BGP Churn Evolution:
A perspective from the core

Ahmed Elmokashfi, **Amund Kvalbein**, Constantine Dovrolis

[simula . research laboratory]

Fast growth in the Internet over the past decade

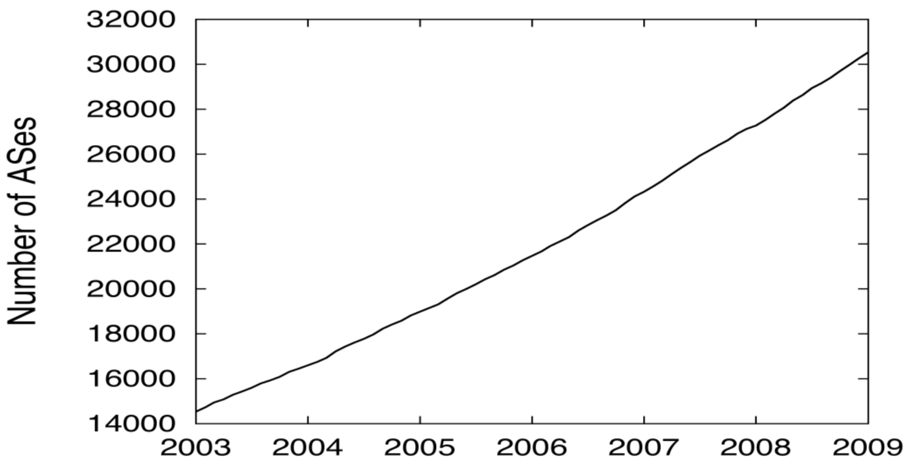

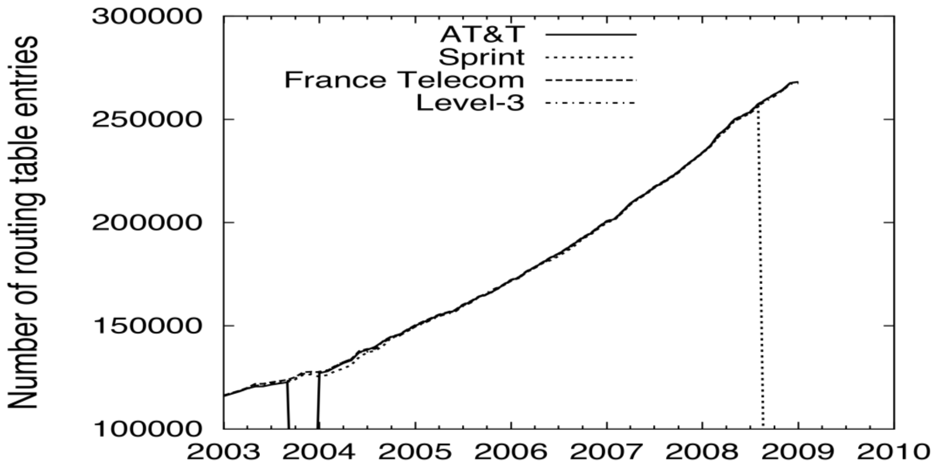

The number of ASes has increased approximately by a factor of 2.

The number of network destinations has increased approximately by a factor of 2.

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

"There is a need to devise a scalable routing and addressing system"

Why?

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

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

2003-2009

4 core networks

RouteViews project

Contrast to yesterday’s presentation by Geoff Huston

Geoff’s work	Our work
Look at both RIB size and churn	Look only at churn
1 monitoring point	4 monitoring points
3 years of data	6 years of data
Perspective from mid-tier AS	Perspective from the core
Full control of monitoring setup	Less control of monitoring setup
Describe trend in churn	Decompose and filter before describing trend

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…

Example

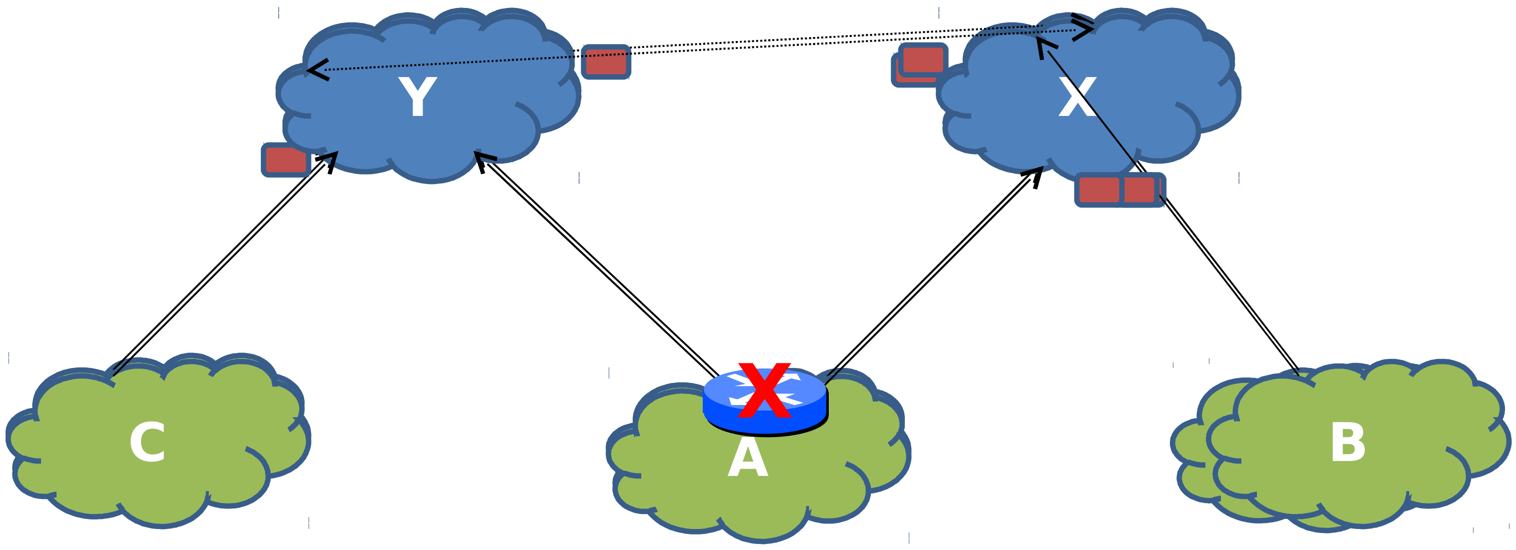

BGP churn timeseries are bursty, and it is difficult to identify a trend in them.

The timeseries is dominated by large frequent spikes.

There are several level shifts

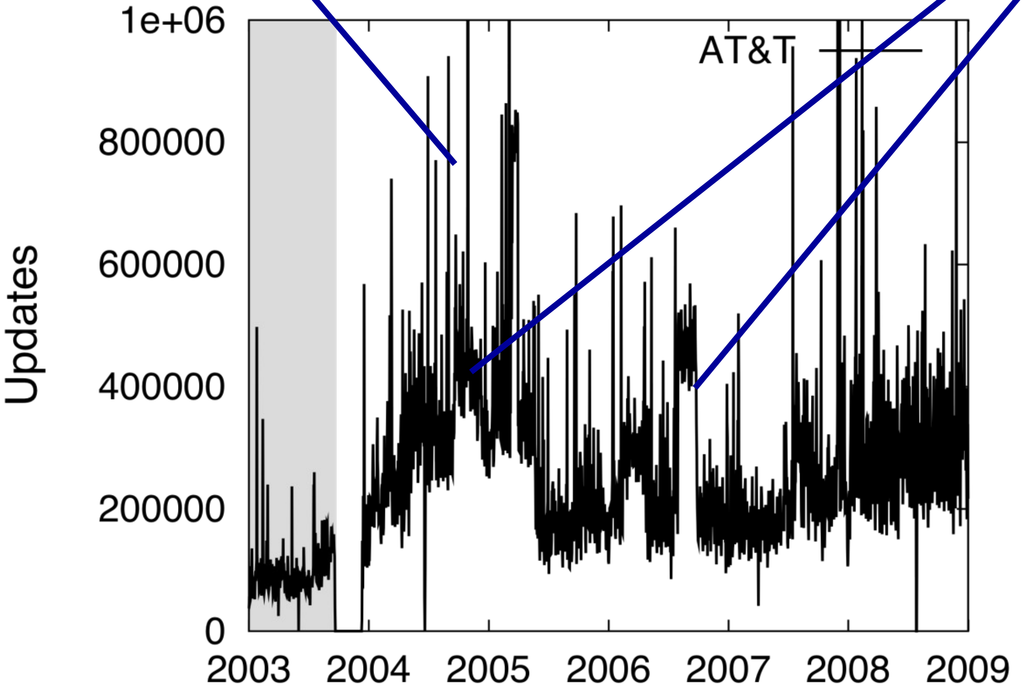

Little or no correlation between monitors

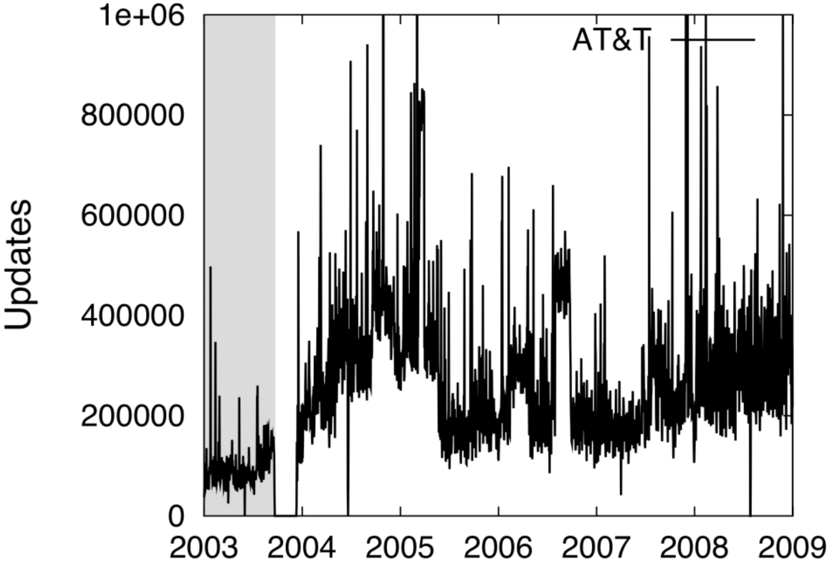

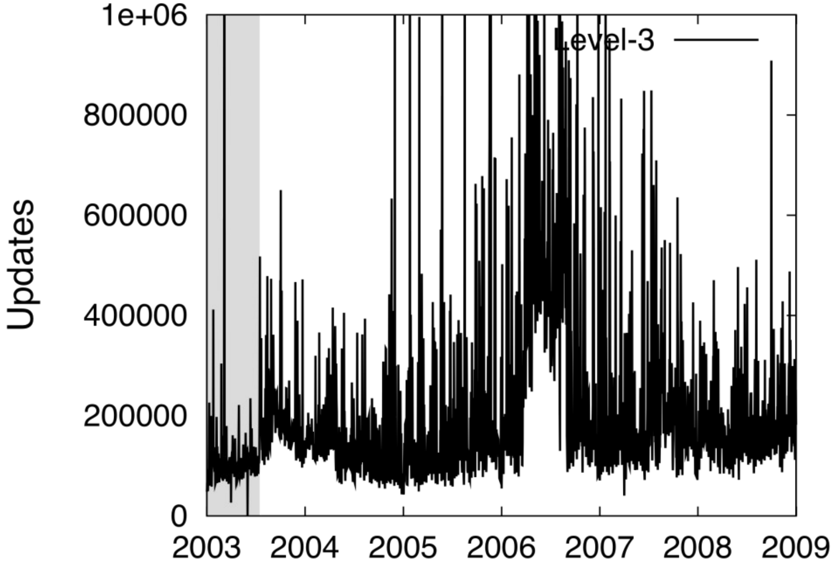

Understanding BGP churn

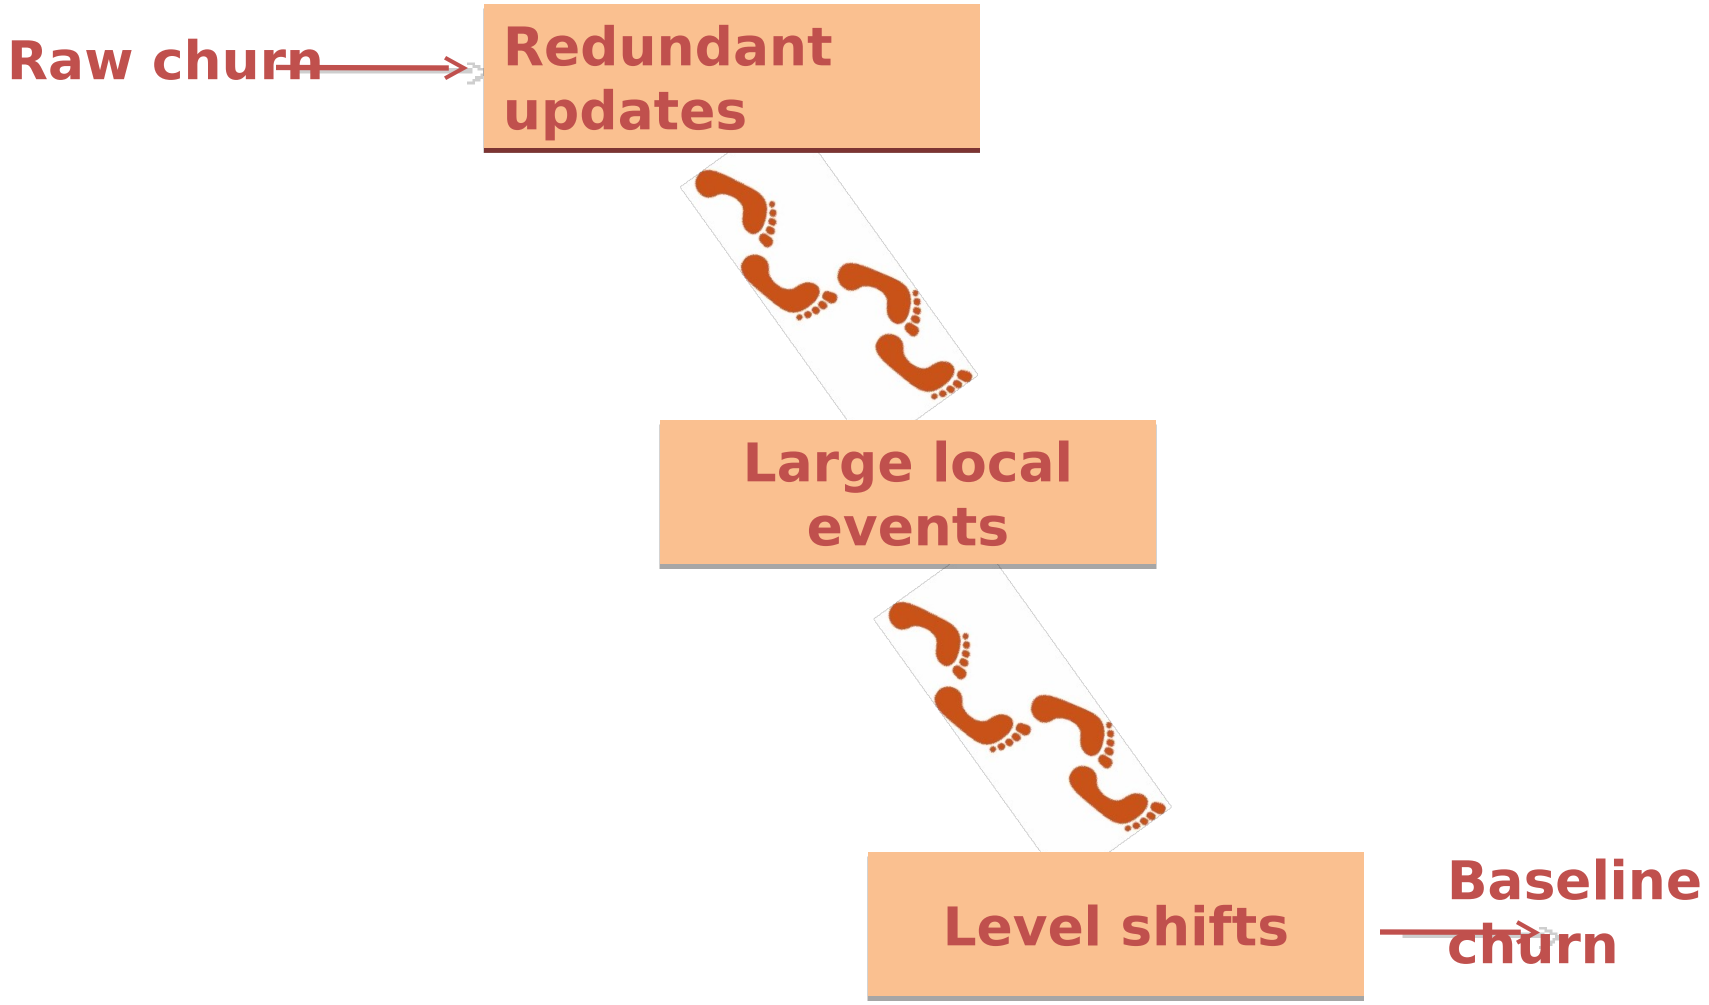

Duplicate updates account for about 40% of churn!

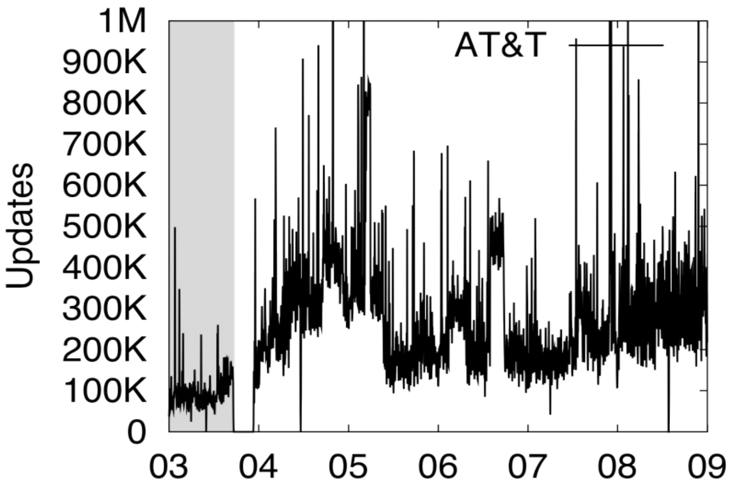

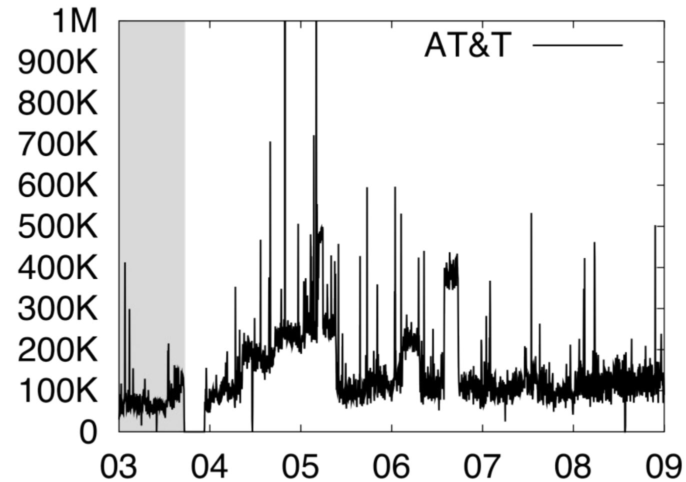

	AT&T	Level-3	FT	Sprint
2003	23.7	40.7	33.0	7.2
2004	47.6	53.8	45.2	23.5
2005	34.8	61.7	52.0	41.1
2006	31.8	46.1	43.5	17.7
2007	52.6	42.3	50.0	14.5
2008	59.6	32.4	43.2	12.9

What causes du

Stateless implemen

E-BGP I-BGP intera

Could they be filtered out?

Sure, but there is a cost wrt processing and state

Understanding BGP churn

Raw churn

Redundant updates

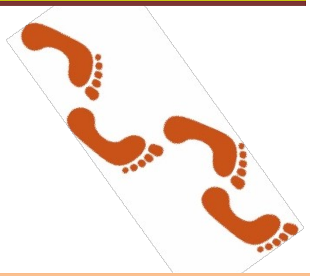

Large local events

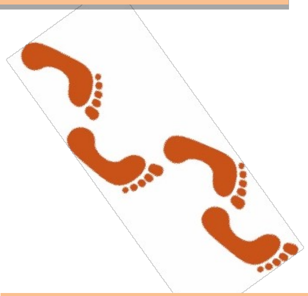

Level shifts

Baseline churn

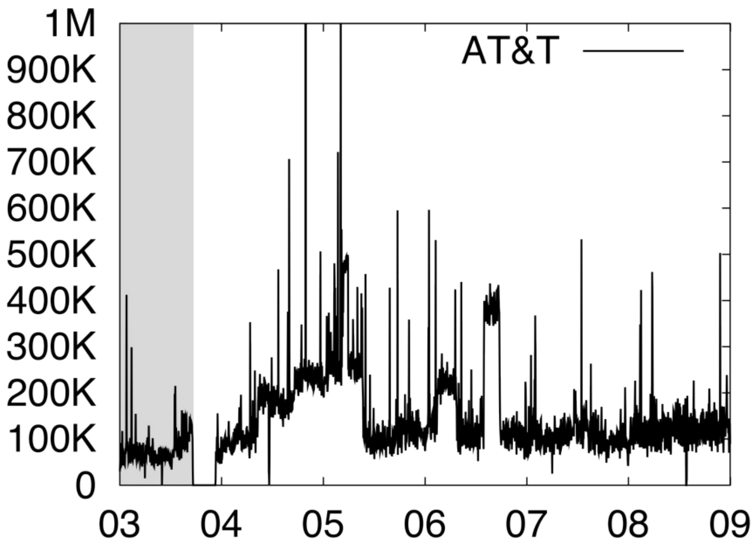

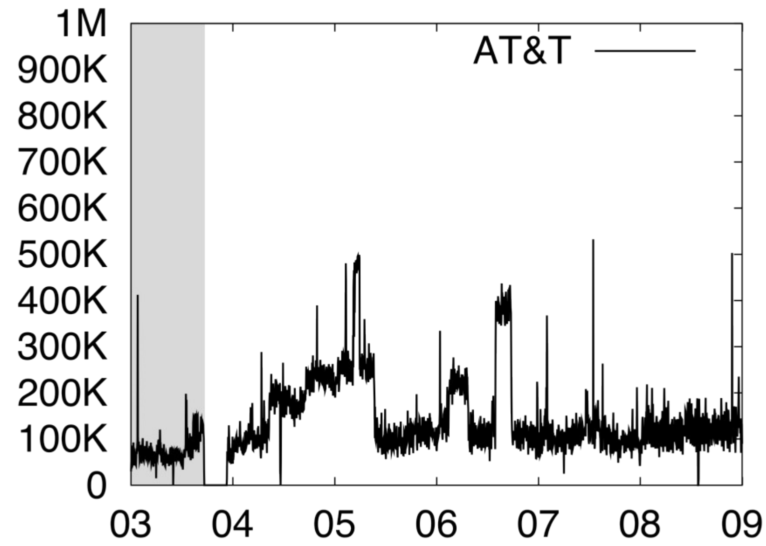

Large events: affect more than 2000 prefixes

- **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

Raw churn

Redundant updates

Large local events

Level shifts

Baseline churn

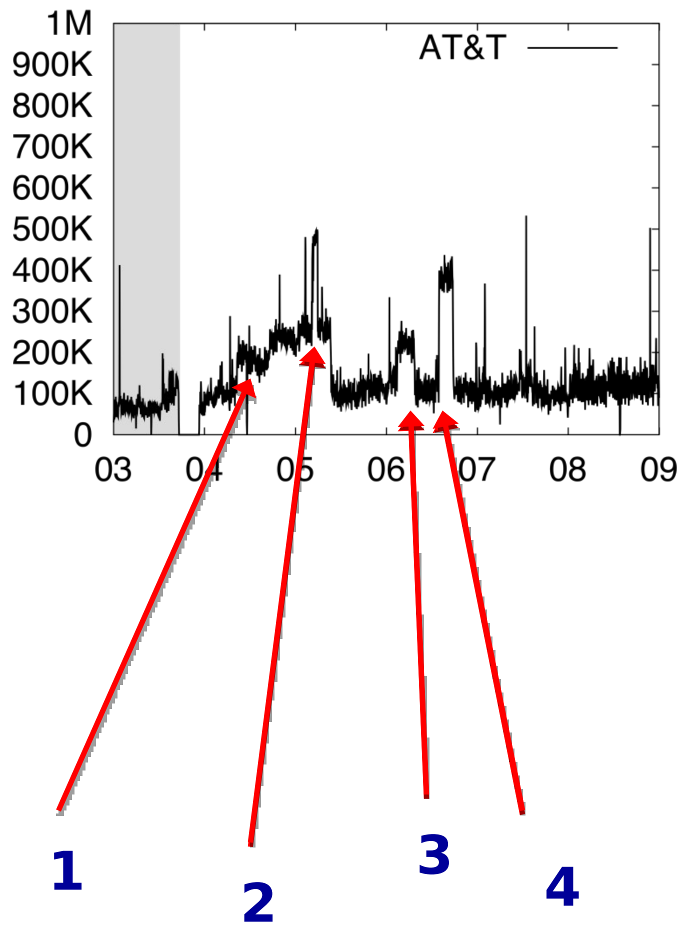

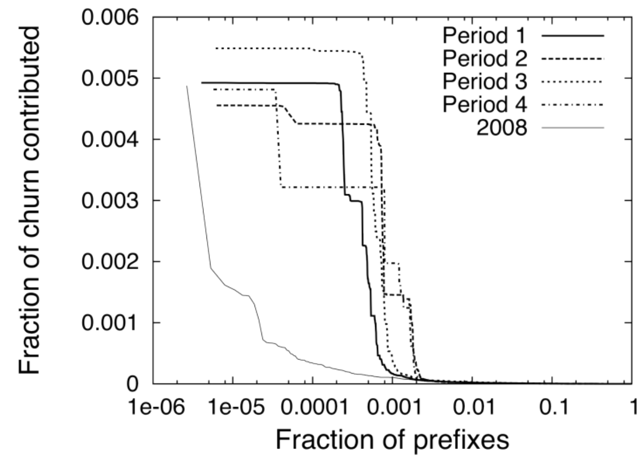

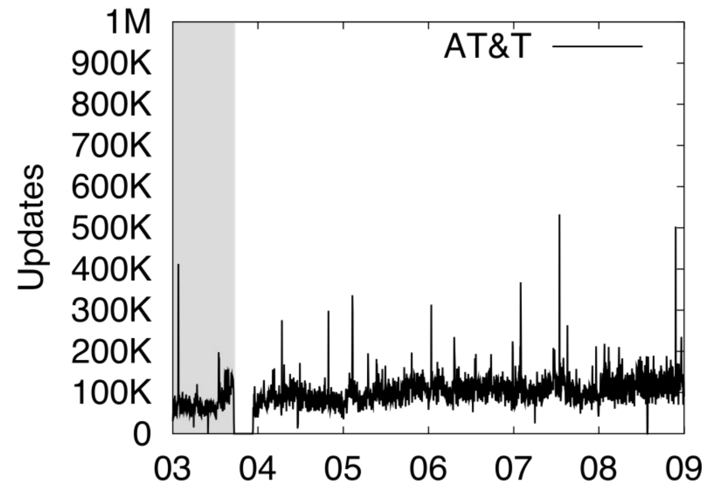

Level shifts are usually caused by specific failures or misconfigurations in or near the monitored AS

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

Raw

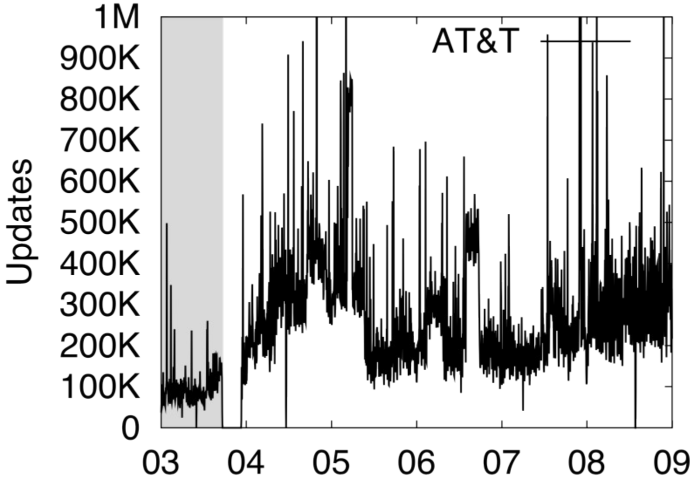

Filtering duplicates

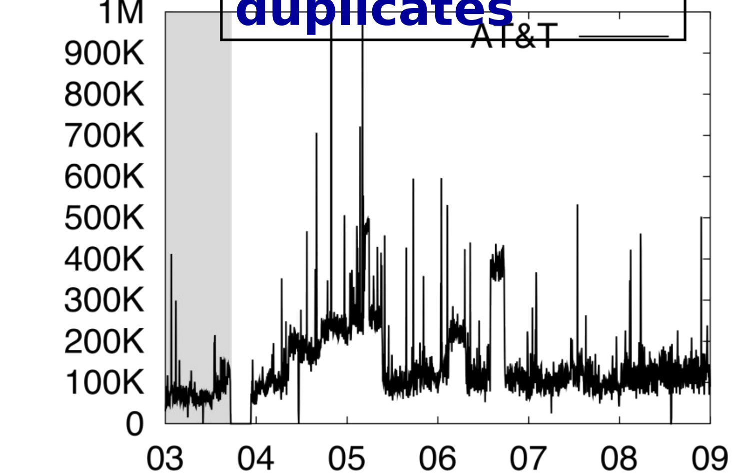

Filtering Large events

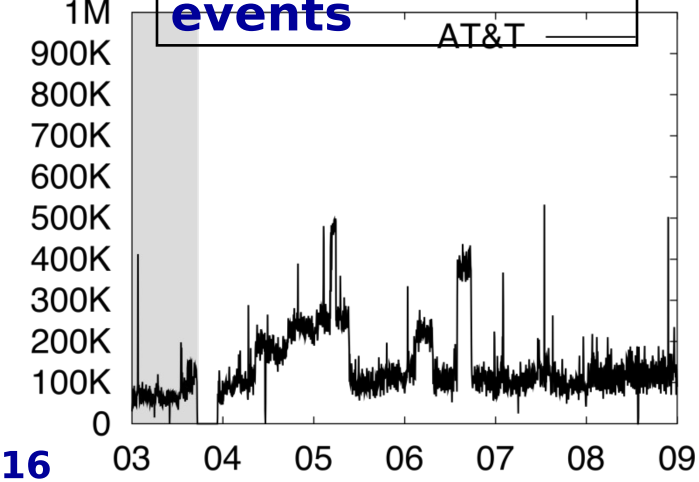

Filtering level shifts

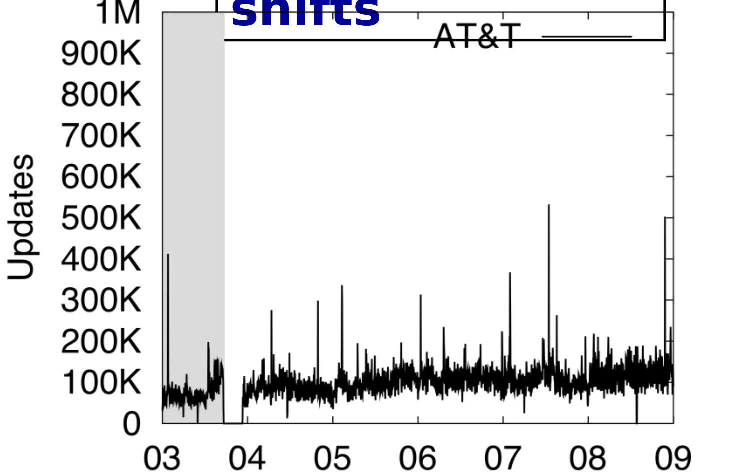

Churn increases at a slower rate than the number of network destinations in the Internet

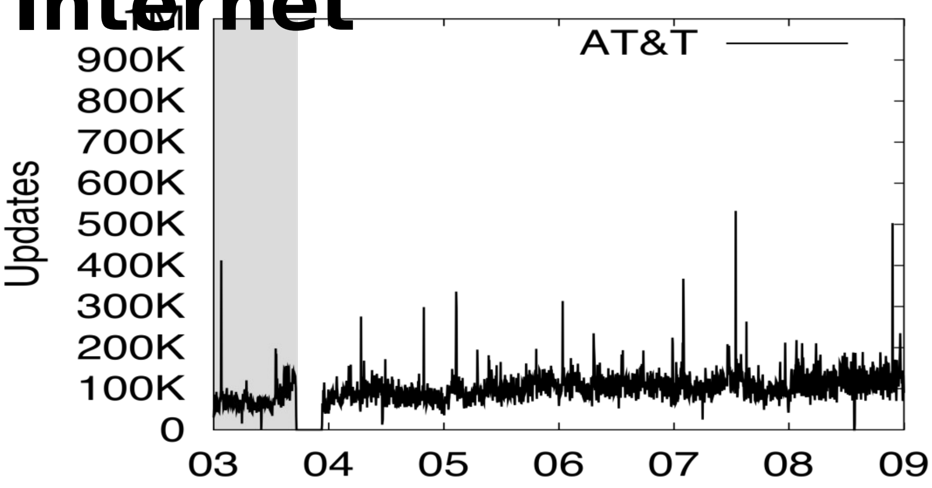

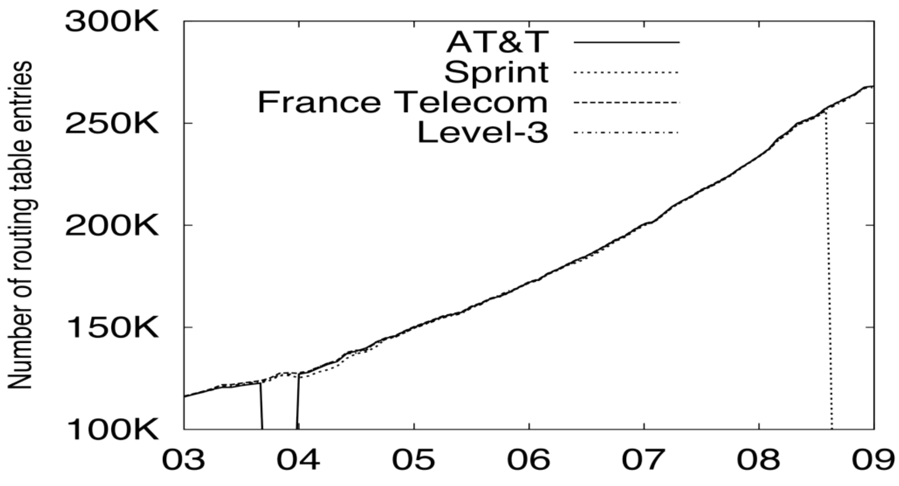

The baseline churn has increased by ~ 20-80% over the past six years depending on monitor

The number of network destinations has increased approximately by a factor of 2

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

Ongoing work:

Why does churn grow so slowly?

Is it because of topological densification?

We know that [Dhamdhere et al 2008]

The Internet is getting denser as it grows

Increasing multihoming gives more paths to each destination

The average path length is constant

How would this influence churn?

+ More paths to explore when a prefix fails
- Some events will not be globally visible - alternative paths are preferred

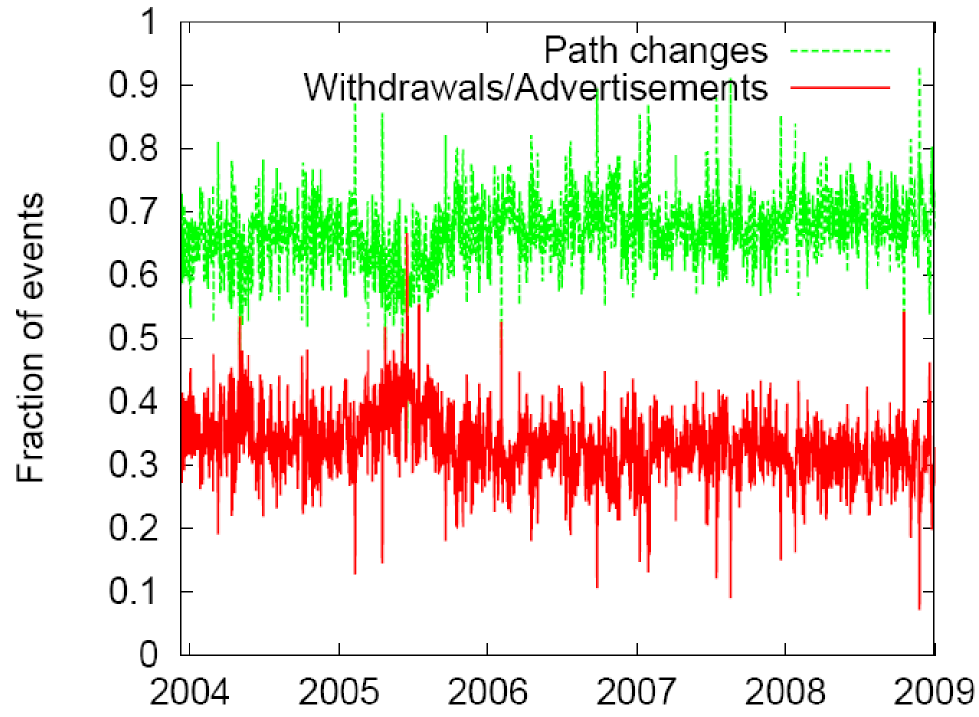

Densification increases path exploration

The use of rate limiting (MRAI timers or OutDelay) masks this

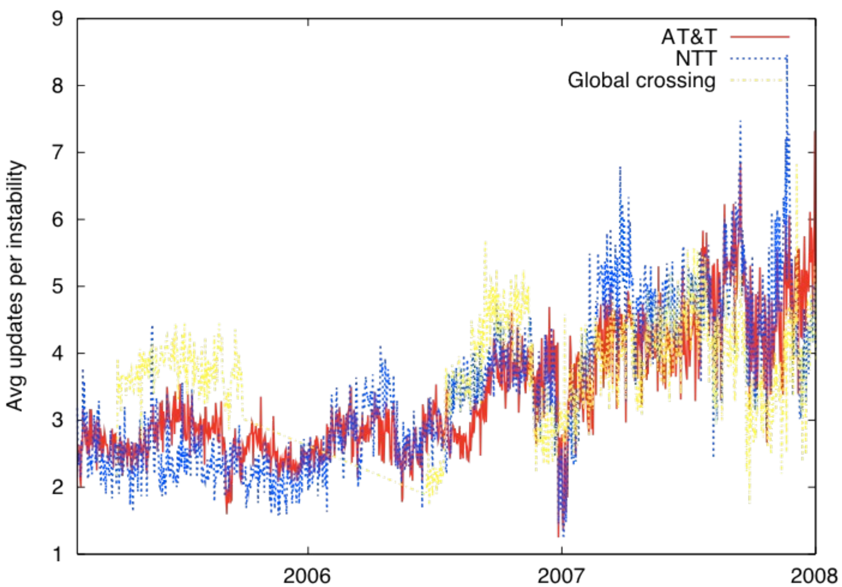

Beacon prefix withdrawal:
nitor sessions without rate limiting

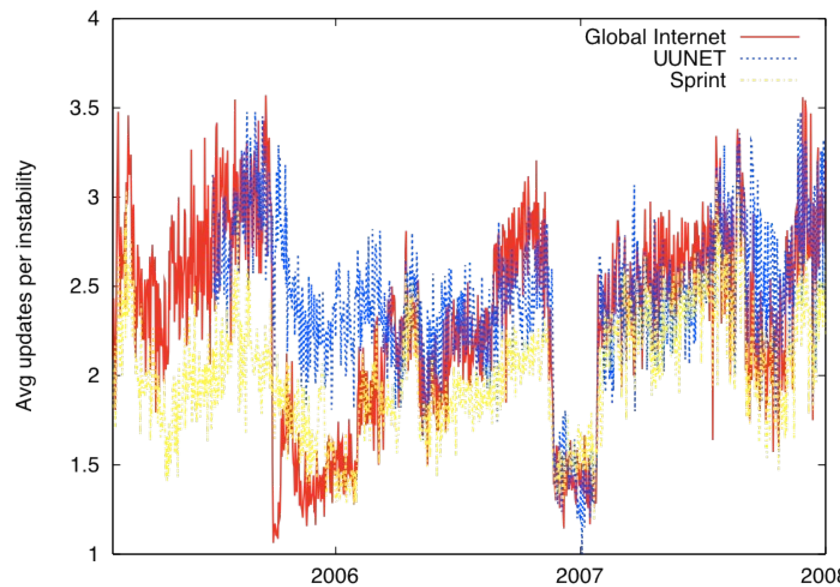

Beacon prefix withdrawal:
Monitor sessions with rate limiti

Densification limits the visibility of routing changes

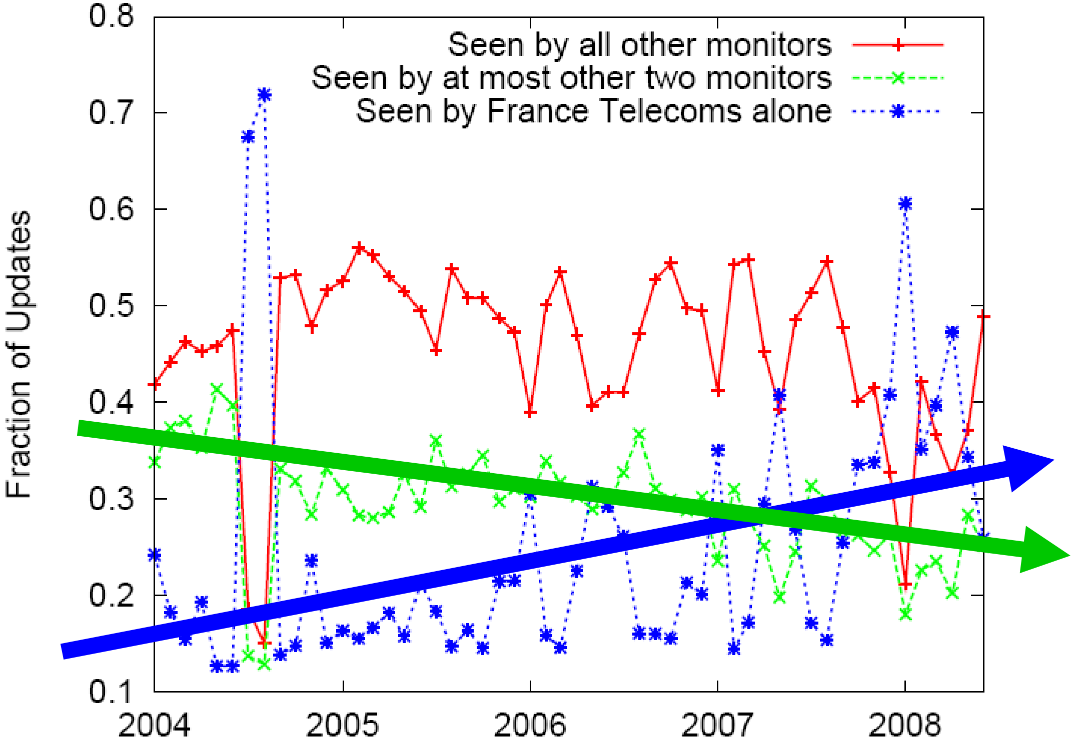

There seems to be a trend that there is less cross-correlation between monitors - events are seen by fewer monitors

Thank you – questions?

http://simula.no/people/amundk

http://simula.no/research/nd/publications/Simula.nd.435