

# BGP Churn Evolution: A perspective from the core


Ahmed Elmokashfi<sup>\*</sup>, Amund Kvalbein<sup>\*</sup>, Constantine Dovrolis<sup>\*\*</sup>

<sup>\*</sup> [ **simula** . research laboratory ]

<sup>\*\*</sup> **Georgia  
Tech**




# Fast growth in the Internet over the past decade



**The number of ASes has increased approximately by 146%**

Number of routing table entries



**The number of network destinations has increased approximately by 168%**

**Report from the IAB Workshop on Routing and Addressing 2007  
RFC 4984**

**“Routing scalability is the most important problem facing the Internet today”**

**“There is a need to devise a scalable routing and addressing system”**

**Why?**

- The rapid growth of the DFZ RIB**
- Increasing BGP churn**

# Increasing BGP churn is a potential threat to routing scalability, but what is the current status?

## We know little about the characteristics of BGP churn

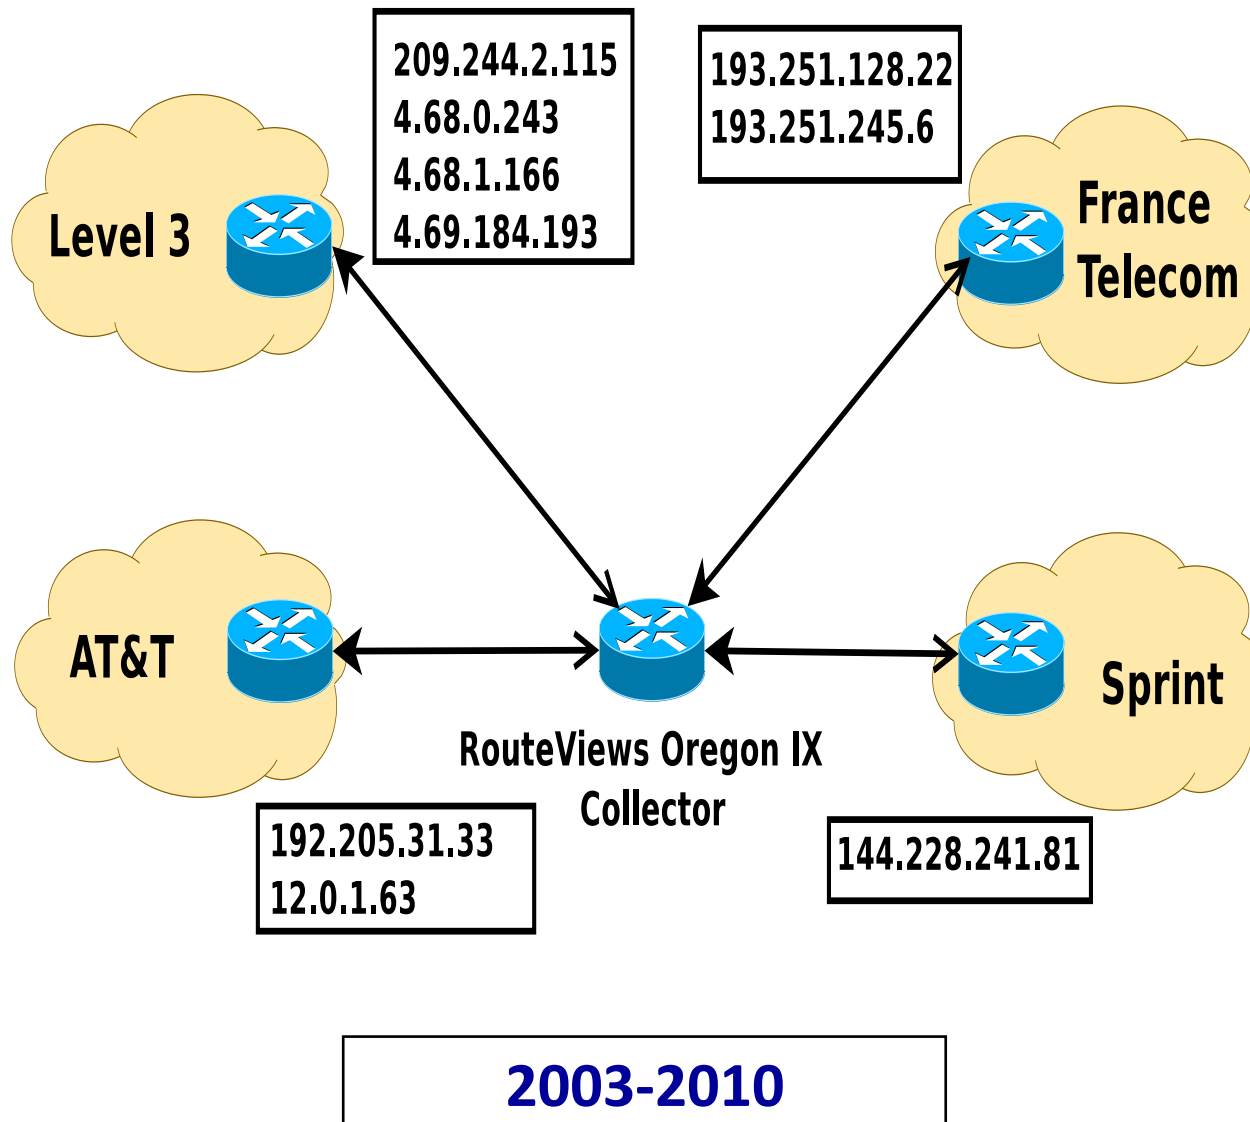
- An earlier study by Huston and Armitage reported an alarming growth in churn

*“G. Huston and G. Armitage. Projecting future IPv4 router requirements from trends in dynamic BGP behaviour. December 2006”*

## Our work

- Measures churn evolution at the core of the Internet in the past few years
- The most comprehensive study of churn evolution

# We focus on understanding BGP churn evolution over time at the core of the Internet



# What determines the observed BGP churn rate?

## The size of the network

- More elements that can fail/change/act

## The structure of the network topology

- Who peers with who?
- How many and which providers does an AS have?
- Depth of Internet hierarchy/path lengths

## Policies and protocol configuration

- MRAI timer
- Route Flap Dampening
- Route filtering and aggregation

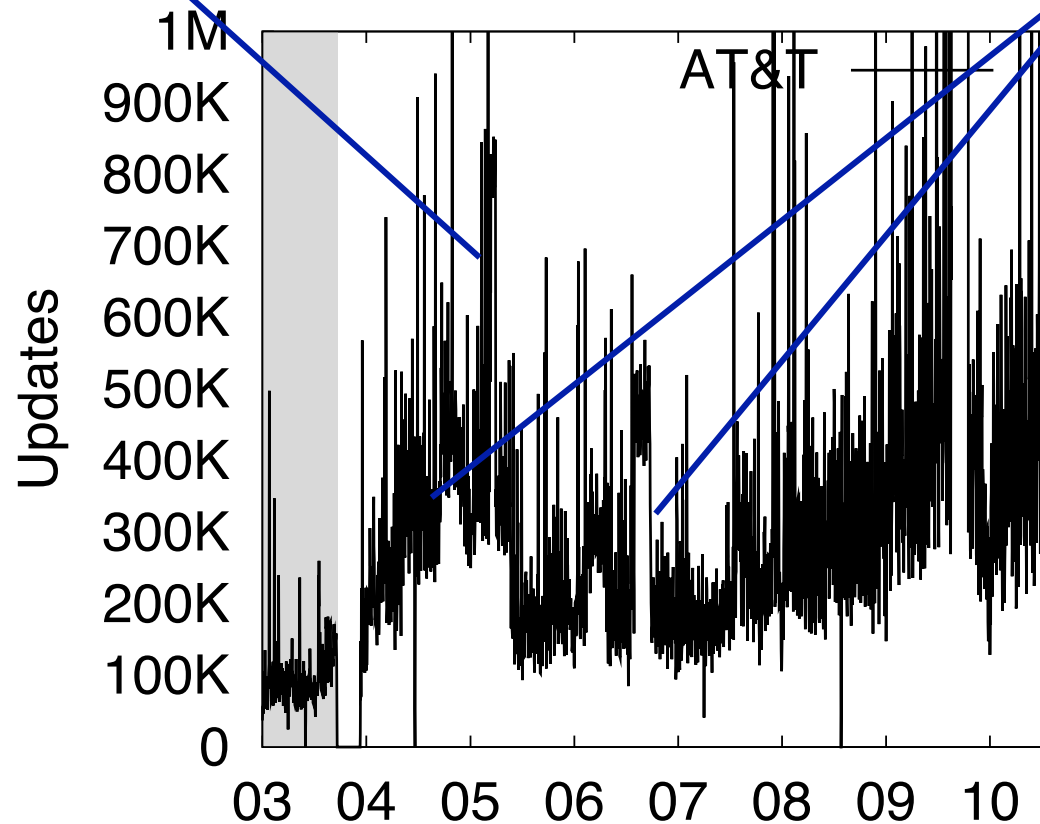
## Event types and frequencies

- Prefix withdrawals, link failures, TE operation...

# BGP churn timeseries are bursty, and it is difficult to identify a trend in them

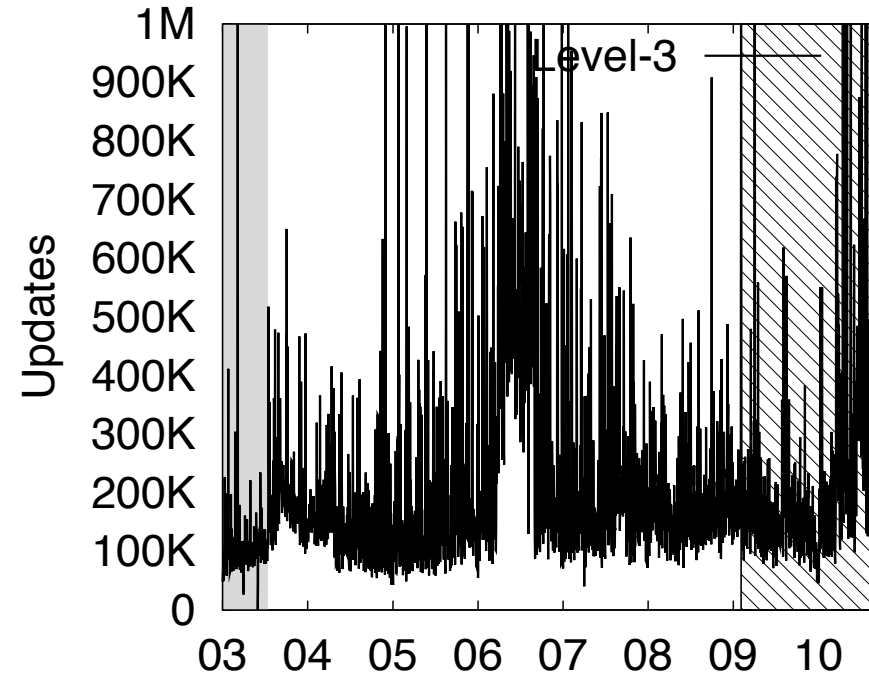

The timeseries is dominated by large frequent spikes

There are several level shifts



**Number of BGP updates per day (2003-2010)**

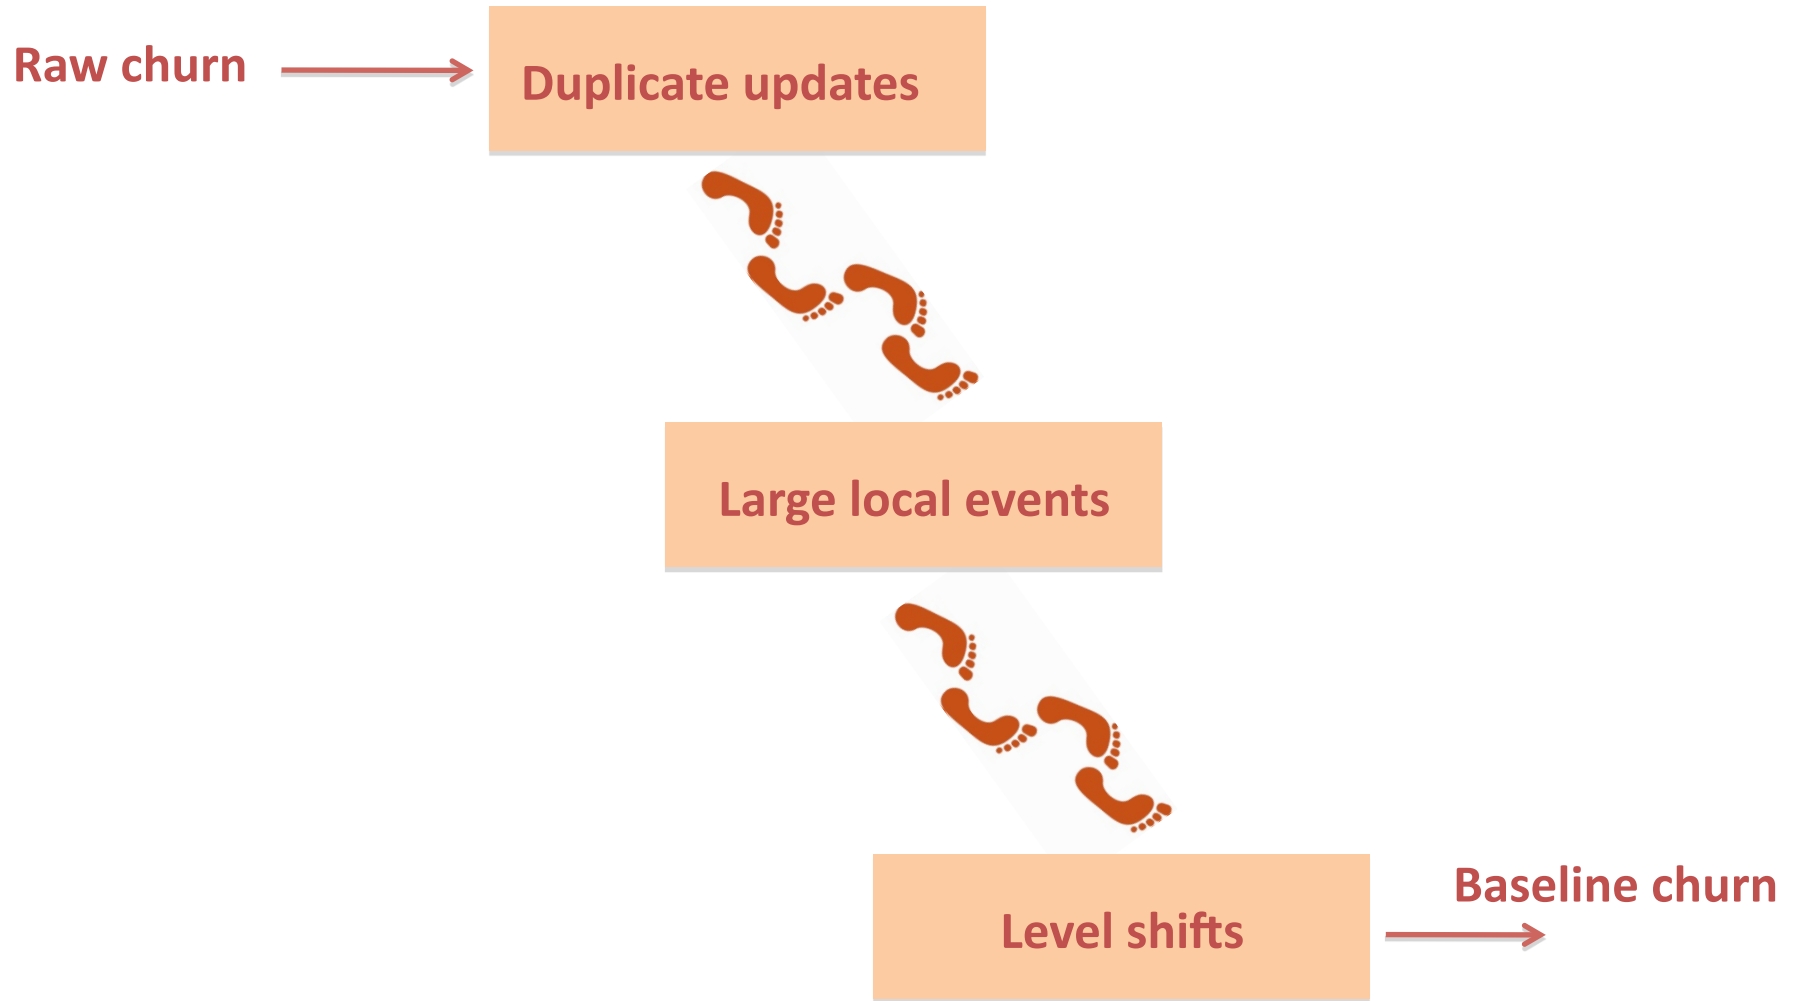
# Little or no correlation between monitors



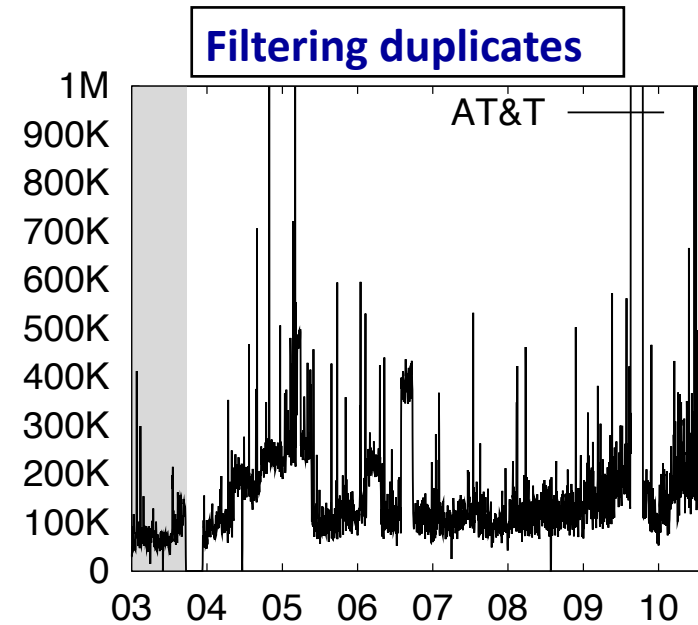

**Number of BGP updates per day (2003-2010)**



# Understanding BGP churn



# Duplicate updates account for about 40% of churn!



**Duplicates %**

**What causes**

**Stateless implem**

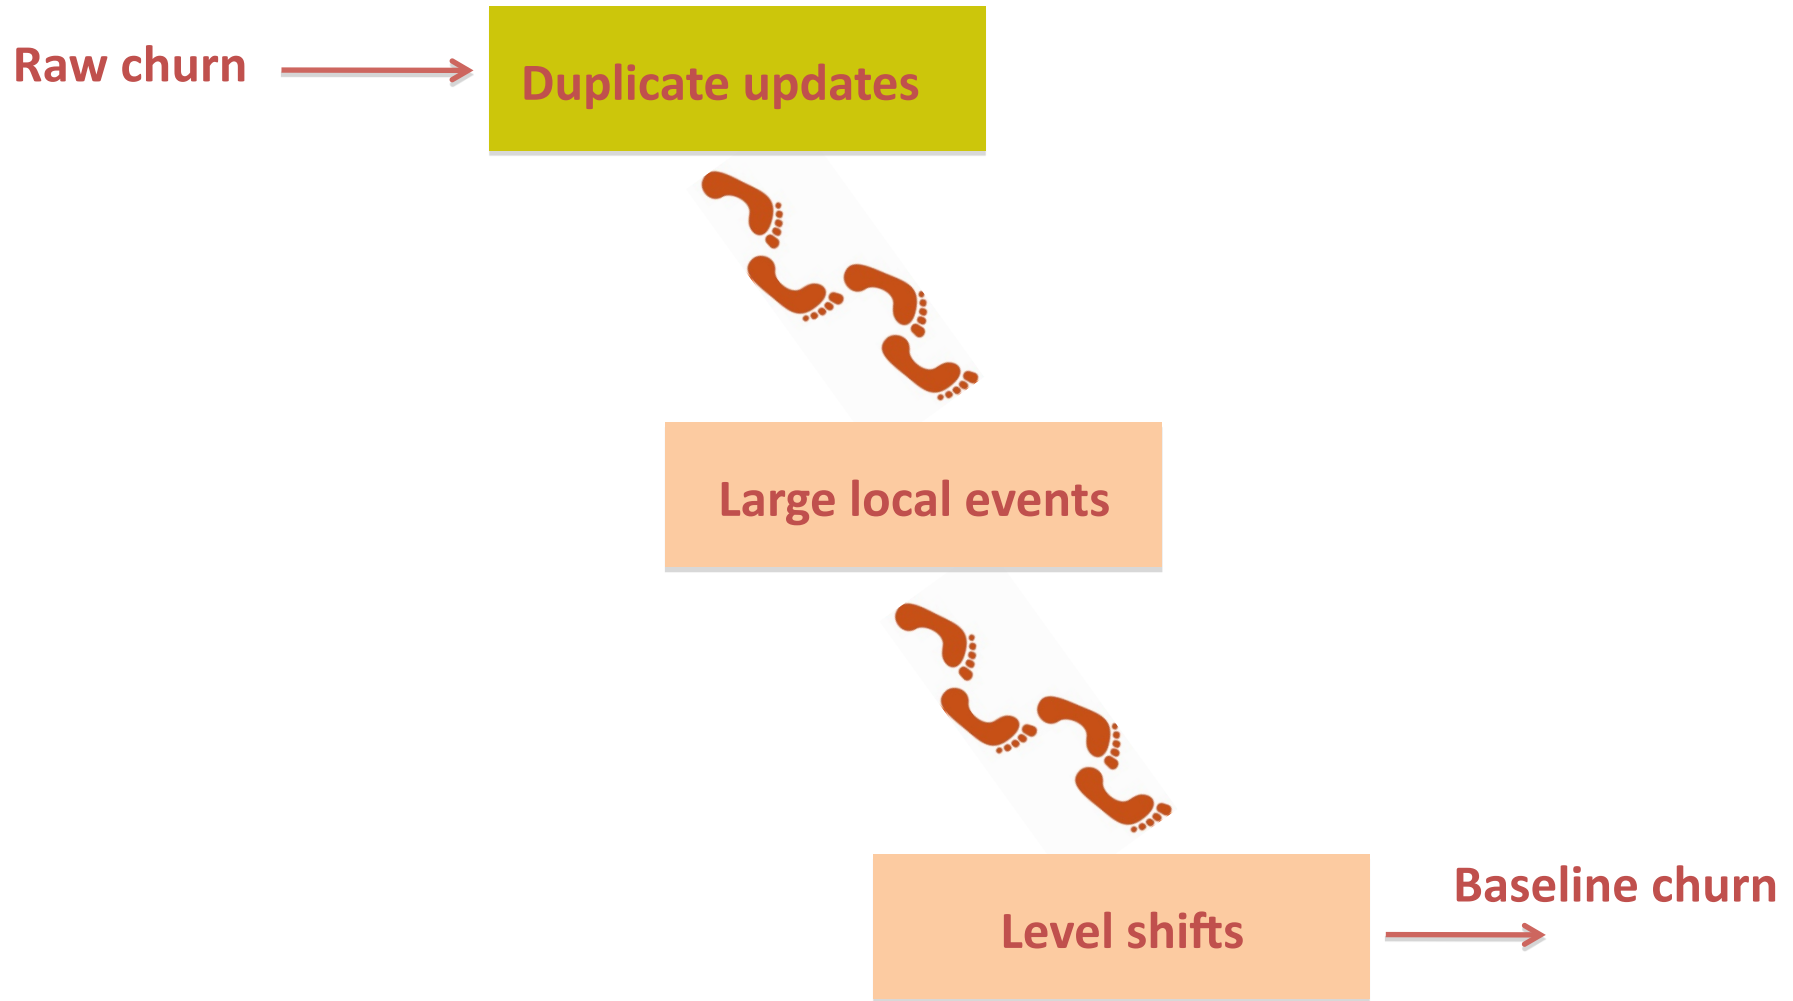
**E-BGP I-BGP int**



	AT&T	Level-3	FT	Sprint
2003	23.7	40.7	33.0	7.2
2004	47.6	53.8	45.2	23.5
2005	34.8	61.7	52.0	41.1
2006	31.8	46.1	43.5	17.7
2007	52.6	42.3	50.0	14.3
2008	59.6	32.4	43.2	12.9
2009	39.7	22.6	-	-
2010	20.2	34.9	-	-

**Could they be filtered out?**

**Sure, but there is a cost wrt processing and state**

# Understanding BGP churn



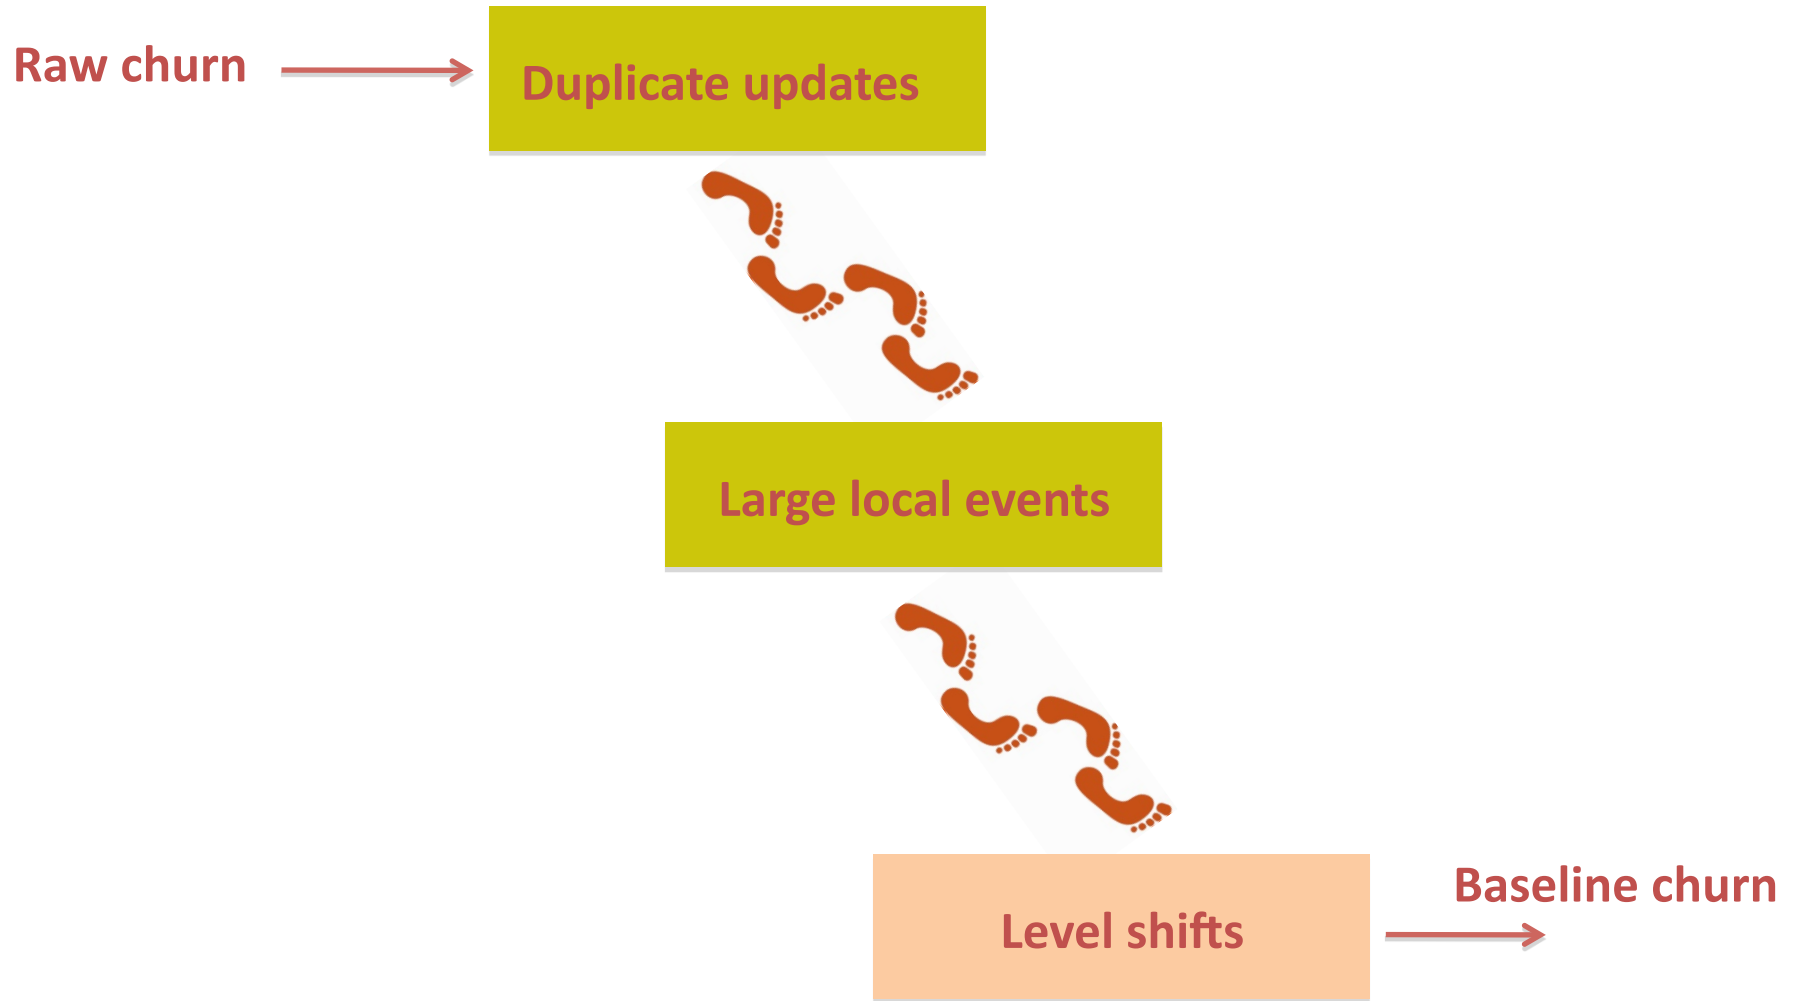


## Large events: affect more than 2000 prefixes

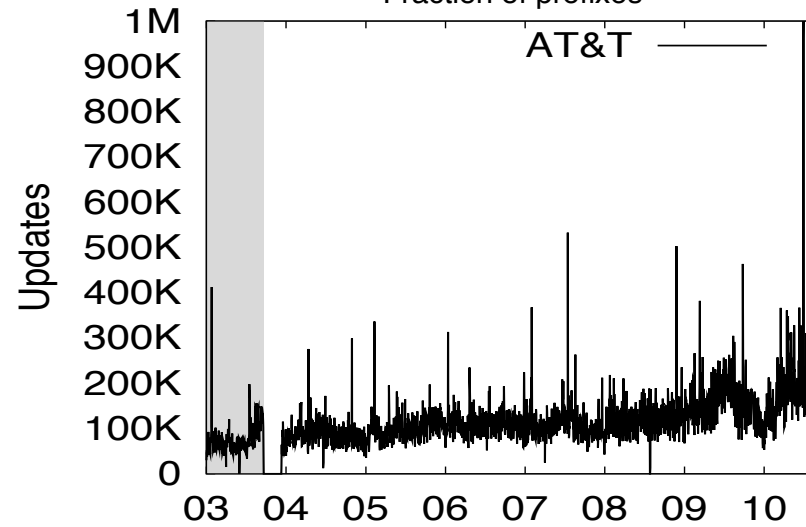
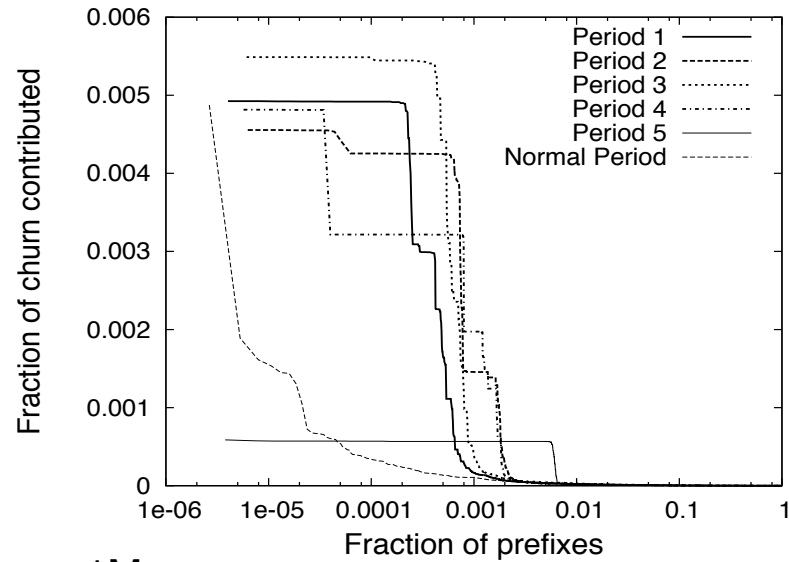
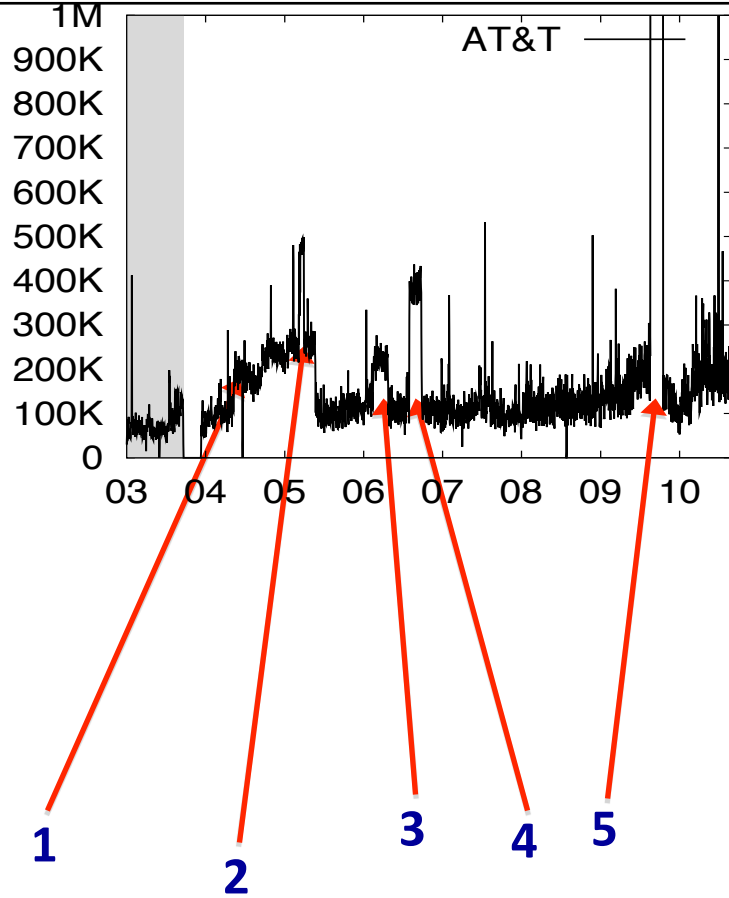
- Almost always caused by events in or close to the monitored AS
- Major causes are MED oscillations, use of Communities for TE and failures in/close to the monitored AS
- Different monitors experience large events with different magnitude and frequency

Most remaining large spikes in the duplicate-free churn are related to large events

# Understanding BGP churn



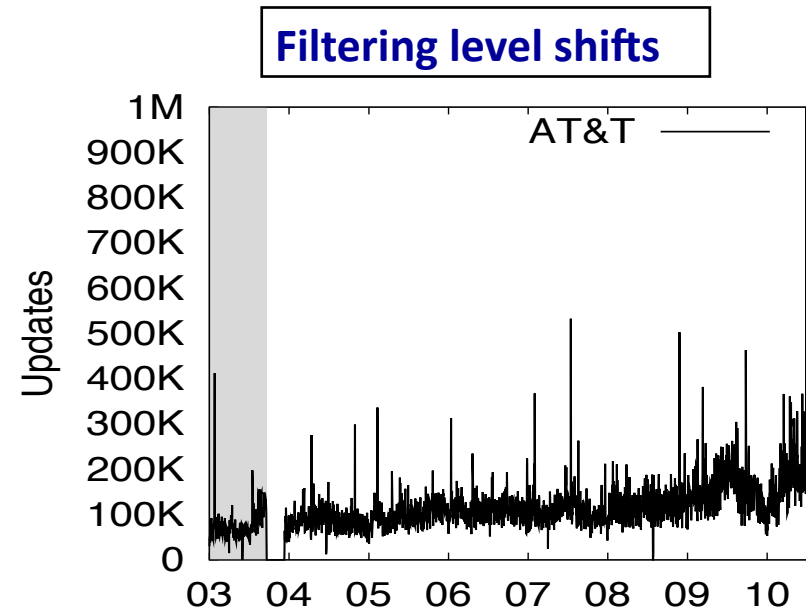
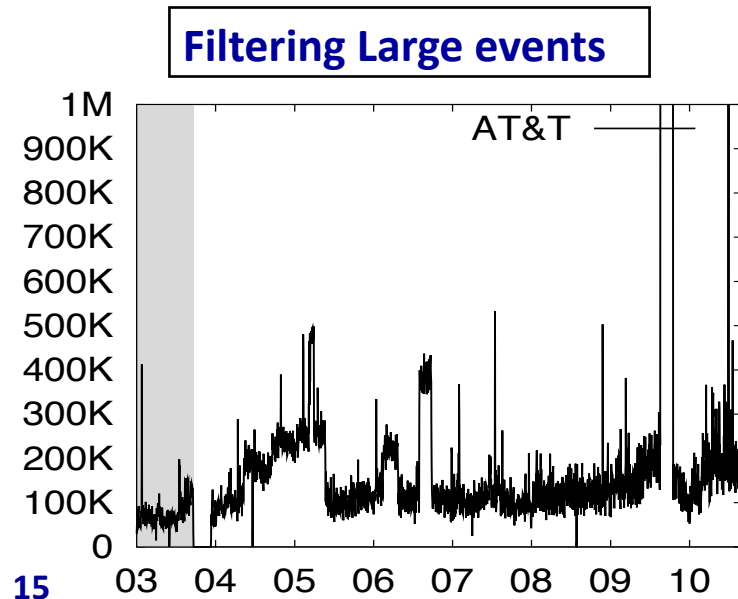

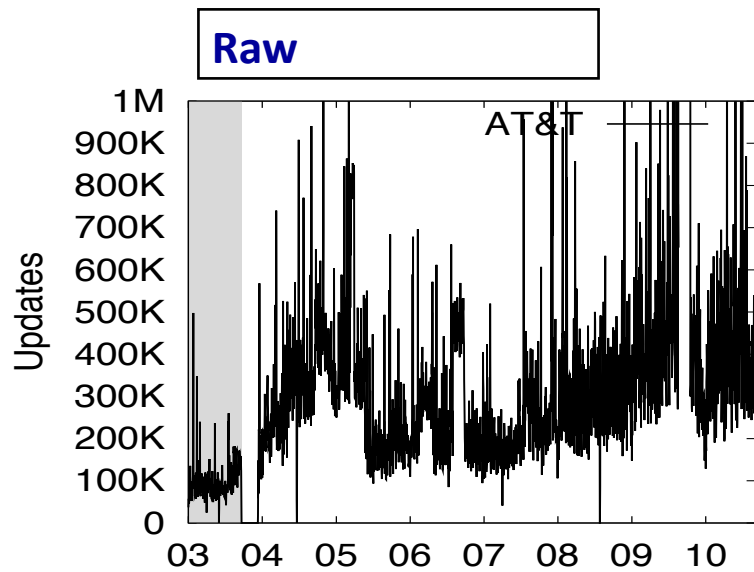
## Duplicates and large events free time series



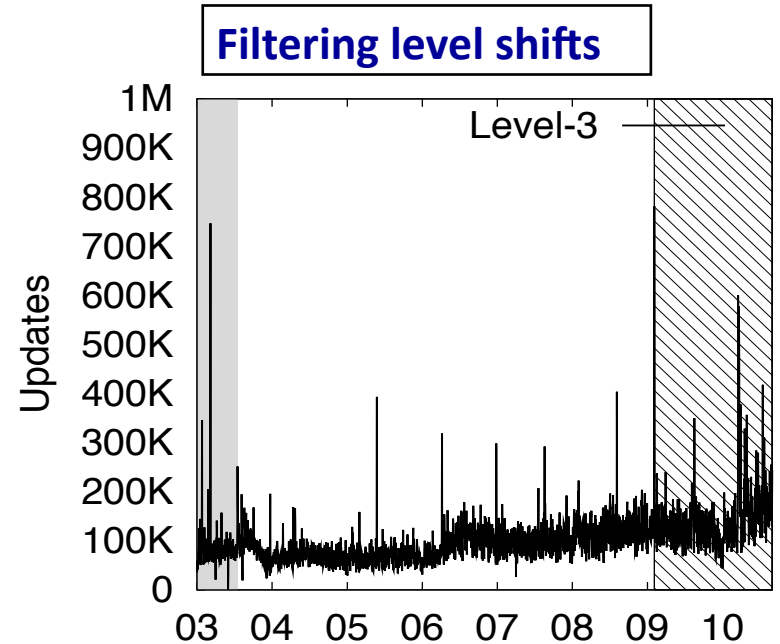
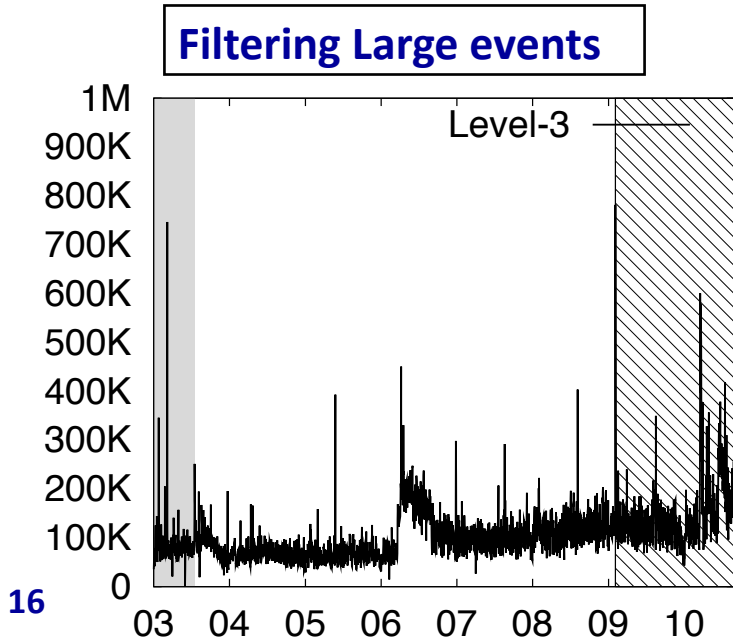

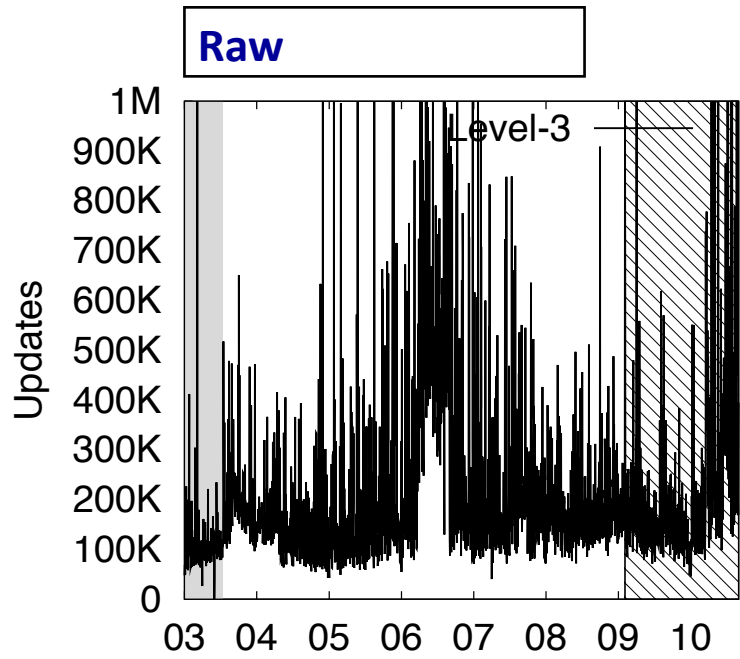
Filtering level shifts

Level shifts are usually caused by specific failures or misconfigurations in or near the monitored AS

# Incidents that are local to the monitored network cause most of the large spikes and level shifts.

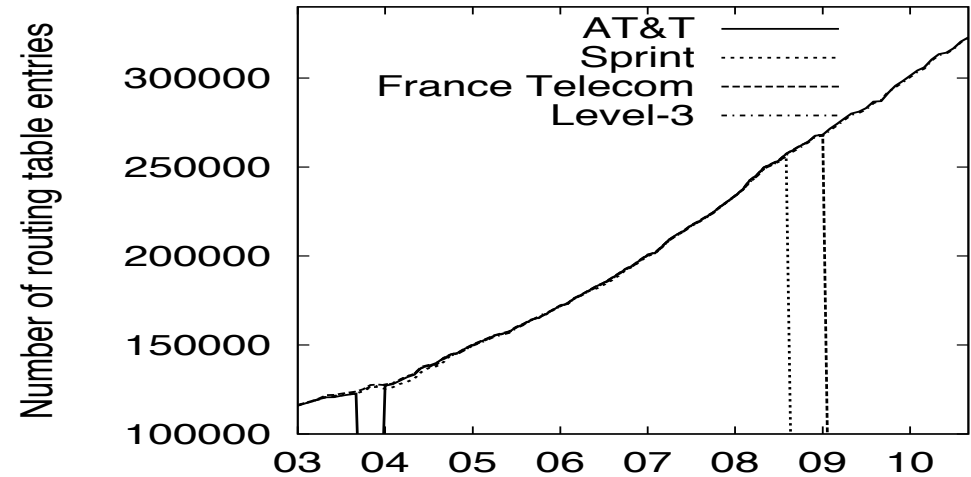
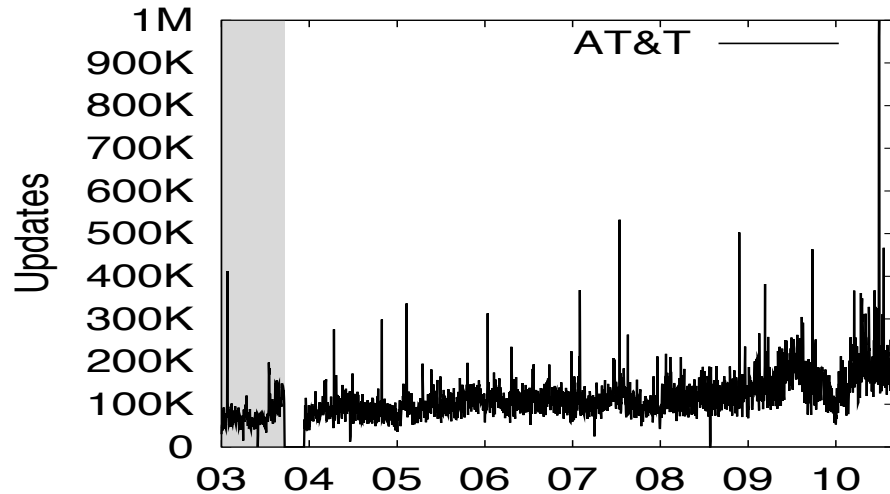


# Level-3 monitor





# Baseline Churn increases at a slower rate than the number of network destinations in the Internet



The baseline churn has increased at most by **103.7%** depending on monitor

The number of network destinations has increased approximately **168%**


The most severe churn bursts are not caused by global effects

The increase in the baseline churn is relatively slow, and will not pose a serious scalability problem in the near future

# Summary

Baseline churn increases more slowly than the number of network destinations in the Internet.

The most severe churn bursts are not caused by global effects



**Decomposing Churn**

**Ongoing work:**

**Why does churn grow so slowly?**

# Questions?

**Ahmed Elmokashfi**

**<http://simula.no/people/ahmed>**

**ahmed@simula.no**

[ **simula** . research laboratory ]

