Time-based Updates in Software Defined Networks

Tal Mizrahi, Yoram Moses
Technion – Israel Institute of Technology

SDNRG, IETF Meeting, Berlin, July 2013

Introduction

- Time has been used for many years in distributed systems, typically not for network configuration.
- SDN: rapid and frequent configuration changes ➔ simple, scalable updates with minimal transient effects.
- Accurate network time synchronization has evolved (NTP, PTP) and has become accessible.

Time-triggered Configuration Updates

- We propose to use time as a tool for network updates / reconfiguration.

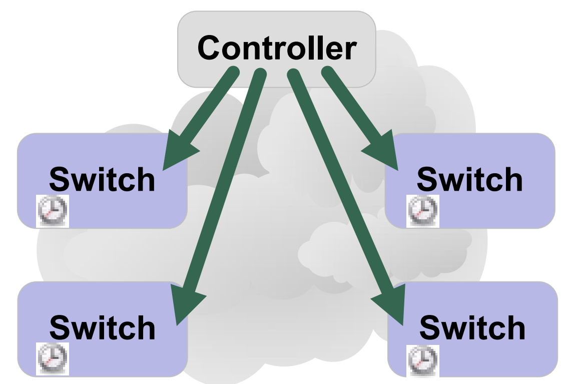

- **TimeConf**: a class of time-triggered configuration scenarios:
  - Coordinated updates.
  - An ordered sequence of updates based on a sequence of scheduled times.

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Example: Path Reconfiguration

- Flow A: path 1 → path 2.

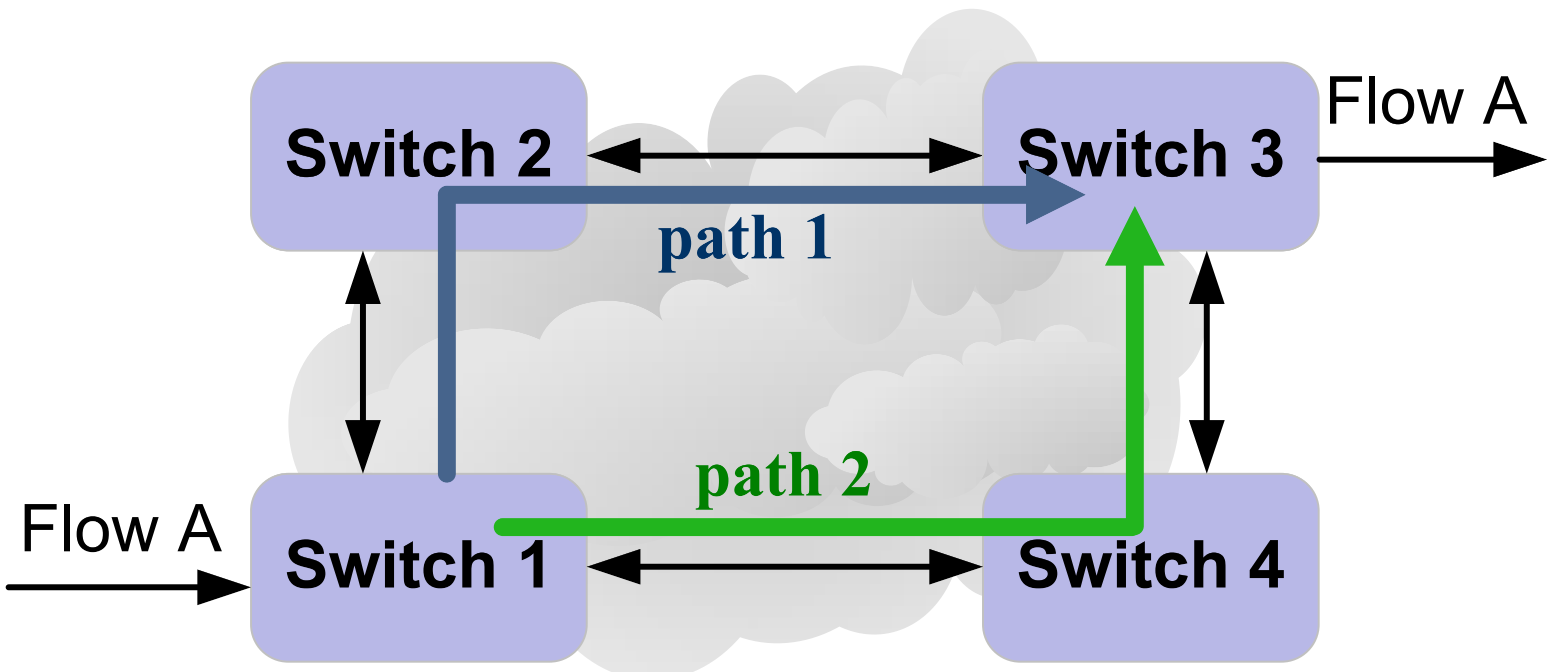

Example: Path Reconfiguration

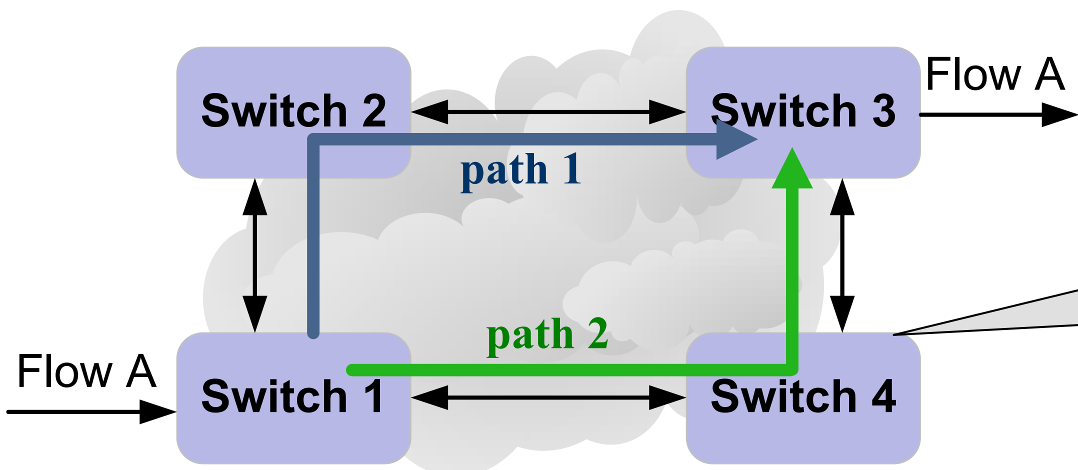

Best-effort approach: update all switches when possible.

Example: Path Reconfiguration

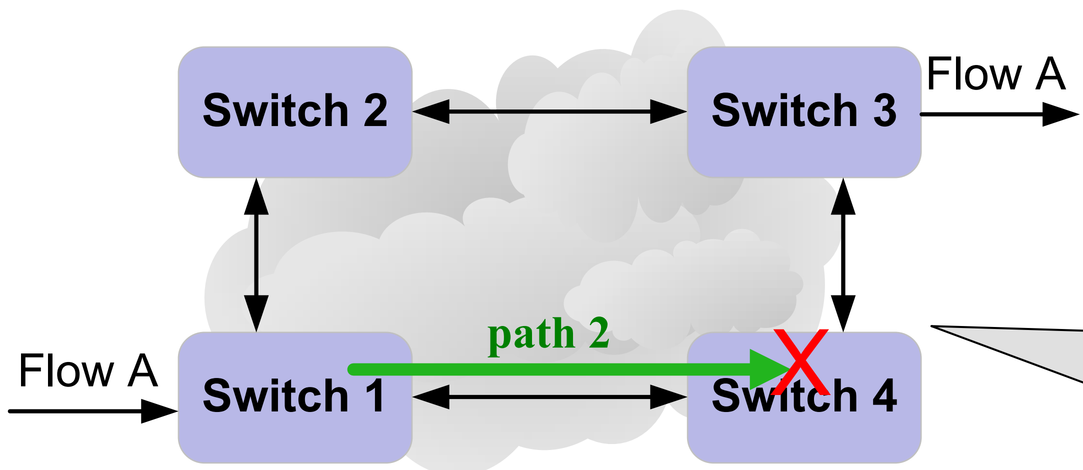

Best-effort approach:
What happens if Switch 1 is updated before Switch 4?

Inconsistency ☹

- Consistency [Reitblatt et al.]: every packet sent through the network is processed according to a single configuration version, either the previous or the current one.

Example: Path Reconfiguration

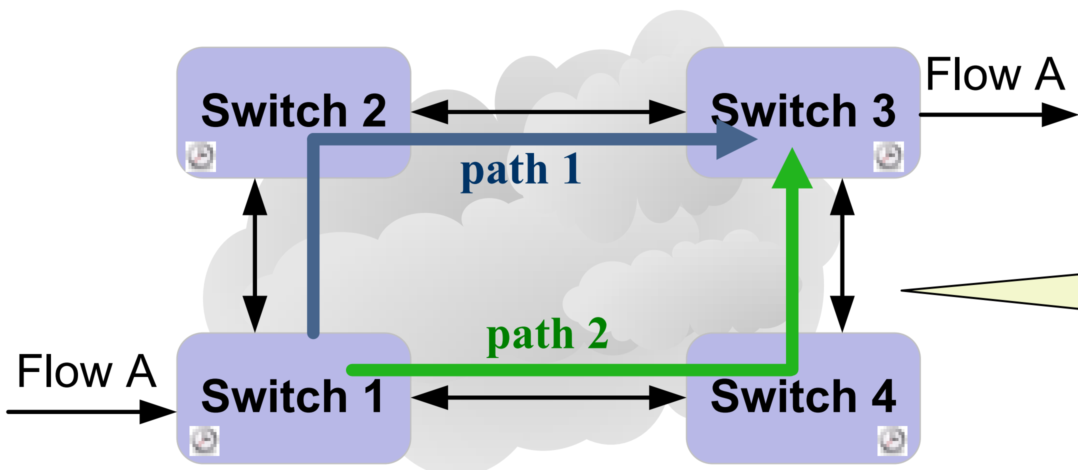

TimeConf: all switches are updated at (T-δ, T+ δ).

Example: Path Reconfiguration

What can we say about consistency with TimeConf?

D = maximal delay from Switch 1 to Switch 3 through path 1

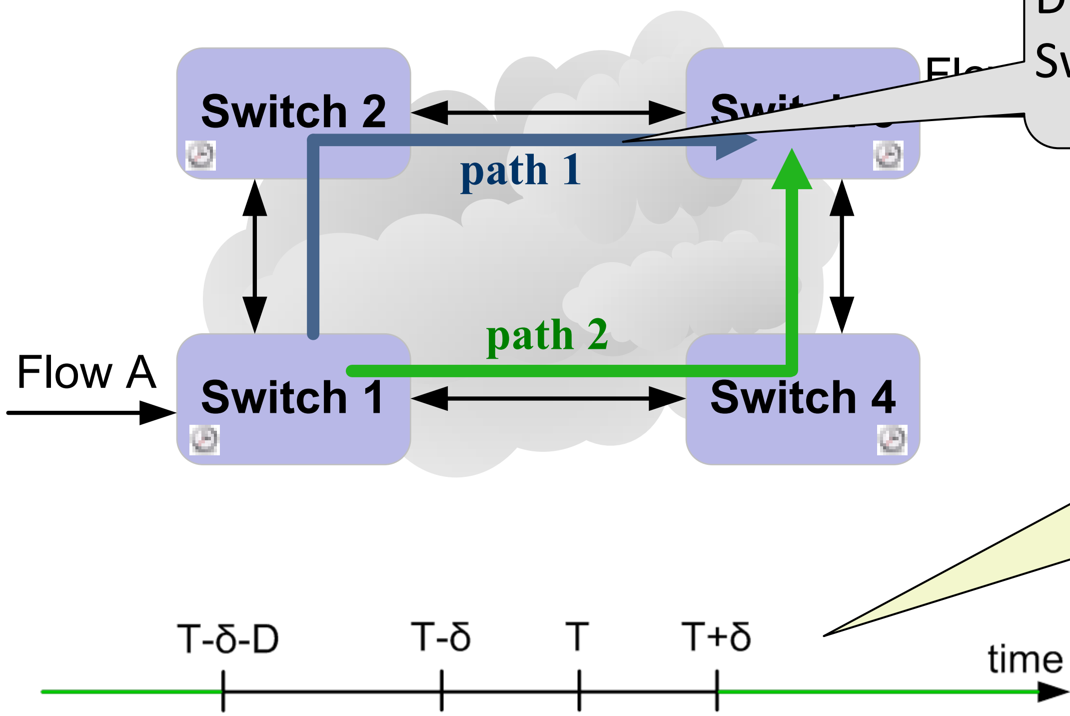

TimeConf: inconsistent for a **short** period (T-δ-D, T+ δ).

Consistent otherwise.

TimeConf: short transition period.

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Example: Path Reconfiguration

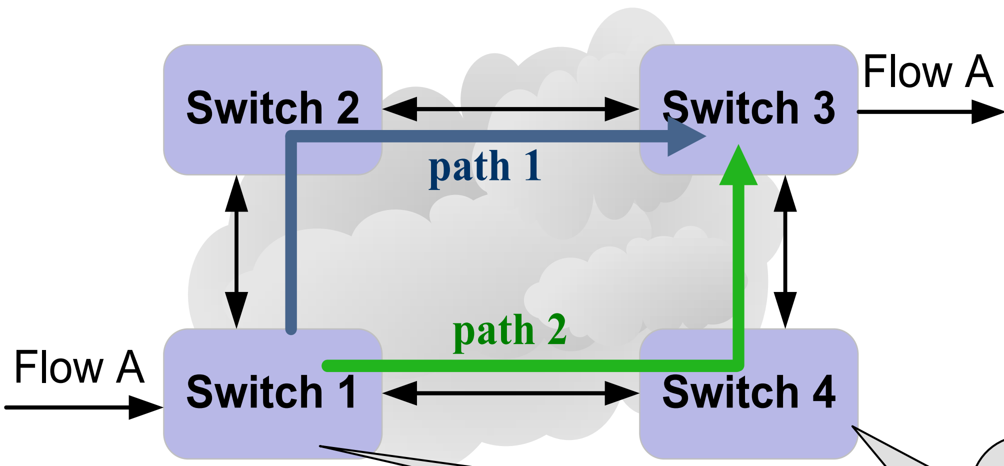

Sequential approach:

1. Update Switch 4.
2. Receive ACK from Switch 4.
3. Update Switch 1.

Example: Path Reconfiguration

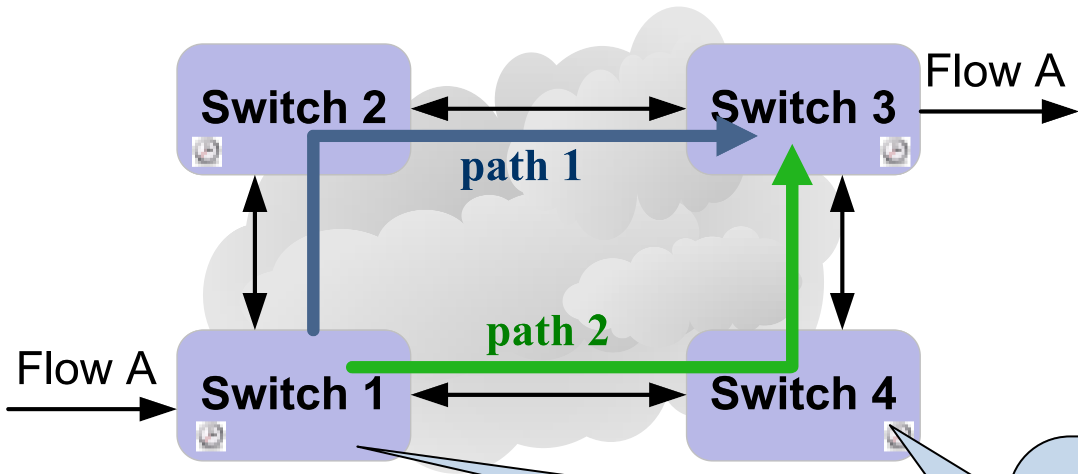

Sequential TimeConf:
Send updates to all switches:
-Switch 4 at time T4
-Switch 1 at time T1
$T_4 < T_1$

Sequential TimeConf: maintains consistency

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Updating Physical Layer Attributes

- Physical attributes such as:
  switch resources, flow rate limit, port transmission rate, …
- Per-packet consistency is irrelevant.
- TimeConf can apply the update to all switches at the same time T.

Example: Port Rate Reconfiguration

What happens when the sequential approach is not possible?

Both switches are configured to change the port rate at time T.

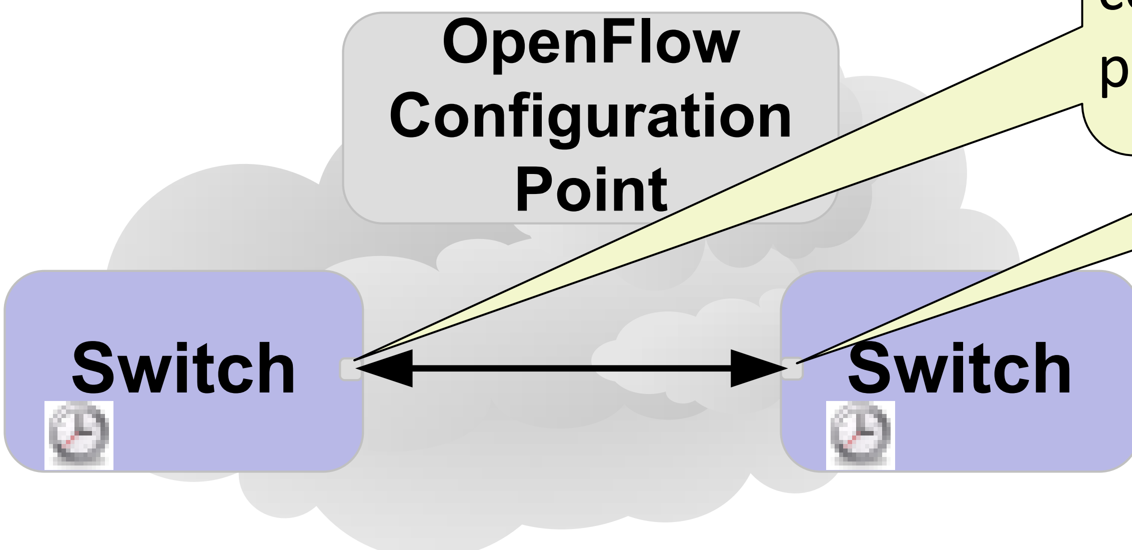

TimeConf allows physical layer updates with short transition period.

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Consistency / Simplicity Tradeoff

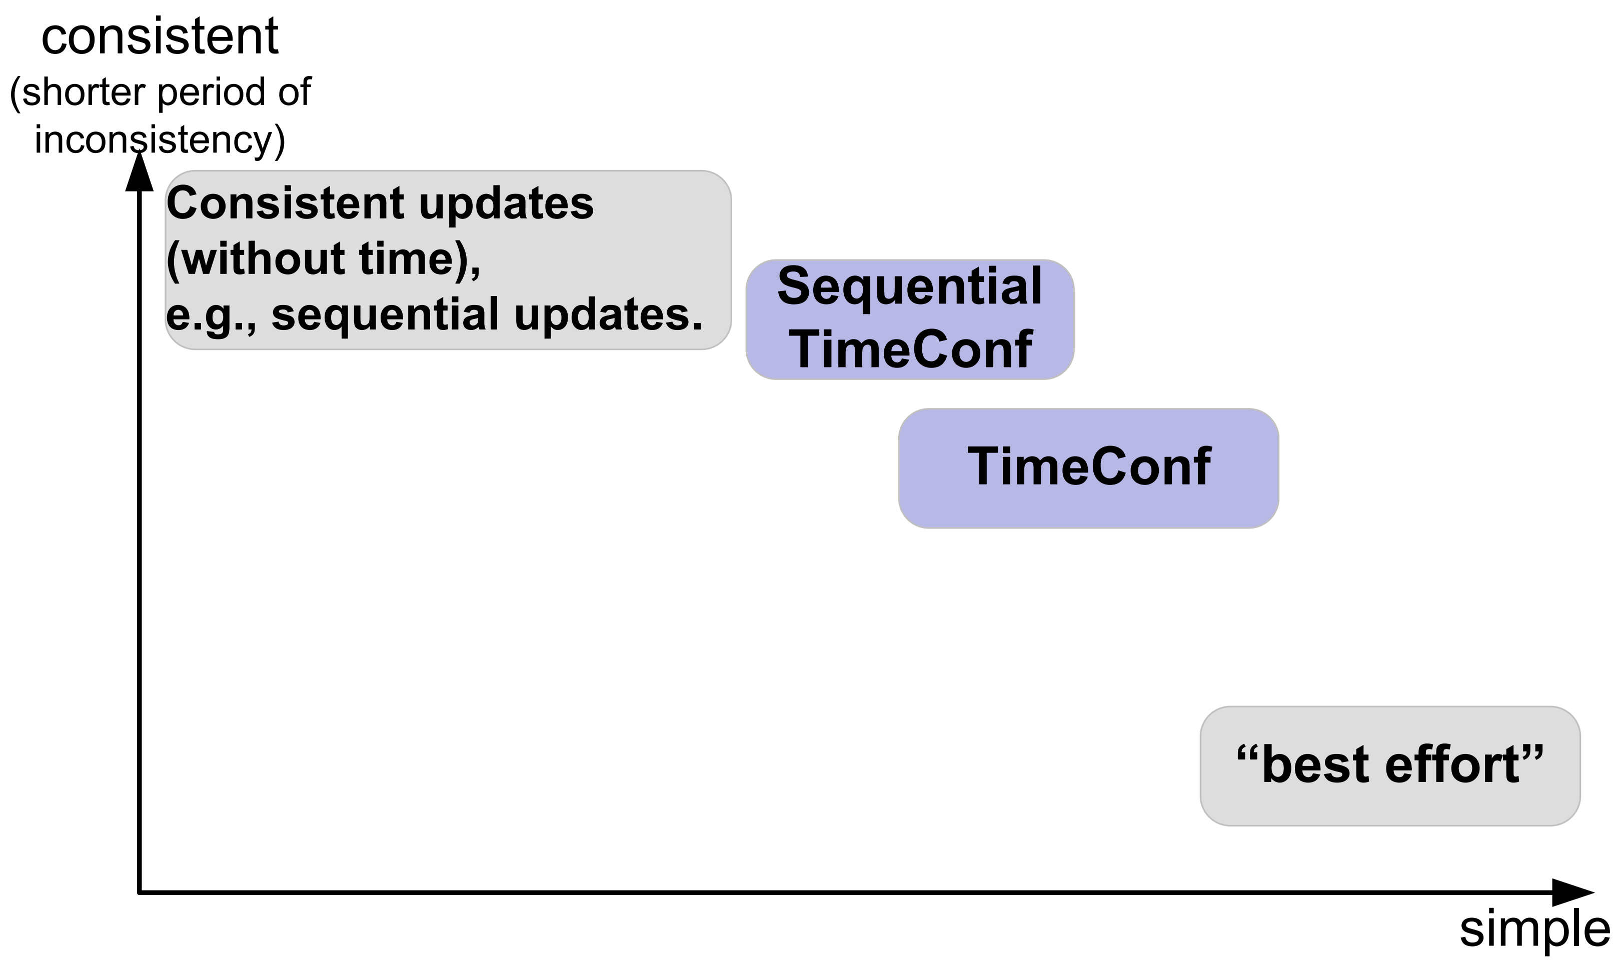

Agenda

- Transition time during updates.
- Sequential updates.
- Updating physical layer configurations.
- Consistency/simplicity tradeoff.
- Time-based updates in practice.

Time-based Updates in Practice

- Proposal to the Open Networking Foundation (ONF):
  - Time-based Updates in OpenFlow: A Proposed Extension to the OpenFlow Protocol
    http://tx.technion.ac.il/~dew/OFTimeTR.pdf
  - Presented in July 2013 to:
    - ONF Extensibility WG
    - ONF Configuration and Management WG.

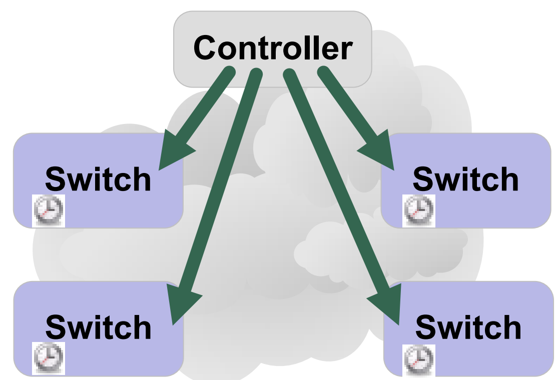

- Proposal to the IETF:
  - http://tools.ietf.org/html/draft-mm-netconf-time-capability
  - To be presented in IETF 87 in Berlin, Jul 2013.

Time-based Updates in Practice: Routing Change

- I2RS: rapid and dynamic routing changes
  (e.g., http://tools.ietf.org/html/draft-atlas-i2rs-problem-statement-00).
- This example: change the route to AS1.

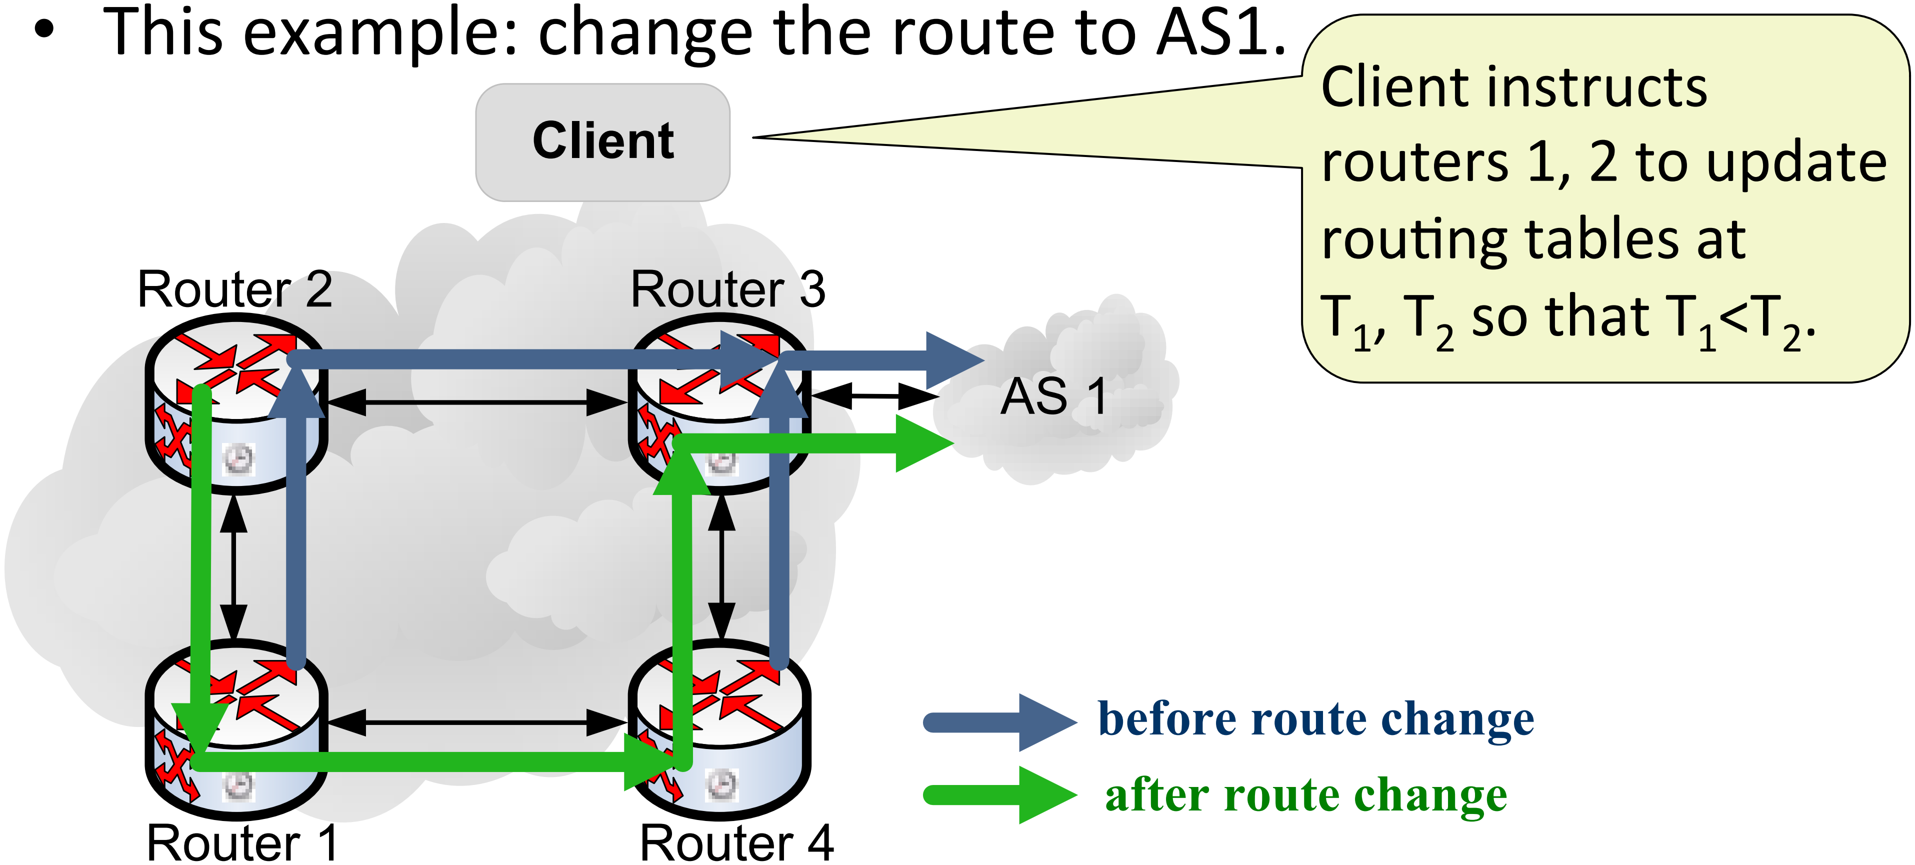

Summary

- We propose to use time as a tool for coordinated configuration updates.
- Can be used for:
  - Reducing transition period.
  - Physical layer updates.
  - Simplifying update procedure.
- This work presents a tradeoff between consistency and simplicity of the update procedure.

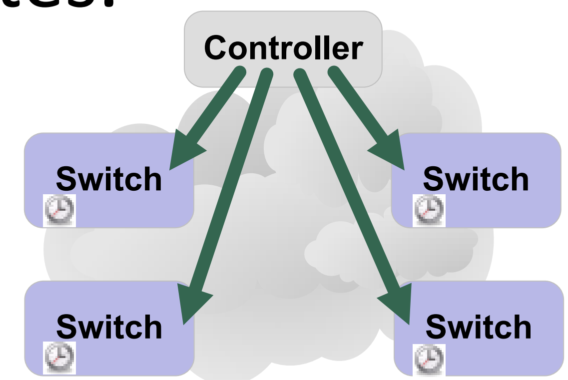

THANKS !

Further Reading

- Mizrahi, T., Moses, Y., "Time-based Updates in Software Defined Networks", the second workshop on hot topics in software defined networks (HotSDN), to appear, 2013.
  http://tx.technion.ac.il/~dew/TimeSDN.pdf
- Mizrahi, T., Moses, Y., "Time-based Updates in OpenFlow: A Proposed Extension to the OpenFlow Protocol", Technion - Israel Institute of Technology, technical report, CCIT Report #835, July 2013, EE Pub No. 1792, 2013.
  http://tx.technion.ac.il/~dew/OFTimeTR.pdf
- Mizrahi, T., Moses, Y., "Time Capability in NETCONF," IETF, draft-mm-netconf-time-capability, work in progress, 2013.
  http://tools.ietf.org/html/draft-mm-netconf-time-capability
- L. Lamport, "Using time instead of timeout for fault-tolerant distributed systems", ACM Trans. Program. Lang. Syst., vol. 6, pp. 254–280, Apr. 1984.

BACKUP SLIDES

Example: Rate Limiting Reconfiguration

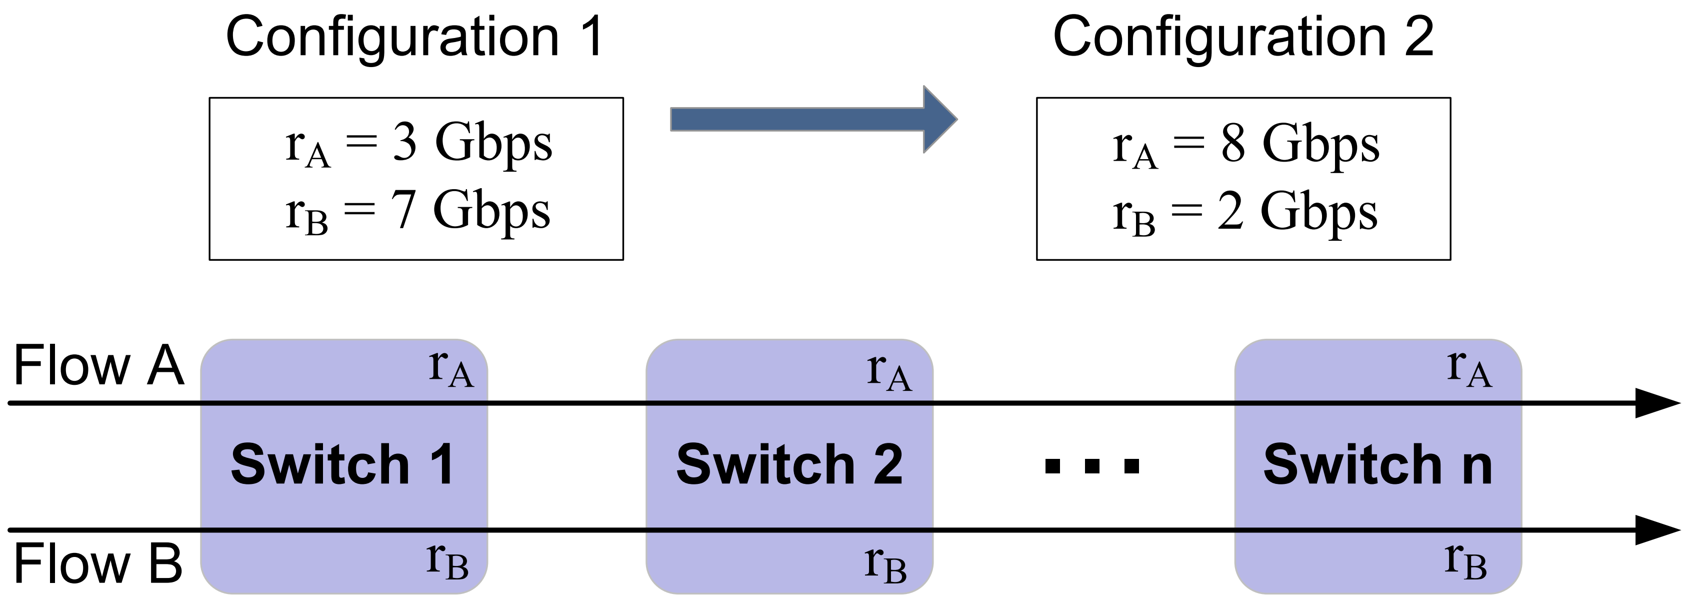

Using time: controller sends update messages to the n switches, scheduled to time T.

Example: Spanning Tree Reconfiguration

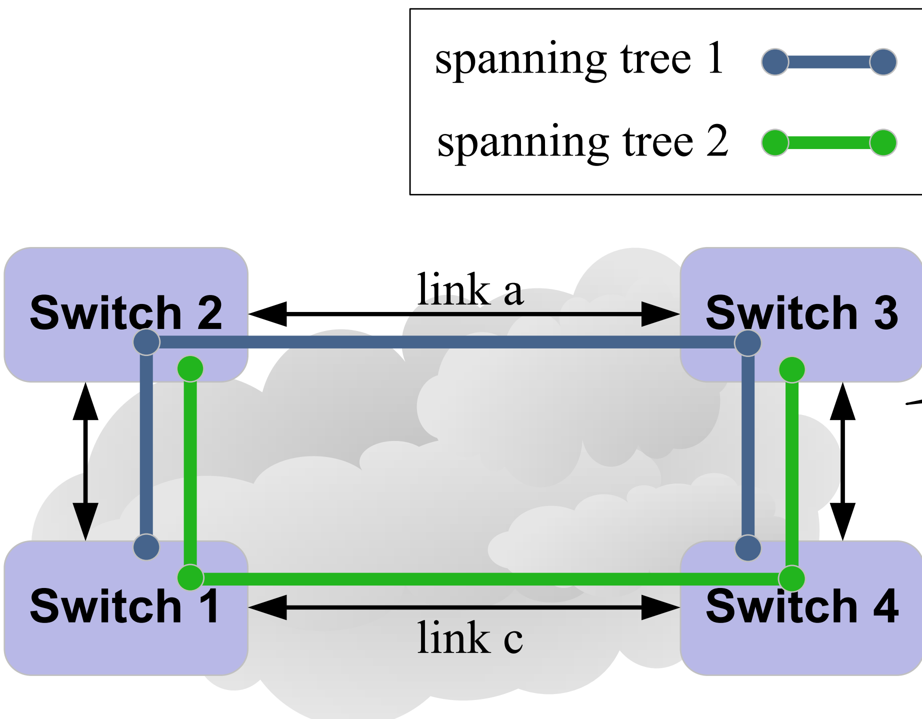

Versioning approach:

0. ST 1 uses VLAN 1 (tagged).
1. Configure VLAN 2 (tagged) with ST 2.
2. Configure switches to start using VLAN 2 (and still accept traffic from VLAN 1).
3. Remove VLAN 1 from all switches.

Example: Spanning Tree Reconfiguration

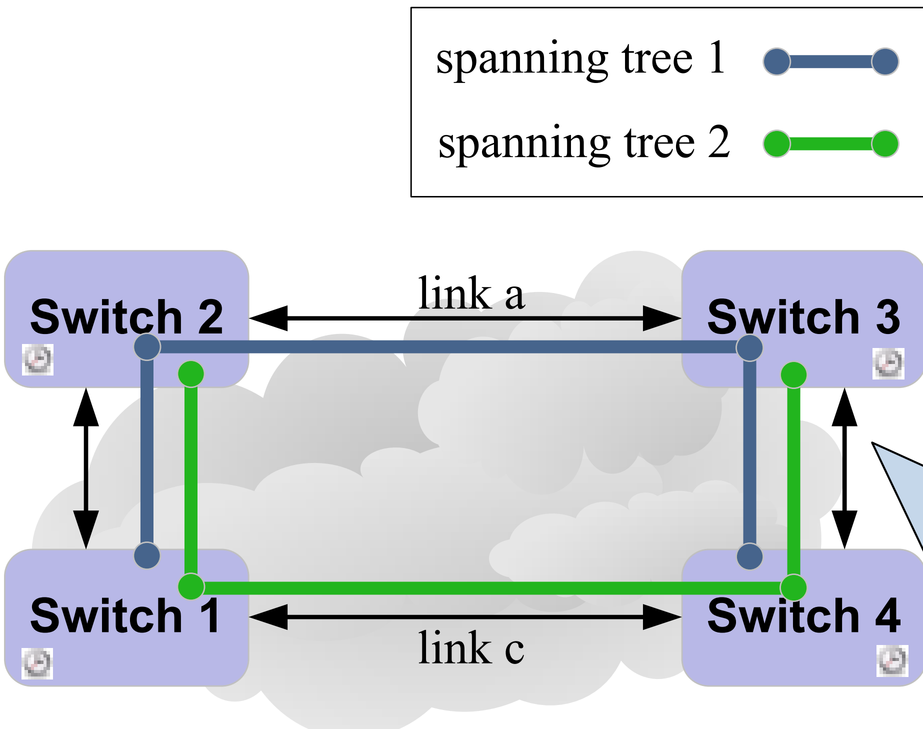

Time-based versioning approach:

0. ST 1 uses VLAN 1 (tagged).

1. Configure VLAN 2 with ST 2 ➔ T_1
   Configure switches to start VLAN 2 ➔ T_2
   Remove VLAN 1 from all switches ➔ T_3

$T_1 < T_2 < T_3$