Observations on redundant BGP traffic and flaps from the RIPE RIS collectors

Ski Ilnicki & Alexander Tudor
Agilent Labs

Agilent Technologies
Innovating the HP Way

Agenda

- Rationale
- Measurement taxonomy
- Recent historical measurements
  - Comparison with previous historical measurements
  - Graphs of recent historical measurements
- Daily Measurements
  - Sample graphs
  - Next steps
- References
- Credits
- Questions

Rationale

- **Identify patterns in BGP traffic measurements**
- **Quantify performance cost of BGP processing**
- **Detect anomalies**

Measurement taxonomy [LA99]

- **Tup & Tdown: Fluctuations in the reachability for a given prefix. An announced route is withdrawn and transitions down (Tdown), or a currently unreachable prefix is announced as reachable and transitions up (Tup).**
- **WWDup: The repeated transmission of BGP withdrawals for a prefix that is currently unreachable.**
- **AADup: A route is implicitly withdrawn and replaced with a duplicate of the original route. [...] a *duplicate route* does not differ in any BGP path attribute information.**
- **AADiff: A route is implicitly withdrawn and replaced by an alternative route as the original route becomes unreachable, or a preferred alternative path becomes available**
- **Flap: Tdown followed by Tup where the prefix has identical attributes when first announced and then re-announced after a withdrawal.**

Recent historical measurements

- **1 full feed peer at LINX and 1 full feed peer at AMSIX from 8/1/2001 to 1/31/2002**
- **measurements:**
  - **daily updates, announcements, duplicate announcements**
  - **daily flaps**
  - **distribution of updates & announcements inter-arrival time and flap duration**

Previous historical measurements [LA99]

- **multiple full-feed and partial-feed peers at MAE-East 1/1996 – 8/1998**
- **measurements:**
  - **AADup**
  - **WWDup**
  - **AADiff**
  - **Tdown**
  - **Tup**

Comparison with previous historical measurements

- [LA99] conclusions
  - AADup 20% of all updates (~50K prefixes)
  - 50% of all updates have 30 second periodicity
  - WWDup 50% of all updates (5% after router software fix)
- [LA99] additional comments
  - AADup behavior is well distributed across studied ISPs
  - AADup and WWDup caused by small service providers

- Current work conclusions
  - AADup 20% of all updates (~100K prefixes)
  - 35% of all updates have 30 second periodicity
  - WWDup < 3% of all updates

Announcements, AADups & Flaps as % of Updates per Day
AMS-IX Peer
8/1/2001 - 1/31/2002

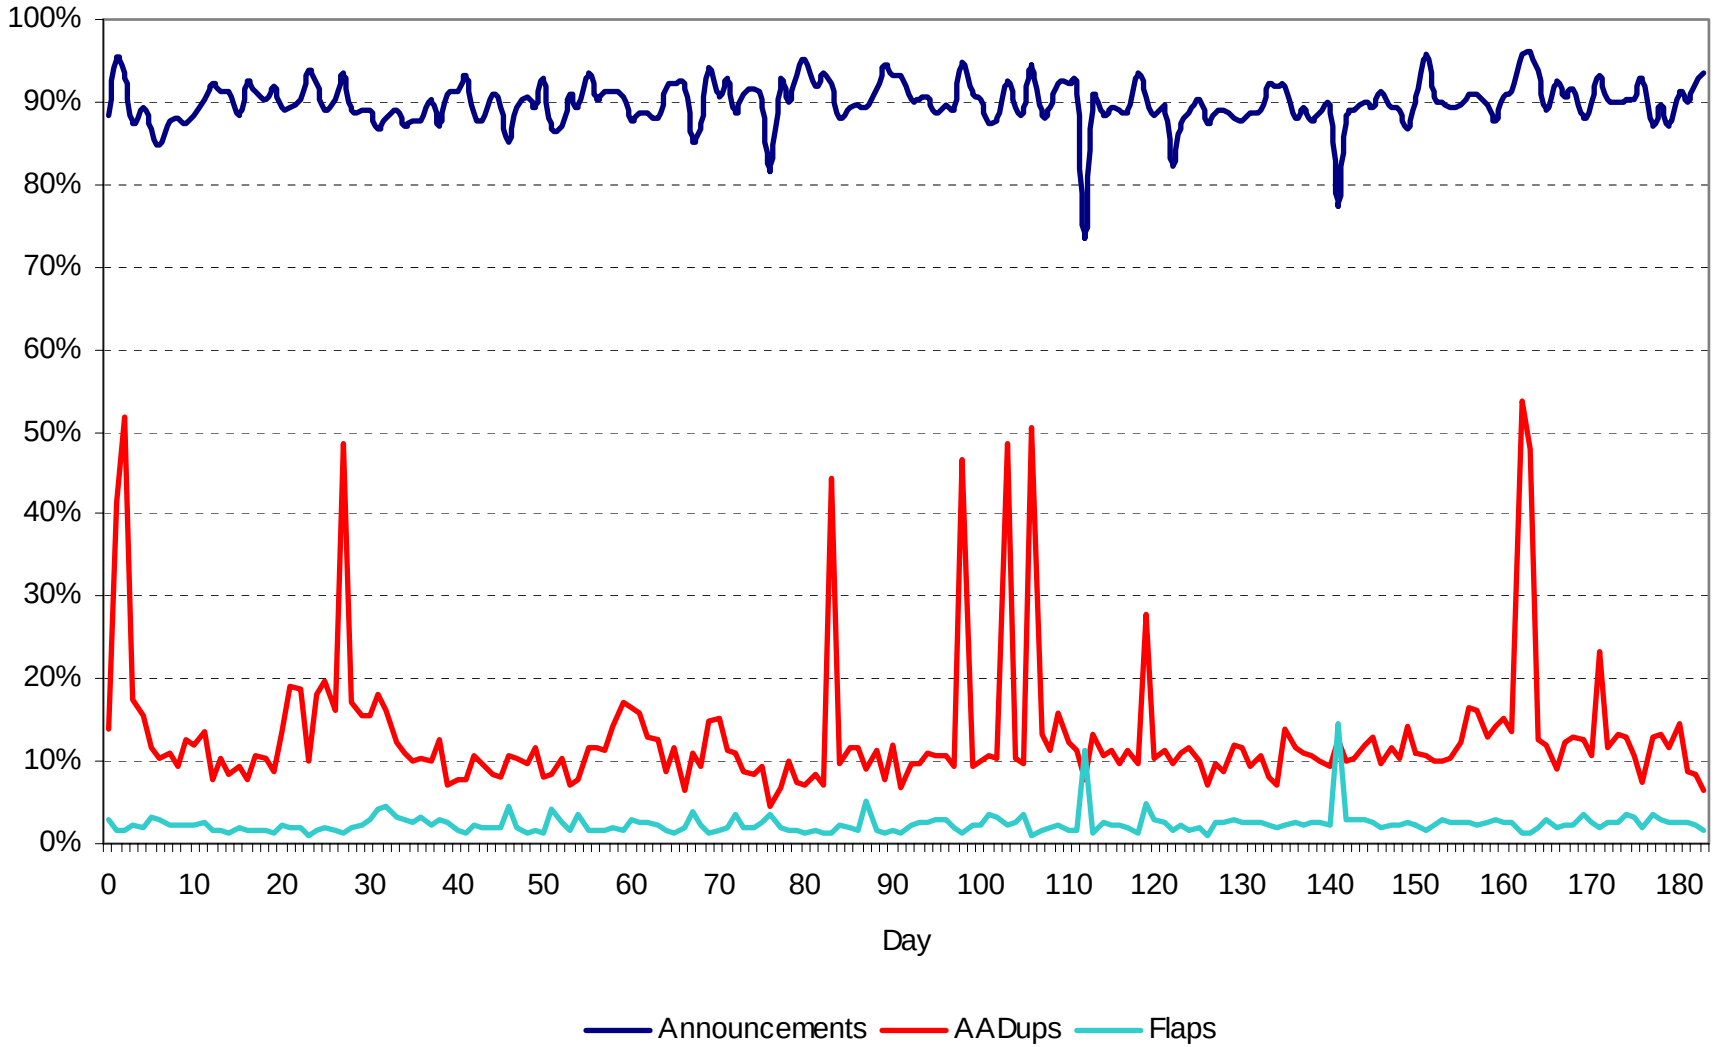

Announcements, AADups & Flaps as % of Updates per Day
LINX Peer
8/1/2001 - 1/31/2002

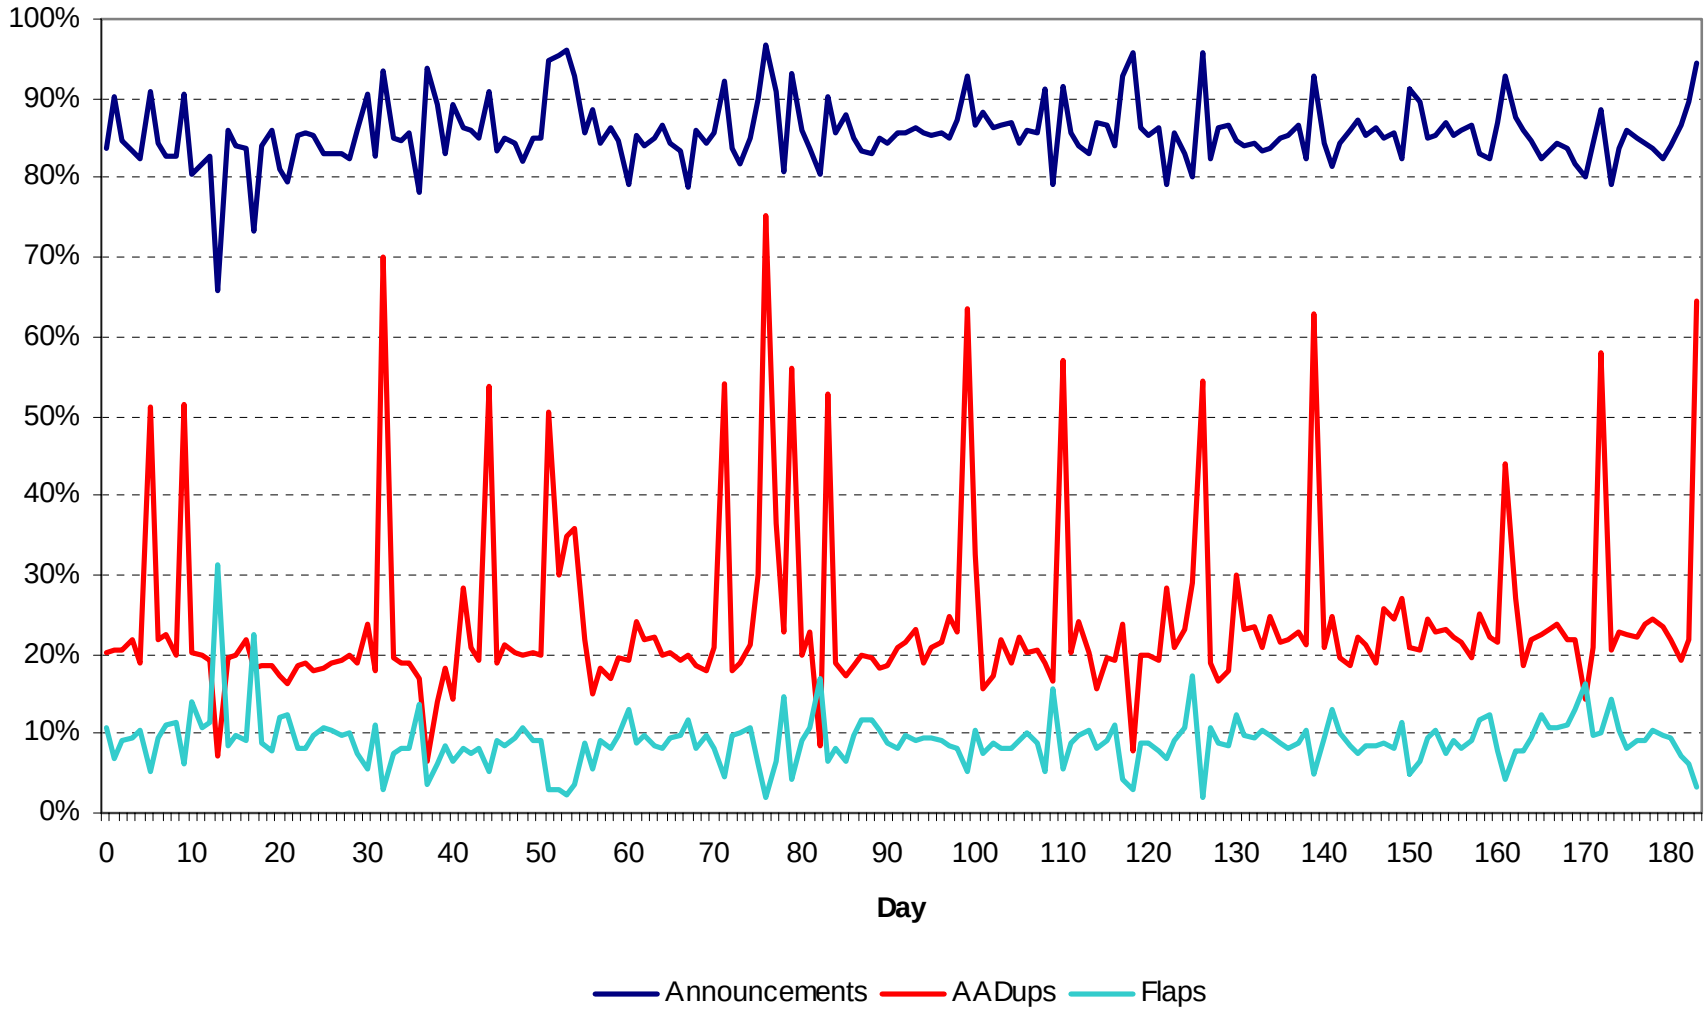

Update Inter-Arrival Time Distribution

AMS-IX

8/1/2001 - 1/31/2002

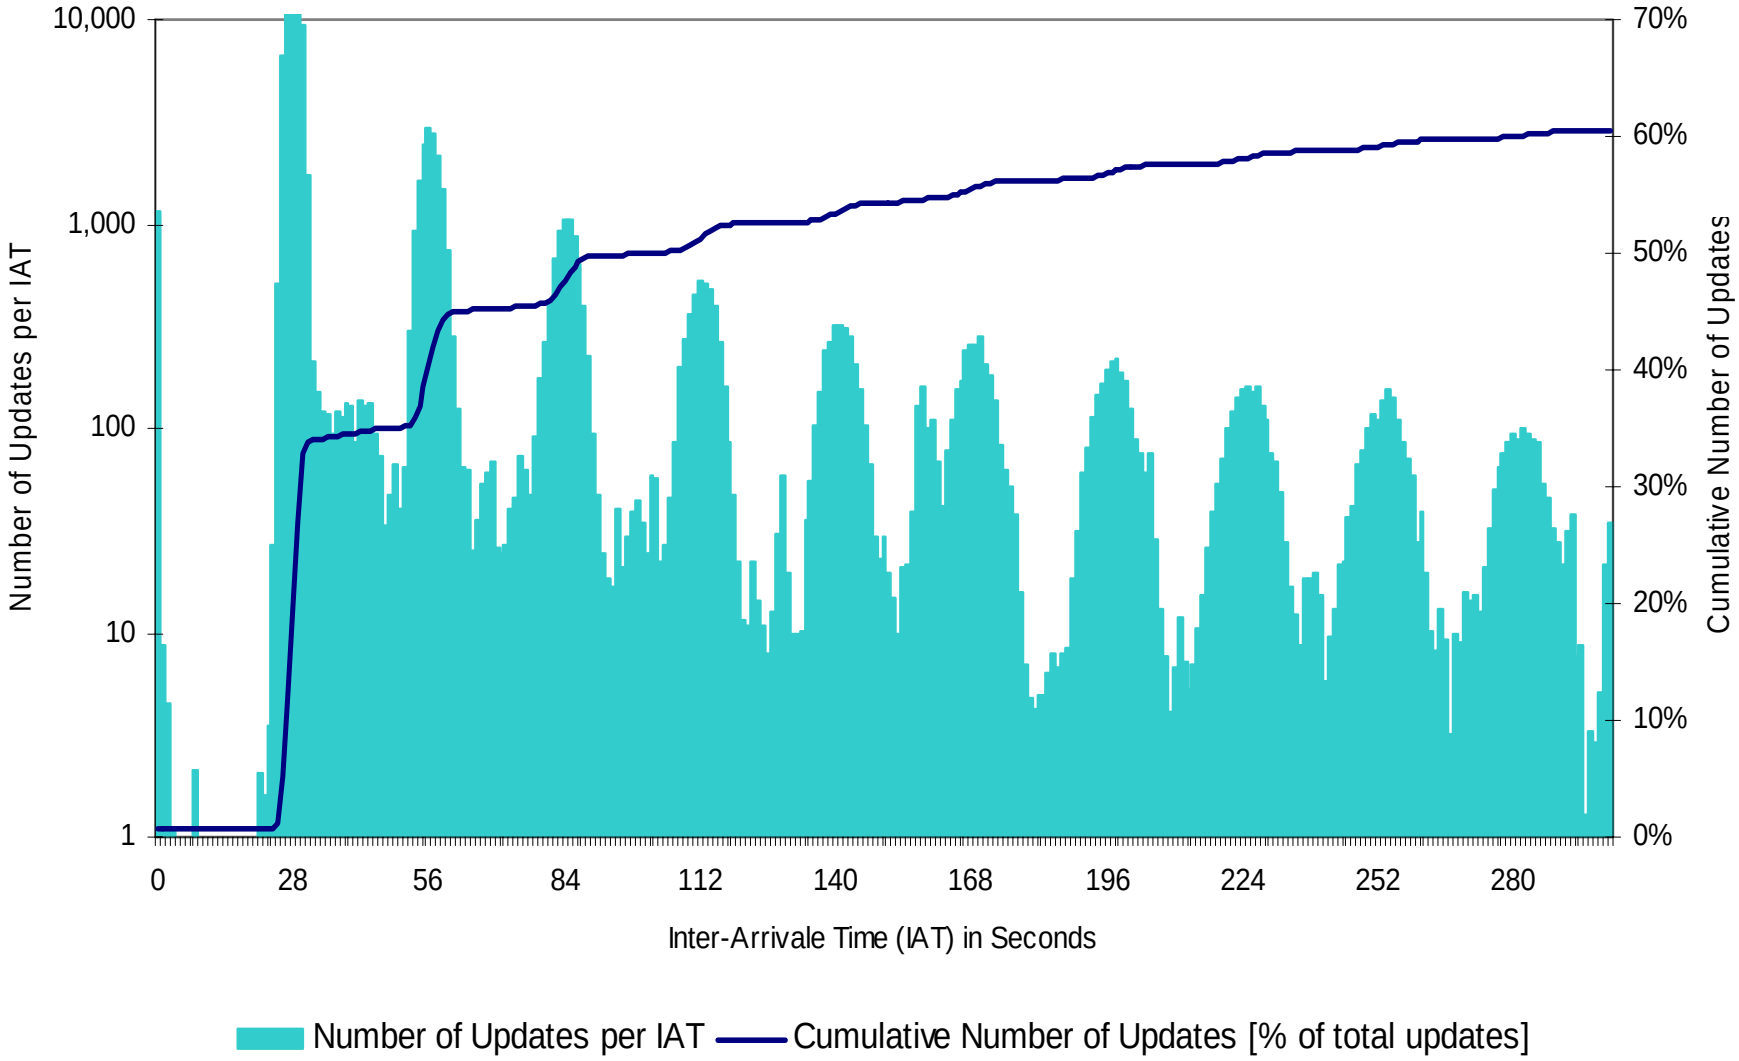

Announcement Inter-Arrival Time Distribution
AMS-IX Peer
8/1/2001 - 1/31/2002

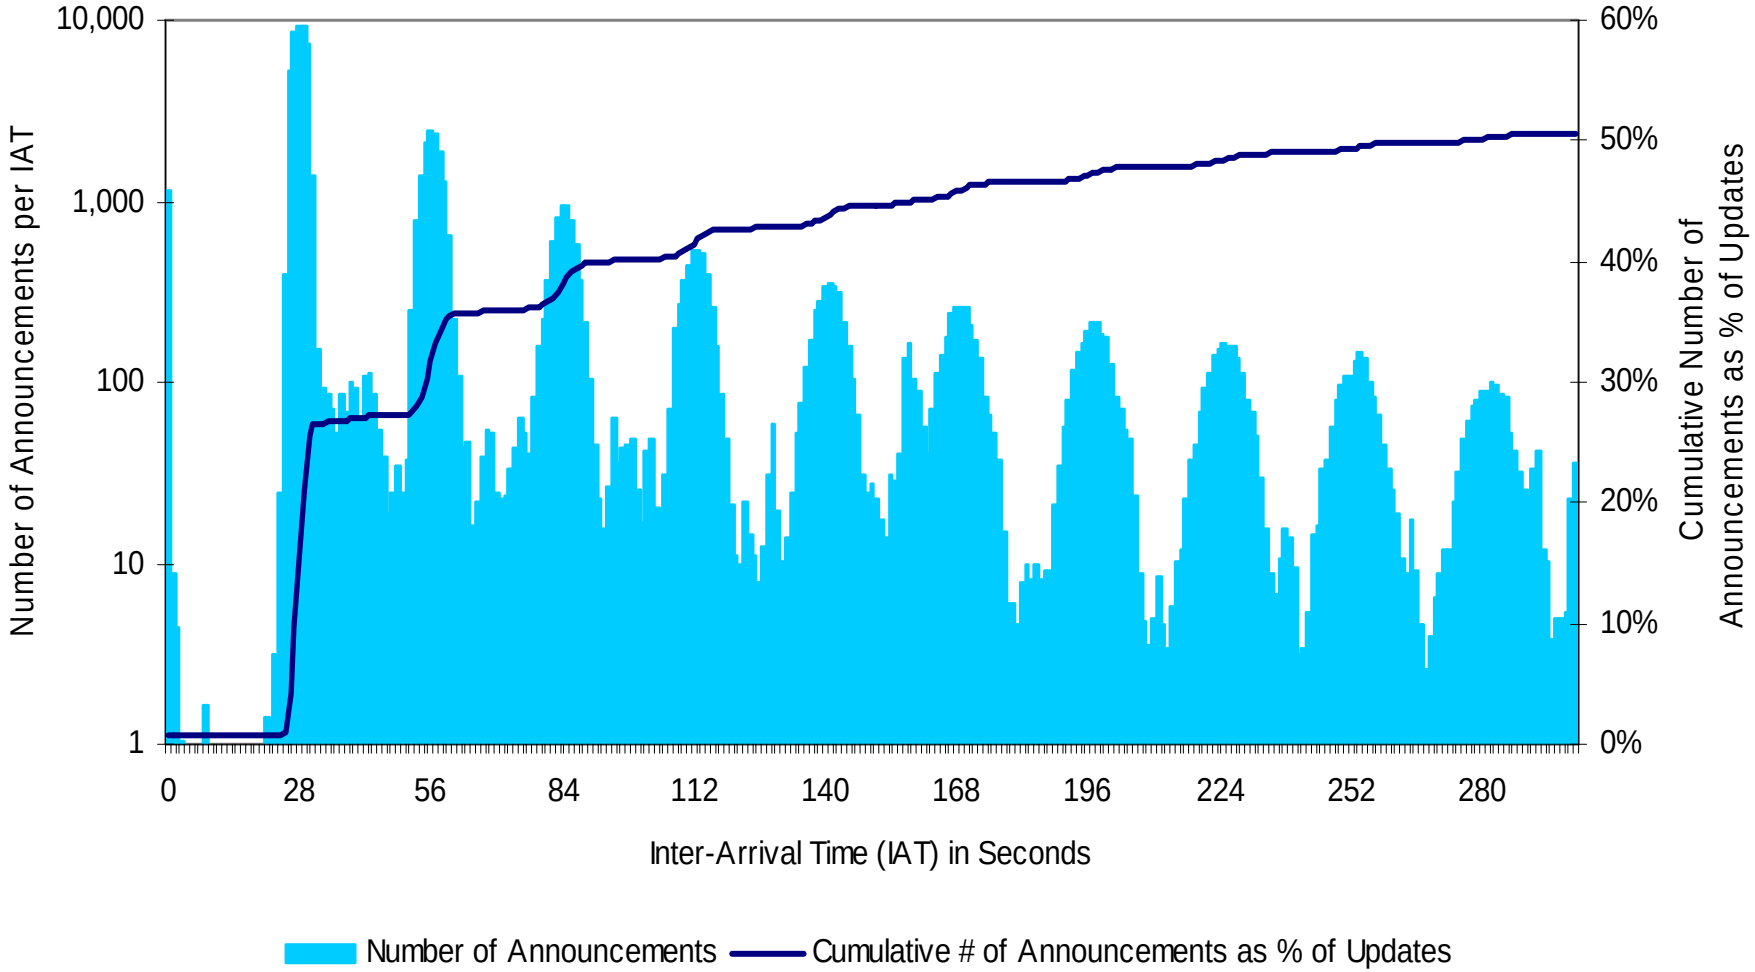

AADup Inter-Arrival Time Distribution
AMS-IX Peer
8/1/2001 - 1/31/2002

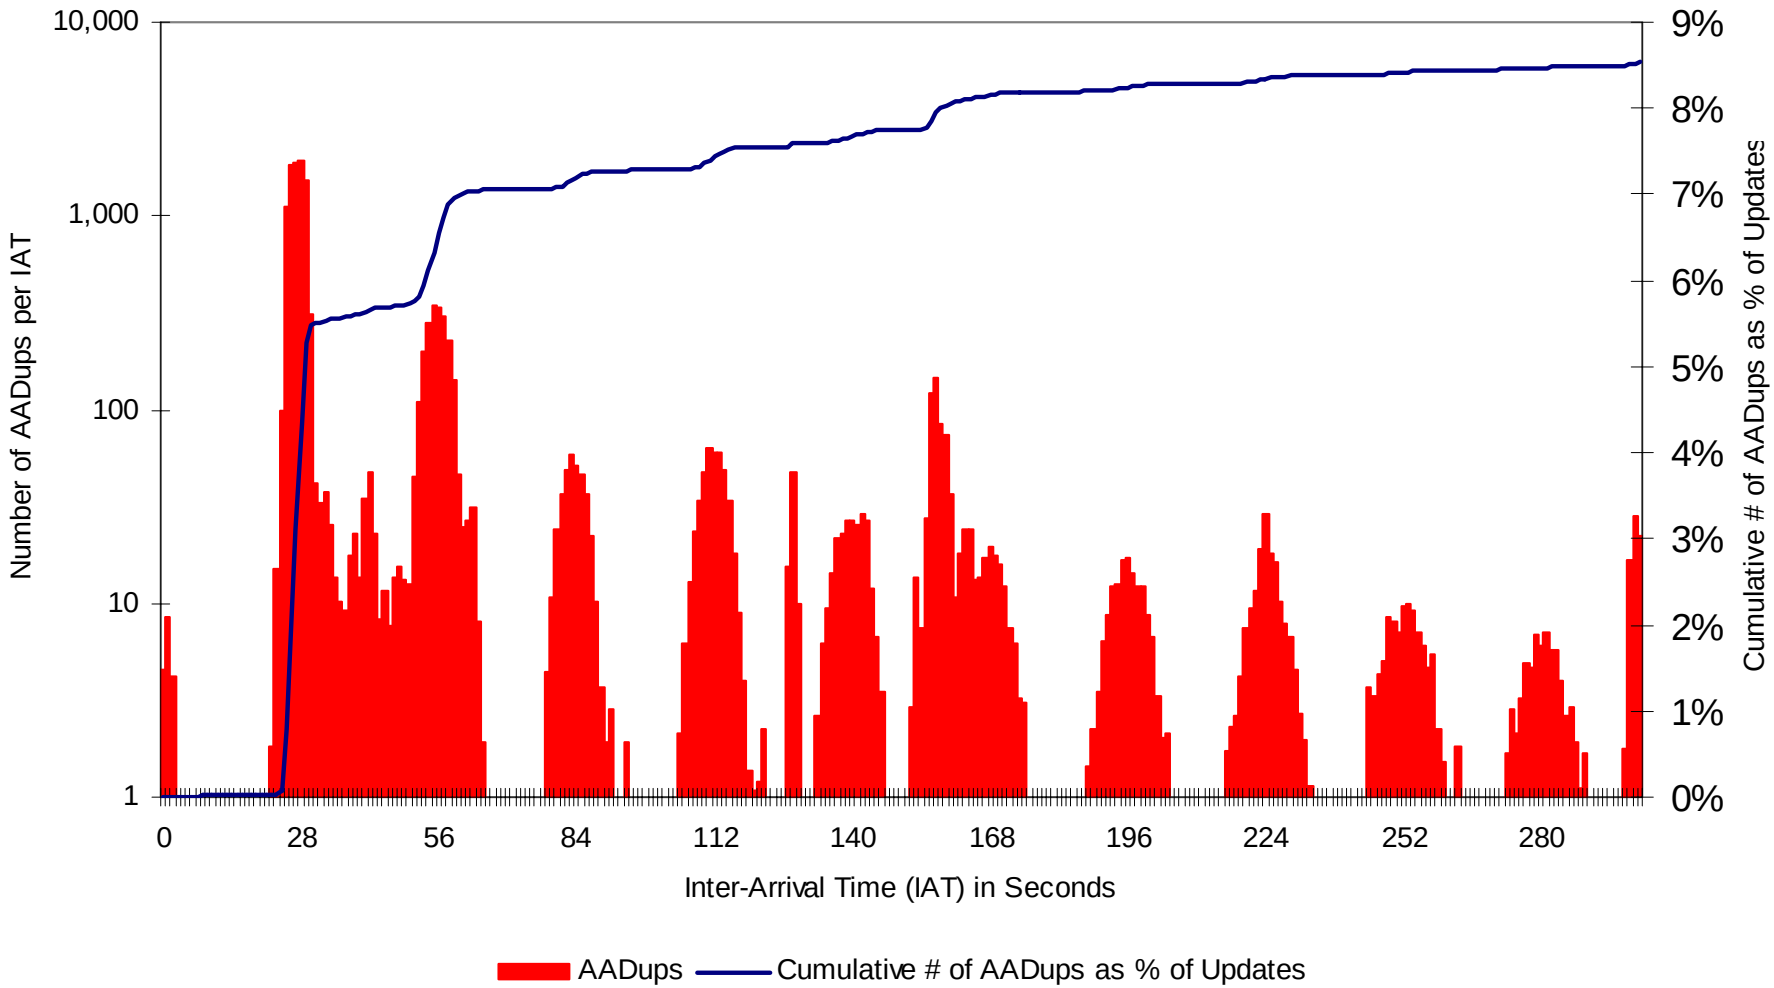

Flap Duration Time Distribution
AMS-IX Peer
8/1/2001 - 1/31/2002

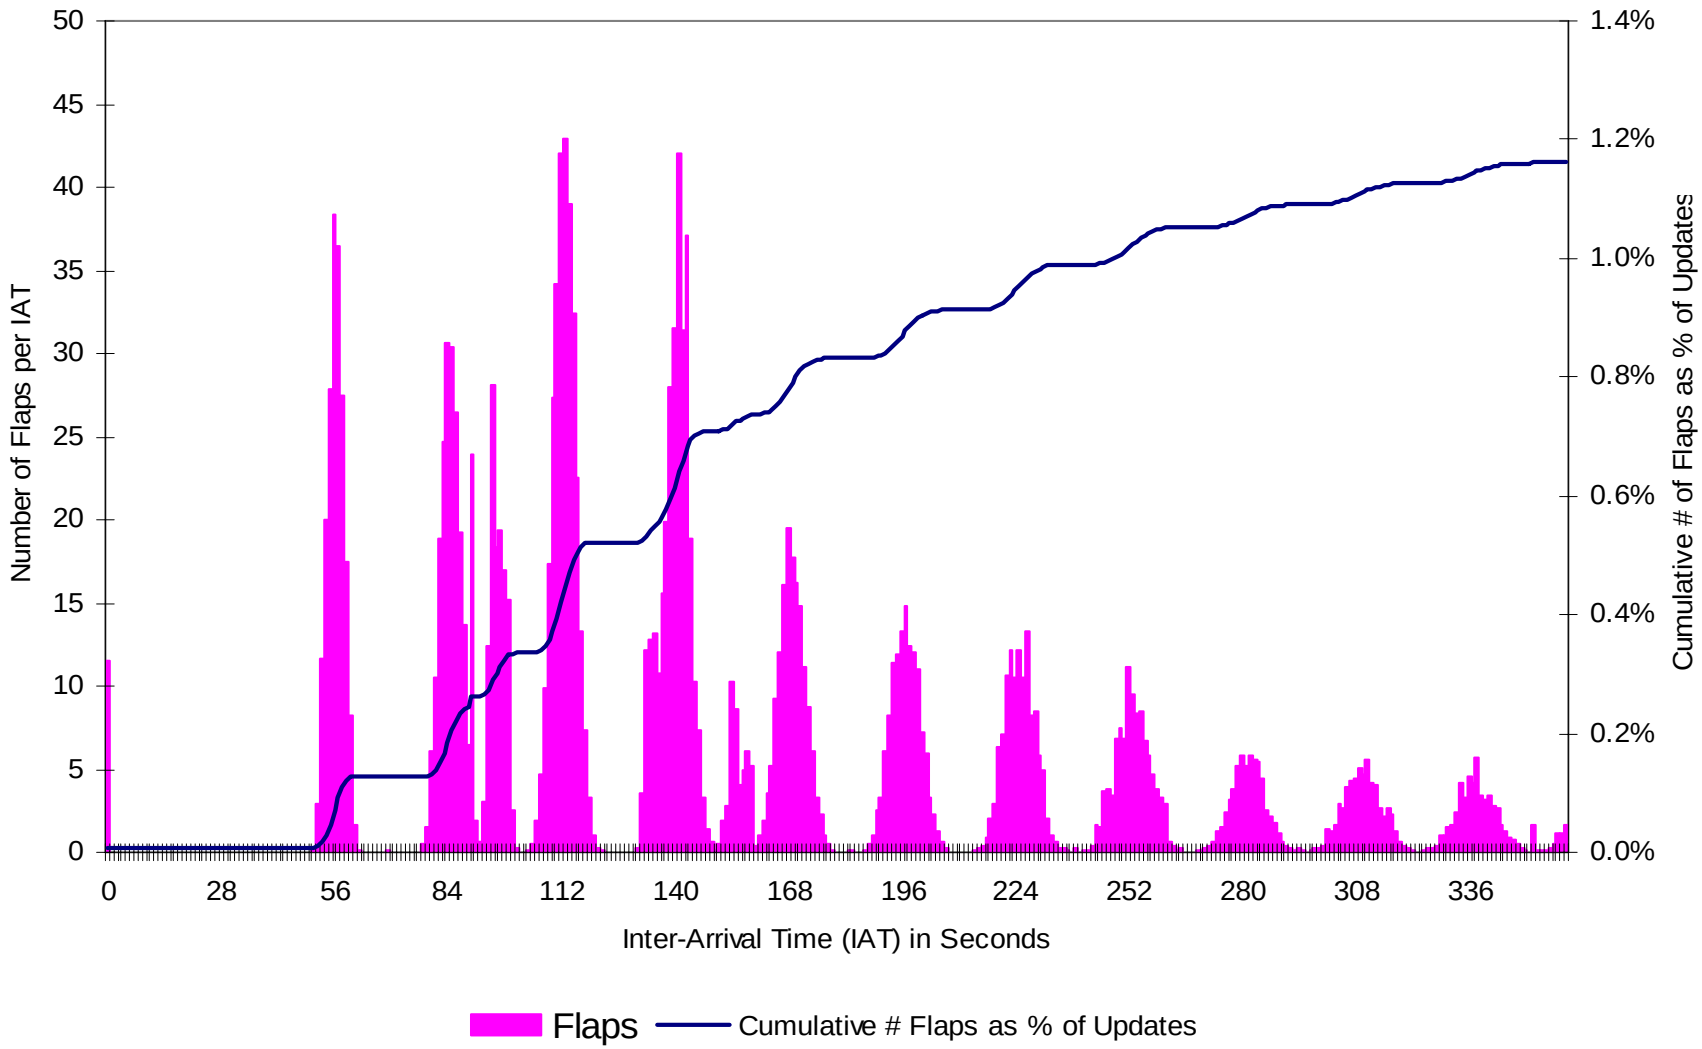

WWDup as % of Updates per Day
LINX Peer
8/1/2001 - 1/31/2002

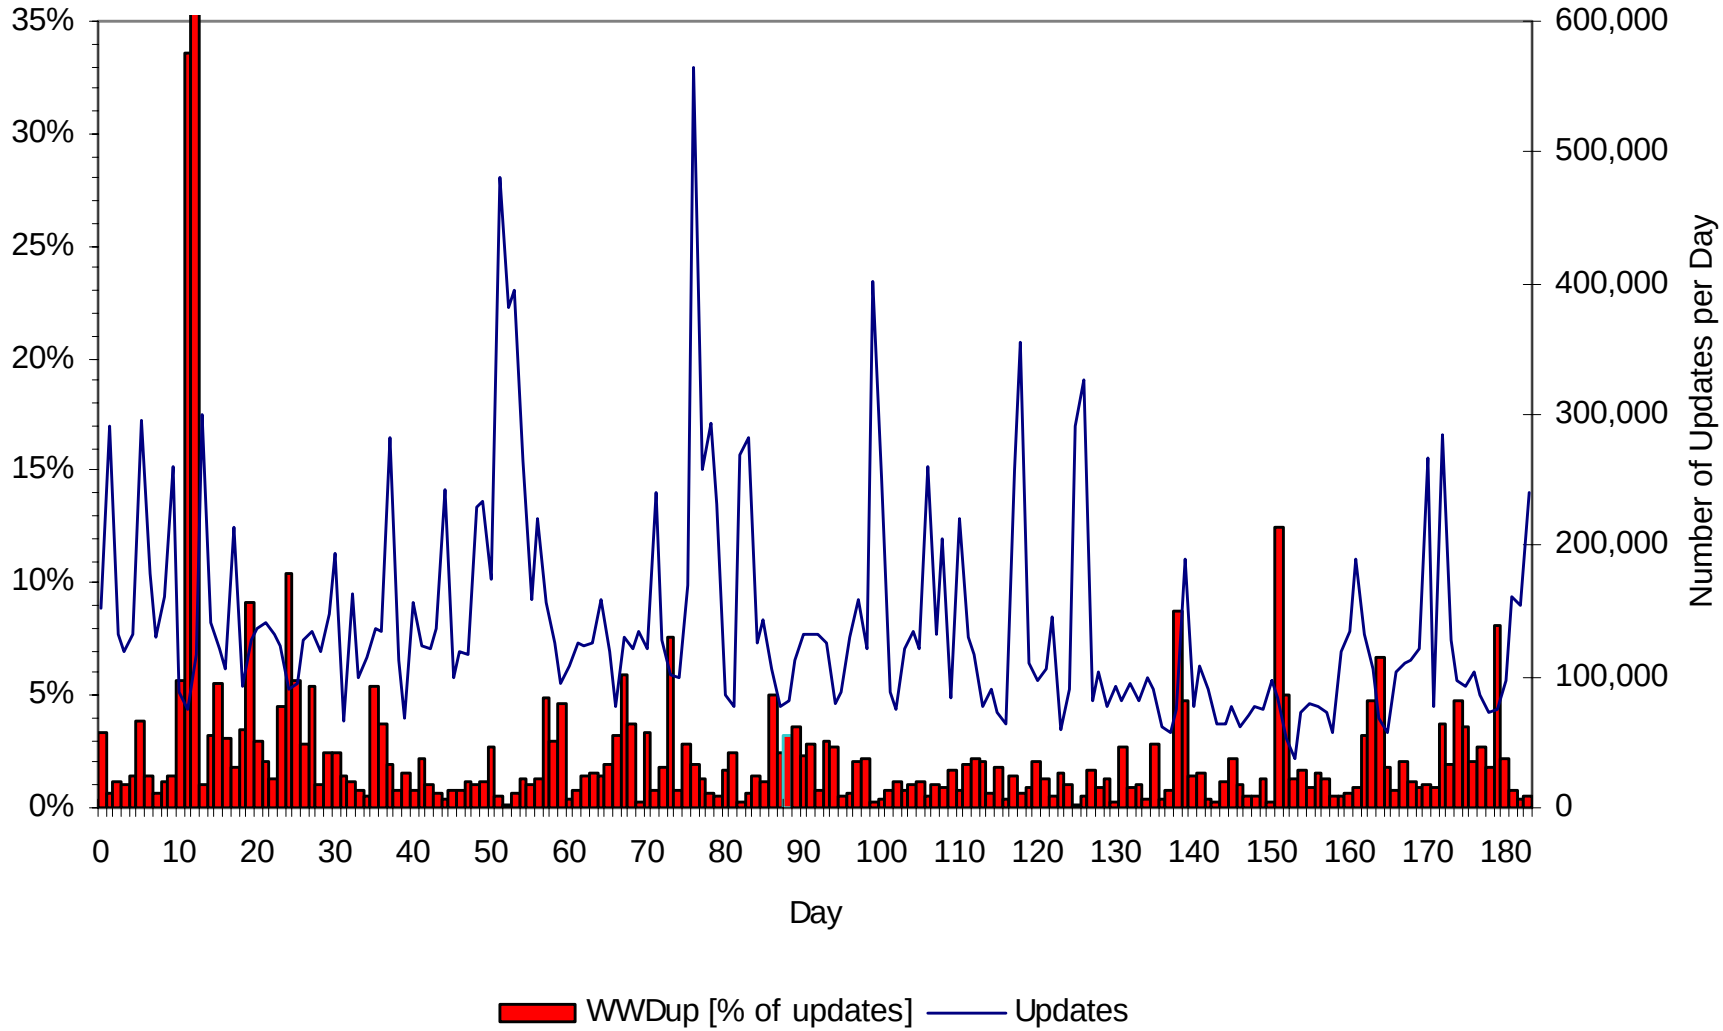

Daily measurements

- **12 RIPENCC peers, 3 AMSIX peers, 1 LINX peer**
- **by peer by hour**
- **measurements:**
  - **A & AADup**
  - **W & WWDup**
  - **flaps**

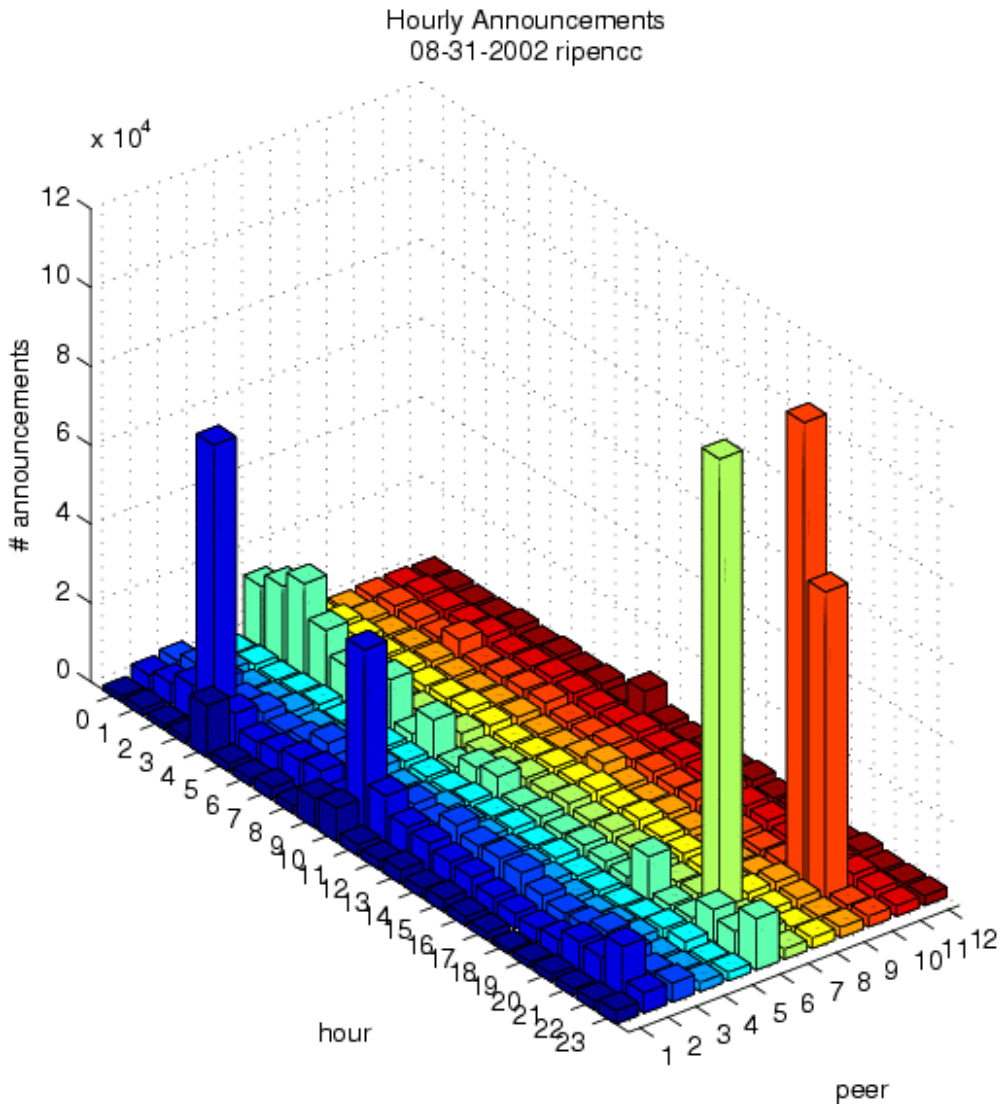
Hourly Announcements
08-31-2002 ripencc
x 10^4
announcements
0
2
4
6
8
10
12
hour
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
peer
1 2 3 4 5 6 7 8 9 10 11 12

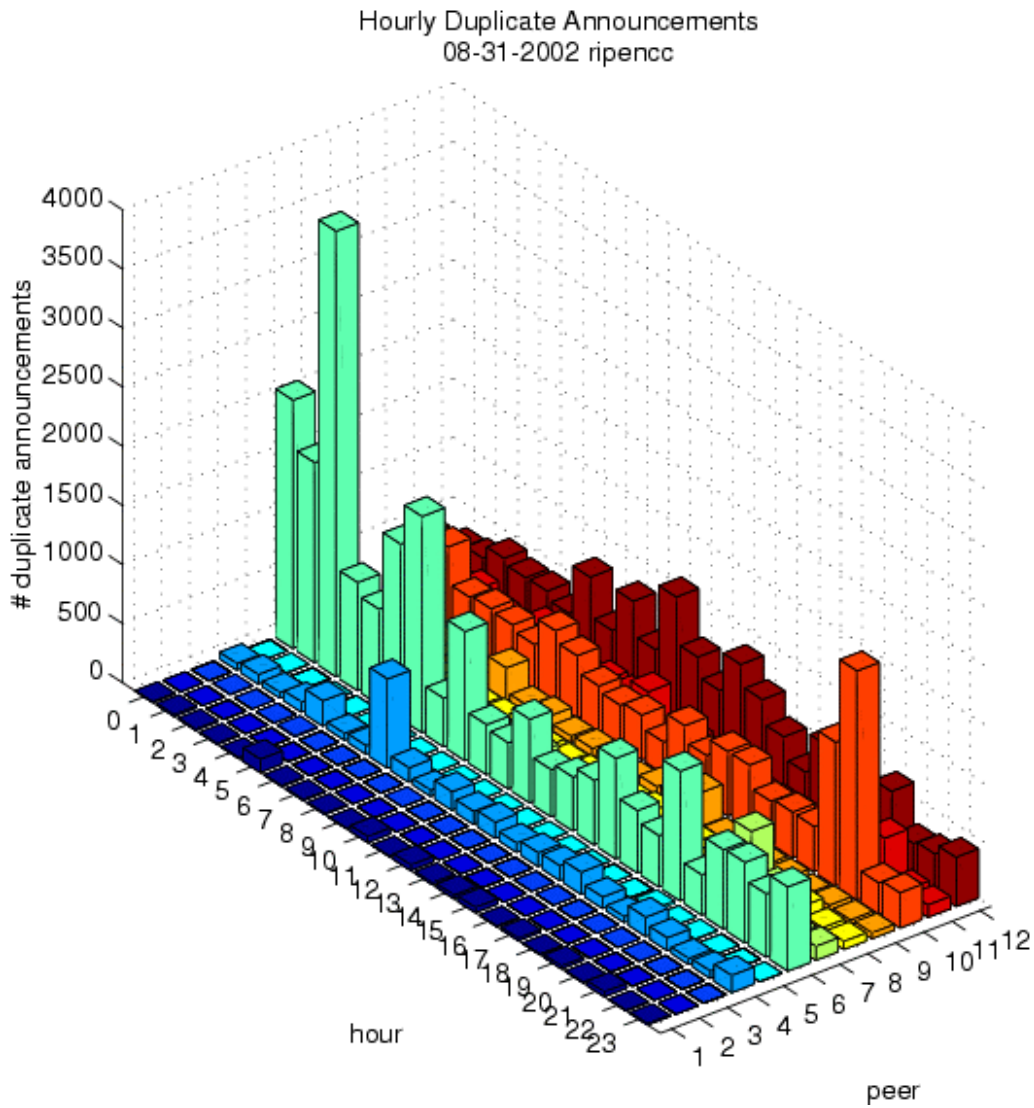
Hourly Duplicate Announcements
08-31-2002 ripencc
duplicate announcements
4000
3500
3000
2500
2000
1500
1000
500
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
hour
1 2 3 4 5 6 7 8 9 10 11 12
peer

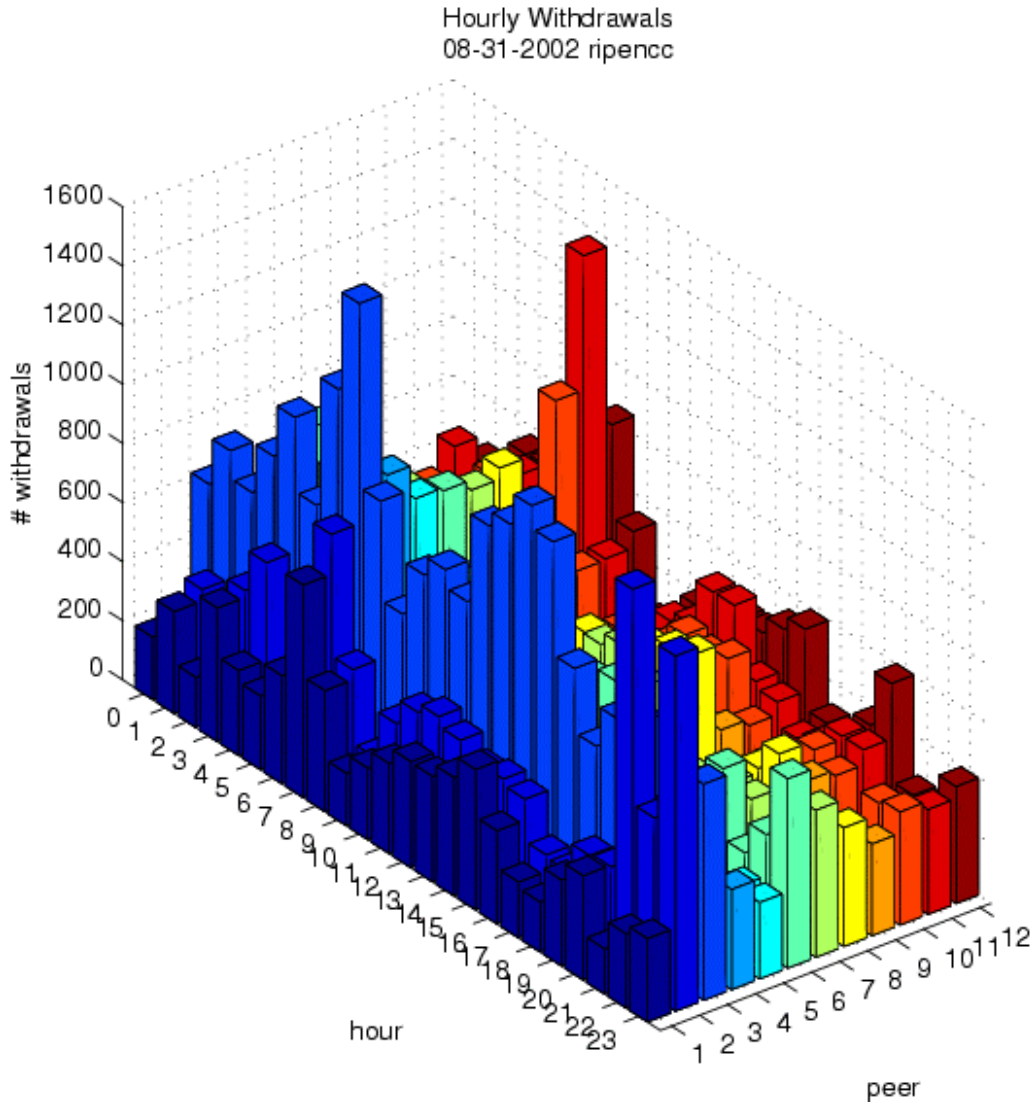
Hourly Withdrawals
08-31-2002 ripencc
withdrawals
1600
1400
1200
1000
800
600
400
200
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
hour
1 2 3 4 5 6 7 8 9 10 11 12
peer

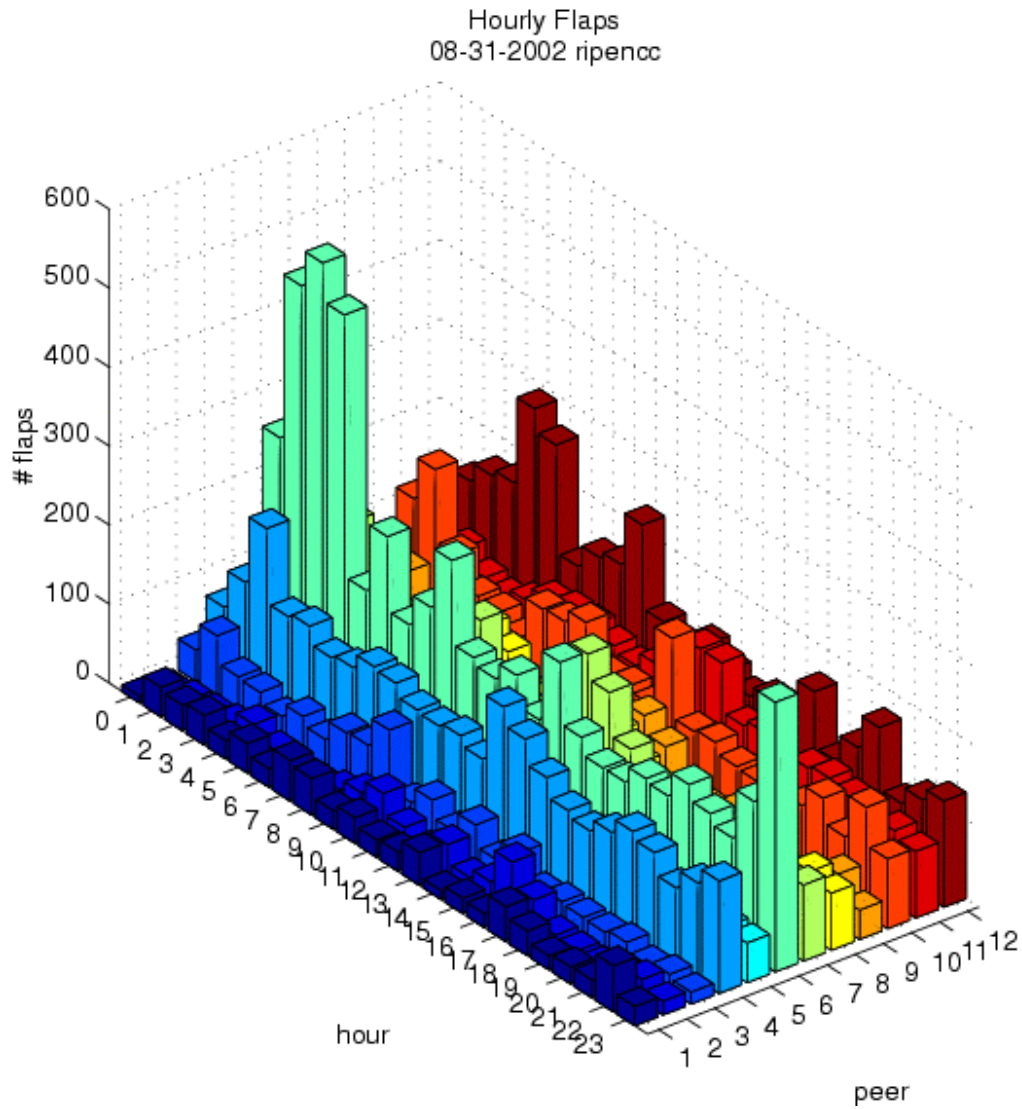
Hourly Flaps
08-31-2002 ripencc
flaps
600
500
400
300
200
100
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
hour
1 2 3 4 5 6 7 8 9 10 11 12
peer

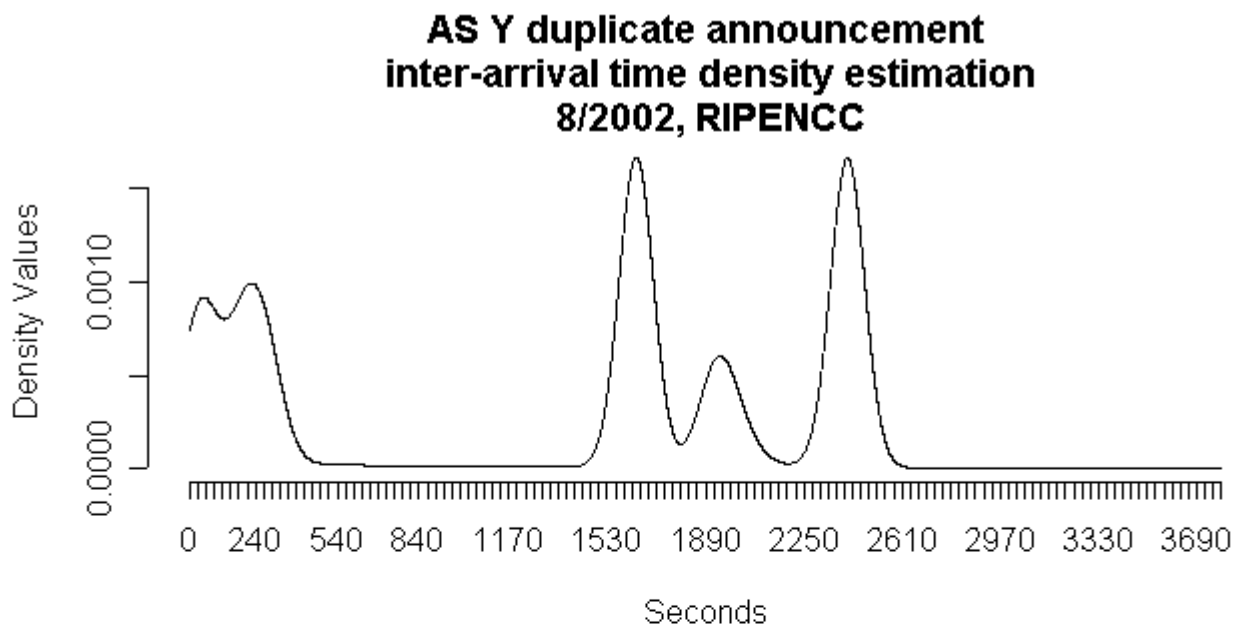
AS Y duplicate announcement
inter-arrival time density estimation
8/2002, RIPENCC
Density Values
0.0000
0.0010
0 240 540 840 1170 1530 1890 2250 2610 2970 3330 3690
Seconds

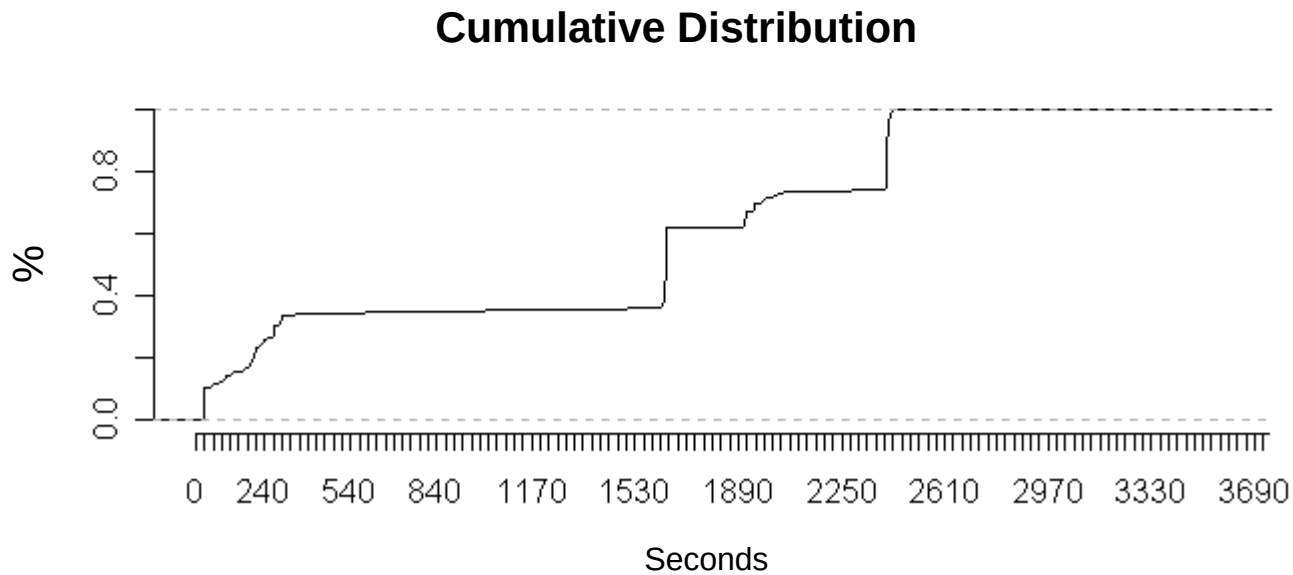
Cumulative Distribution
%
0.0
0.4
0.8
0 240 540 840 1170 1530 1890 2250 2610 2970 3330 3690
Seconds

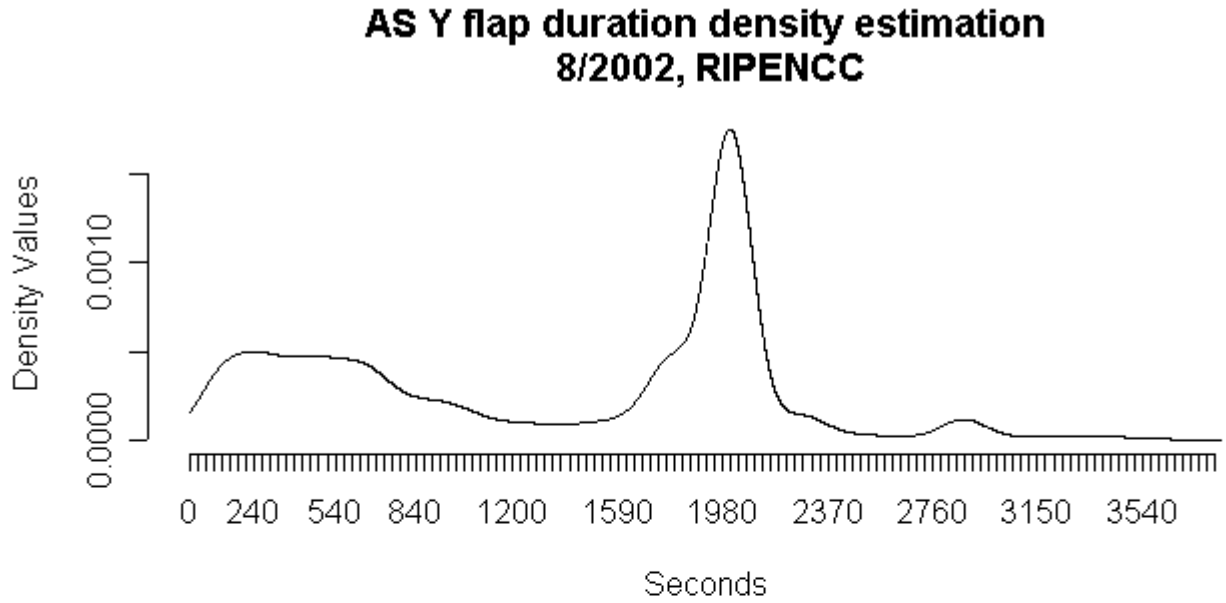
AS Y flap duration density estimation
8/2002, RIPENCC
Density Values
0.0000
0.0010
0 240 540 840 1200 1590 1980 2370 2760 3150 3540
Seconds

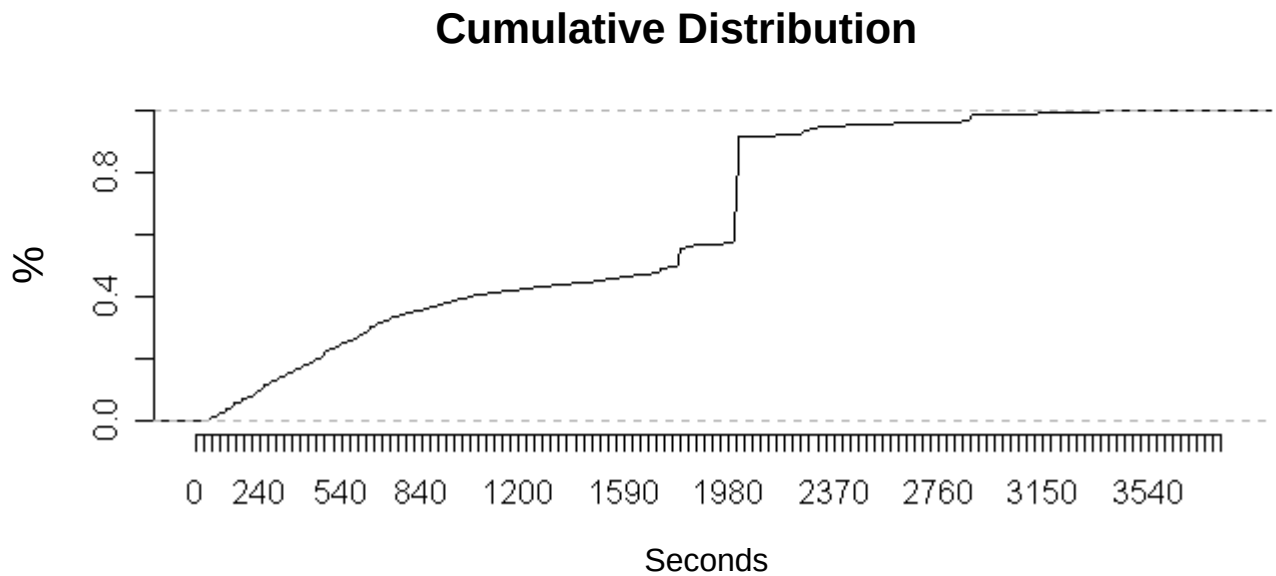
Cumulative Distribution
%
0.0
0.4
0.8
0 240 540 840 1200 1590 1980 2370 2760 3150 3540
Seconds

Next steps

- **attempt to identify spatial and temporal patterns using defined measurements**
- **is there a baseline?**
- **develop methodology to quantify performance cost of BGP processing**

References & Related work

- **[LA99] C. Labovitz, “Scalability of the Internet Backbone Routing Infrastructure”, Ph.D. Thesis, University of Michigan, 1999**
- **MRTd http://www.mrtd.net/**
- **Matlab http://www.mathworks.com/**
- **R http://www.r-project.org/**
- **Density estimation http://www.maths.uwa.edu.au/~duongt/seminars/intro2kde/**
- **Related work:**
  - **http://www.sprintlabs.com/Department/IP-Interworking/Routing/PyRT/index.html**
  - **http://bgp.lcs.mit.edu/bgpview.cgi**
  - **http://www.renesys.com/projects/bgp_instability/**

Credits

- **RIPENCC's RIS Team**
- **Oregon Routeviews**
- **Jonathan Li, Lance Tatman**

Questions

- **over what time period should peers be compared?**
- **are we taking the *right* measurements?**
- **what additional measurements would you like to see?**
- **flap & dup counts are cut off at 1 hr; is this reasonable?**
- **we need volunteers to correlate observed behavior with trouble tickets**

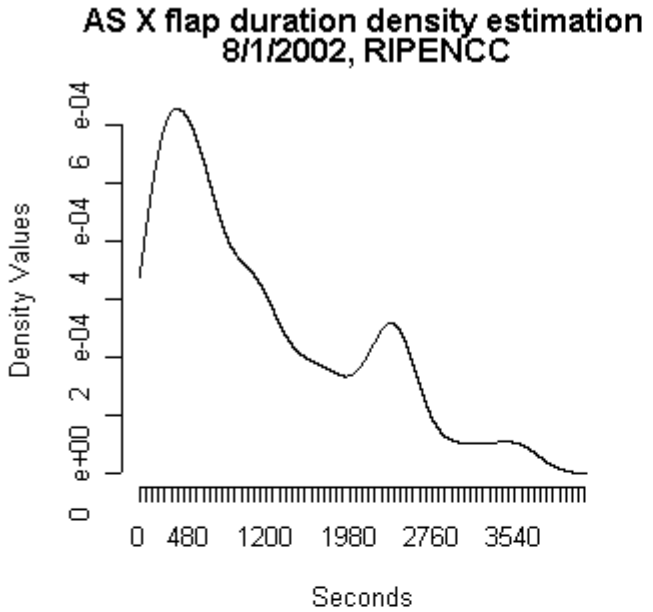
AS X flap duration density estimation
8/1/2002, RIPENCC
Density Values
0 e+00 2 e-04 4 e-04 6 e-04
0 480 1200 1980 2760 3540
Seconds

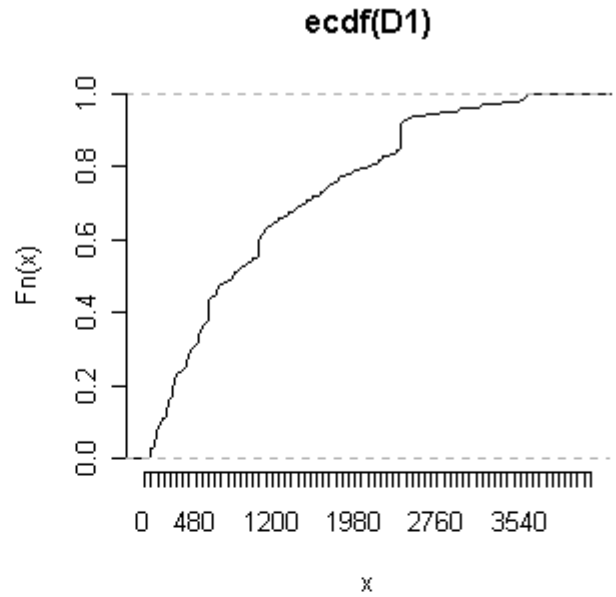
ecdf(D1)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1200 1980 2760 3540
x

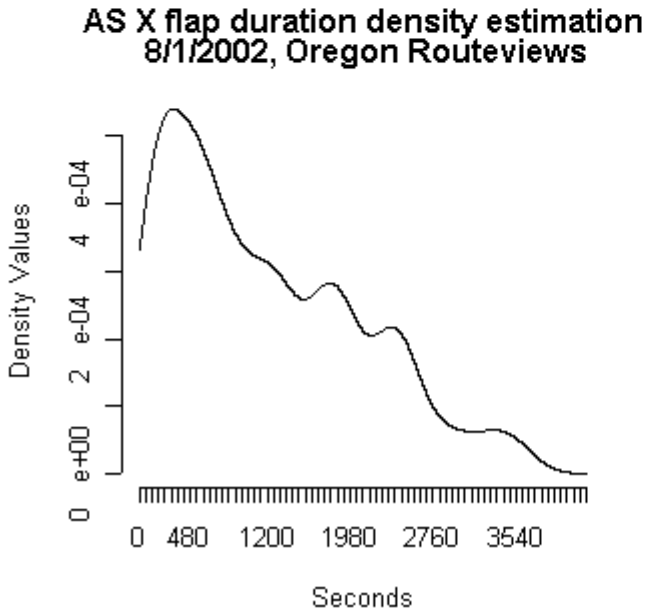
AS X flap duration density estimation
8/1/2002, Oregon Routeviews
Density Values
0 e+00 2 e-04 4 e-04
0 480 1200 1980 2760 3540
Seconds

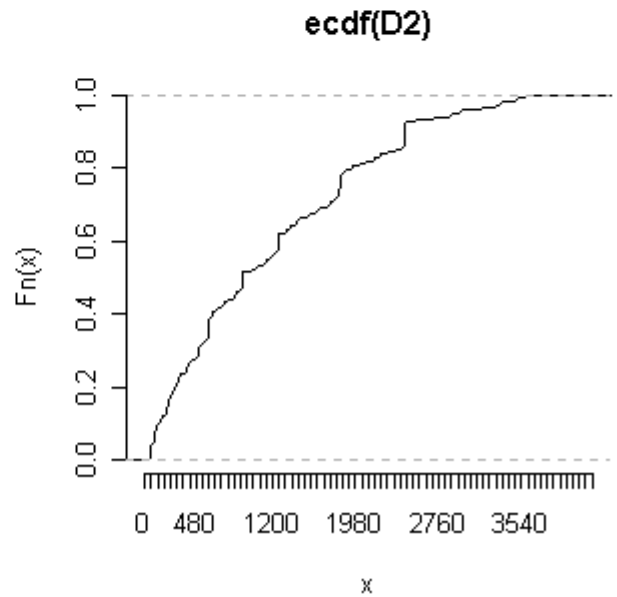
ecdf(D2)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1200 1980 2760 3540
x

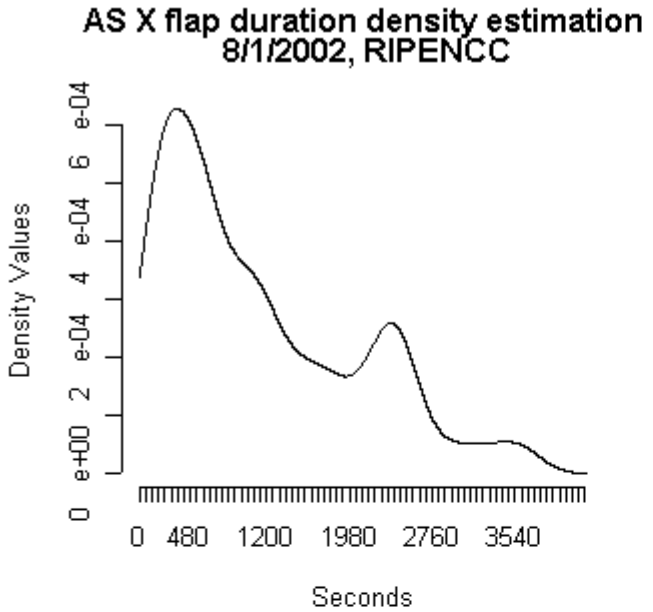
AS X flap duration density estimation
8/1/2002, RIPENCC
Density Values
0 e+00 2 e-04 4 e-04 6 e-04
0 480 1200 1980 2760 3540
Seconds

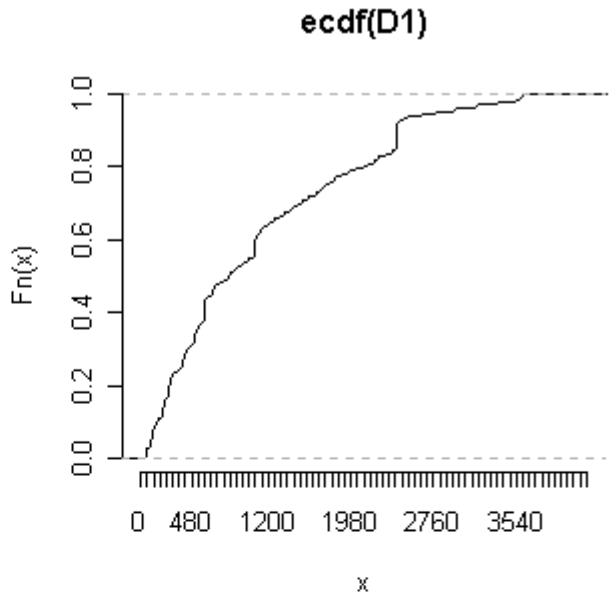
ecdf(D1)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1200 1980 2760 3540
x

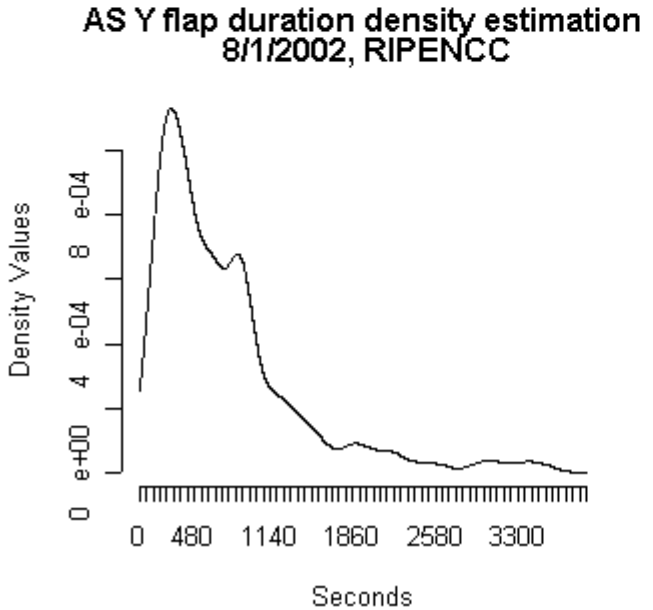
AS Y flap duration density estimation
8/1/2002, RIPENCC
Density Values
0 e+00 4 e-04 8 e-04
0 480 1140 1860 2580 3300
Seconds

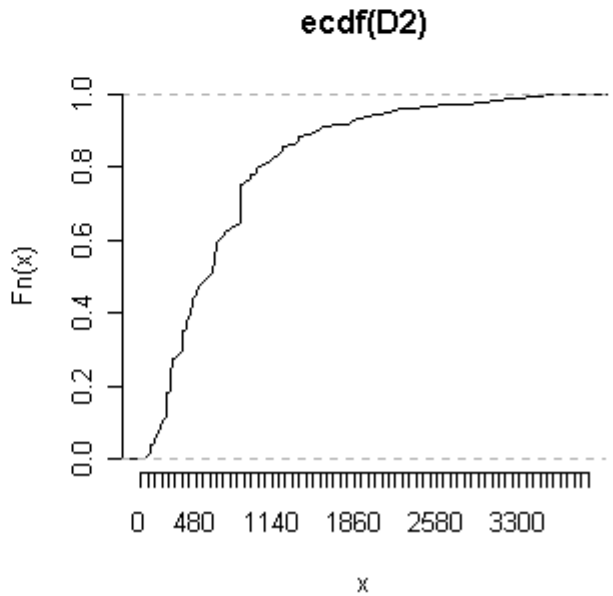
ecdf(D2)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1140 1860 2580 3300
x

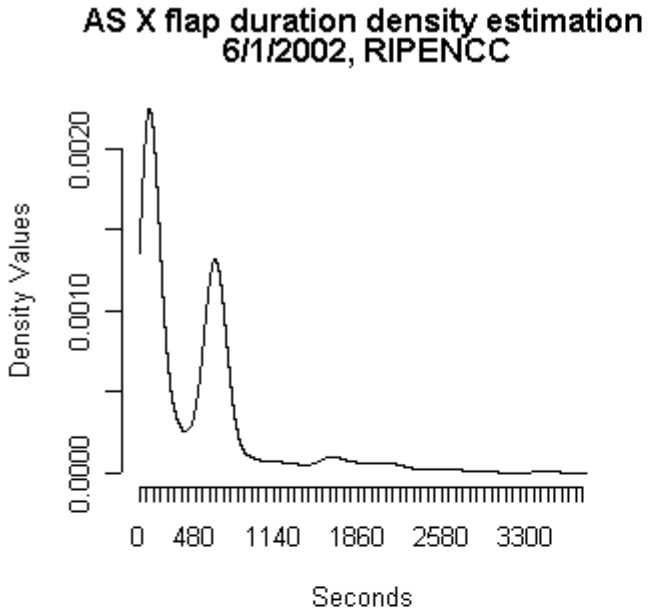
AS X flap duration density estimation
6/1/2002, RIPENCC
Density Values
0.0000
0.0010
0.0020
0 480 1140 1860 2580 3300
Seconds

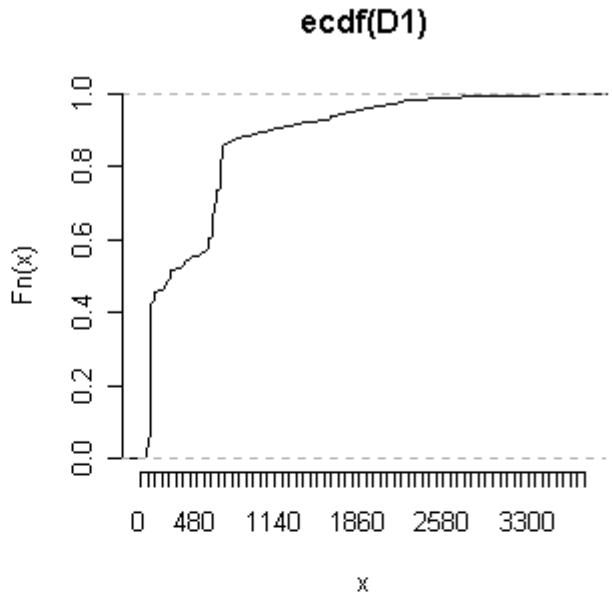
ecdf(D1)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1140 1860 2580 3300
x

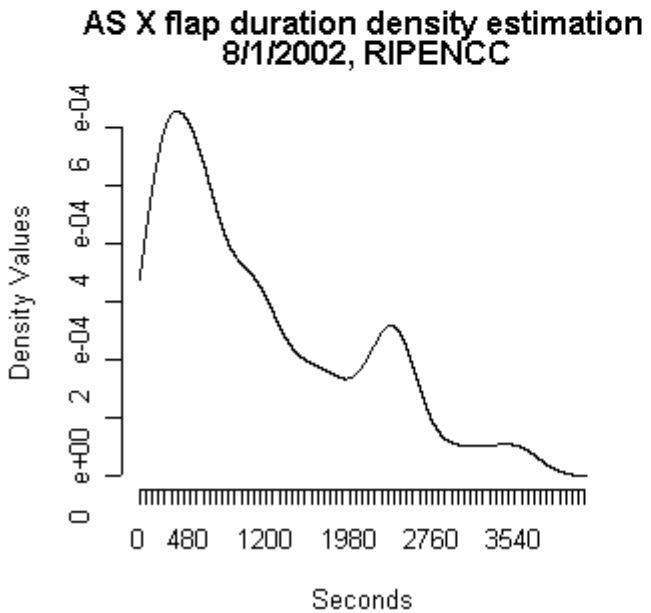
AS X flap duration density estimation
8/1/2002, RIPENCC
Density Values
0 e+00 2 e-04 4 e-04 6 e-04
0 480 1200 1980 2760 3540
Seconds

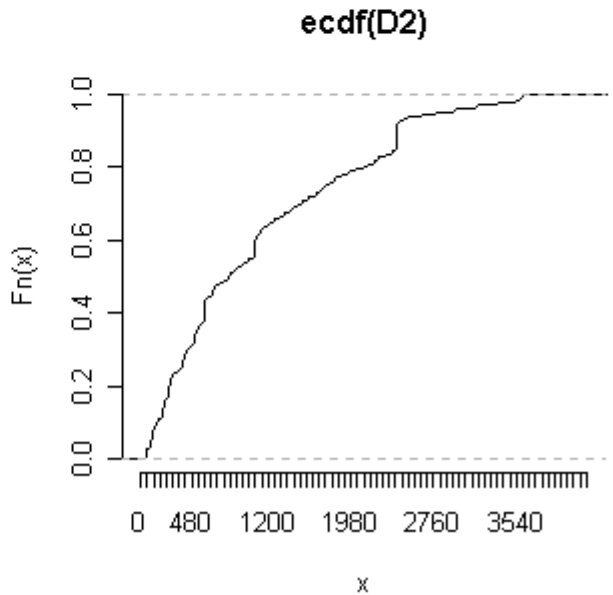
ecdf(D2)
Fn(x)
0.0 0.2 0.4 0.6 0.8 1.0
0 480 1200 1980 2760 3540
x

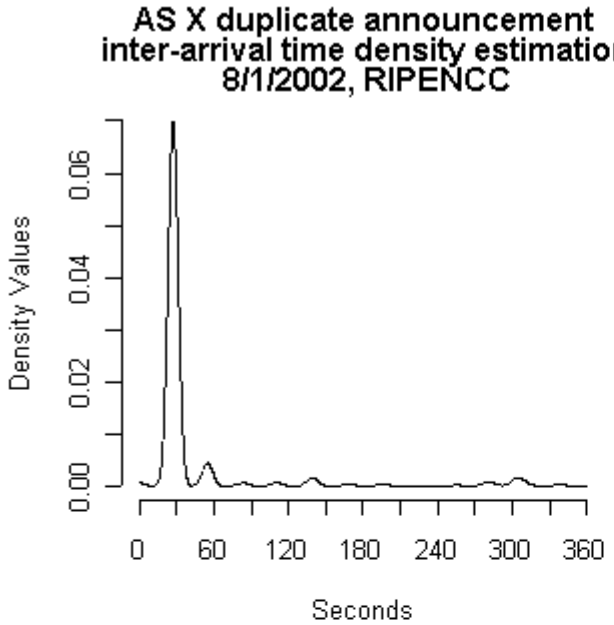
AS X duplicate announcement
inter-arrival time density estimation
8/1/2002, RIPENCC
Density Values
0.00
0.02
0.04
0.06
0
60
120
180
240
300
360
Seconds

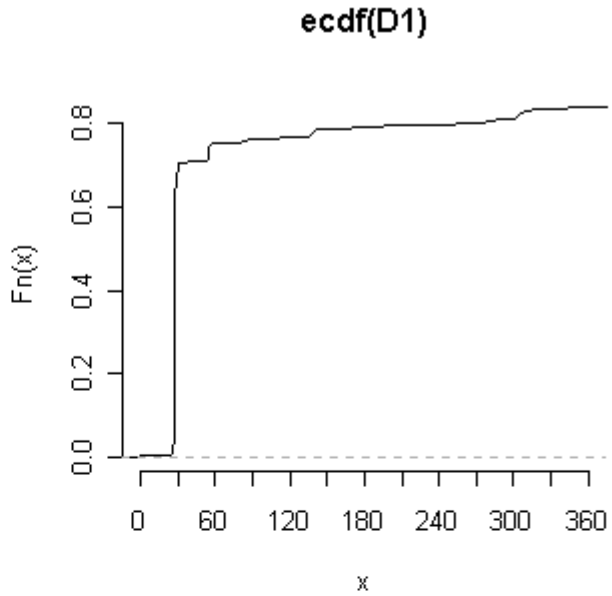
ecdf(D1)
Fn(x)
0.0
0.2
0.4
0.6
0.8
0
60
120
180
240
300
360
x

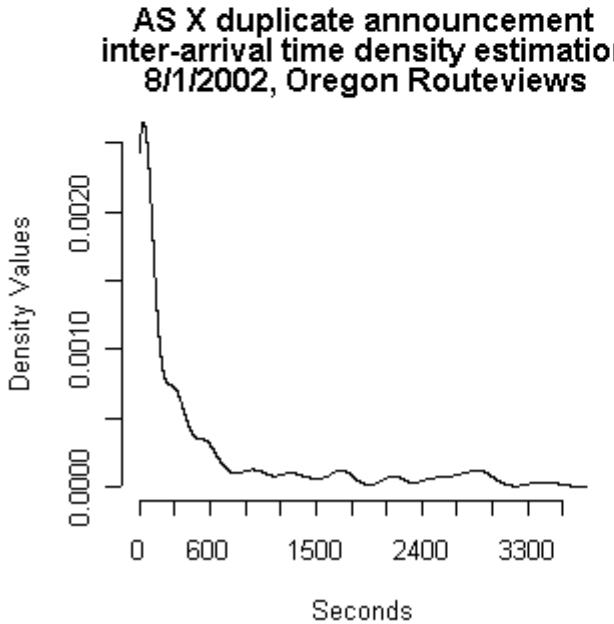
AS X duplicate announcement
inter-arrival time density estimation
8/1/2002, Oregon Routeviews
Density Values
0.0000
0.0010
0.0020
0
600
1500
2400
3300
Seconds

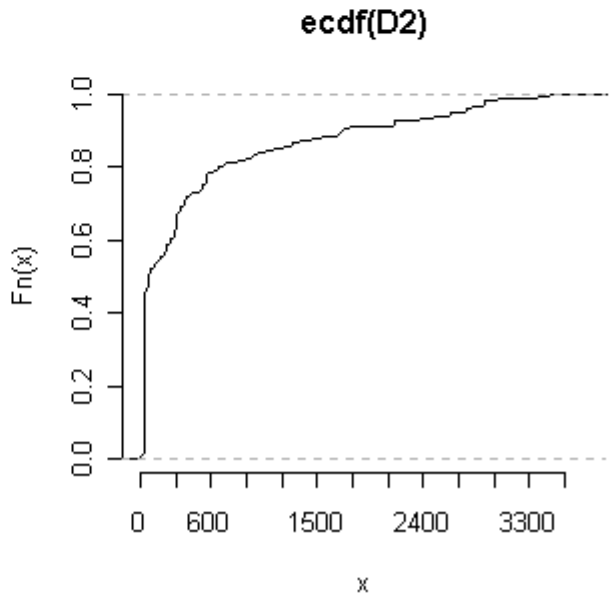
ecdf(D2)
Fn(x)
0.0
0.2
0.4
0.6
0.8
1.0
0
600
1500
2400
3300
x

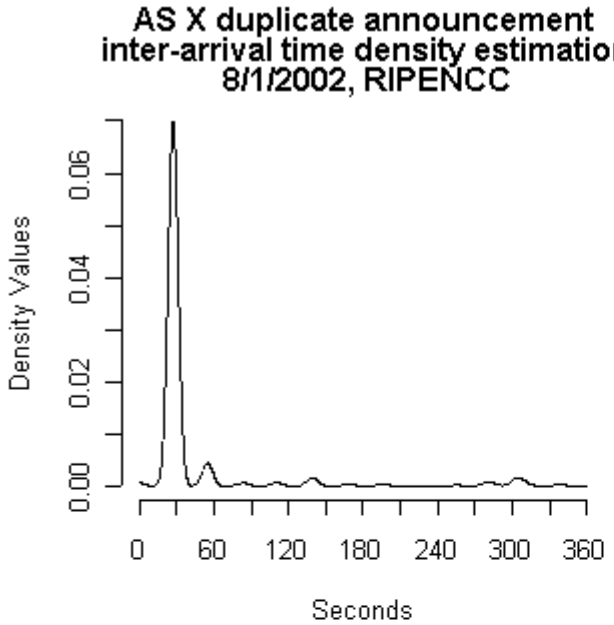
AS X duplicate announcement
inter-arrival time density estimation
8/1/2002, RIPENCC
Density Values
0.00
0.02
0.04
0.06
0
60
120
180
240
300
360
Seconds

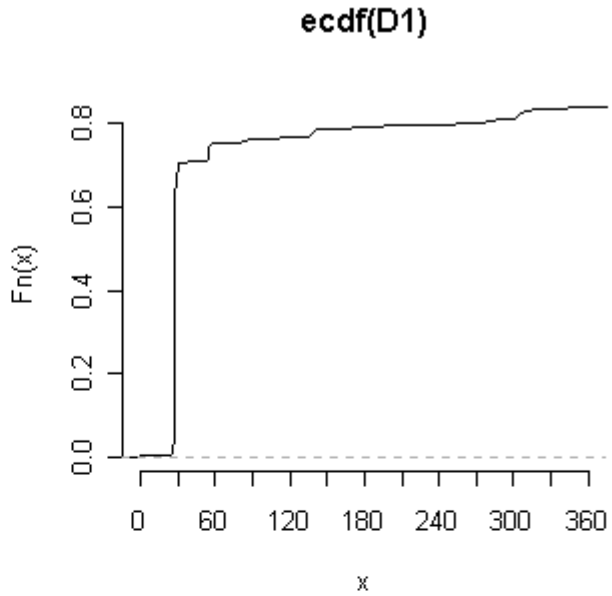
ecdf(D1)
Fn(x)
0.0
0.2
0.4
0.6
0.8
0
60
120
180
240
300
360
x

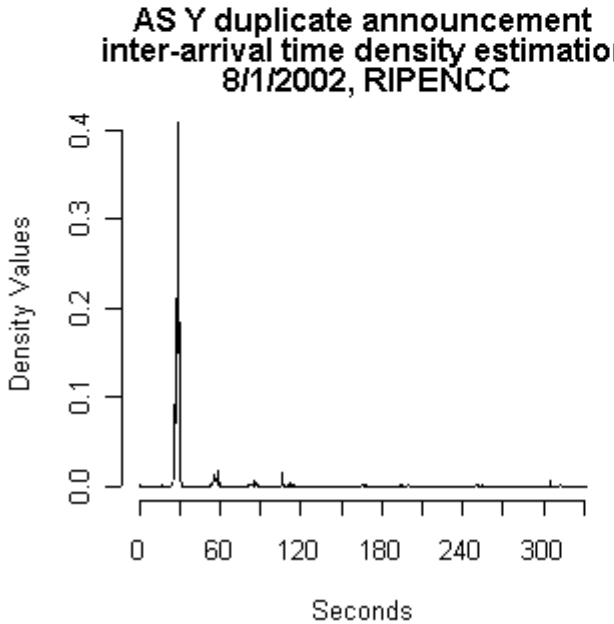
AS Y duplicate announcement
inter-arrival time density estimation
8/1/2002, RIPENCC
Density Values
0.0
0.1
0.2
0.3
0.4
0
60
120
180
240
300
Seconds

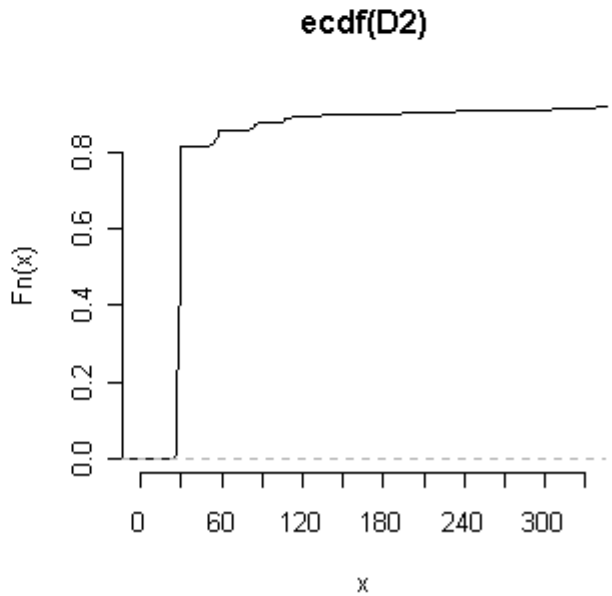
ecdf(D2)
Fn(x)
0.0
0.2
0.4
0.6
0.8
0
60
120
180
240
300
x

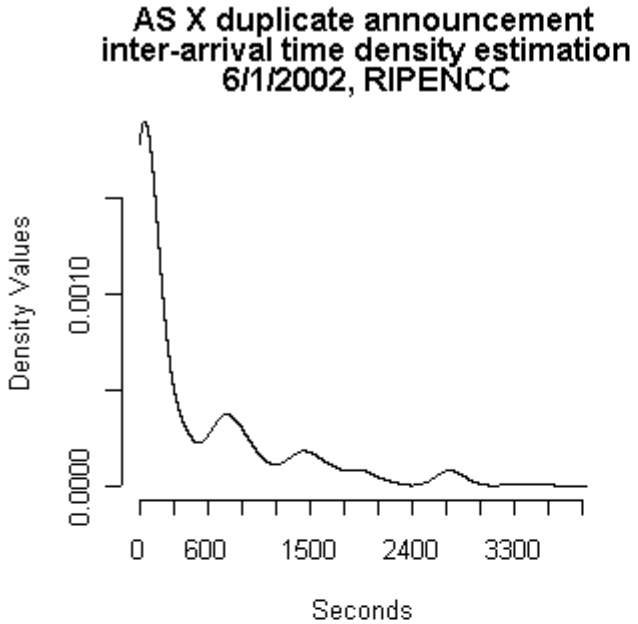
AS X duplicate announcement
inter-arrival time density estimation
6/1/2002, RIPENCC
Density Values
0.0000
0.0010
0
600
1500
2400
3300
Seconds

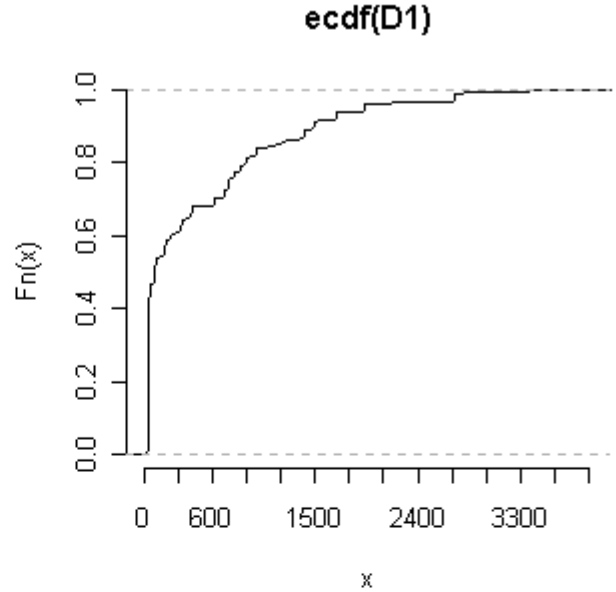
ecdf(D1)
Fn(x)
0.0
0.2
0.4
0.6
0.8
1.0
0
600
1500
2400
3300
x

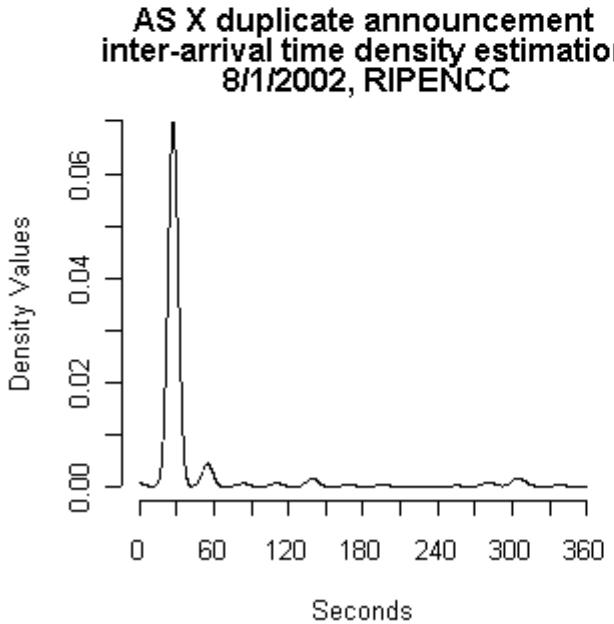
AS X duplicate announcement
inter-arrival time density estimation
8/1/2002, RIPENCC
Density Values
0.00
0.02
0.04
0.06
0
60
120
180
240
300
360
Seconds

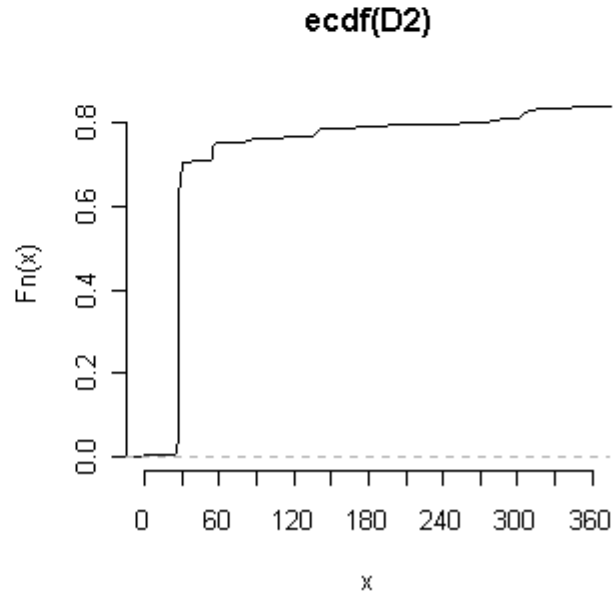
ecdf(D2)
Fn(x)
0.0
0.2
0.4
0.6
0.8
0
60
120
180
240
300
360
x

Updates, Announcements, AADups & Flaps per Day
AMS-IX Peer
8/1/2001 - 1/31/2002

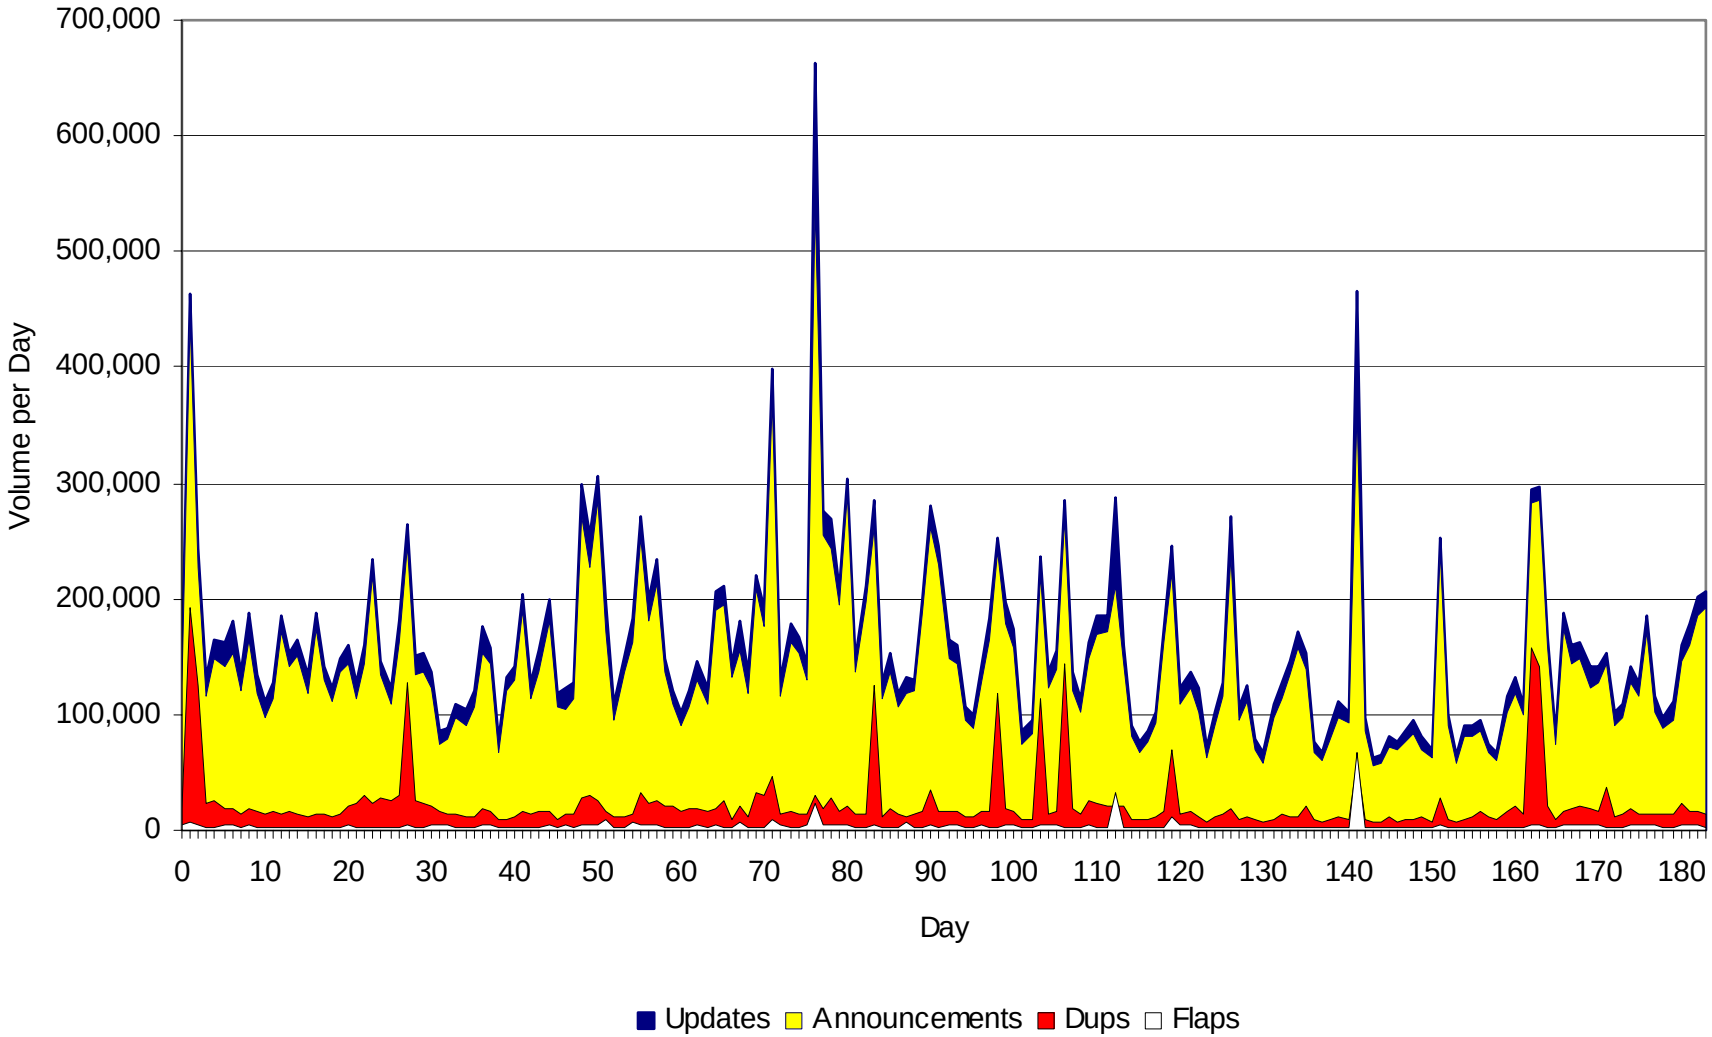

Updates, Announcements, AADups, Flaps per Day
LINX Peer
8/1/2001 - 1/31/2002

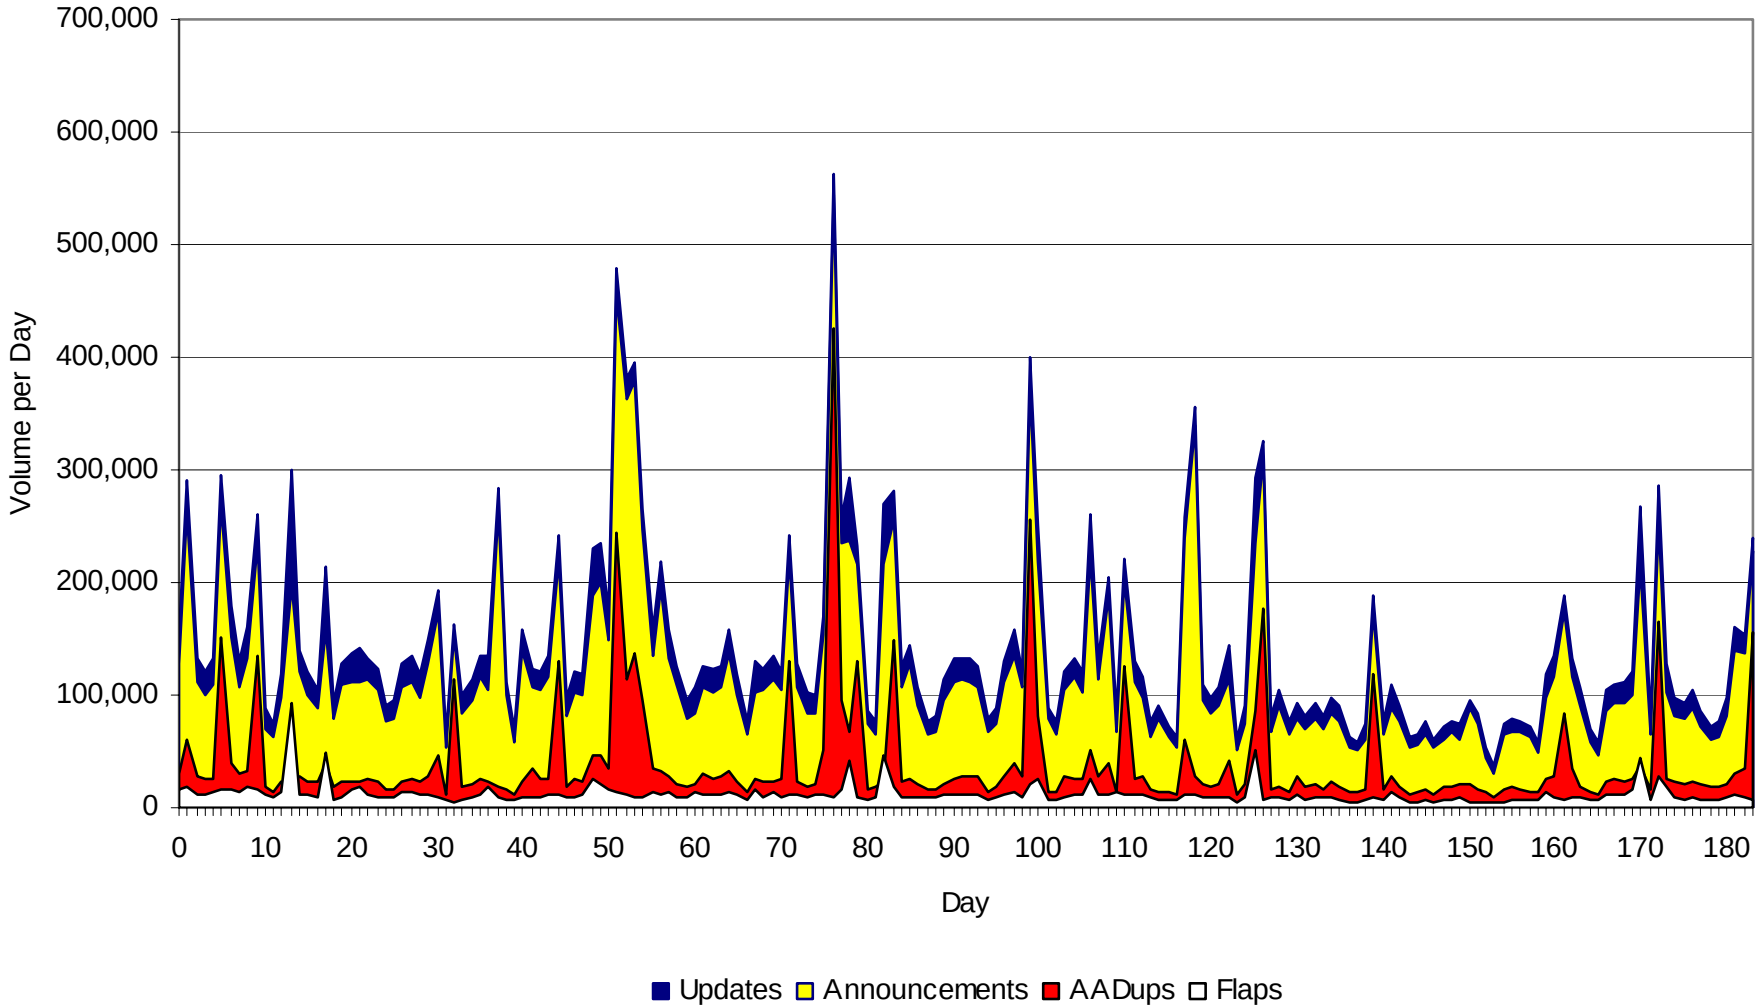

Update Inter-Arrival Time Distribution
LINX Peer
8/1/2001 - 1/31/2002

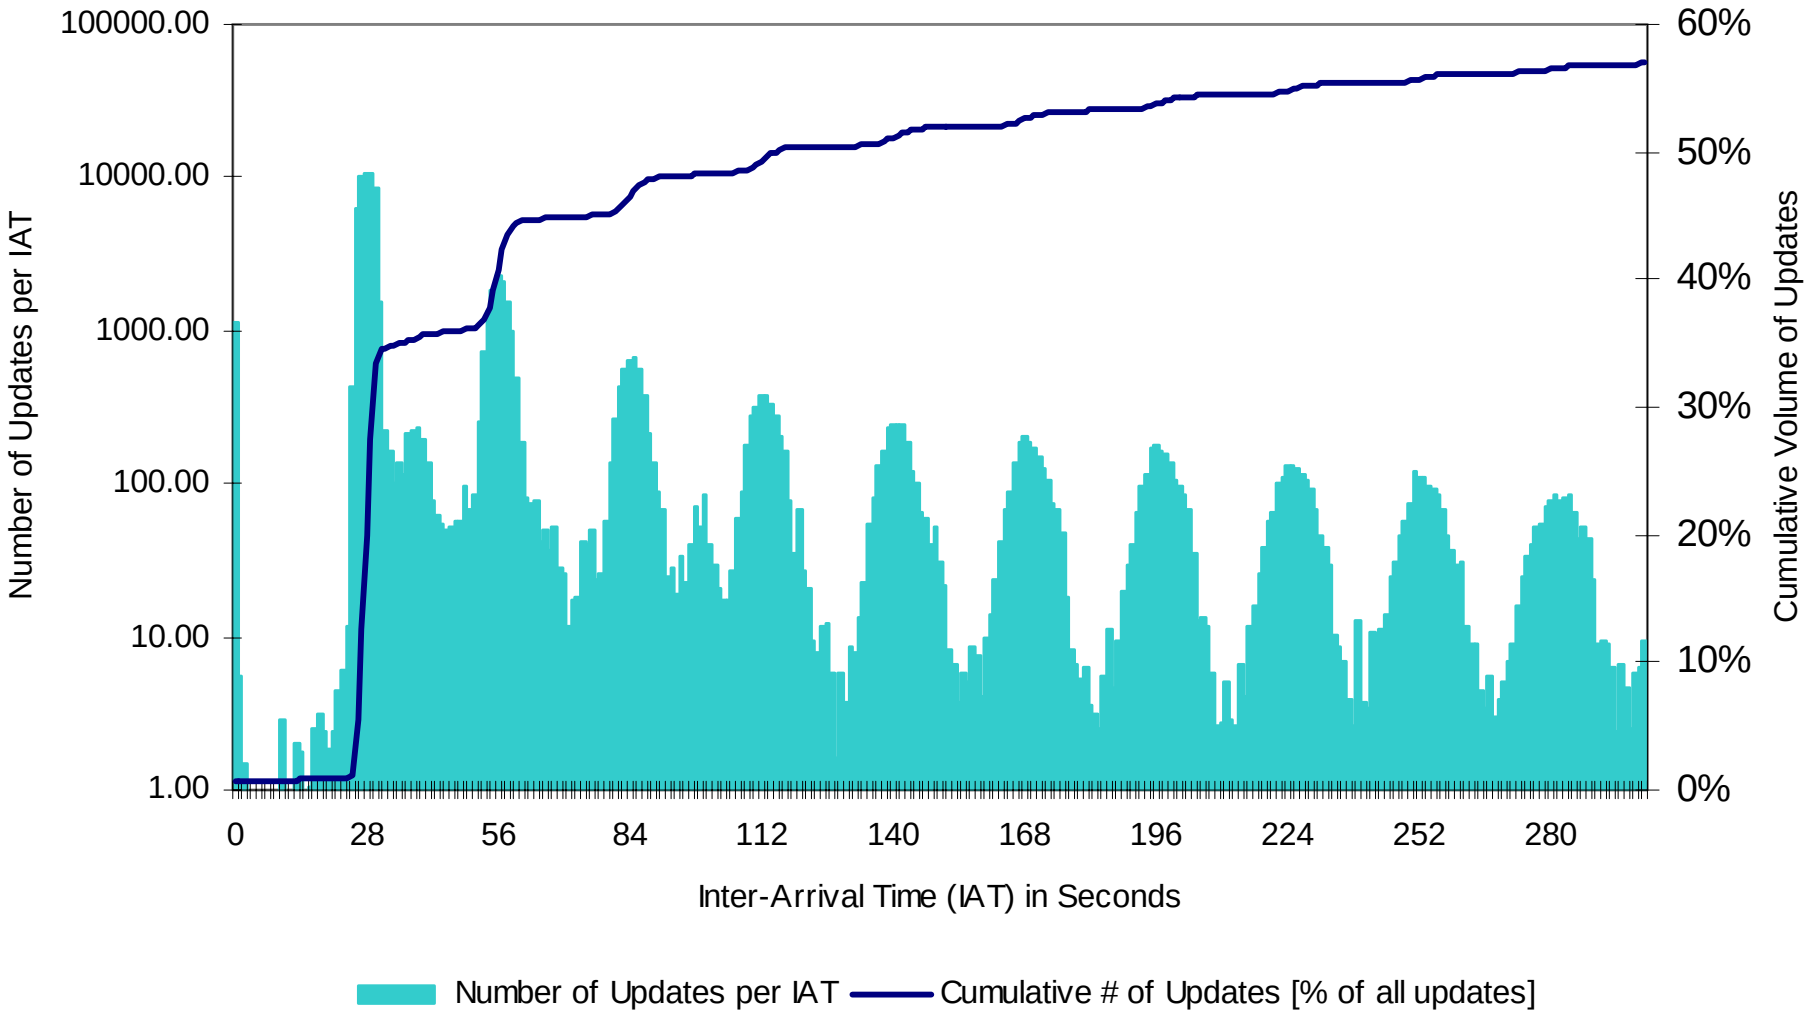

Announcement Inter-Arrival Time Distribution
LINX Peer
8/1/2002 - 1/31/2002

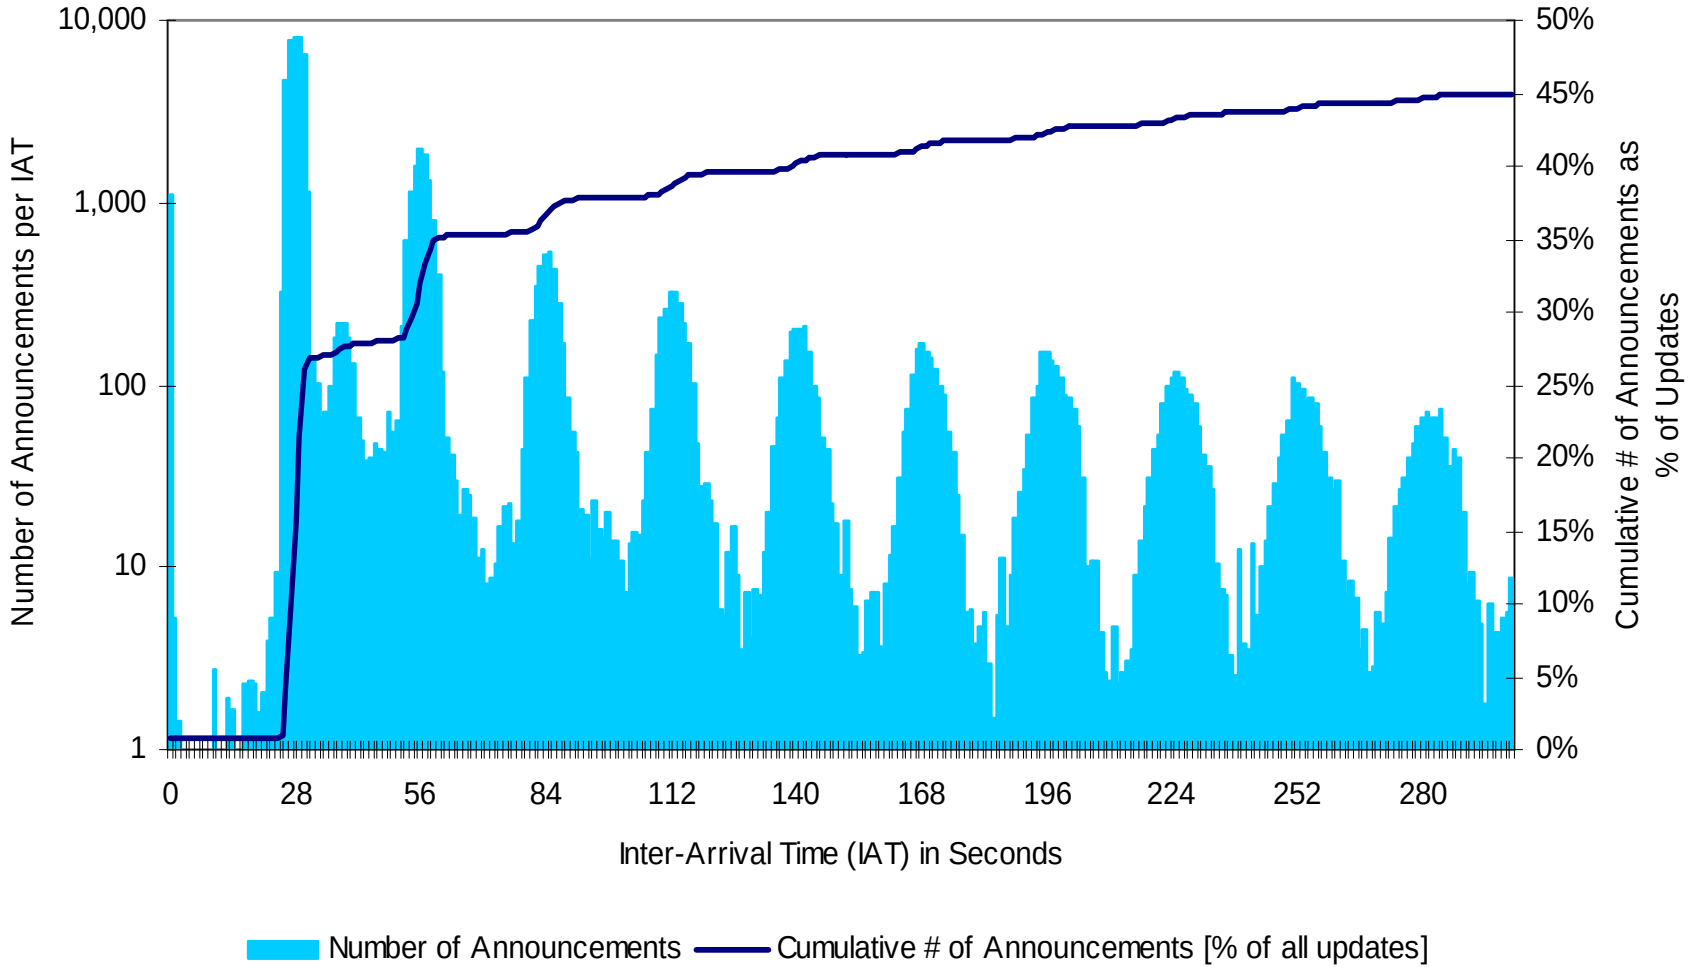

Flap Duration Distribution
LINX Peer
8/1/2001 - 1/31/2002