

Generating realistic BGP traffic

Anja Feldmann, Olaf Maennel
Saarland University, Saarbrücken, Germany

Motivation

- Internet evolution
 - New features, components, software versions
- Correctness, effectiveness???
- Test lab
- Real network
- Test lab limitations:
 - Small number of components
 - Realistic workloads
 - Traffic, e.g., temporal and address variability
 - Routing

Internet in a lab

Routing:


Tool: "RTG"
this talk

Benutzerverhalten:

Tool: "Surge"
"Surge++"
BU
in work

Topologie:

Tool: "Dummynet"
Pisa



BGP workload model

- Identify structure of BGP traffic
- Basis for BGP workload model
- Tool, RTG, realization of workload model
- Characterize structure of BGP traffic
- Verification of tool

Benefit of RTG

- Explore BGP in a test-lab
 - Basic BGP implementation regression tests
 - Test BGP implementation features
 - Test settings of BGP parameters
 - Test BGP's scalability
 - Test interactions routing vs. Forwarding
 - Experiment with changes in BGP workload
 - Understand reasons for BGP's limited scalability

Outline

- BGP
- BGP workload model: characteristics
- RTG: workload realization
- BGP traffic characterization
- Summary

BGP

- Updates

Timestamp				Updated Prefix					
1011363829	A	195.66.224.112	3549	80.96.15.0/24	3549	3300	702	8708	
1011387198	W	195.66.224.112	3549	80.96.15.0/24					
1011387339	A	195.66.224.112	3549	80.96.15.0/24	3549	701	702	8708	
1011387369	A	195.66.224.112	3549	80.96.15.0/24	3549	3300	702	8708	
1010976980	W	195.66.224.112	3549	80.96.150.0/24					
1010977007	A	195.66.224.112	3549	80.96.150.0/24	3549	209	1755	15471	

- Protocol elements limiting BGP updates

- Min-Route Advertisement Interval (30s)
- Route dampening (typically 10-30 minutes)

BGP workload model: Central notions

- Updates
 - Instability creators
 - Instability bursts
- Changes
 - BGP attributes
- Baseline
 - Routing table \leftrightarrow prefix forest
- Correlation
 - AS-Path

BGP updates

- Routing instabilities
 - Session establishment/teardown/reset
 - Parameter change
 - Addition/deletion of prefixes
 - Prefix policy changes
- Instability creator
 - Two peering ASs
 - Prefix

BGP updates (cont.)

- Instability event
 - update bursts for set of prefixes
- Update burst:
 - Single prefix
 - Set of updates
 - Close in time

Metrics: BGP updates

- Updates
 - Interarrival time, attribute change
- Update bursts
 - Interarrival time, duration, # of updates
- Session resets (possible)
 - Interarrival time, duration, # of prefixes

BGP routing table

- Prefix forest

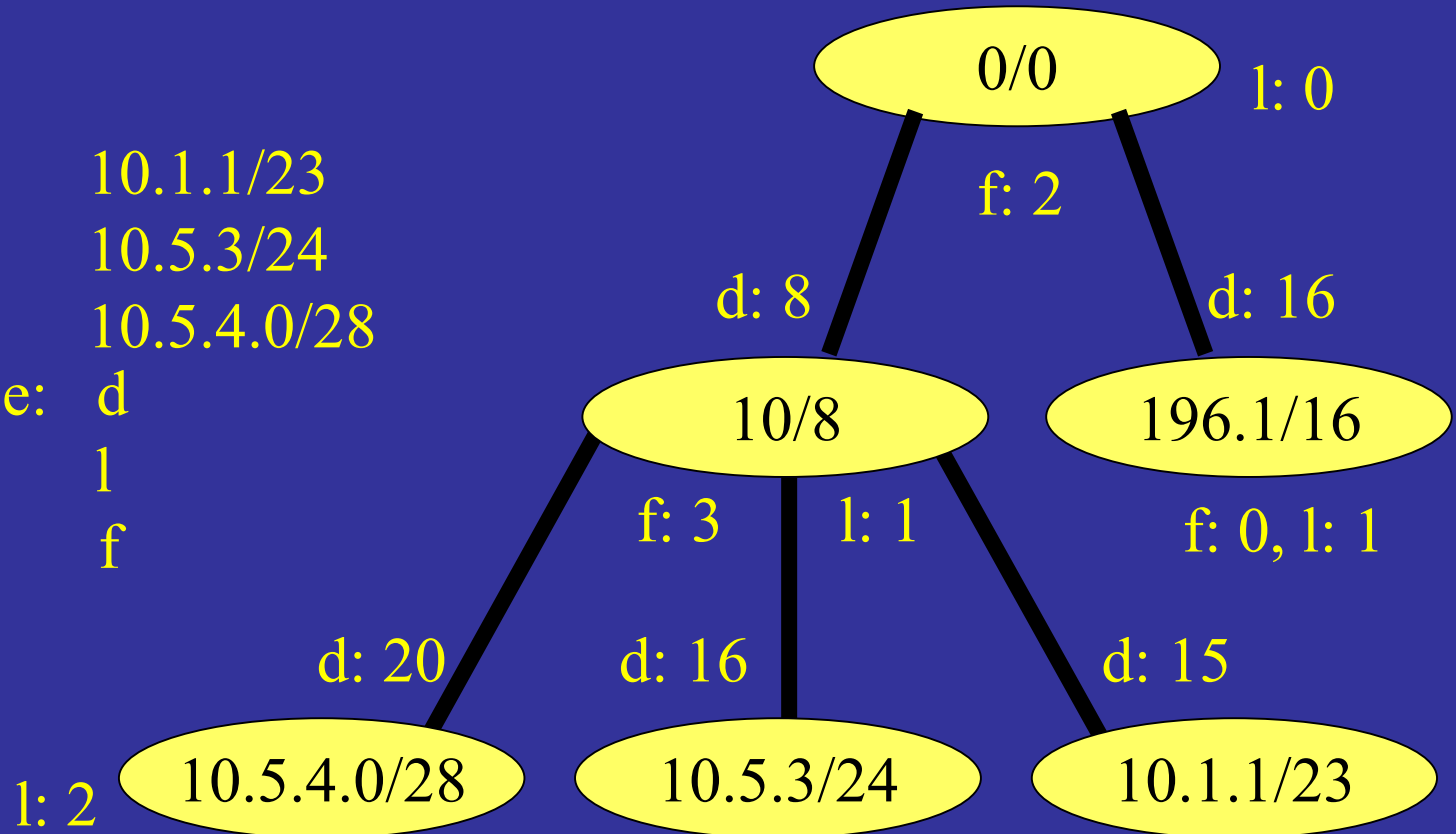
Prefix:

0/0 10.1.1/23
10/8 10.5.3/24
196.1/8 10.5.4.0/28

Distance: d

Depth: l

Fanout: f



Metrics: BGP routing table

- Routes within IP address range
- Prefix length
- Prefix depth
- Prefix fanout
- Prefix distance

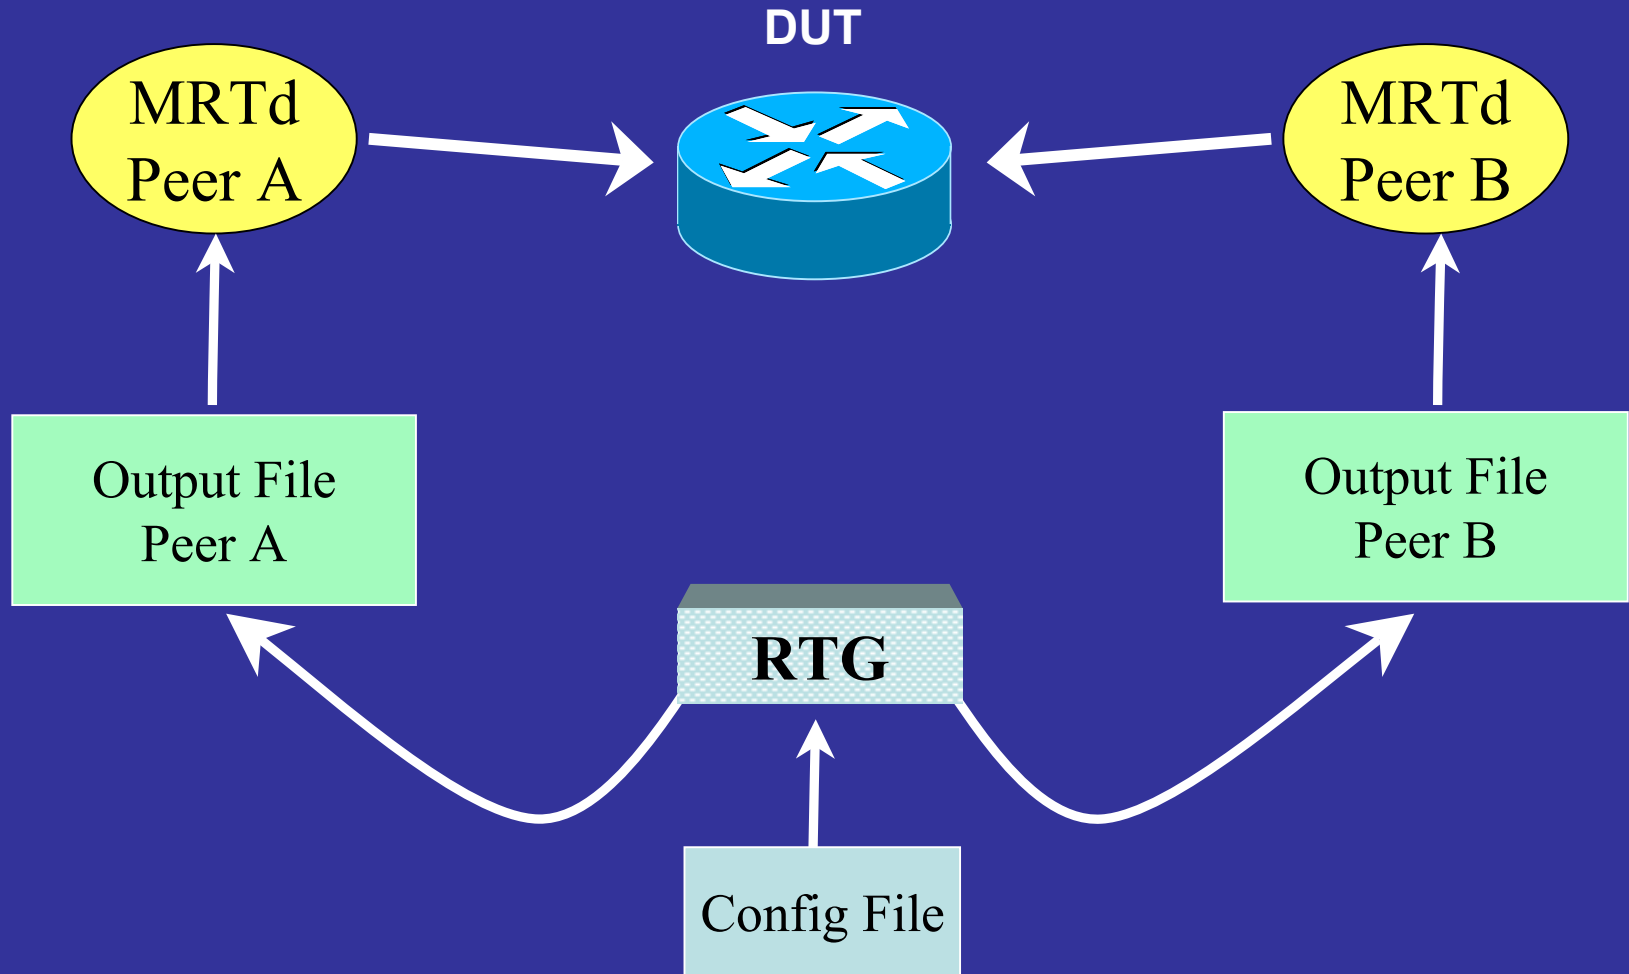
BGP AS

- Correlation: routing instabilities and prefix structure
 - AS, AS-path
- Characterization AS
 - # of originating routes
 - # of transiting routes
 - Distance of AS to peer
- Characterization AS-path
 - Length
 - Number of unique ASs
 - Position of replicated ASs
 - Number of replicated ASs

RTG: workload realization

- **Idea**
 - Generate updates off-line (stored in file): RTG
 - Feed them to system under test: e.g., MRTd
- **RTG**
 - Build routing table
 - Routing table size and characteristics
 - Generate BGP attributes
 - AS, AS-path characteristics
 - Create BGP updates
 - Event log: session reset, update burst, single update
- **Parameters**
 - Configuration files (automatic, semi-manual, manual)

Test Bed Methodology



RTG: BGP updates

```
sub updates {
  foreach $event ( event_log ) {
    ($time, $type, $as, $prefix) = split($event);
    if ( $type == session_reset ) {
      if ($as == 0) {select_as()}
      $p = select_fraction();
      foreach ($prefix = select_prefix( $as )) {
        updates( $time, update_burst, $as, $prefix);
      } elseif ( $type == update_burst ) {
        if ($prefix == 0) {select_prefix()}
        (@updates, @times) = select_burst();
        foreach (@updates, @times) {
          updates( $timeu, single, $as, $prefix)
        } else { # single update
          @a = select_attr_change( $as, $prefix);
          print FILE $time, $prefix, @a;
        }
      }
    }
  }
}
```

BGP traffic characterization

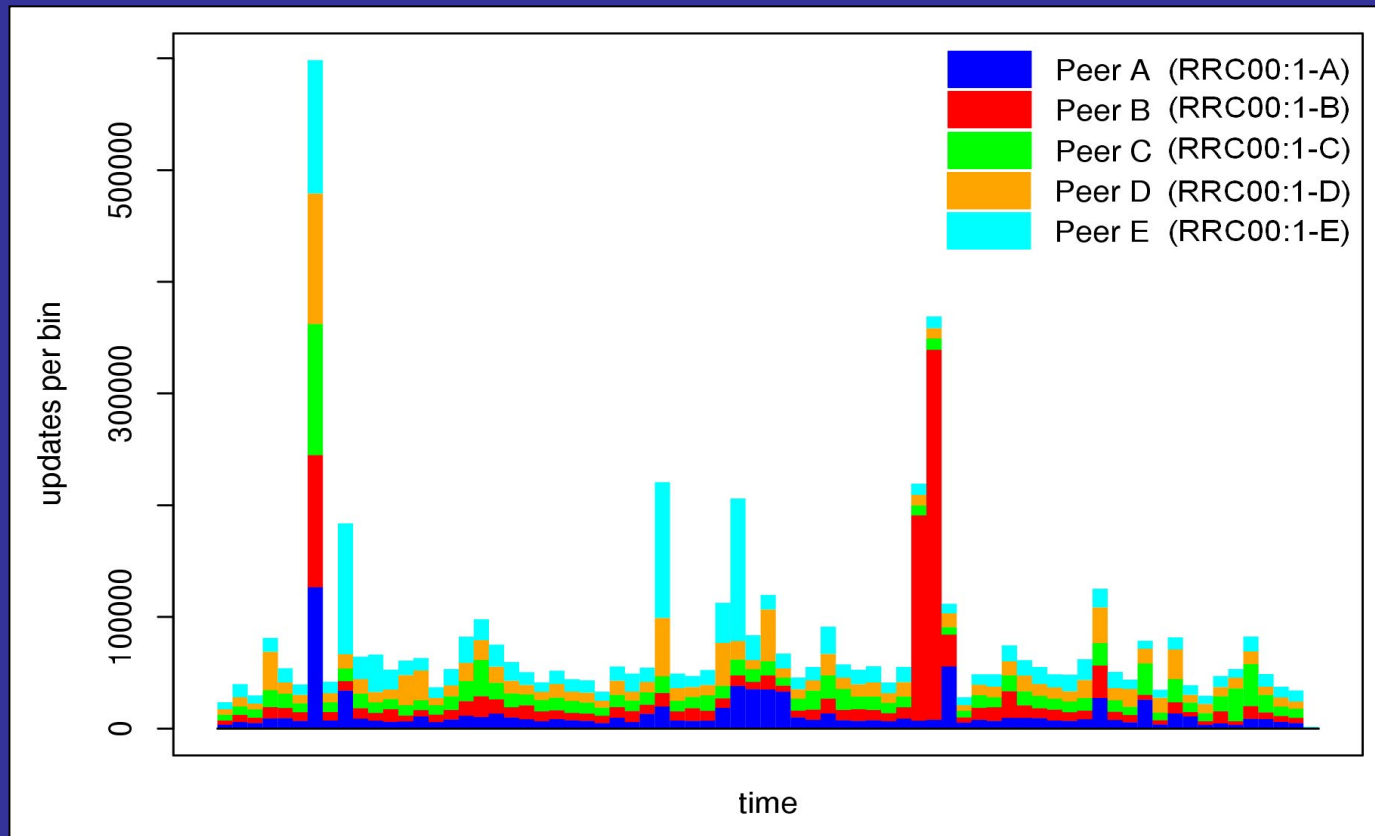
- **Tables**
 - Prefix forest, fanout, nesting, ...
- **Attributes**
 - AS-Path, Community, ...
- **Updates (TOOL: "Character")**
 - Updates
 - Update bursts
 - Session resets

TOOL: "Character"



- Data munching
 - automatic processing of raw data
 - providing an intermediate level
- Characterizing BGP updates
 - Processing of updates in the context of
 - Related (same prefix)
 - Surrounding (near in time)
 - Identification of update events
 - Type of changes, flapping, update burst, session resets, ...

Example data set

- RIPE's RRC00:
 - Jan 14, 2002 01:00 - Jan 20, 2002 01:10



What has changed



Sets of updates for a prefix with same attributes

1.

new change

2.

duplicate

3.

flapping

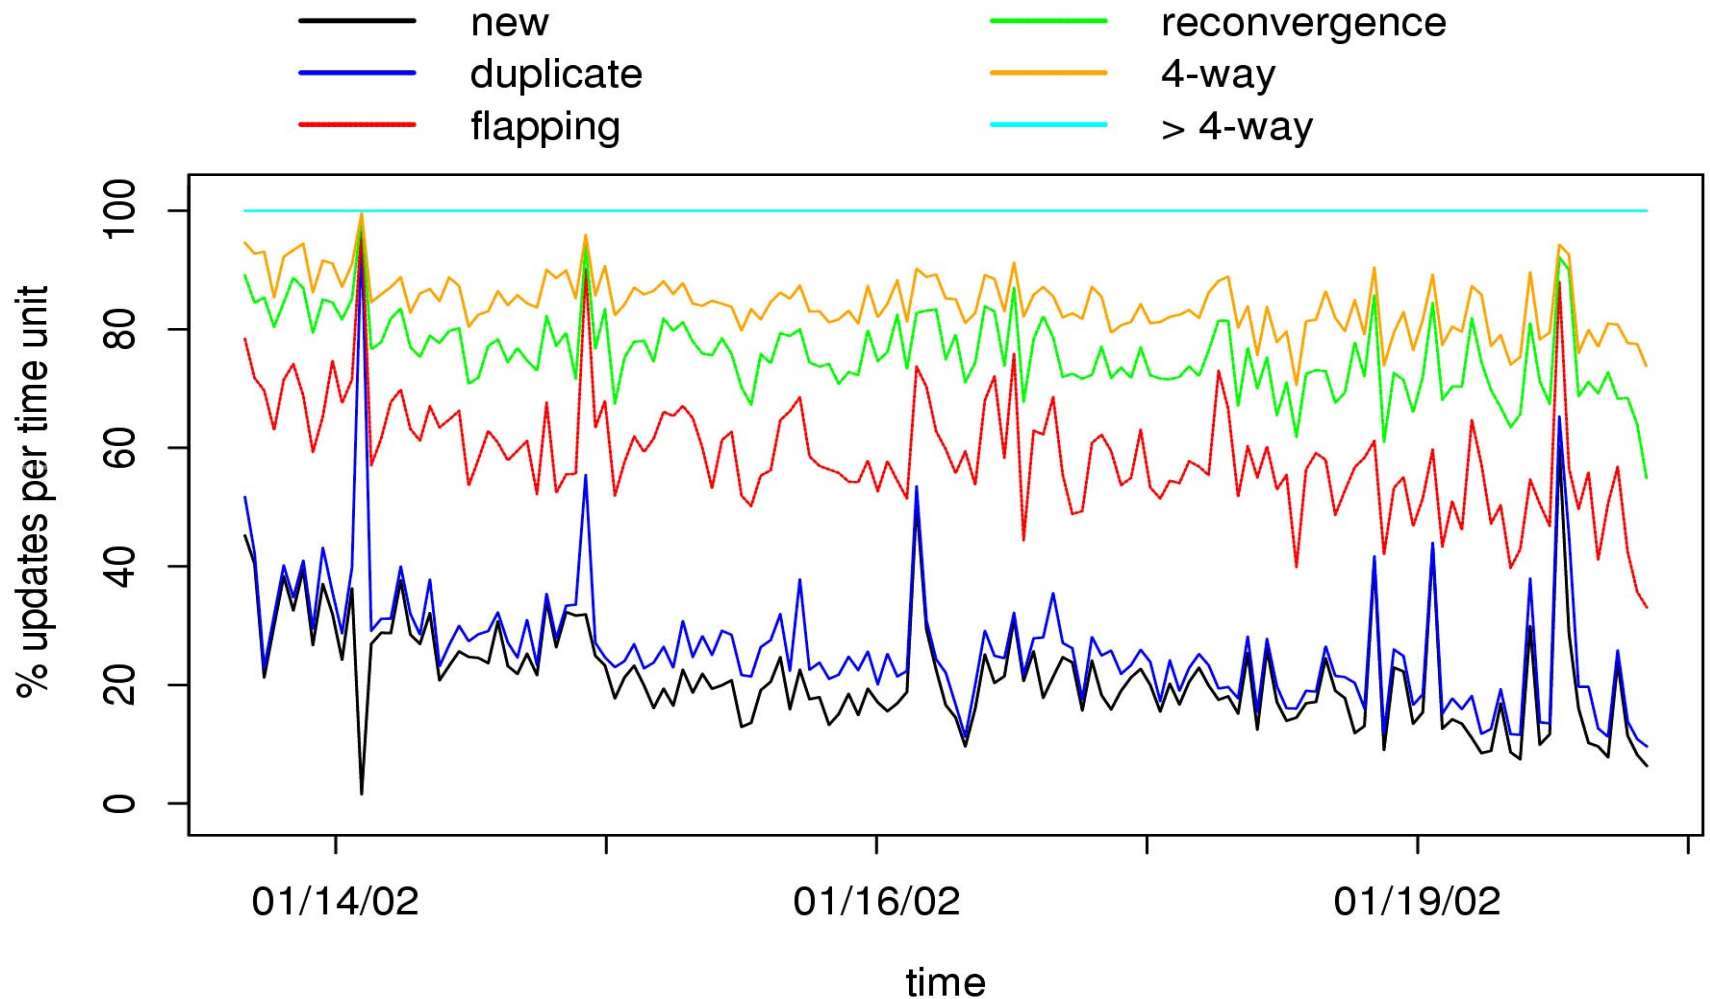
4.

reconvergence

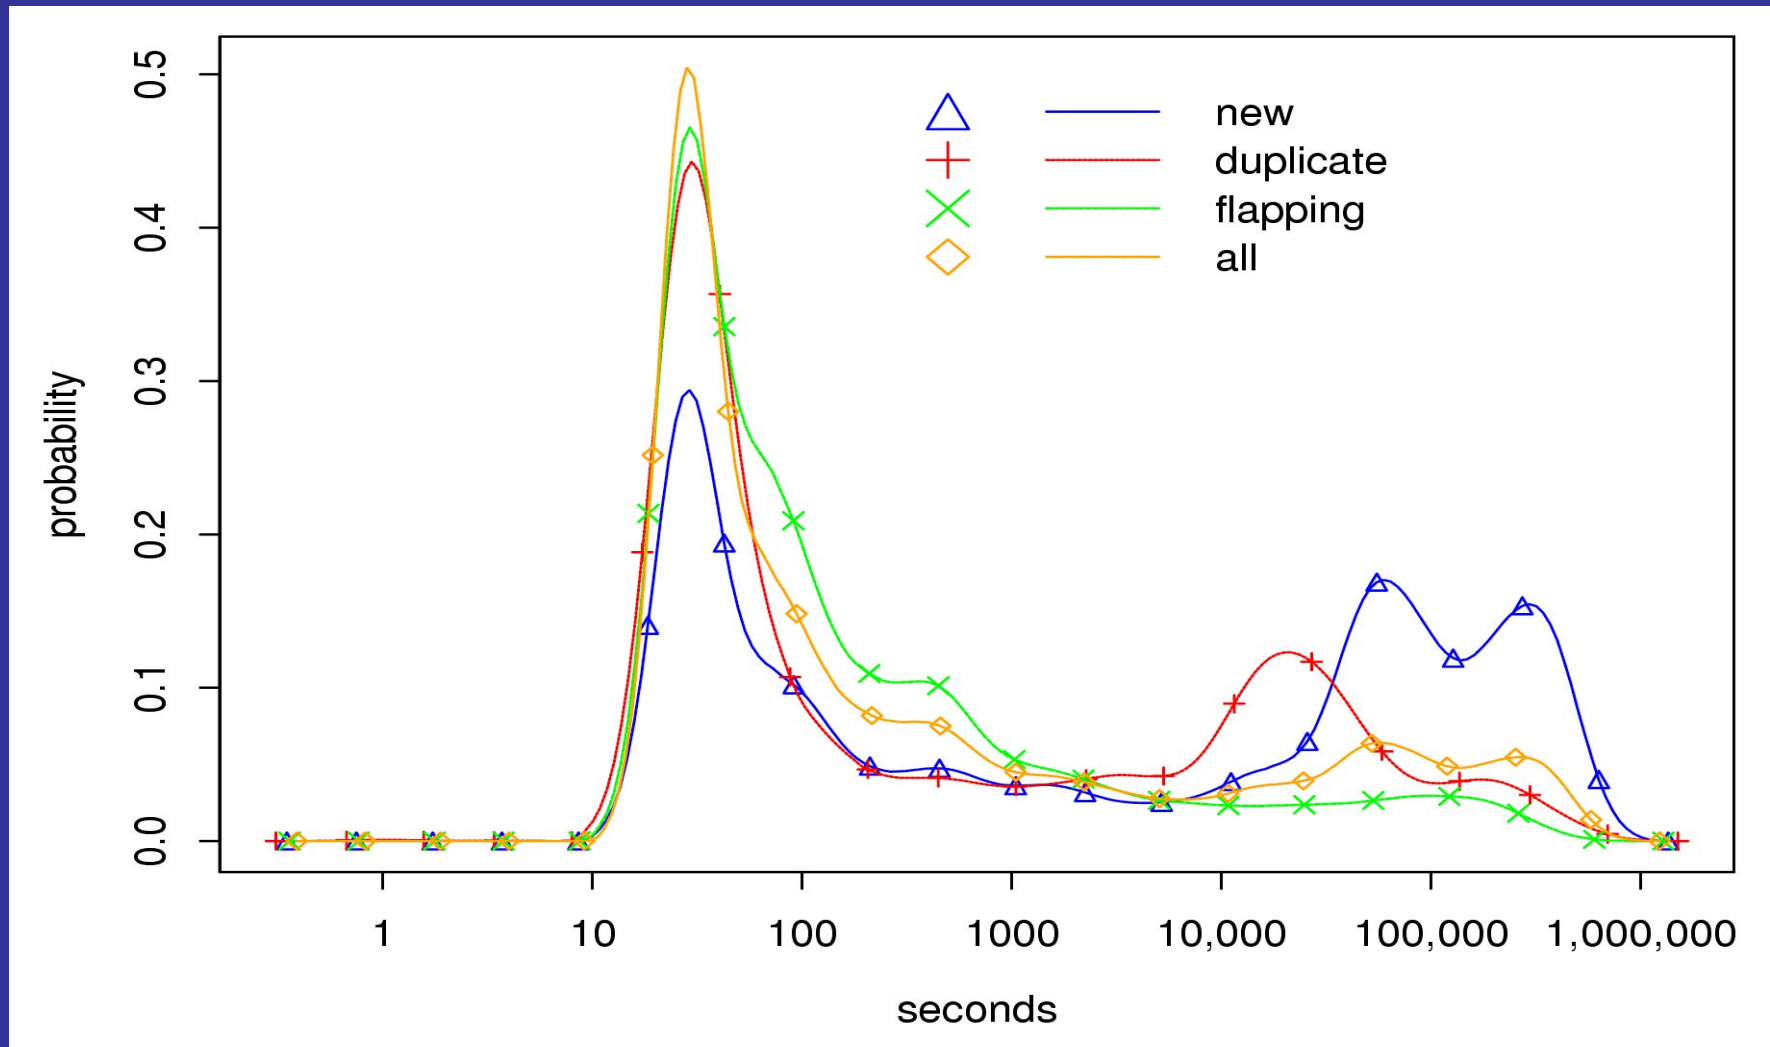
>4

n-way change

Categorization of changes



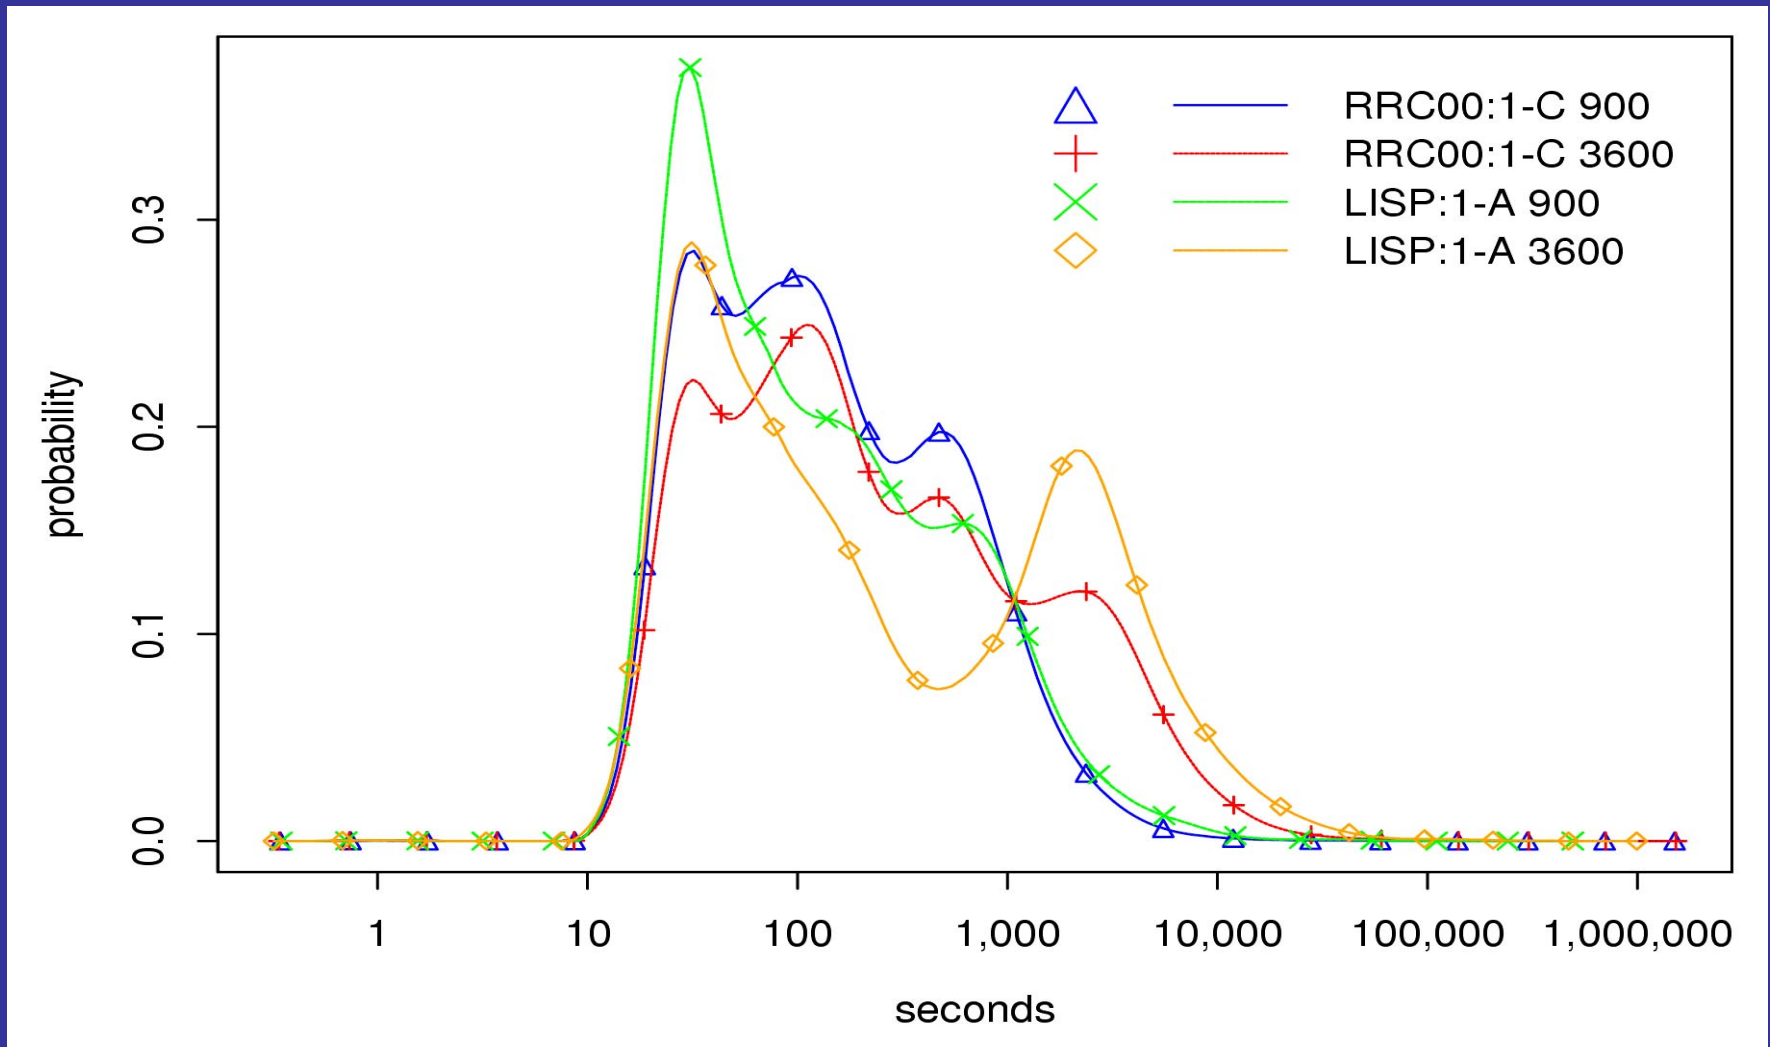
Update interarrival time



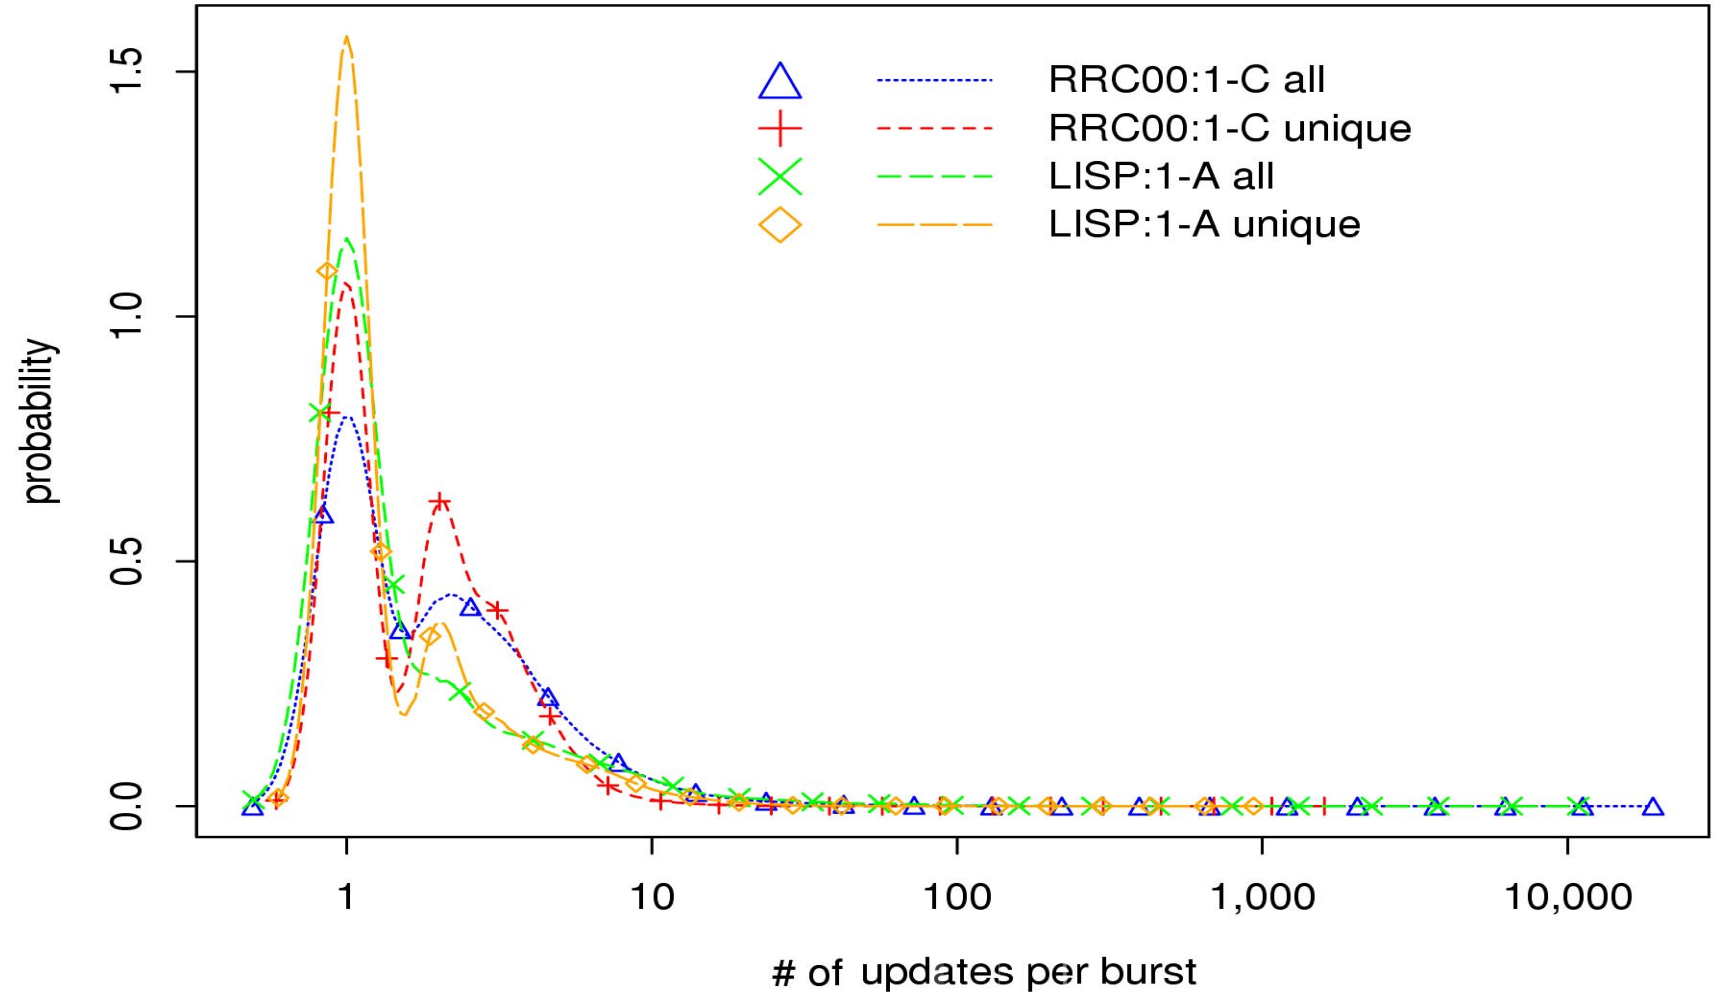
Update burst

- Bursts consists of several updates
 - Same prefix
 - Short time window
- Like packet flows

Burst duration



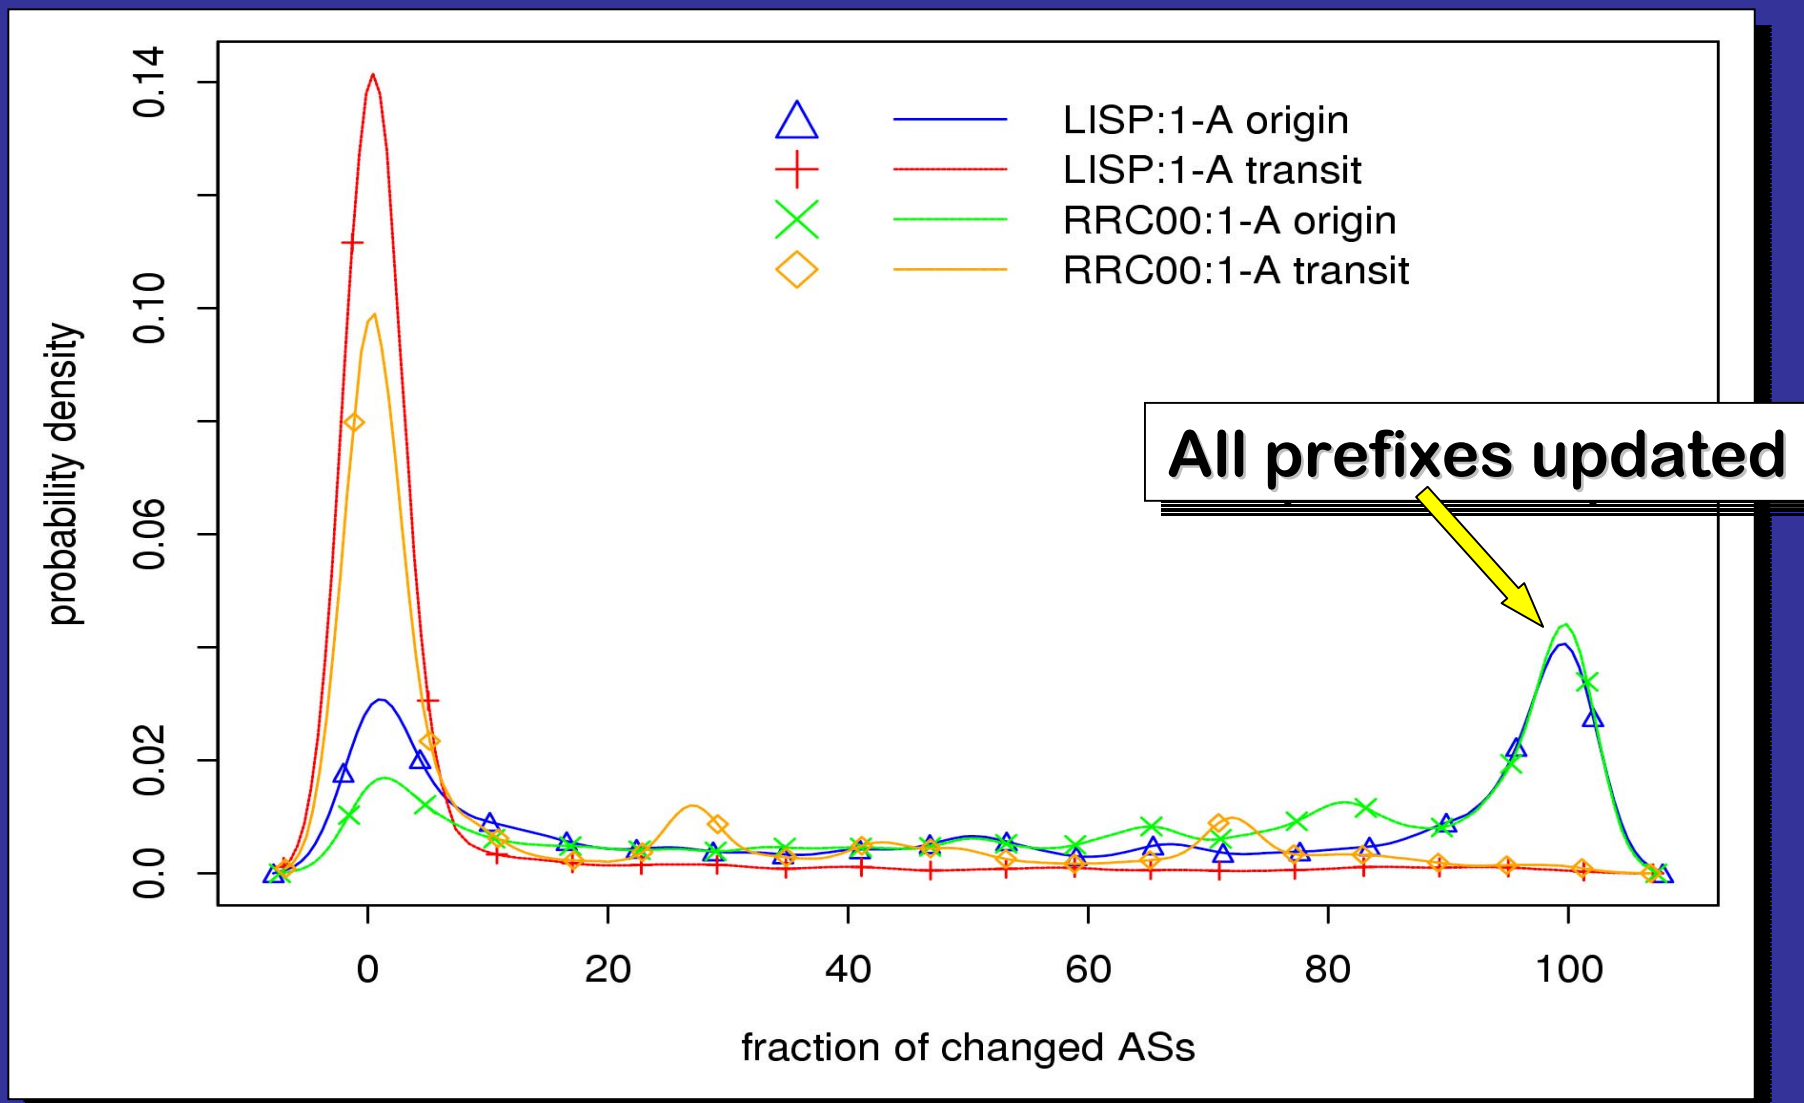
Updates in burst



Session resets

- Peering connection breakdown - a whole table must be exchanged
- Update storms are propagated through the internet...
- How big is the problem?

Identification of session resets



Summary

- BGP workload model
 - Identify structure of BGP traffic
 - Tool, RTG, realization of workload model
 - Characterize structure of BGP traffic
 - Verification of tool
- One more component for an Internet Lab
- Towards a better testing methodology