Using Netflow data for forecasting

Les Cottrell[SLAC] and Fawad Nazir[NIIT],

Presented at the CHEP06 Meeting, Mumbai India, February 2006

www.slac.stanford.edu/grp/scs/net/talk06/icfa-chep06.ppt

Partially funded by DOE/MICS for Internet End-to-end Performance Monitoring (IEPM)

Why Netflow

- Traceroute dead for dedicated paths
- Some things continue to work
  - Ping, owamp
  - Iperf, thrulay, bbftp … but

- Packet pair dispersion needs work, its time may be over

- Passive looks promising with Netflow
- SNMP needs AS to make accessible - perfSONAR
- Capture expensive
  - ~$100K (*Joerg Micheel*) for OC192Mon

Netflow

- Router/Switch identifies flow by sce/dst ports, protocol
- Cuts record for each flow:
  - src, dst, ports, protocol, TOS, start, end time
- Collect records and analyze
- Can be a lot of data to collect each day, needs lot cpu
  - Hundreds of MBytes to GBytes
- No extra traffic injected, & real: traffic, collaborators, applications
- No accounts/pwds/certs/keys
- No reservations etc
- Characterize traffic: top talkers, applications, flow lengths etc.
- Internet 2 backbone
  - http://netflow.internet2.edu/weekly/
- SLAC:
  - www.slac.stanford.edu/comp/net/slac-netflow/html/SLAC-netflow.html

Typical day's flows

- Very much a work in progress …
- Look at SLAC border
- Typical day:
  - For >100KB flows
  - ~ 28K flows/day
  - ~ 75 sites with > 100KByte bulk-data flows
  - Few hundred flows > GByte

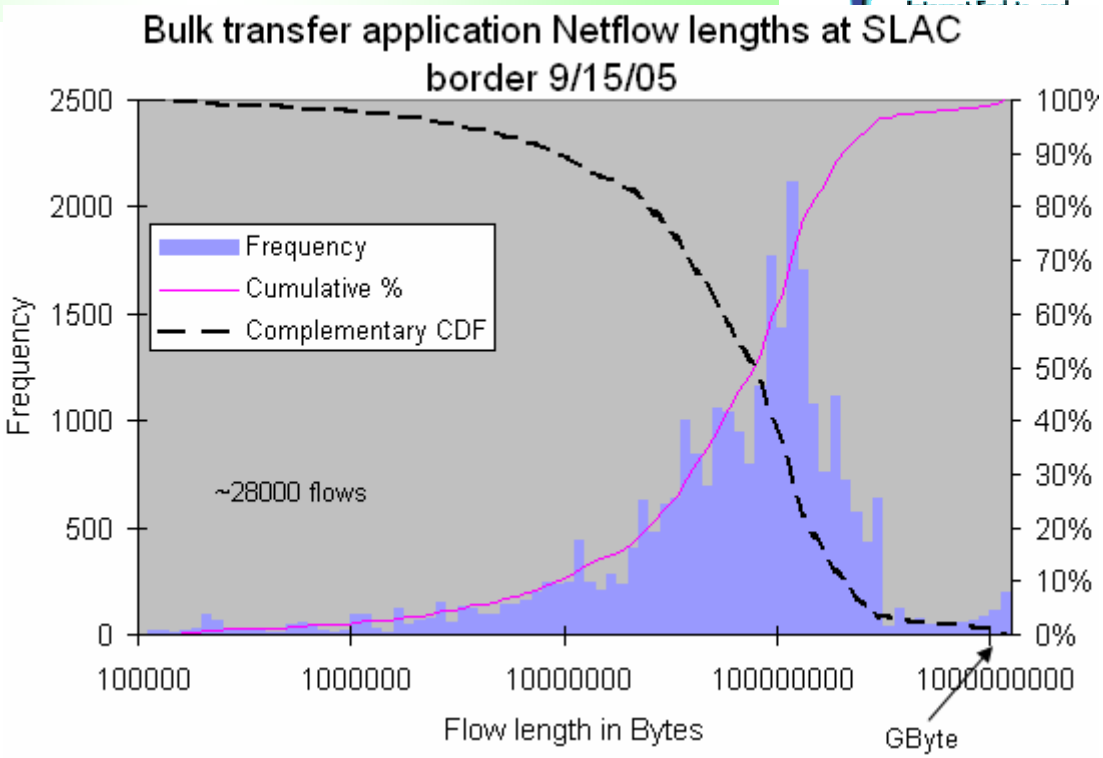

Forecasting?

- Collect records for several weeks
- Filter 40 major collaborator sites, big (> 100KBytes) flows, bulk transport apps/ports (bbcp, bbftp, iperf, thrulay, scp, ftp
- Divide by remote site, aggregate parallel streams
- Fold data onto one week, see bands at known capacities and RTTs

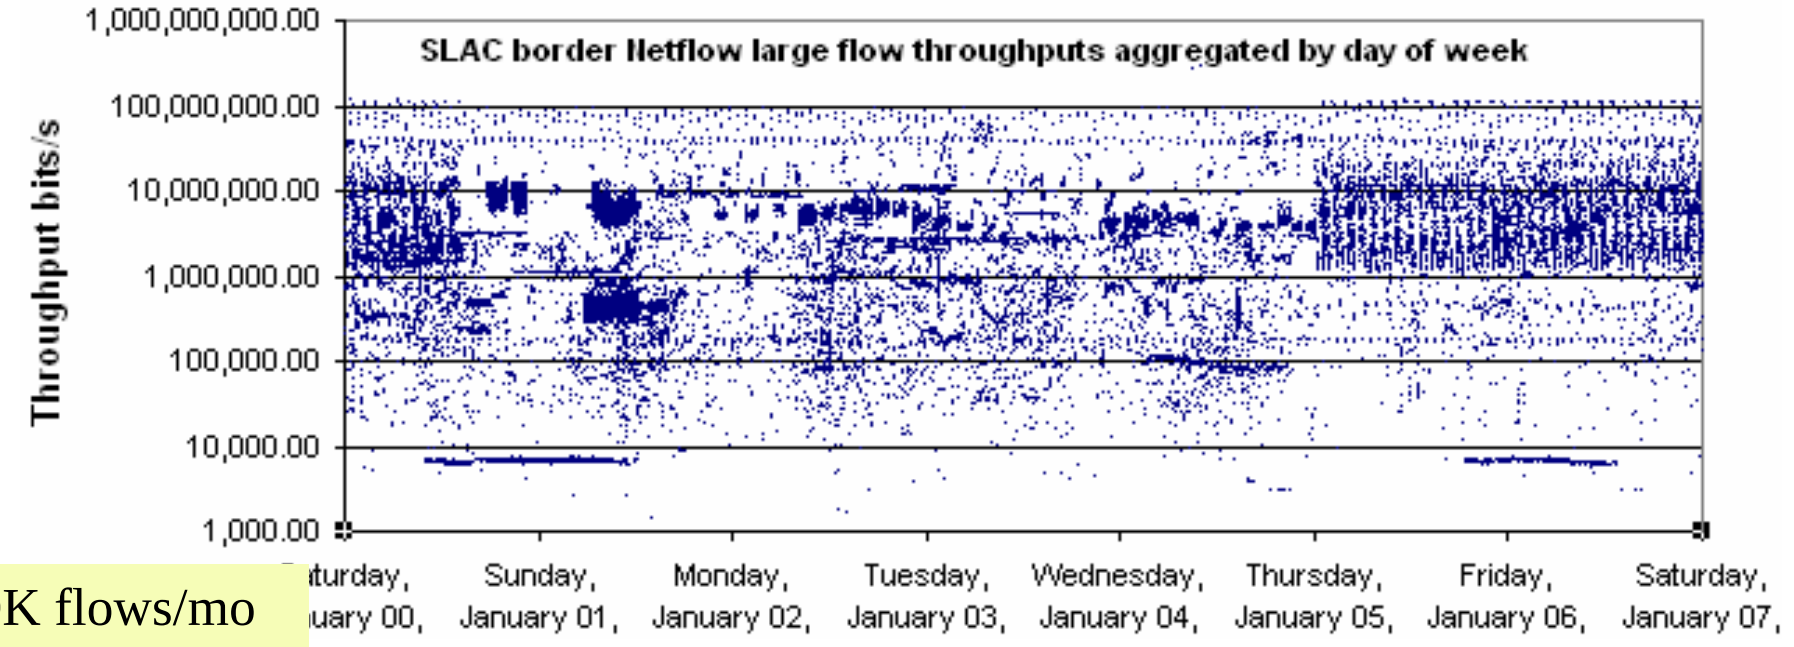

~ 500K flows/mo

Netflow et. al.

Peaks at known capacities and RTTs

RTTs might suggest windows not optimized

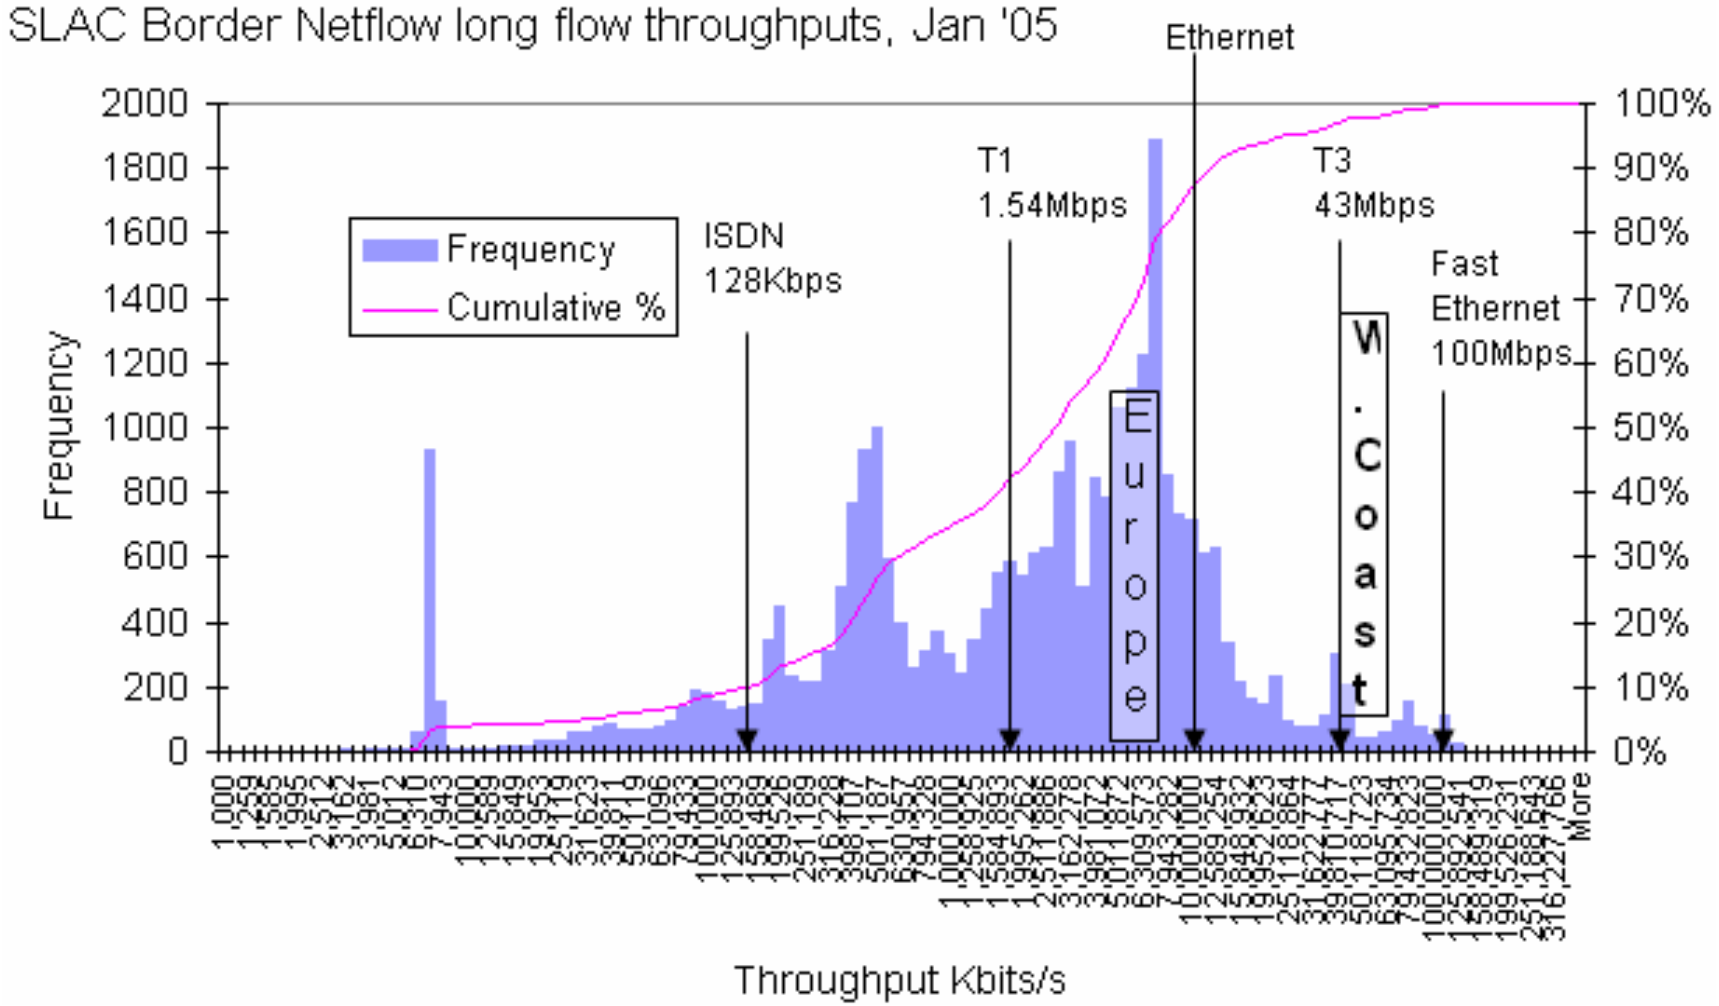

How many sites have enough flows?

- In May ’05 found 15 sites at SLAC border with > 1440 (1/30 mins) flows
  - Maybe enough for time series forecasting for seasonal effects
- Three sites (Caltech, BNL, CERN) were actively monitored
- Rest were “free”
- Only 10% sites have big seasonal effects in active measurement
- Remainder need fewer flows
- So promising

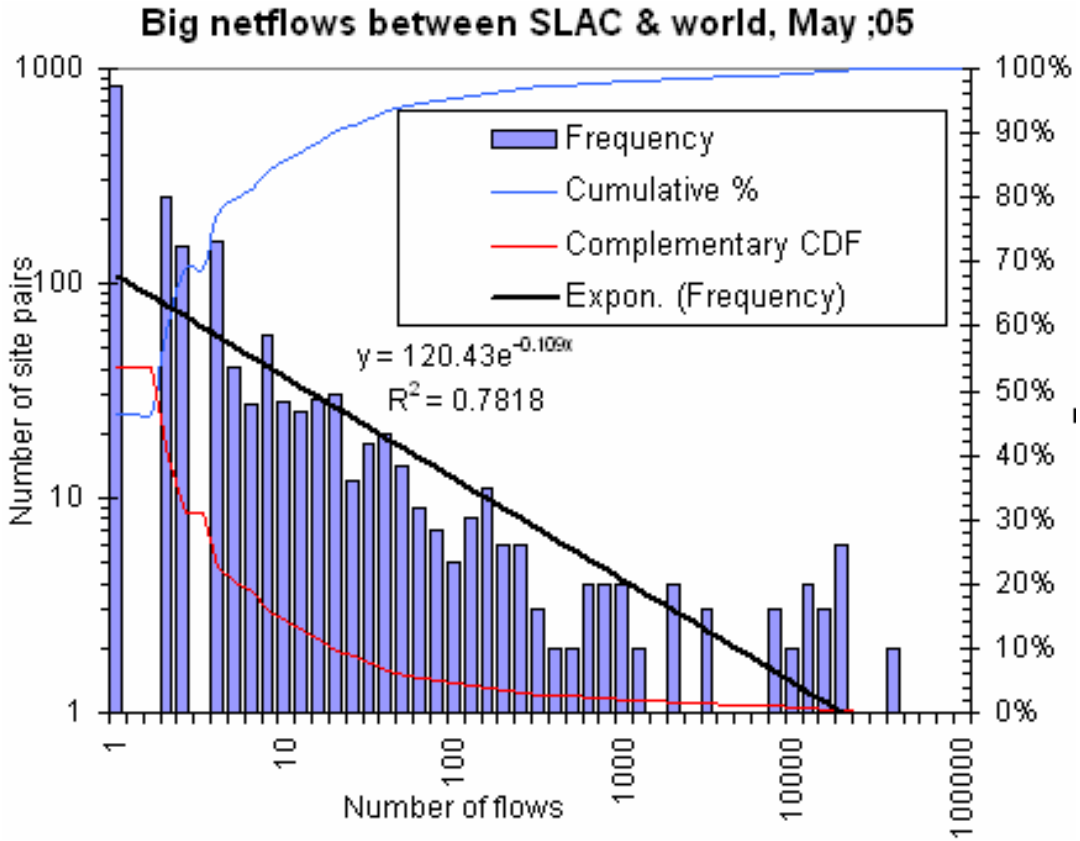

Compare active with Passive

Scatter plot: thru_active vs. thru_passive has strong correlation

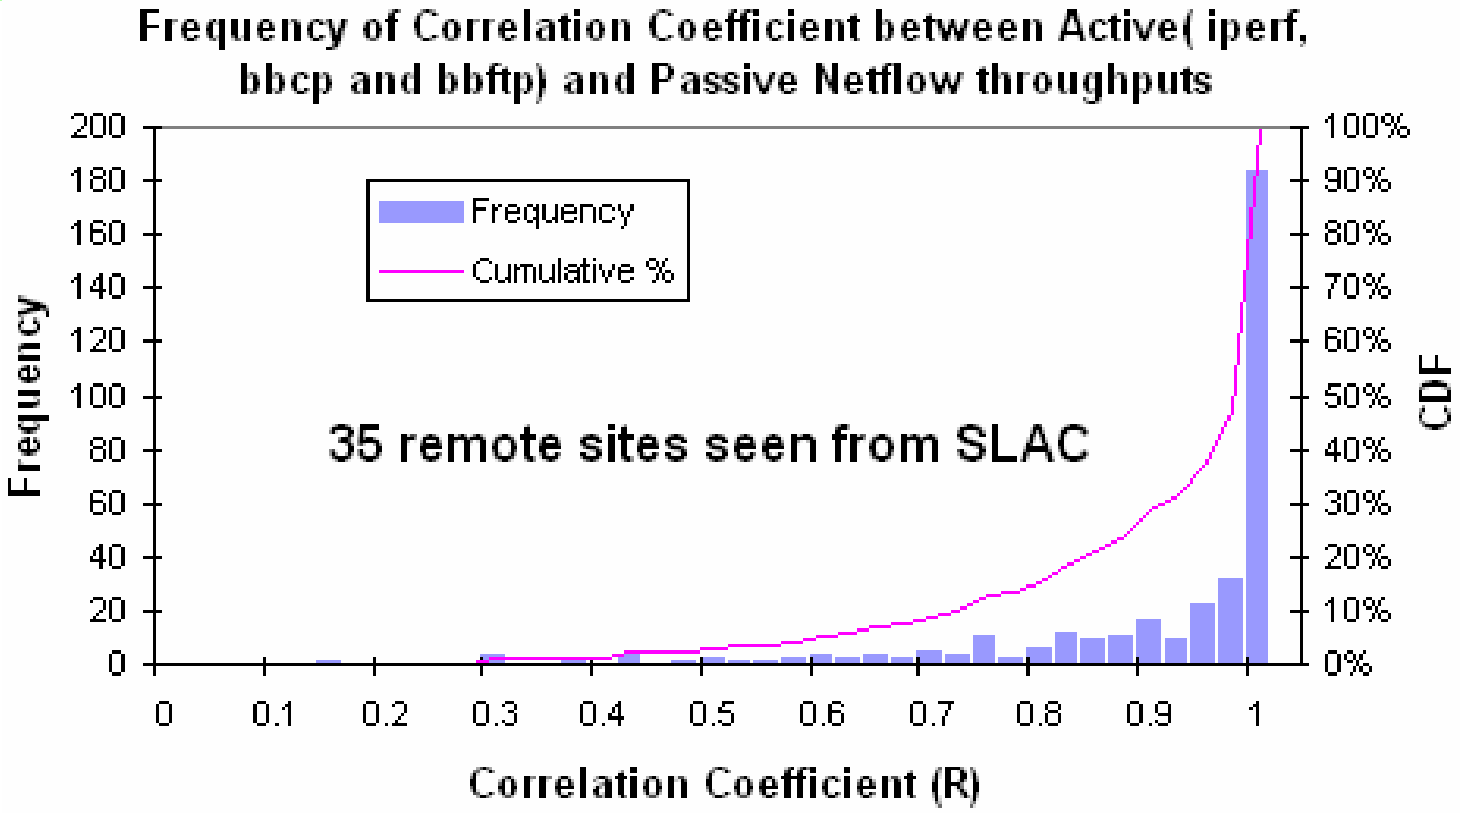

- Poor correlation usually caused by long flows
  - i.e. one stream of parallel flows lingers well after others
- See
  - www.slac.stanford.edu/comp/net/bandwidth-tests/web100/

Compare active with passive

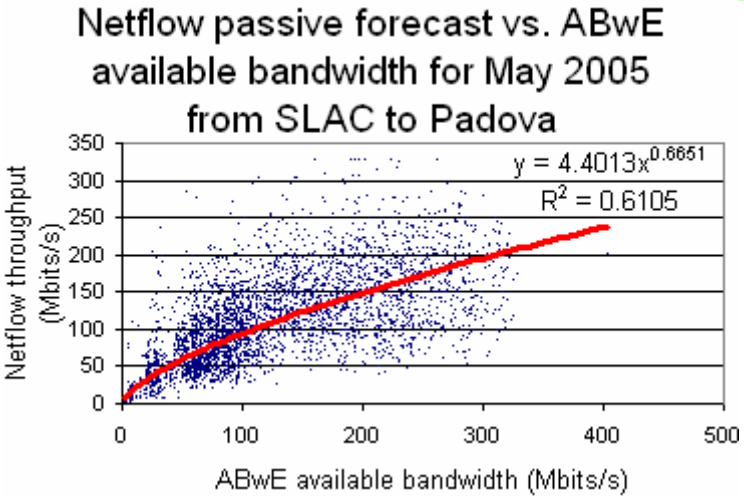

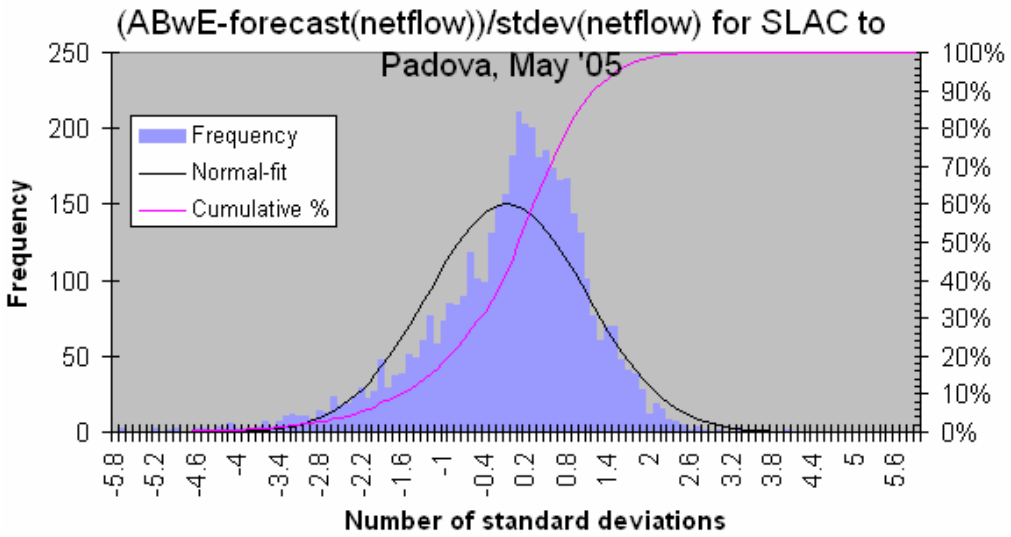

- Predict flow throughputs from Netflow data for SLAC to Padova for May '05
- Compare with E2E active ABwE measurements

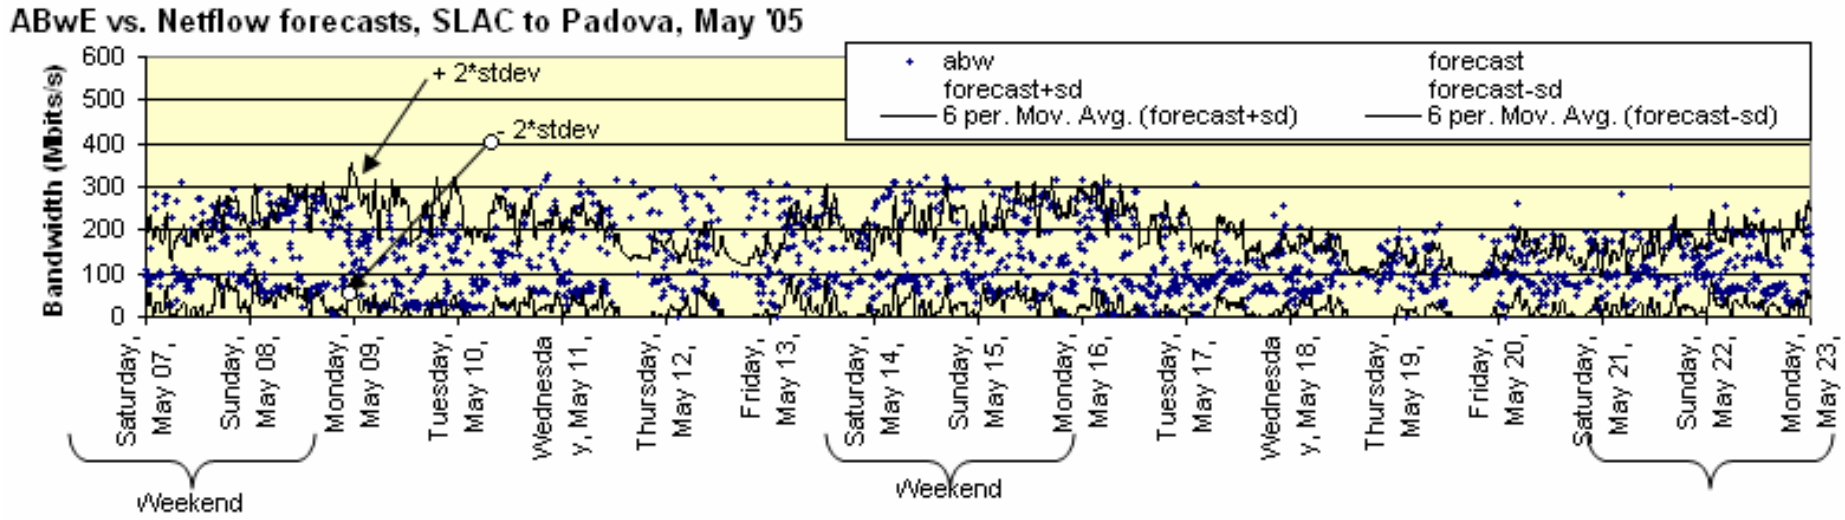

Mining data for sites

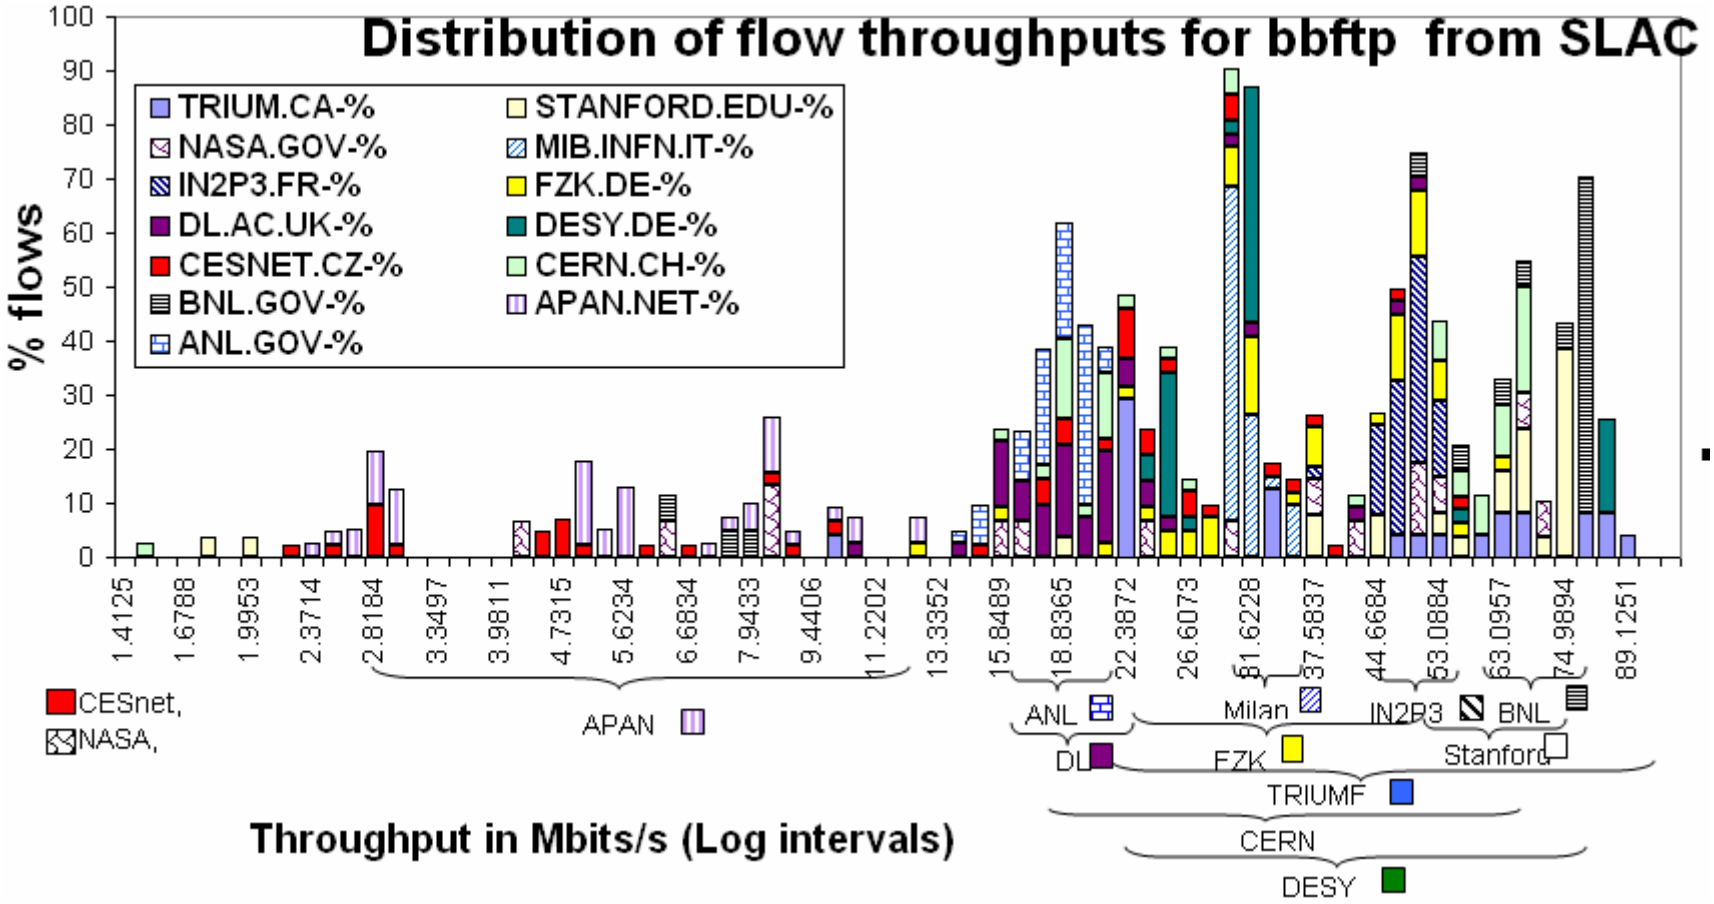

One month for one site

- Bbcp SLAC to Padova

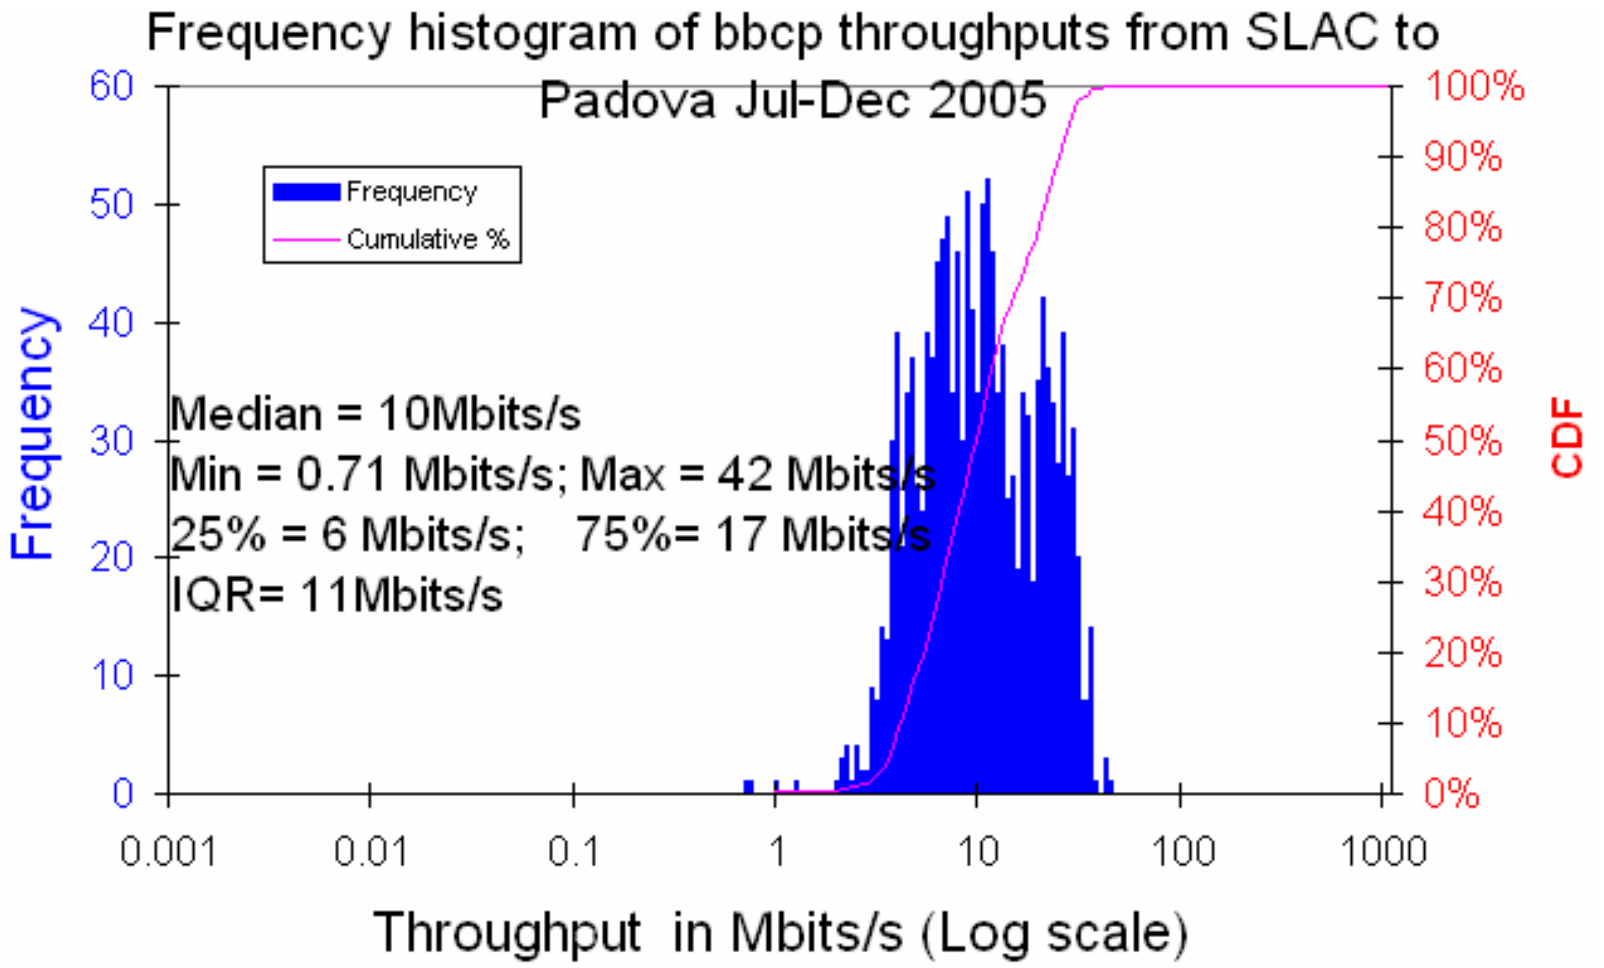

Multi months

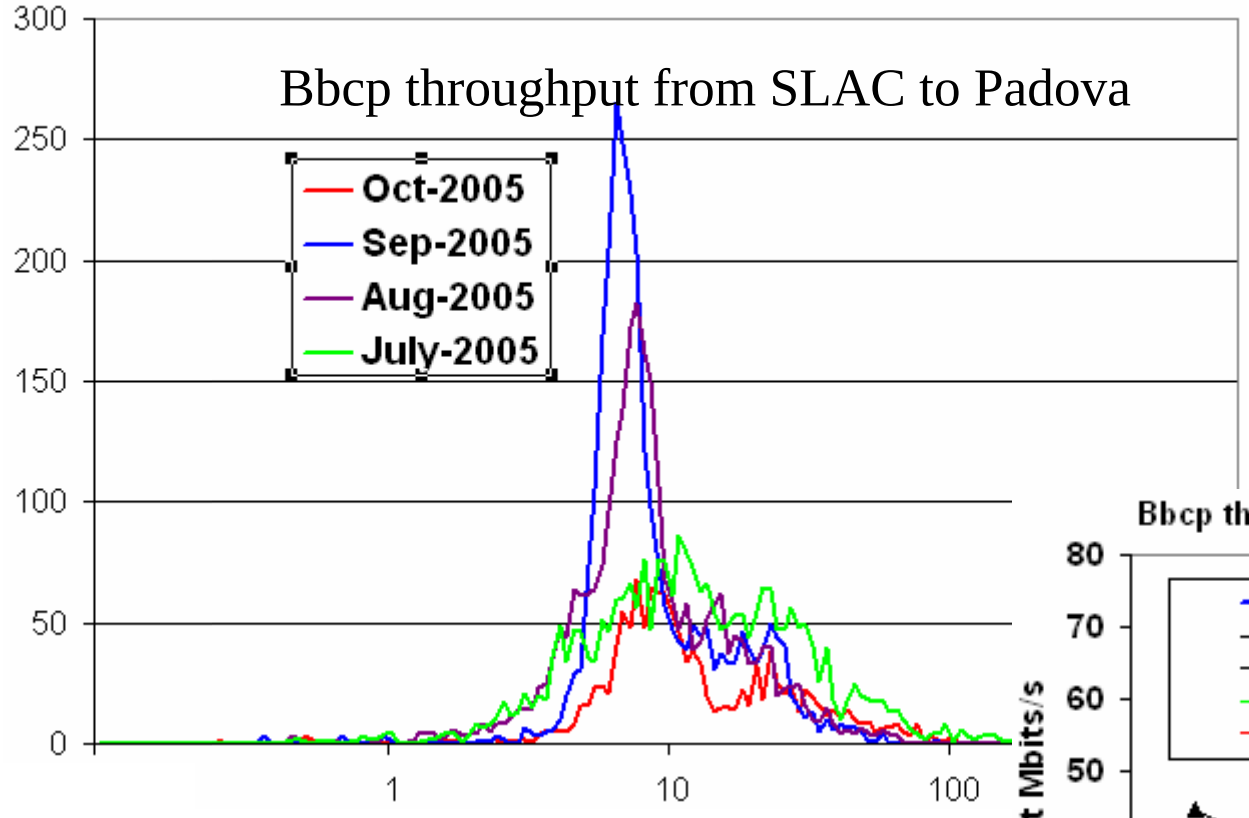

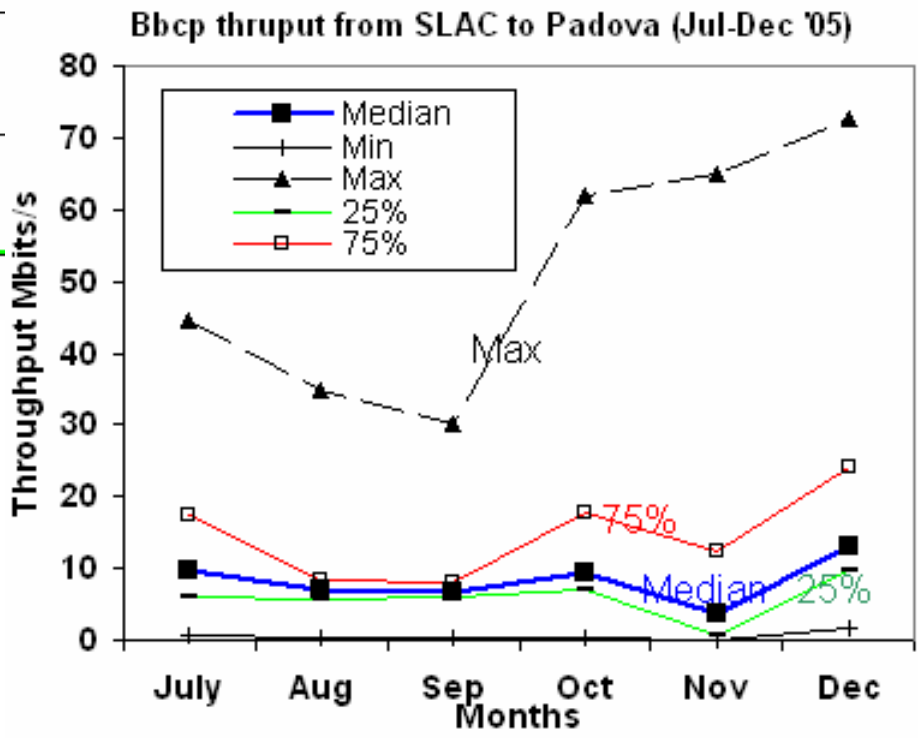

Diurnal behavior

- Some evidence of diurnal behavior

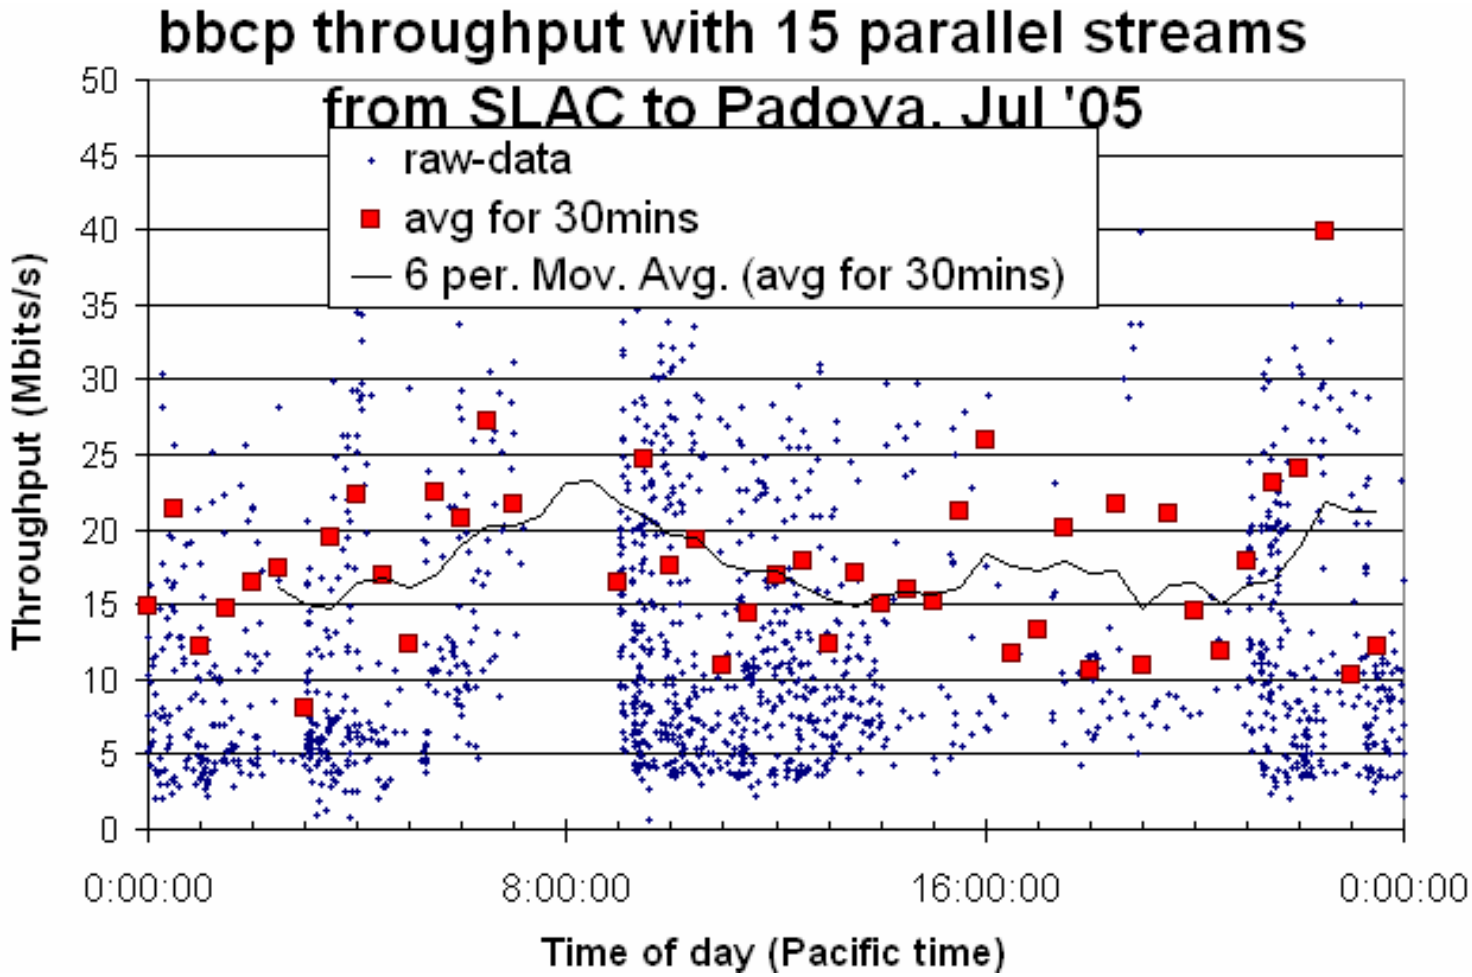

Effect of multiple streams

- Dilemma what do you recommend:
  - Maximize throughput but unfair, pushes other flows aside
  - Use another TCP stack, e.g. BIC-TCP, H-TCP etc.

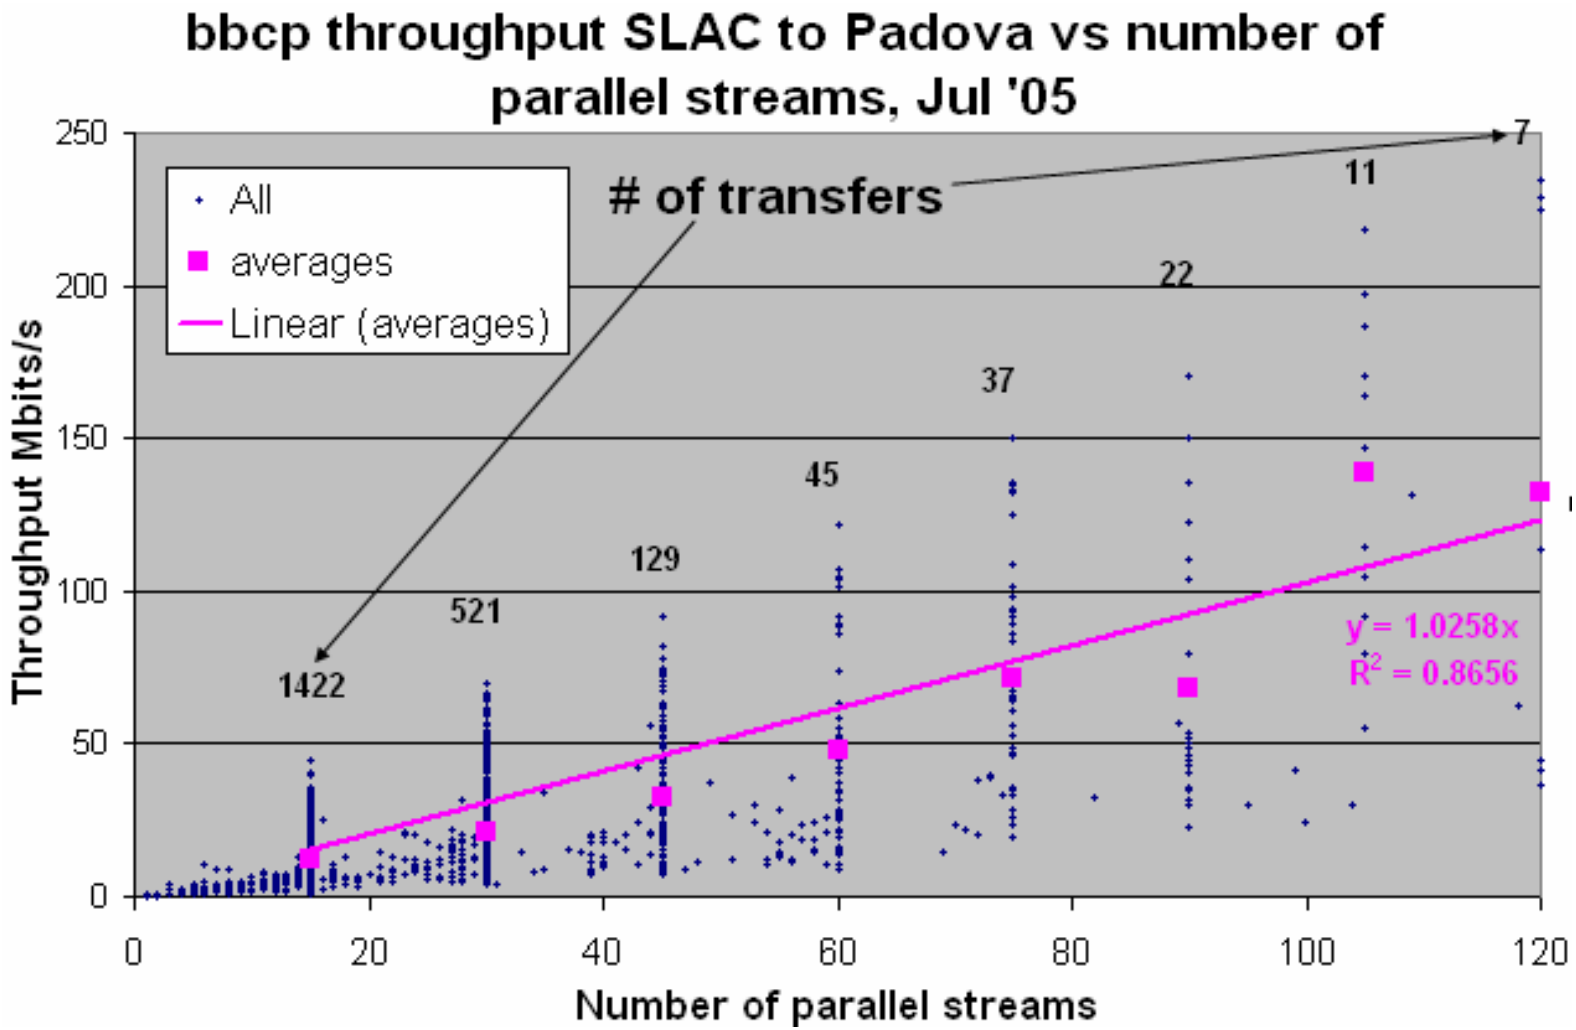

Netflow limitations

- Use of dynamic ports.
  - GridFTP, bbcp, bbftp can use fixed ports
  - P2P often uses dynamic ports
  - Discriminate type of flow based on headers (not relying on ports)
    - Types: bulk data, interactive …
    - Discriminators: inter-arrival time, length of flow, packet length, volume of flow
    - Use machine learning/neural nets to cluster flows
    - E.g. http://www.pam2004.org/papers/166.pdf
- Aggregation of parallel flows (needs care, but not difficult)
- SCAMPI/FFPF/MAPI allows more flexible flow definition
  - See www.ist-scampi.org/
- Use application logs (OK if small number)

More challenges

- Throughputs often depend on non-network factors:
  - Host interface speeds (DSL, 10Mbps Enet, wireless)
  - Configurations (window sizes, hosts)
  - Applications (disk/file vs mem-to-mem)
- Looking at distributions by site, often multi-modal
- Predictions may have large standard deviations
- How much to report to application

Any questions?

- Comparisons of Active, web100 & Netflow measurements
  - http://www.slac.stanford.edu/comp/net/bandwidth-tests/web100/