


AF-TM 94-0987

Ordered BECN




Raj Jain, Shiv Kalyanaraman, Ram Viswanathan

**Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
<http://www.cse.wustl.edu/~jain/>**

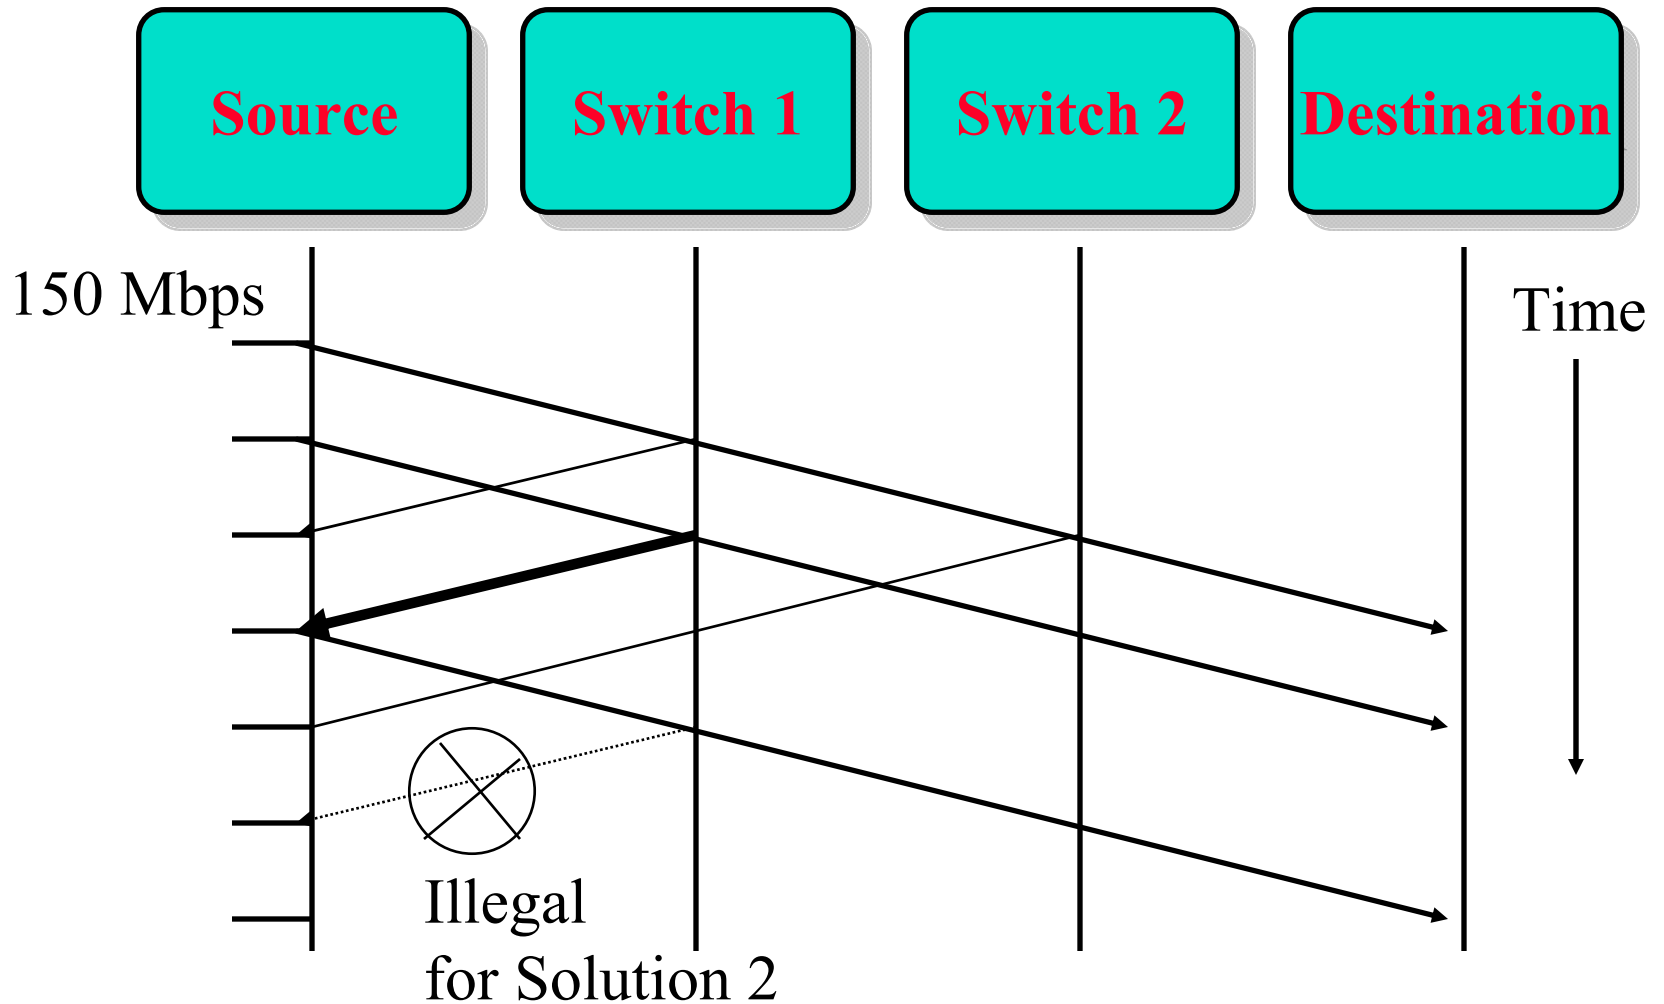


- ❑ Why BECN's cause confusion
- ❑ How to Fix it

Backward Explicit Congestion Notification



Backward Explicit Congestion Notification



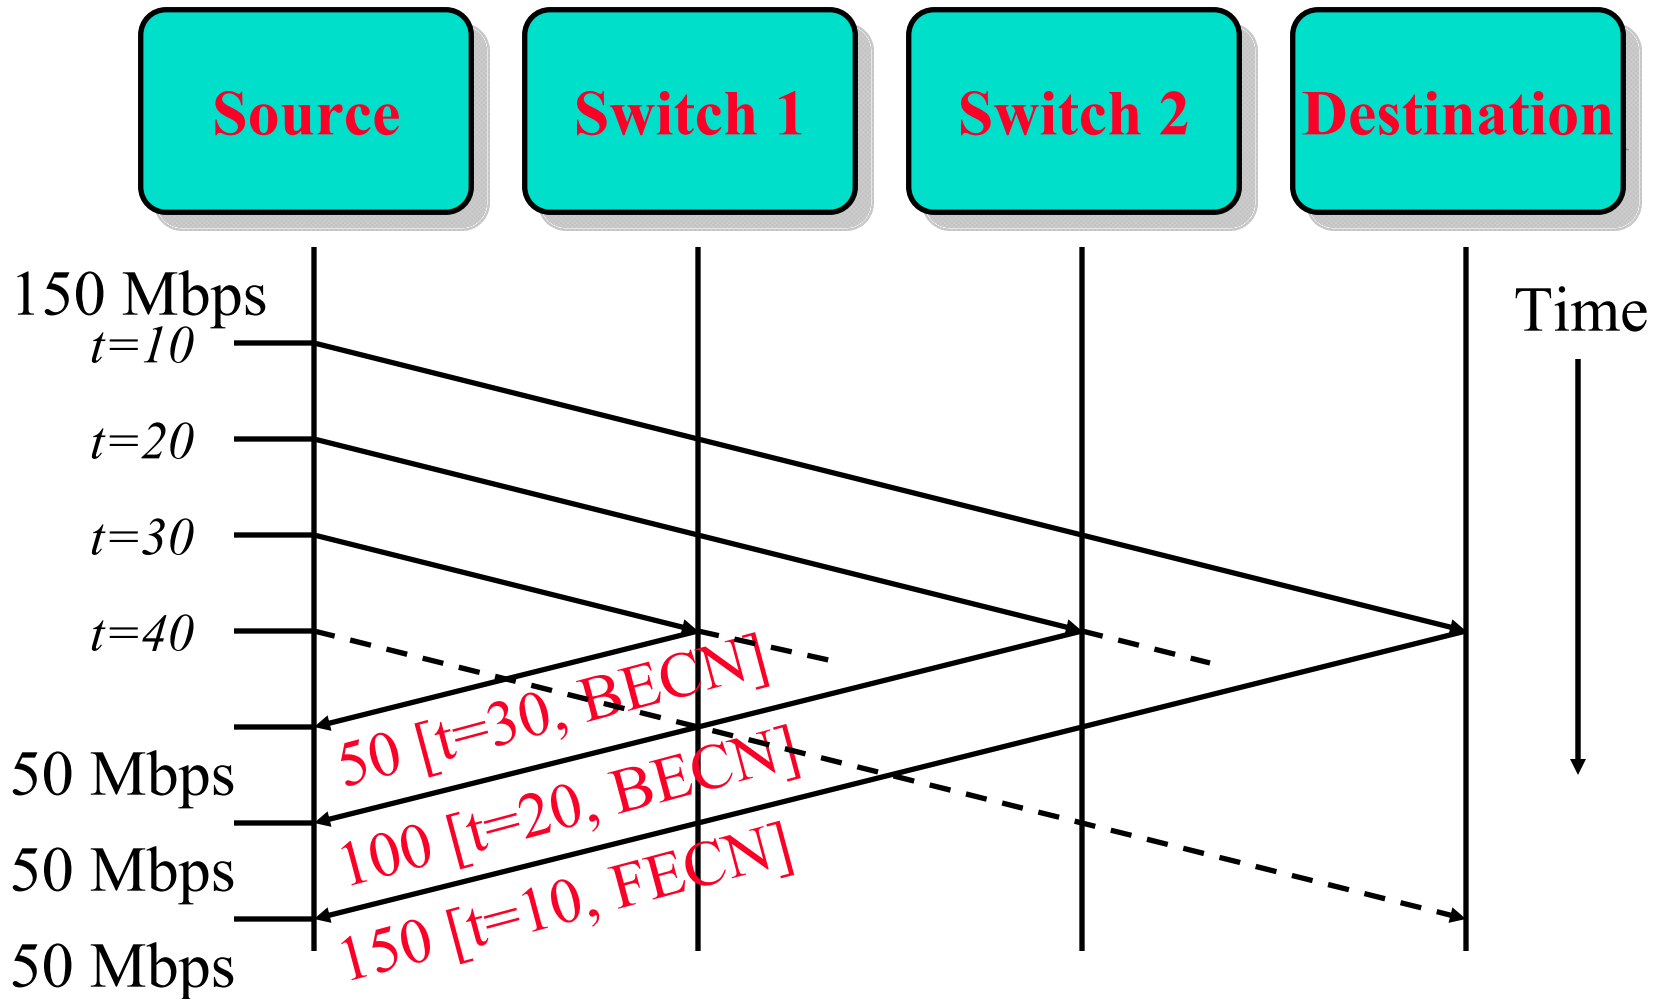
BECN: Problems

- ❑ BECN's from multiple switches can come out of order
- ❑ One BECN can overrule the effect of other
- ❑ BECN can actually increase the rate

Ordered BECN

- ❑ Each RM Cell contains a timestamp from the source
- ❑ Source ignores out-of-date information. Remembers the time up to which it has already acted
- ❑ BECN can not increase
FECN can increase/decrease
- ❑ RM Cell contains BECN/FECN bit
- ❑ Out-of-date FECN's are also ignored

Ordered BECN: Example



OSU Ordered BECN Rules

- ❑ BECN is optional for switches
- ❑ BECN should be sent only if overloaded
- ❑ Source includes a timestamp in each RM cell
The timestamp field is ignored by other nodes
- ❑ RM cells include BECN/FECN bit initialized to FECN
- ❑ The original RM cell is forwarded towards the destination.
- ❑ A copy of the RM cell with BECN/FECN bit changed to BECN is returned to the source
- ❑ The source remembers the timestamp of the last BECN/FECN RM cell received
- ❑ All BECN/FECNs with earlier timestamp are ignored
- ❑ All BECN that would result in rate increase are ignored

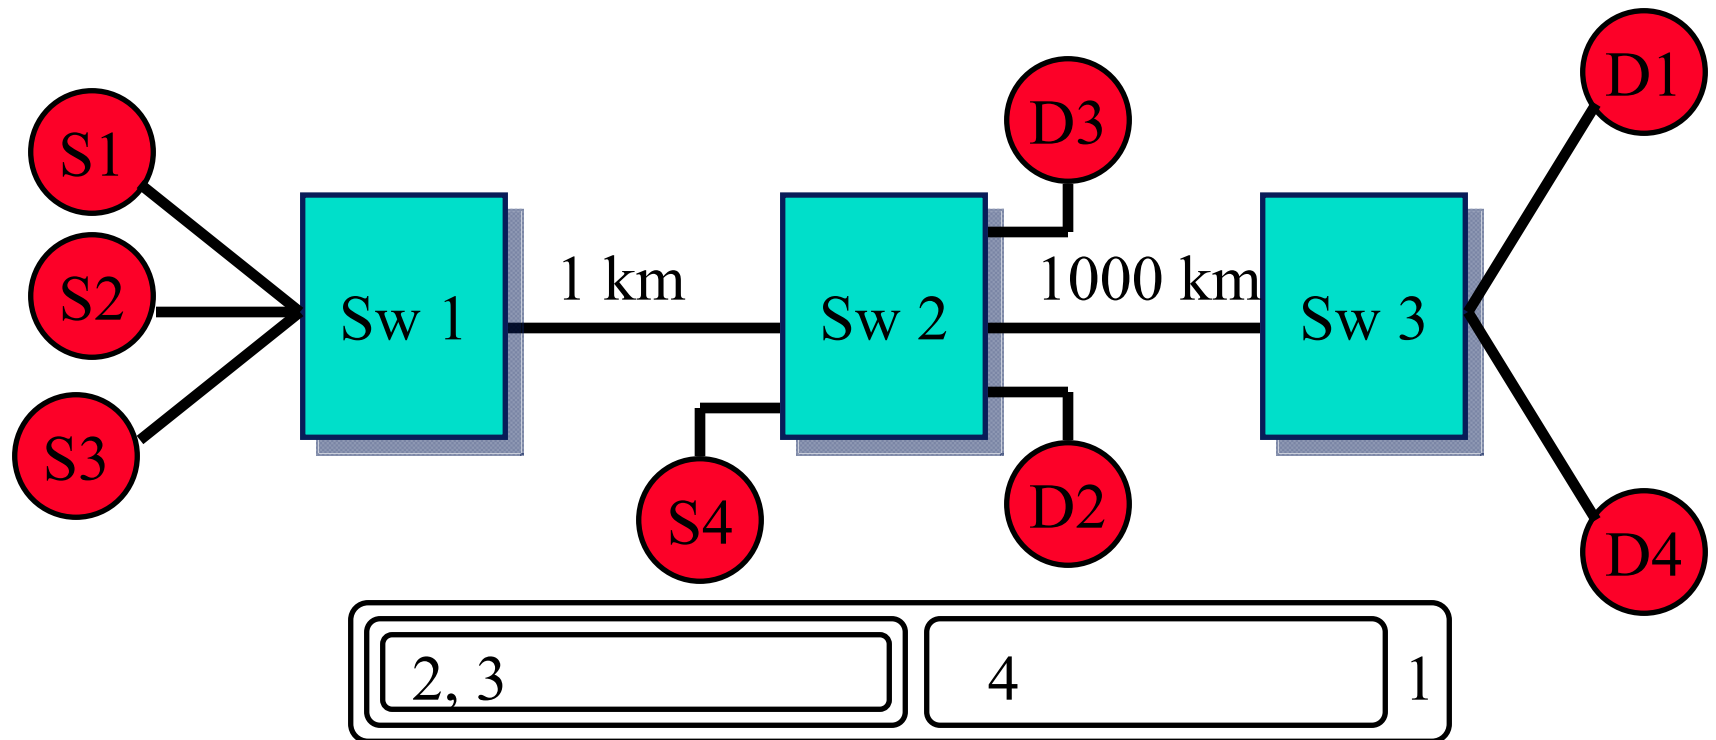
BECN Rules

- ❑ BECN :When a switch generates an RM cell destined for the source when the switch is not acting as a virtual destination
- ❑ RM cells include BECN/FECN bit initialized to FECN
- ❑ All BECN that would result in rate increase are ignored
- ❑ BECN is optional for switches
- ❑ BECN should be sent only if overloaded
- ❑ The original RM cell is forwarded towards the destination.
- ❑ A copy of the RM cell with BECN/FECN bit changed to BECN is returned to the source
- ❑ BECN can be sent on the “First RM” cell
- ❑ If the source has received an RM cell from the destination, then no switch should generate a BECN
- ❑ => Need “Not yet heard from destination” bit in RM cells

Subsequent BECN Solutions

- ❑ Source ignores all BECNs that result in rate increase.
- ❑ Switch does not generate a subsequent BECN if it has seen other reverse RM cells for this VC.
=> Need bit per active VC in the switch
- ❑ The source includes a timestamp in each RM Cell and remembers the timestamp of the last BECN/FECN RM cell acted. Ignores BECN/FECNs with earlier timestamps are ignored.

Upstream Configuration



- ❑ All links 155 Mbps. Link 1: 1 km, Link 2: 1000 km
- ❑ Max-min optimal: 51.3, 51.3, 51.3, 103.6 Mbps
- ❑ Goal: To check that all unused capacity is allocated to S4

Summary



- ❑ Without a timestamp field, BECN's can cause confusion
- ❑ Need a timestamp field and BECN/FECN bit in the RM Cell
- ❑ Need to specify rules for BECN/FECN usage
- ❑ Sequence number can be used in place of timestamps. But, timestamp can also be used for round-trip delay estimate