On the Impact of Bursting on TCP Performance

PAM 2005

Boston, MA

Ethan Blanton (Purdue Univ.)
Mark Allman (ICIR)

Background

- Much conversation over the years about a *micro-burst* limiter for TCP

- A number of proposals on how to handle micro-bursts
  - are these required?

- No real measurement data on:
  - how often micro-bursts occur
  - how often micro-bursts cause problems
  - what causes micro-bursts

Background (cont.)

- Micro-bursts:
  - data segments liberated by a single ACK packet
  - presumably sent into the network back-to-back

- TCP's expected micro-burstiness is 1-3 segments per ACK
  - due to window sliding and window growth

- TCP *can* send larger micro-bursts due to network dynamics
  - e.g., ACK loss

Example #1

```
0     ACK 100
1     DATA 200
2     DATA 201
3     ACK 102
4     DATA 202
5     DATA 203
6     ACK 104
7     DATA 204
8     DATA 205
9     DATA 206
      [...]
```

Example #2

```
0     ACK 100
1     DATA 200
2     DATA 201
3     (ACK 102)  (dropped)
4
5
6     ACK 104
7     DATA 202
8     DATA 203
9     DATA 204
10    DATA 205
11    DATA 206
      [...]
```

Micro-Bursting Causes

- ACK loss
- Application sending patterns
- ACK reordering
- Stretch ACK generation
- Advertised window limit during loss recovery

Alternate Bursting

- TCP can also send *macro-bursts*
  - e.g., due to slow start
  - e.g., due to ACK compression

- We do not consider these bursts in this paper

Scope

- We also scope on considering the impact of micro-bursting on the connection itself

- I.e., our data sheds no light on the impact of bursting on cross-traffic

Measurement Methodology

- Collect sender-side packet traces of web traffic from 3 networks at 4 different times
  - Anon: 7/24/03, 26 hours, 295K connections
  - LBNL: 10/22/03, 11 hours, 196K connections
  - ICSI1: 1/4/04, 14 days, 224K connections
  - ICSI2: 9/18/04, 14 days, 199K connections

- At most 2% of connections not analyzed due to *craziness*

Measurement Methodology (cont.)

- Micro-burst definition:
  - at least 4 data segments sent between two subsequent ACK arrivals

- Custom TCP analysis software
  - assesses the causality between ACKs and data segments
    - *actually seemingly impossible!*
  - tracks the size of micro-bursts
  - tracks whether the micro-bursts experience loss
  - determines the root cause of each micro-burst

Example #1 (revisited)

```
0     ACK 100
1     DATA 200
2     DATA 201
3     ACK 102
4     DATA 202
5     DATA 203
6     ACK 104
7     DATA 204
8     DATA 205
9     DATA 206
      [...]
```

Example #3

```
0     ACK 100
1     ACK 102
2     DATA 200
3     DATA 201
4     DATA 202
5     DATA 203
6     ACK 104
7     DATA 204
8     DATA 205
9     DATA 206
    [...]
```

Measurement Methodology (cont.)

- We cope by *accumulating ACKs*
  - each ACK arrival allows the transmission of up to 3 segments
  - *this is potentially bogus*

- No sound way to assess the size of bursts if and when they occur
  - so, we do not try
  - *this is potentially bogus*

- Rather, we wave our hands ...

Measurement Methodology (cont.)

- Measurement is about *insight not numbers*
- The *numbers* from our analysis are nearly surely not quite right
- The implications of bursts [N-m,N+m] are roughly the same
  - therefore, we believe the *insights* derived from our data are sound
- Not an exact science here

Burst Characteristics

Burst length in packets:

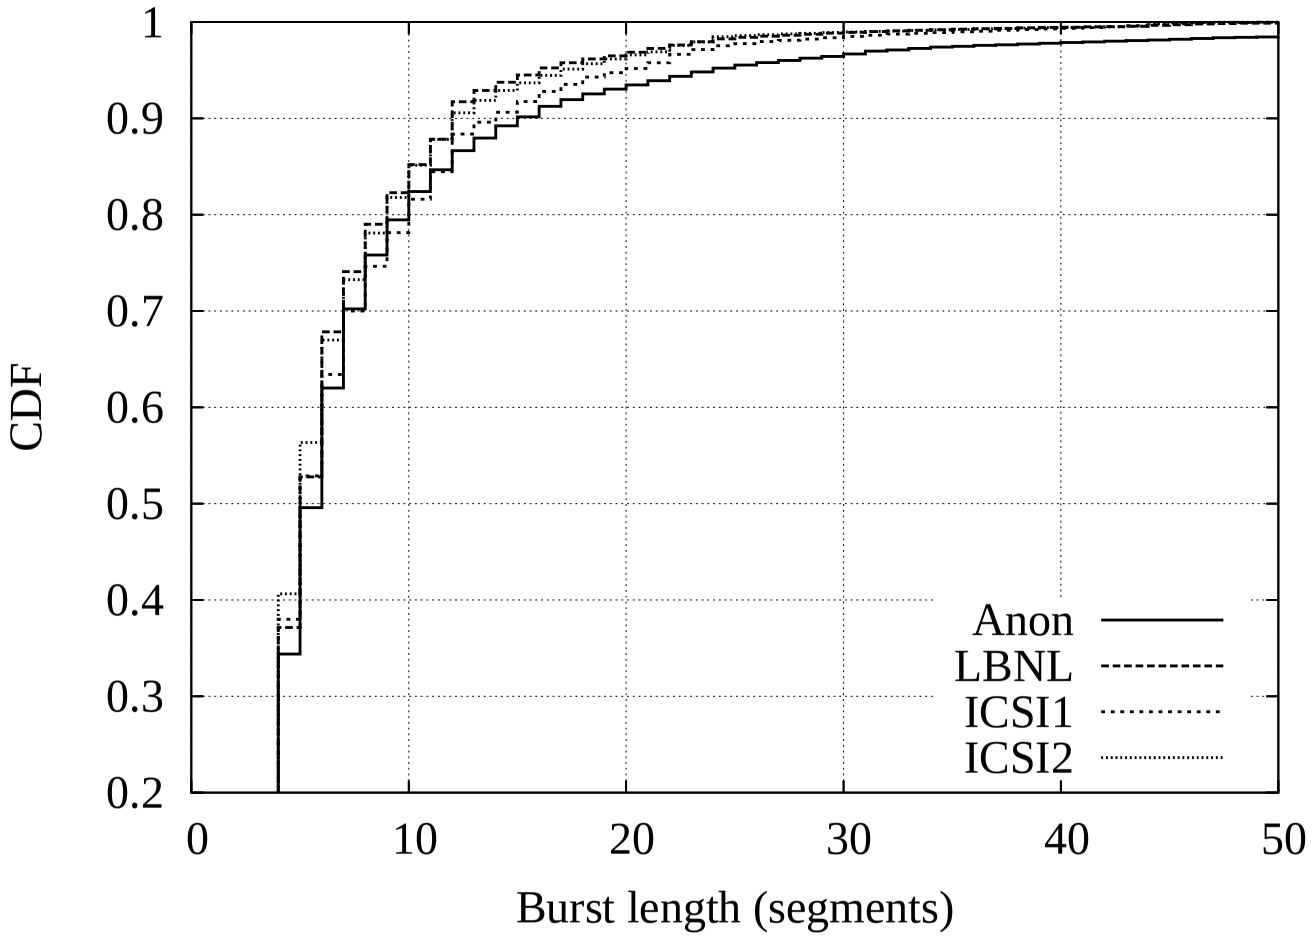

Burst Characteristics (cont.)

Burst length in bytes:

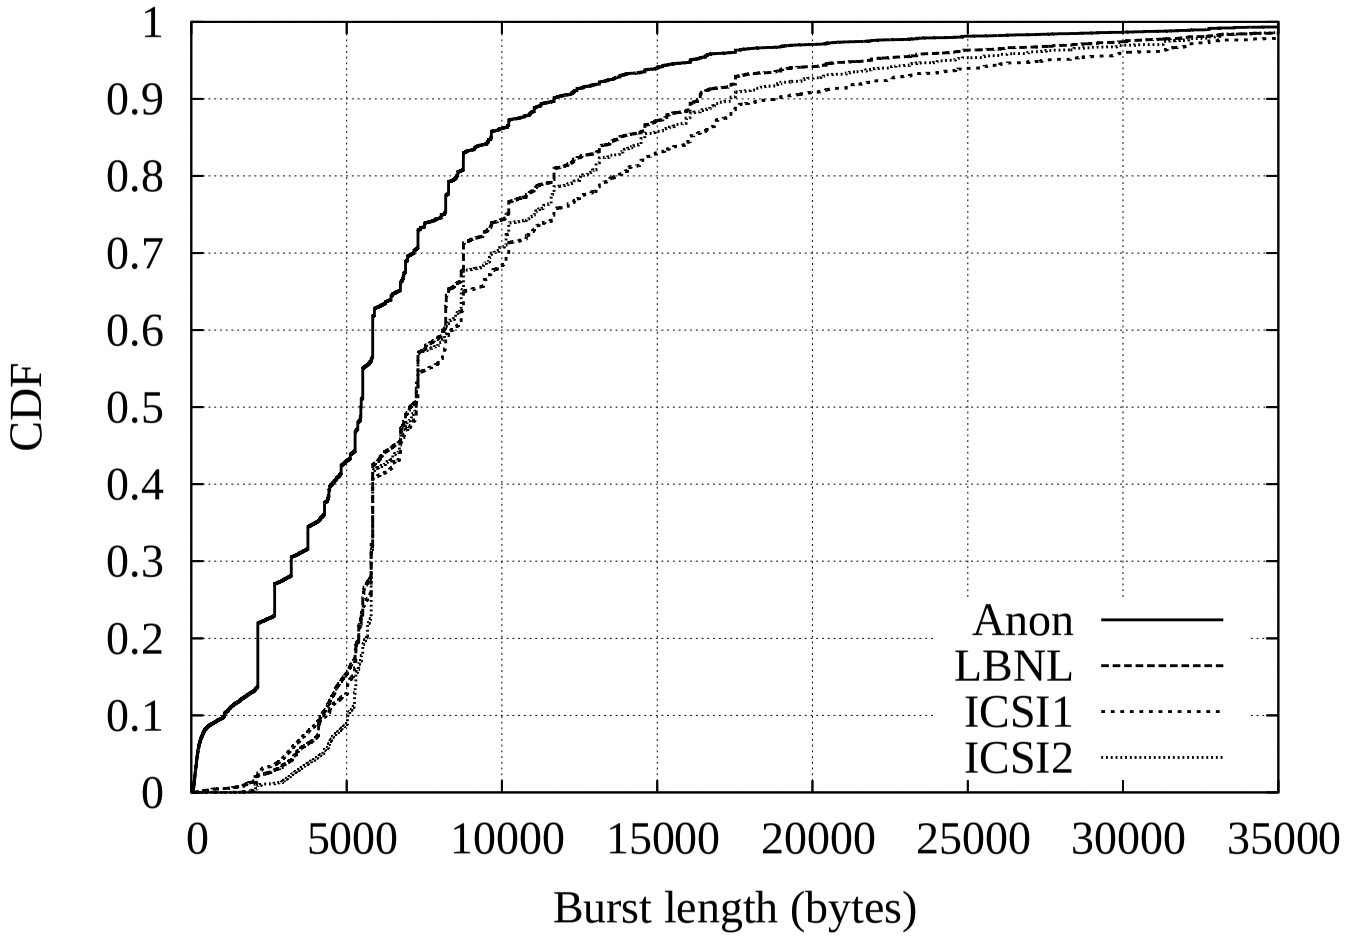

Burst Characteristics (cont.)

Bursts per connection:

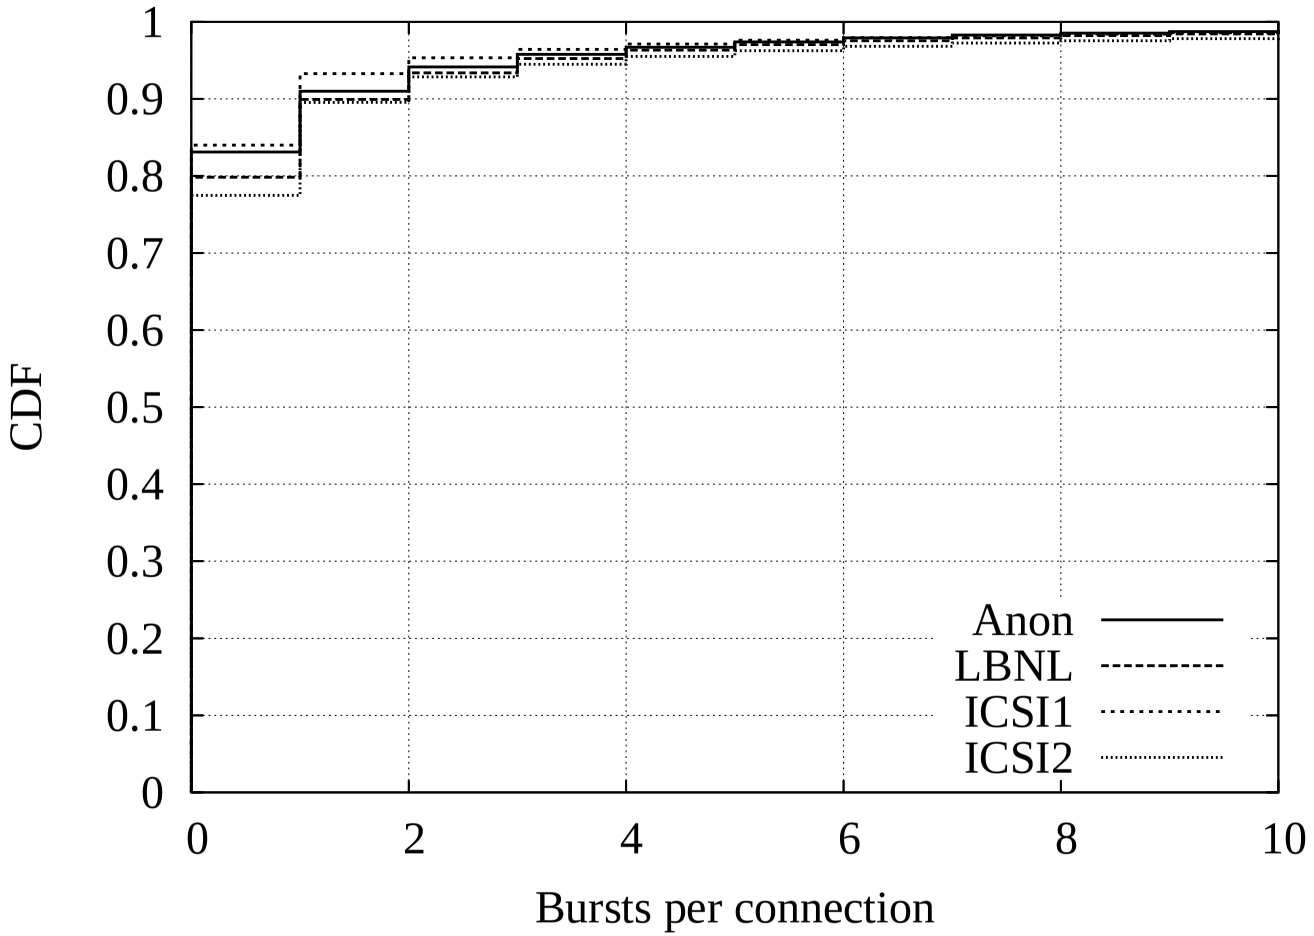

Implications

Chance of losing at least 1 segment vs. burst size (LBNL):

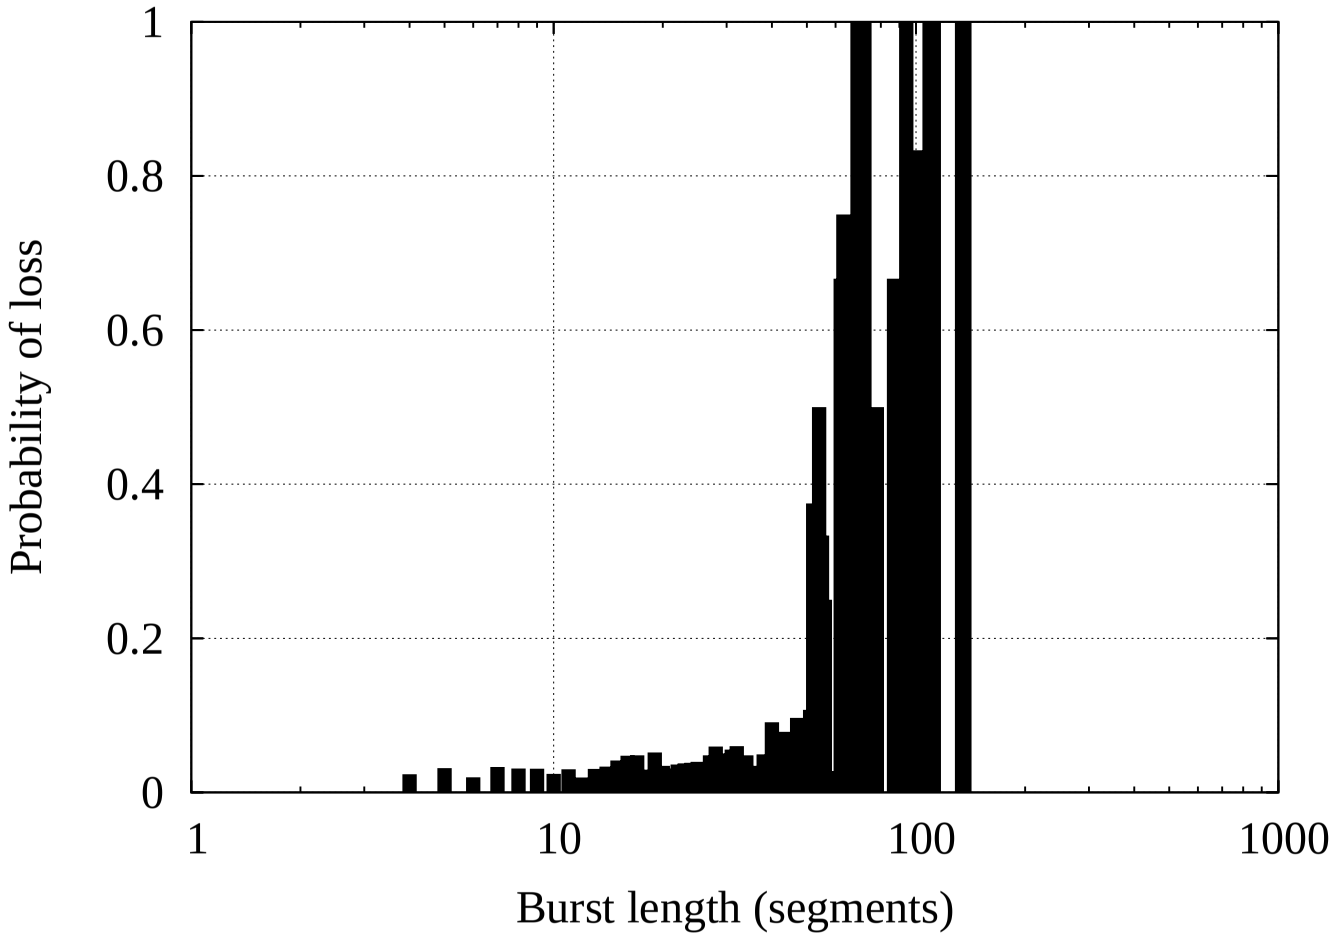

Implications (cont.)

Chance of losing at least 1 segment vs. burst size (Anon):

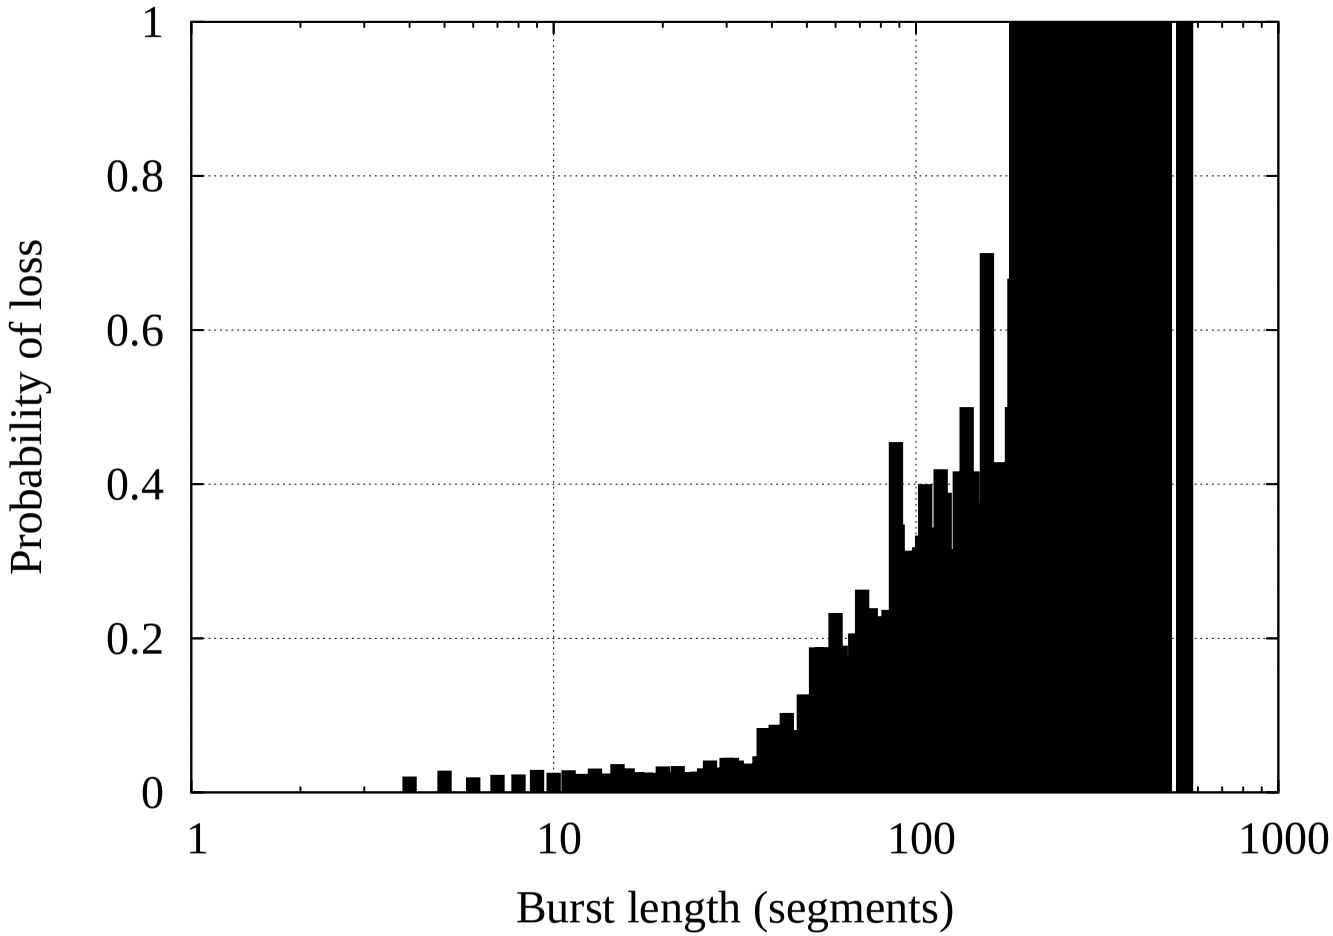

Implications (cont.)

- Loss rates inside and outside bursts
  - scoped on connections with both bursts and loss

Dataset	Conns.	Inside (%)	Outside (%)
Anon	4,233	70.9	41.2
LBNL	5,685	23.0	22.9
$ICSI_1$	4,805	16.1	14.5
$ICSI_2$	8,201	26.6	20.5

- Loss rates appear high because of our winnowing

Implications (cont.)

- Root causes of bursts:

Dataset	Bursts	Initial Window	Exit Loss Recovery	Stretch ACKs	Window Opening	App. Pattern	Unknown
Anon	274,880	1.8	0.2	26.3	5.0	17.0	49.6
LBNL	187,176	0.9	0.3	22.9	3.1	32.8	40.0
$ICSI_1$	165,023	6.4	0.7	23.5	4.8	24.0	40.6
$ICSI_2$	228,063	4.2	5.1	22.4	4.5	23.3	45.1

- Predominant causes:
  - stretch ACKs
  - application patterns

- A large number of bursts are of undetermined cause ....

Conclusions

- Micro-bursting does not happen very often (with 75% of connections not bursting)
- When bursting happens the burst sizes are modest (<= 15 segments)
- The chance of modest bursts causing loss is small
- However, the chance of large bursts causing loss is quite high
- So, one wonders whether burst mitigation techniques are useful
  - Except that they may be useful for aggregate network traffic

Acknowledgments

- Andrew Moore and Vern Paxson provided some of the data used in our study
  - Thanks!

- Questions?