: if a BDR receives an LSA on another interface, it floods the LSA to all nodes (i.e., it sends the LSA to the *all-rtrs* address)

51

## Four Cases



## Four Cases



## Eliminating Duplicate LSA Traffic

	Case1	Case 2	Case 3	Case 4
<b>Duplicate LSA traffic</b>	High	None	High	None
<b>Deterministic via configuration</b>	Yes	No	No	Yes
<b>Area 2</b>		X		X configuration change
<b>Area 3</b>			X	X configuration change

## Summary

- Categorize and baseline LSA traffic
  - Refresh LSAs: constitute bulk of overall LSA traffic
    - No evidence of synchronization between different routers
    - Refresh LSA traffic predictable from configuration information
- Detect, diagnose and act on anomalies
  - Change LSAs: can indicate persistent yet partial failure modes
    - Internal LSA spikes → hardware router problems → preventive router maintenance
    - External LSA spikes → customer congestion problems → “preventive” customer care
- Propose changes to improve performance
  - Duplicate LSAs: can arise from control plane asymmetries
    - Simple configuration changes can eliminate duplicate LSAs and improve performance

55

## Other Problems Caught

- Configuration problem
  - Identified assignment of same router-id to two routers in enterprise network
- OSPF implementation bug
  - Caught a bug in type-3 LSA generation code of a router vendor in ISP network
    - Faster refresh of LSAs than standards-mandated rate

56

## LSA aGregator (LSAG)

- Analyzes “reflected” LSAs from LSARs in real-time
- Generates console messages:
  - Change in OSPF network topology
    - ADJACENCY COST CHANGE: rtr 10.0.0.1 (intf 10.0.0.2)  
→ rtr 10.0.0.5 old\_cost 1000 new\_cost 50000 area 0.0.0.0
  - Node flaps
    - RTR FLAP: rtr 10.0.0.12 no\_flaps 7 flap\_window 570 sec
  - LSA storms
    - LSA STORM: lstype 3 lsid 10.1.0.0 advrt 10.0.0.3 area 0.0.0.0 no\_lsas 7 storm\_window 470 sec
  - Anomalous behavior
    - TYPE-3 ROUTE FROM NON-BORDER RTR: ntw 10.3.0.0/24 rtr 10.0.0.6 area 0.0.0.0
- Dumps snapshots of network topology

57

## OSPFScan

- Tools for off-line analysis of LSA archives
  - Parse, select (based on queries), and analyze
- Functionality supported by OSPFScan
  - Classification of LSA traffic
    - Change LSAs, refresh LSAs, duplicate LSAs
  - Emulation of OSPF Routing
    - How OSPF routing tables evolved in response to network changes
    - How end-to-end path within OSPF domain looked like at any instance
  - Modeling of topology changes
    - Vertex addition/deletion and link addition/deletion/change\_cost
  - Playback of topology change events
  - Statistics and report generation

58

## Deployment

- Tier-1 ISP network
  - Area 0, 100+ routers; point-to-point links
  - Deployed since January, 2003
  - LSA archive size: 8 MB/day
  - LSAR connection: partial adjacency mode
- Enterprise network
  - 15 areas, 500+ routers; Ethernet-based LANs
  - Deployed since February, 2002
  - LSA archive size: 10 MB/day
  - LSAR connection: host mode

59

## LSAG in Day-to-day Operations

- Generation of alarms by feeding messages into higher layer network management systems
  - Grouping of messages to reduce the number of alarms
  - Prioritization of messages
- Validation of maintenance steps and monitoring the impact of these steps on network-wide OSPF behavior
  - Example:
    - Network operators use cost-out/cost-in of links to carry out maintenance
    - A "link-audit" web-page allows operators to keep track of link costs in real-time

60

## Long Term Analysis by OSPFScan

- LSA traffic analysis
  - Identified excessive duplicate LSA traffic in some areas of Enterprise Network
    - Led to root-cause analysis and preventative steps
- Statistics generation
  - Inter-arrival time of change LSAs in ISP network
    - Fine-tuning configurable timers related to route calculation (= SPF calculation)
  - Mean down-time and up-time for links and routers in ISP network
    - Assessment of reliability and availability

61