Measurement and Analysis of Traffic in a Hybrid Satellite-Terrestrial Network

Qing (Kenny) Shao

Communication Networks Laboratory
http://www.ensc.sfu.ca/cnl
School of Engineering Science
Simon Fraser University

Road map

- Introduction and motivation
- Traffic collection
- Traffic analysis
- Traffic prediction
- Conclusion
- References

Introduction of network traffic measurements

- Focus of networking research:
  - mid to late 1980's
  - early 1990's

Motivation for traffic measurements:

- Understand traffic characteristics of existing networks
- Develop traffic models for research
- Performance evaluation of protocols and applications
- Useful for simulation setup

Measurement methodology

- Traffic collection
- Traffic analysis and modeling
- Traffic prediction

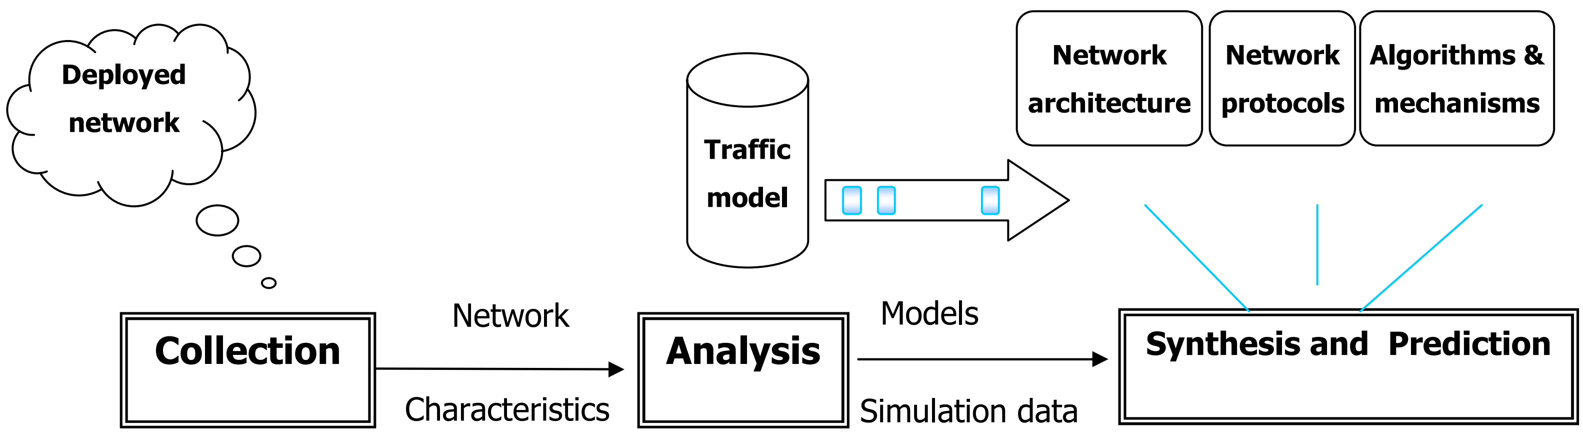

Motivation for this research

- Most traffic traces were collected from research communities:
  - Bellcore, LBNL, Auckland University
- Very few traces were collected from commercial or wireless networks
- Various factors affect Internet traffic patterns:
  - Web, Proxy, Napster, MP3, Web mail
- Evaluate wavelets based approach to data prediction

Road map

- Introduction and motivation
- Traffic collection
- Traffic analysis
- Traffic prediction
- Conclusion
- References

Introduction to DirecPC system

- DirecPC is a satellite one-way broadcast system offering three types of services:
  - turbo Internet access
  - digital package delivery
  - multimedia broadcast
- Manufactured by Hughes Network System (HNS), the largest satellite communication vendor
- DirecPC systems are operating worldwide: Eastern Canada (Bell Express-Vu)
- ChinaSat provides turbo Internet service to 400 Internet cafés nationwide

Turbo Internet service

- TCP splitting:
  - terrestrial links use standard TCP
  - long delay space links use the modified TCP version (increased TCP window size) to improve efficiency
- IP spoofing:
  - customer’s requests are not directly sent to the website
  - they are rerouted to the satellite network operation center (NOC)
  - NOC resends the request to the website
  - website sends to the NOC data to be downloaded

Turbo Internet service

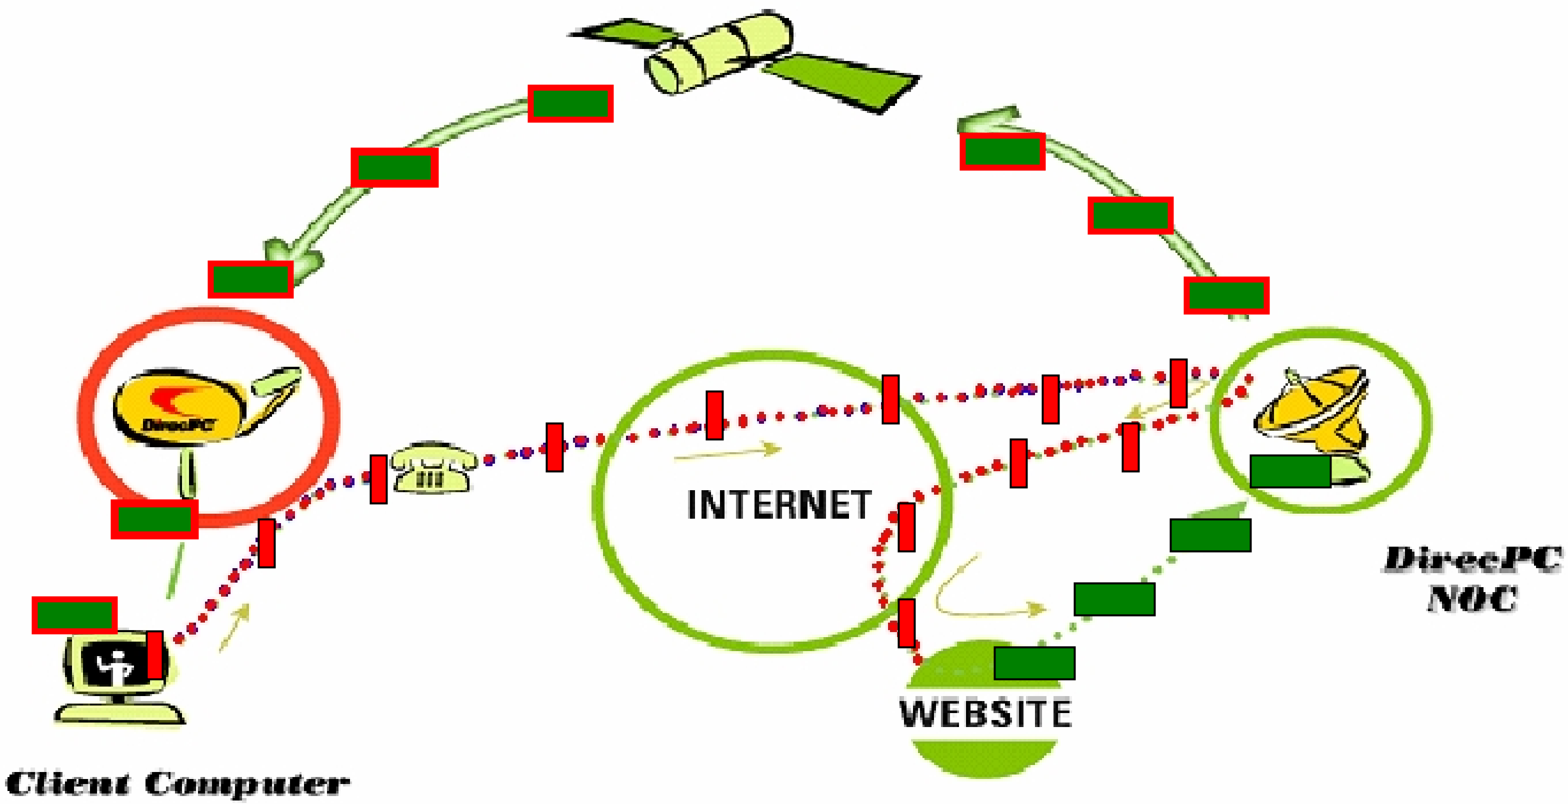

Red: uploaded traffic,
Green: downloaded traffic

Data collection

- One-month trace:
  2002/12/14 11:30 AM - 2003/1/10 11:00 AM
  - each dump file is 500 MBytes large
  - total: 128 dump files, ~63 GBytes
  - granularity: several milliseconds
- Billing records:
  2002/11/1 0:00 AM - 2003/1/10 11:00 AM
  - DirecPC billing system creates one file per hour
  - total: 3,380 files, each ~1-4 Kbytes
  - granularity: 1 hour

tcpdump trace format

- timestamp src > dst: flags data-seqno ack window urgent options
  - 19:12:45.660701 61.159.59.162.12800 > 192.168.1.169.62246: udp 52
  - 19:12:45.672959 192.168.1.242.40849 > 210.51.17.67.9065: P 6541284:6541321(37) ack 1479344110 win 8192 (DF)
  - 19:12:45.674709 192.168.2.30.39042 > 202.101.165.124.4220: . ack 807850998 win 8192
  - 19:12:45.676255 61.152.249.71.55901 > 192.168.1.242.40770: P 2627573783:2627573791(8) ack 5795719 win 63343 (DF)
  - 19:12:45.676256 61.152.249.71.55901 > 192.168.1.242.40846: P 2775973525:2775973533(8) ack 11622145 win 64102 (DF)
  - 19:12:45.688514 192.168.1.242.40770 > 61.152.249.71.55901: . ack 8 win 8192
  - 19:12:45.688843 192.168.1.242.40846 > 61.152.249.71.55901: . ack 8 win 8192
  - 19:12:45.689095 192.168.1.169.63644 > 202.103.69.103.3010: P 1969195:1969259(64) ack 2995916216 win 8192 (DF)
  - 19:12:45.692475 202.101.165.134.80 > 192.168.2.3.45585: . ack 3153903 win 6432
  - 19:12:45.699193 207.46.104.20.80 > 192.168.1.239.4912: R 2405276149:2405276149(0) win 0

- Red: uploaded traffic
- Green: downloaded traffic

Billing records format

RecLen	RecTyp	SiteID	Start	Stop	Cmin	Bill	CTxByt	CRxByt	CTxPkt	CRxPkt
00100	001	0007AA6701	20021220140005	20021220150005	053	2	**0002427264**	0000743362	**0000015045**	0000014686
00100	001	0004553002	20021220140005	20021220150005	060	2	**0069564405**	0005303976	**0000071747**	0000093743
00100	001	0004492702	20021220140005	20021220150005	003	2	**0000016966**	0000009543	**0000000055**	0000000082
00100	001	00045BBA07	20021220140005	20021220150005	059	2	**0008957320**	0001809152	**0000011501**	0000015741
00100	001	00045BEA02	20021220140005	20021220150005	059	2	**0011940914**	0001234279	**0000014215**	0000017098

- Red: uploaded traffic
- Green: downloaded traffic

Data processing approach

- Parse the data:
  - Awk scripts, tcp-con, tcp-slice
- Upload to the data server:
  - mysql
- Extract the data from data server
- Analysis tools:
  - Splus, R language, Matlab

Road map

- Introduction and motivation
- Traffic collection
- Traffic analysis
  - background
  - results
- Traffic prediction
- Conclusion
- References

Self-similarity

- Self-similarity implies a “fractal-like” behavior: examining data on certain time scales produces similar patterns
- A wide-sense stationary process X(n) is called (exactly second order) self-similar:
  - $r^{(m)}(k) = r(k), k \geqslant 0, m = 1, 2, \ldots, n$
- Implications:
  - no natural length of bursts
  - bursts exist across many time scales
  - traffic does not become “smoother” when aggregated (unlike Poisson traffic)

The Hurst parameter

- Properties:
  - slow decaying variance
  - long-range dependence
  - Hurst parameter
- Models with only short range dependence (Poisson model): $H = 0.5$
- Self-similar processes: $0.5 < H < 1.0$
- As the traffic volume increases, the traffic becomes more bursty, more self-similar, and the Hurst parameter increases

Estimate the self-similarity

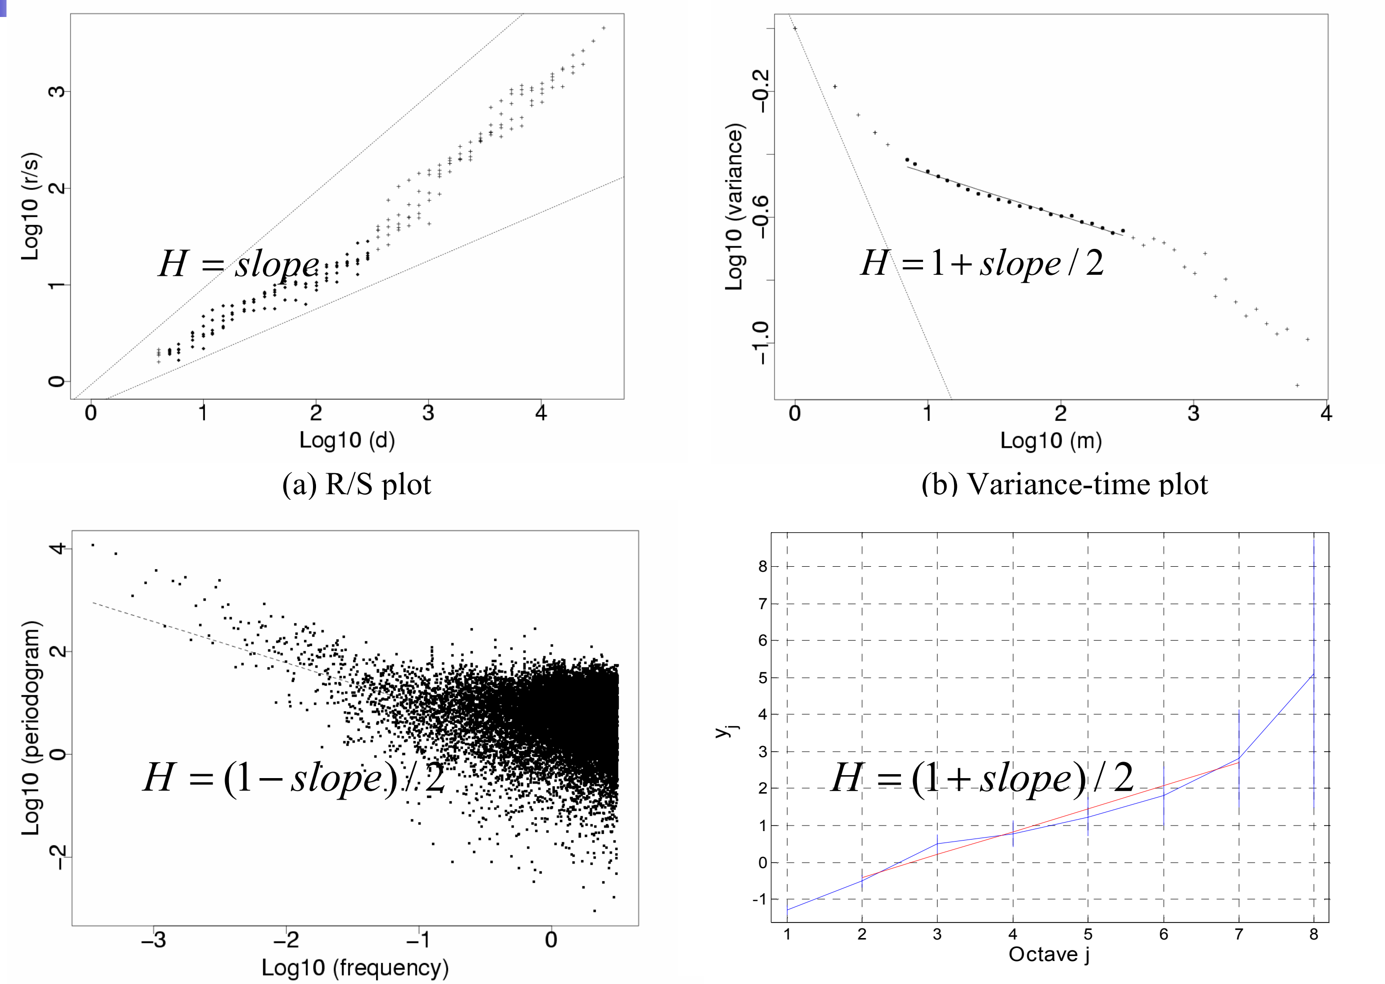

(a) R/S plot

(b) Variance-time plot

(c) Periodogram plot

(d) Wavelet plot

Why self-similar?

- Self-similar process can be generated by aggregation of ON/OFF sources
- The ON/OFF periods are heavy-tailed distributed with infinite variance
- Web and ftp file sizes are heavy-tailed
- A probability distribution X is heavy-tailed if:

$$P[X > x] \sim cx^{-\alpha}, 0 < a < 2, x \to \infty$$

Reference: Mark E. Crovella and Azer Bestavros, "Self-similarity in world wide web traffic: evidence and possible causes," in *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835 - 846, December 1997.

Road map

- Introduction and motivation
- Traffic collection
- Traffic analysis
  - background
  - results
- Traffic prediction
- Conclusion
- References

Analysis of billing records

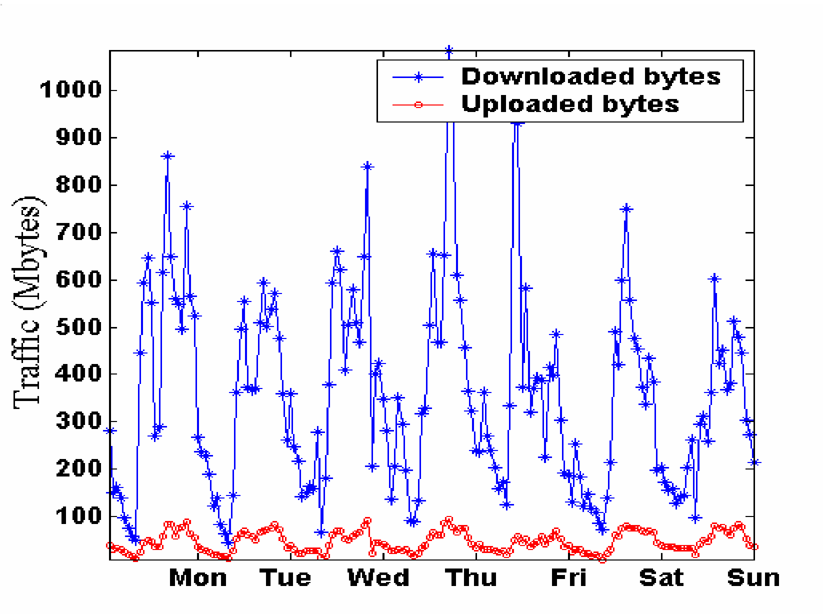

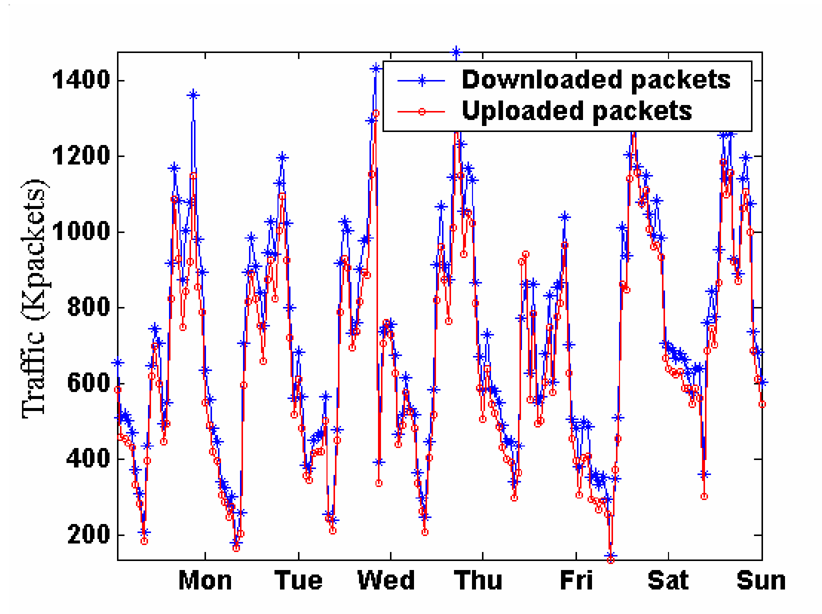

- Weekly traffic volume measured in packets (left) and bytes (right)
- Data was collected from 2002-12-09 to 2002-12-15

Analysis of billing records

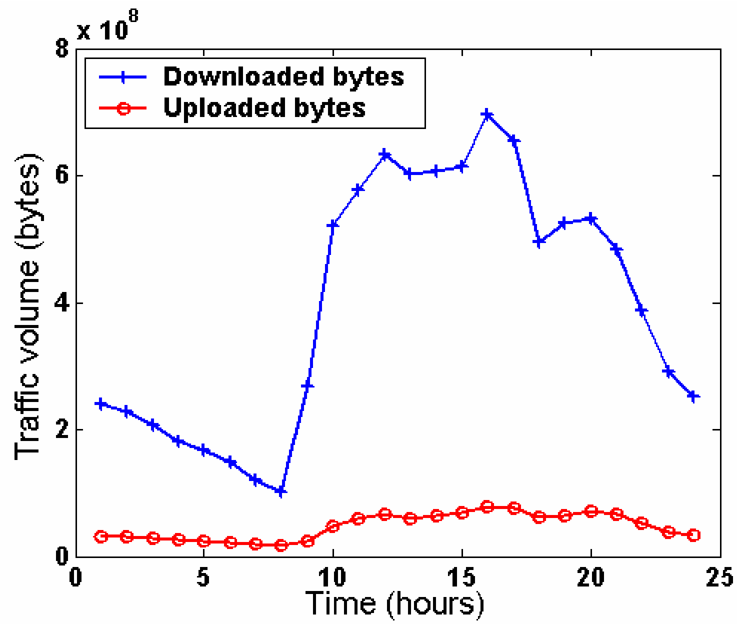

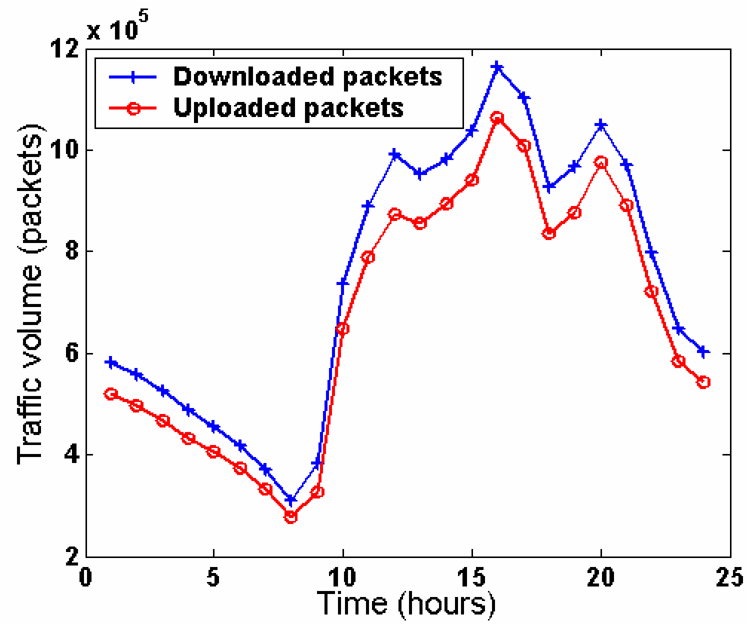

- Average traffic volume over a day in packets (left) and bytes (right)
- Data was collected from 2002-12-09 to 2002-12-15

Applications and protocols distribution

Applications	Connections	% of connections	Bytes	% of bytes
WWW	304,243	90.06	10,203,267,005	75.79
FTP-data	636	0.19	1,440,393,008	10.7
IRC	2,324	0.69	945,965	0.008
SMTP	562	0.17	2,326,373	0.01
POP-3	115	0.03	2,326,373	0.02
Telnet	70	0.02	280,286	0.002
Other	651	8.84	238,099,412	13.47
Total	**308,601**	**100**	**11,885,432,923**	**100**

Protocol	Packet	% of packets	Bytes	% of bytes
TCP	36,737,165	84.32	11,231,147,530	94.49
UDP	6,202,673	14.24	601,157,016	5.06
ICMP	630,528	1.44	53,128,377	0.45
Total	**43,570,366**	**100**	**11,885,432,923**	**100**

- Data was collected from 2002-12-21 22:08 to 2002-12-23 3:28

Packet size: analysis

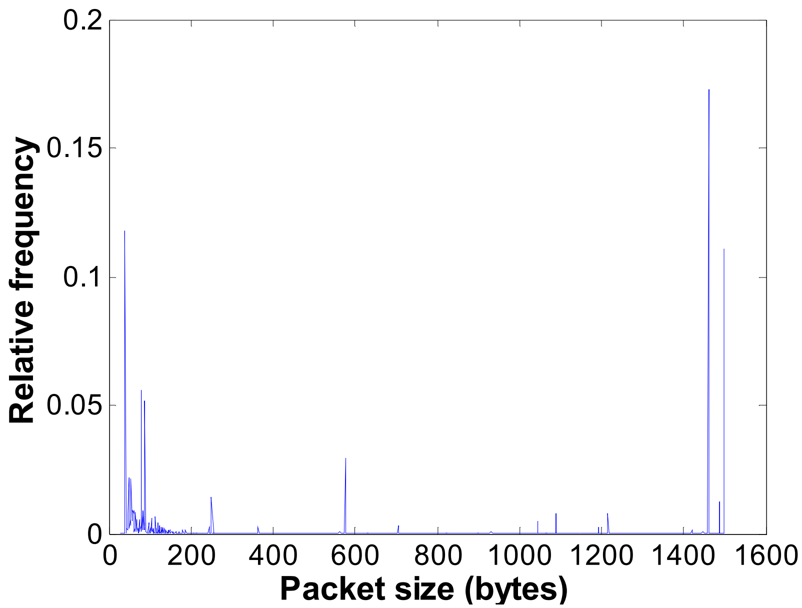

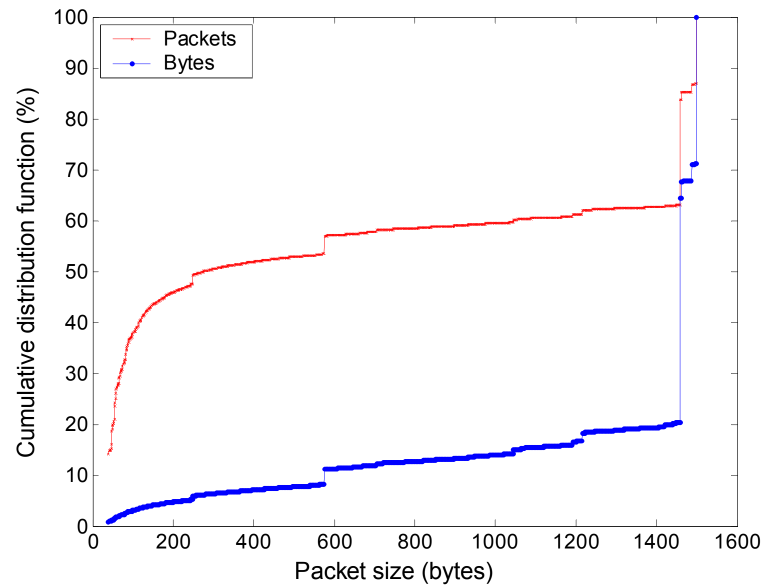

Data was collected from 2002-12-21 22:08 to 2002-12-23 3:28

- Packet size distribution is bimodal:
  - 50% packets less than 300 bytes
  - 40% packets greater than 1400 bytes
- Most bytes are transferred in large packets

Web traffic distribution on the TCP connection level

- Zipf-like distribution fr ~ 1/r$^{\beta}$
  - the frequency of popularity is inversely proportional to rank of popularity

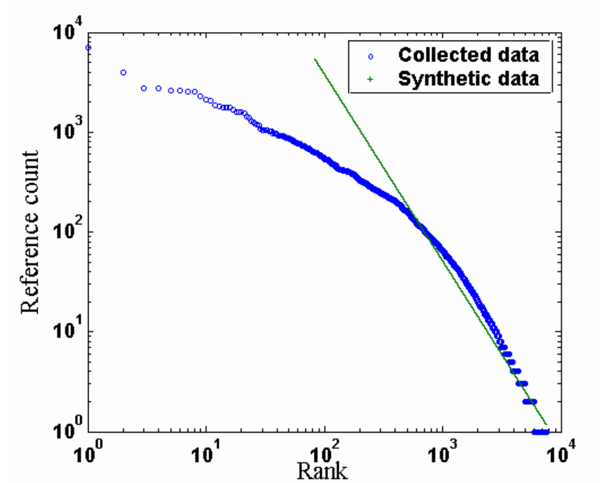

- DGX (discrete lognormal)

$$p(x=k)=\frac{A(\mu,\sigma)}{k}\exp[-\frac{(\ln k-\mu)^2}{2\sigma^2}]$$

$$A(\mu,k)=\{\sum_{k=1}^{\infty}\frac{1}{k}[-\frac{(\ln k-\mu)^2}{2\sigma^2}]\}^{-1}$$

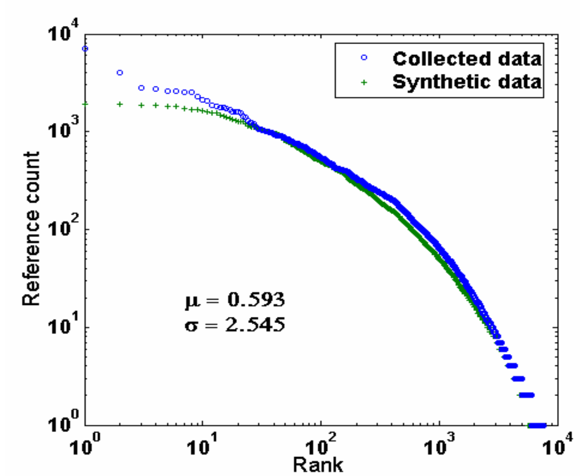

- DGX distribution fits better than the Zipf-like distribution

Web traffic distribution on the TCP connection level

- Traffic is non-uniformly distributed among the hosts on the Internet
  - top 10 web servers account for 62.3 % of the traffic:
    - all registered under the Asia Pacific Network Information Centre (APNIC)
    - the heaviest loaded website is a Chinese search engine website
- Language, geography, and popular sites (commercial influence) greatly affect the traffic distribution
- Important for ISPs when configuring proxy servers and content delivery networks

Estimate the self-similarity

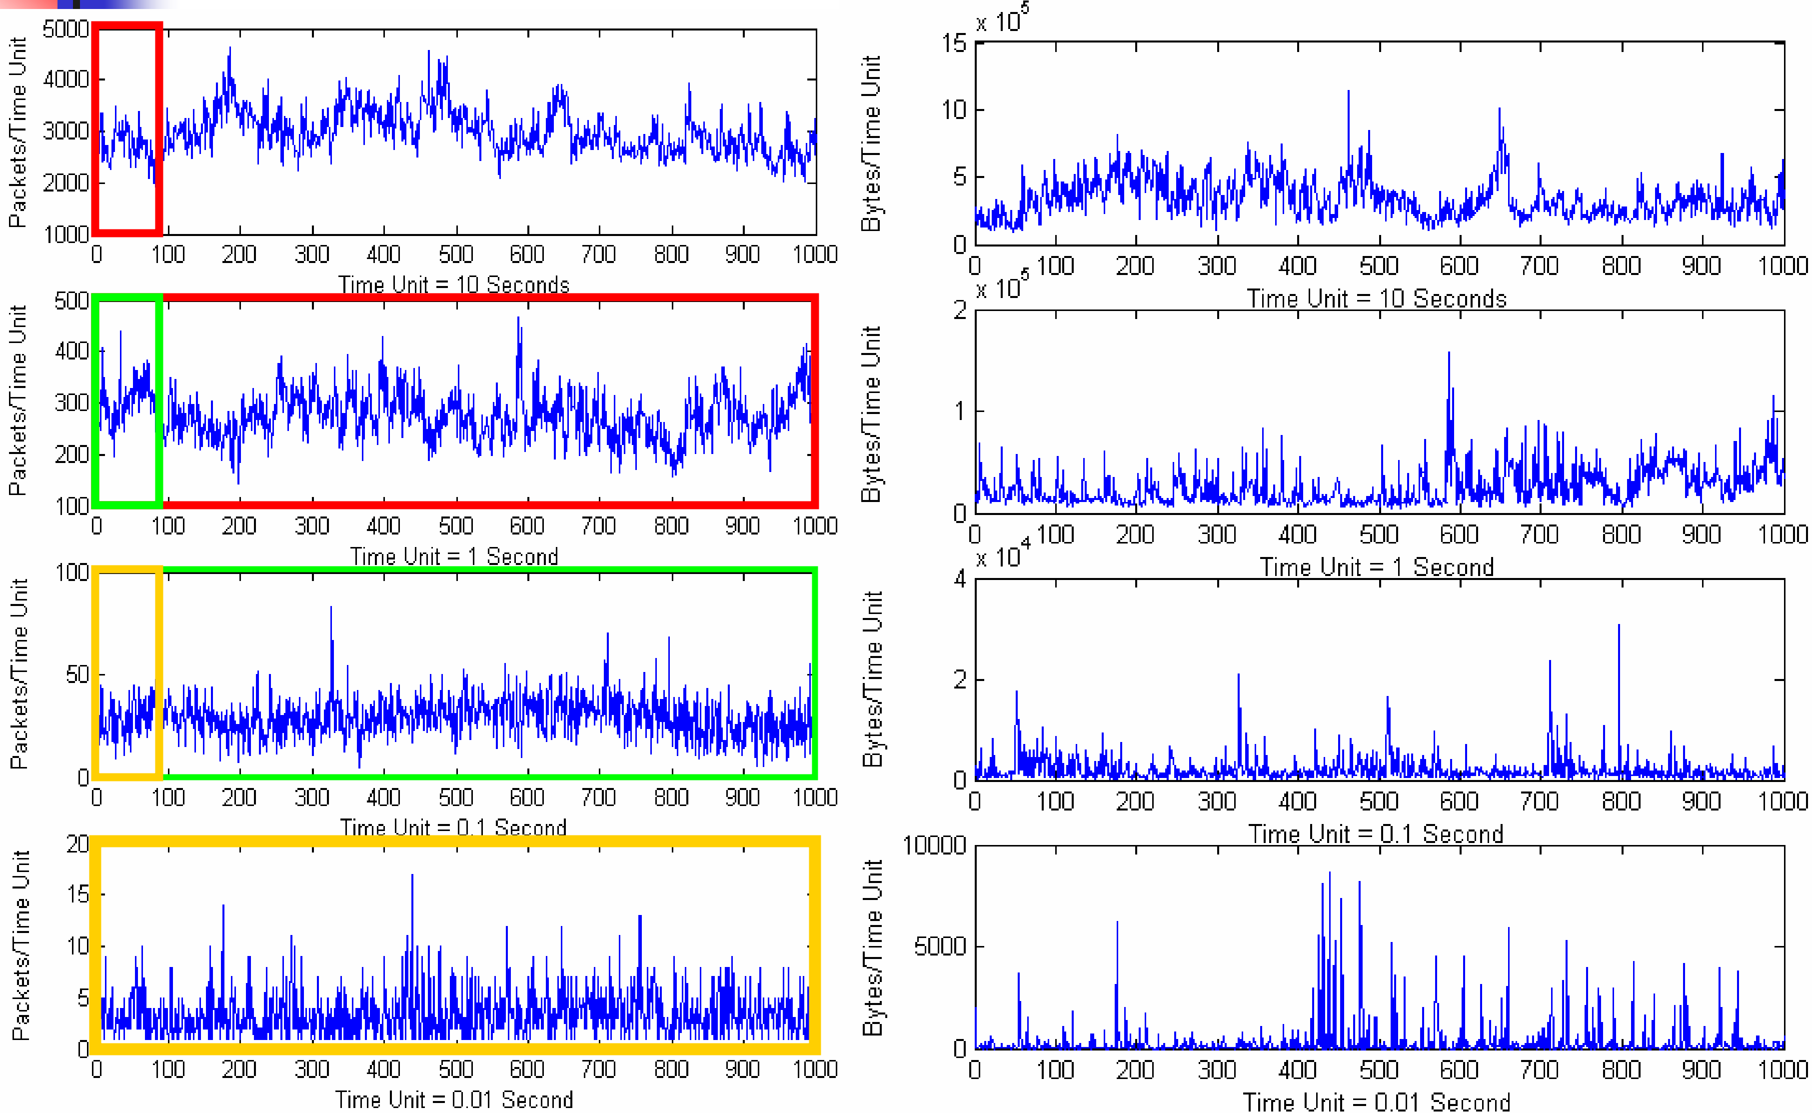

Estimate the self-similarity

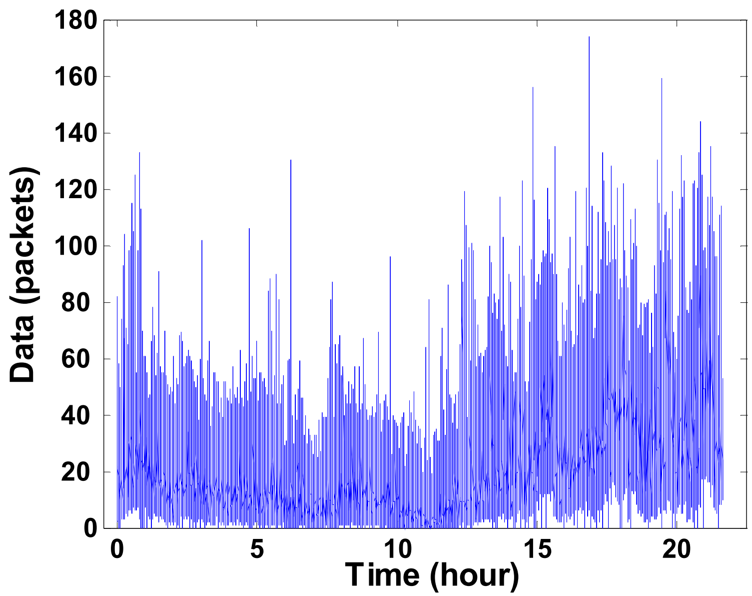

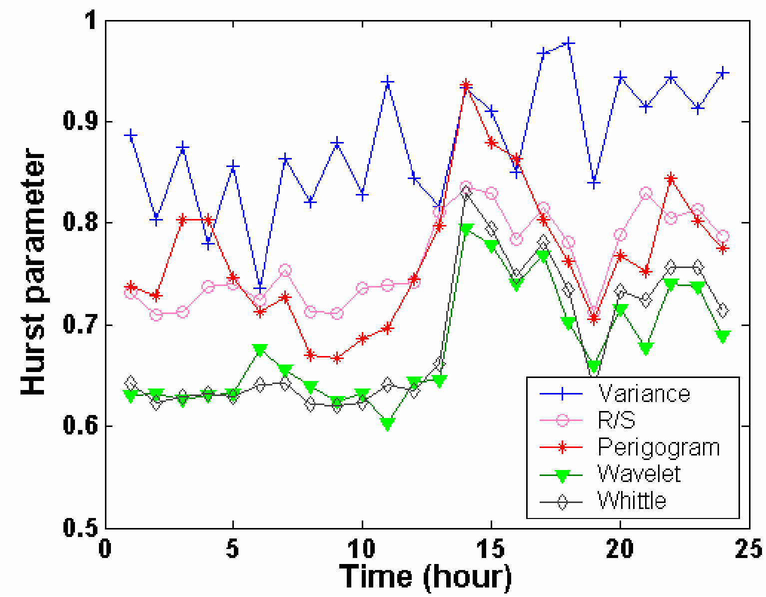

- Traffic data was collected on 2002-12-09

TCP connection model

- Two parameters of TCP connection are:
  - inter-arrival times
  - number of bytes transferred per connection
- Four distributions:

Distribution	Probability density	Cumulative probability
Exponential	$f(x) = \frac{1}{\rho} e^{-x/\rho}$	$F(x) = 1 - e^{-x/\rho}$
Weibull	$f(x) = \frac{1}{a}\left(-\frac{x}{a}\right)^{c-1} \mathrm{e}^{-(x/a)^c}$	$F(x) = 1 - e^{-(x/a)^c}$
Pareto $(k > 0, a > 0; x \geq k)$	$f(x) = \frac{ak^a}{(x)^{k+1}}$	$F(x) = 1 - \left(\frac{k}{x}\right)^a$
Lognormal	$f(x) = \frac{1}{x\sqrt{2\pi\sigma}} e^{-[\log(x)-\xi]^2/2\sigma^2}$	No closed form

TCP connection model

- Fit four distributions :
  - step 1: estimate the parameter of the distribution
    - maximum likelihood estimator
  - step 2: find the best distribution:
    - visualization of the cumulative distribution curve
    - goodness-of-fit measures (discrepancy estimate)

Reference: Vern Paxson, "Empirical derived analytic models of wide-area TCP connections," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316-336, February 1994.

TCP connection model

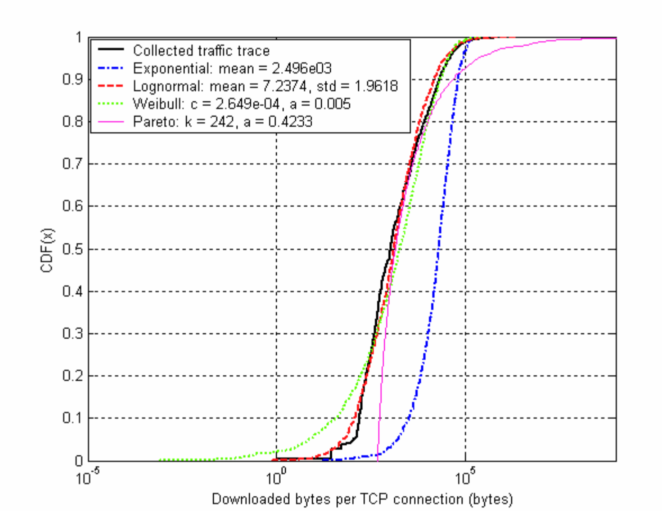

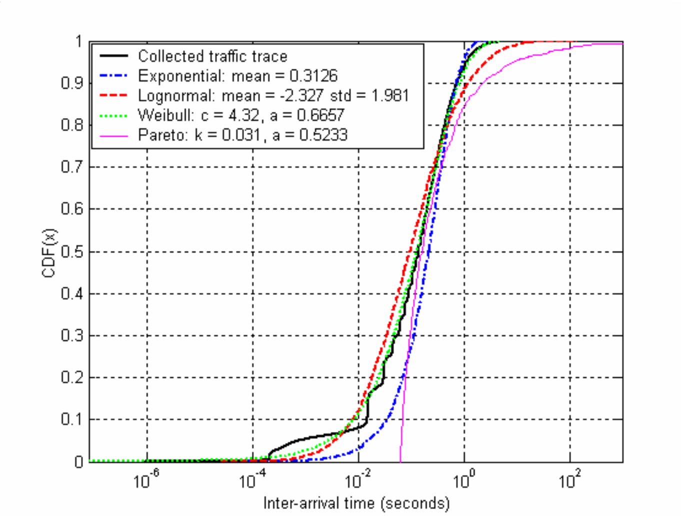

- Discrepancy test:

$$\chi^2 = \frac{X^2}{n} \qquad X^2 = \sum_{i=1}^{N} \frac{(Y_i - np_i)^2}{np_i}$$

Model	Exponential	Lognormal	Weibull
Downloaded bytes	1,261,263	113,620	1,371,293
Inter-arrival	4,103,846	3,779,695	680,707

- Best fit:
  - Lognormal: downloaded bytes per connection
  - Weibull: inter-arrival time

TCP connection model

- Kolmogorov-Smirnov goodness of fit test:
  - hypothesis test
  - significance level
  - P-value

Downloaded bytes per connection

Model	1	2	3	4	5
Exp	2.5×10^{-35}	6.1×10^{-47}	2.7×10^{-40}	5.1×10^{-48}	3.5×10^{-35}
lognormal	0.16	0.2910	0.3685	0.2744	0.2045
Weibull	5.5×10^{-3}	8.4×10^{-3}	3.32×10^{-3}	8.43×10^{-4}	1.54×10^{-3}

Inter-arrival time

Model	1	2	3	4	5
Exp	6.2×10^{-5}	7.2×10^{-6}	2.5×10^{-7}	1.1×10^{-3}	3.7×10^{-5}
lognormal	6.1×10^{-4}	6.2×10^{-3}	7.7×10^{-3}	4.8×10^{-3}	5.9×10^{-4}
Weibull	0.175	0.734	0.616	0.753	0.242

Road map

- Introduction and motivation
- Traffic collection
- Traffic analysis
- Traffic prediction
  - long-term prediction (ARIMA model)
  - short-term prediction (multi-resolution analysis)
- Conclusion
- References

Long-term prediction: ARIMA model

- "Time series analysis - forecasting and control"
  - G. E. P. Box and G. M. Jenkins (1976)
- AutoRegressive Integrated Moving-Average

$$X(t) = u + \phi_1 X(t-1) + ...\phi_p X(t-p) + e(t) + \theta_1 e(t-1) + ...\theta_q e(t-q)$$

  - past values
    - AutoRegressive (AR) structure
  - past random fluctuant effect
    - Moving Average (MA) process

$$(p,d,q) \times (P,D,Q)_s$$

Partial-autocorrelation (PACF) and differencing

- PACF is the correlation between:

$y_t - E(y_t \mid y_{t-1},\dots y_{t-k+1})$ and $y_{t-k} - E(y_{t-k} \mid y_{t-1},\dots y_{t-k+1})$

- Consecutive differencing (order d):
  - $d(Y) = Y_t - Y_{t-1}$
  - usually d = 0, 1
- Seasonal differencing (order D):
  - $D(Y) = Y_t - Y_{t-168}$
  - usually D = 1

Order identification $(p,d,q)\times(P,D,Q)_s$

- Once stationarity has been achieved through differencing, we employed ARMA structure

Process	Autocorrelation plot	Partial autocorrelation plot
AR(p)	Dominated by either damped exponentials or sine waves	The lag beyond which the value cuts off is order p
MA(q)	The lag beyond which the cuts off is the order q	Dominated by damped exponentials and sine waves
ARMA(p,q)	Exponential or sine wave decay after lag (q-p)	Exponential or sine wave decay after lag (p-q)

- Finding P, Q is similar to finding p and q, but the plot should be examined s units apart

One week ahead prediction

- We employed the identification method on six weeks of billing records and identified the following parameters:
  - d=0, D=1, s=168, p=1, q=0, P=0, Q=1
  - collected traffic fits the model $(1,0,0)\times(0,1,1)_{168}$

- Normalized mean squared error (NMSE) is used to measure the quality of prediction:

$$NMSE = \frac{1}{\sigma^2 N}\sum_{k=1}^{N}(x(k)-\hat{x}(k))^2$$

Predictability evaluation

Traffic type	Uploaded (Mbytes)	Downloaded (Mbytes)	Uploaded (packets)	Downloaded (packets)
NMSE	0.3653	0.5988	0.5268	0.5244

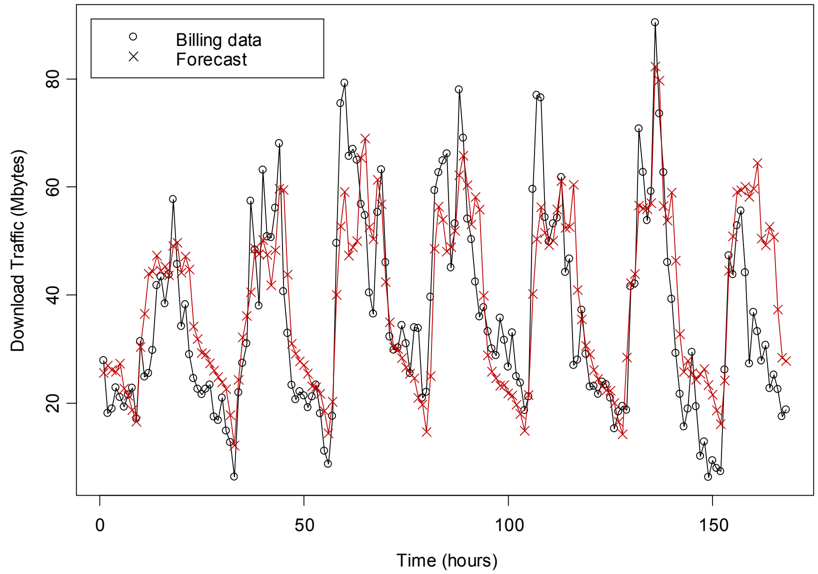

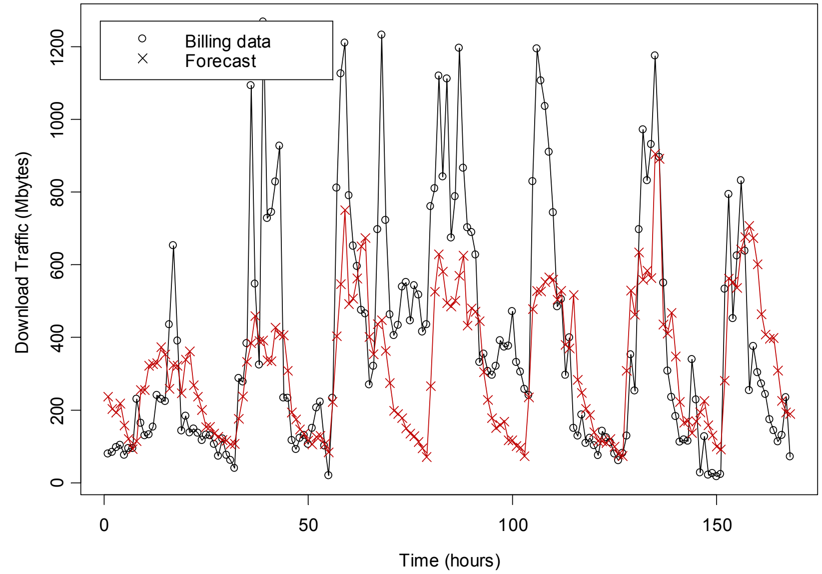

- Downloaded traffic (bytes) is more difficult to predict than uploaded traffic (bytes)

Short-term prediction

- Multi-resolution analysis
- Prediction combined wavelets and AR model

Multi-resolution analysis (MRA)

- From Fourier (1807) to wavelets (1980s)

Redundant wavelet transform

- Use the redundant information to fill the gap

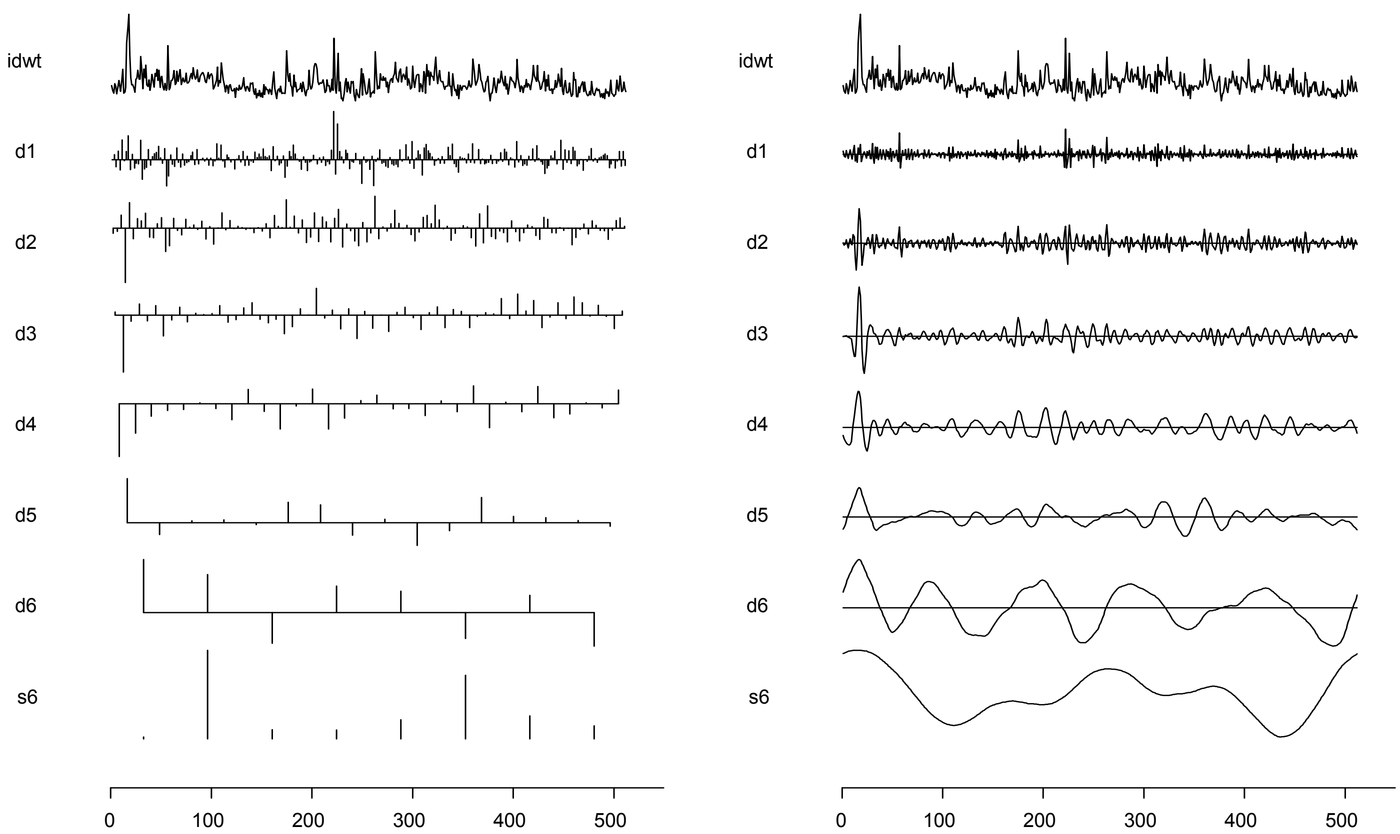

Atrous wavelet transform

- Step 1: First initialize i to 0 and start with a signal:

$$s_0(k) = y(k) \qquad s_i(k) = \sum_{l=-\infty}^{l=\infty} h(l) s_{i-1}(k + 2^{i-1} l)$$

- Step 2: Calculate the detail coefficients:

$$d_i(k) = s_{i-1}(k) - s_i(k)$$

- Step 3: Continue Step 1 and 2 until i reaches the specified scale level

Atrous inverse transform

- The inverse transform is:

$$y(k) = s_p(k) + \sum_{j=1}^{p} d_j(k)$$

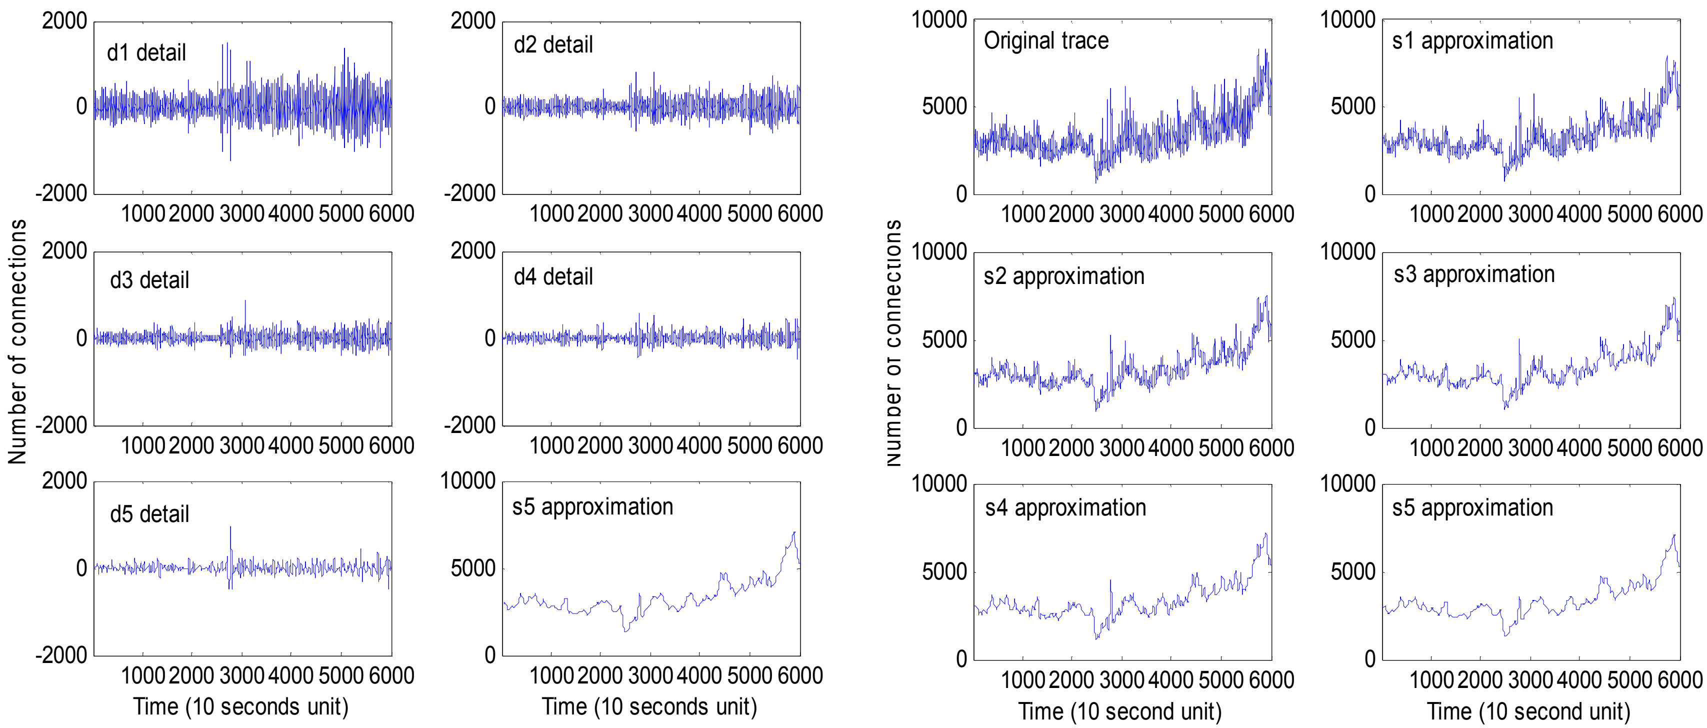

Combined wavelet method and AR model for short-term prediction

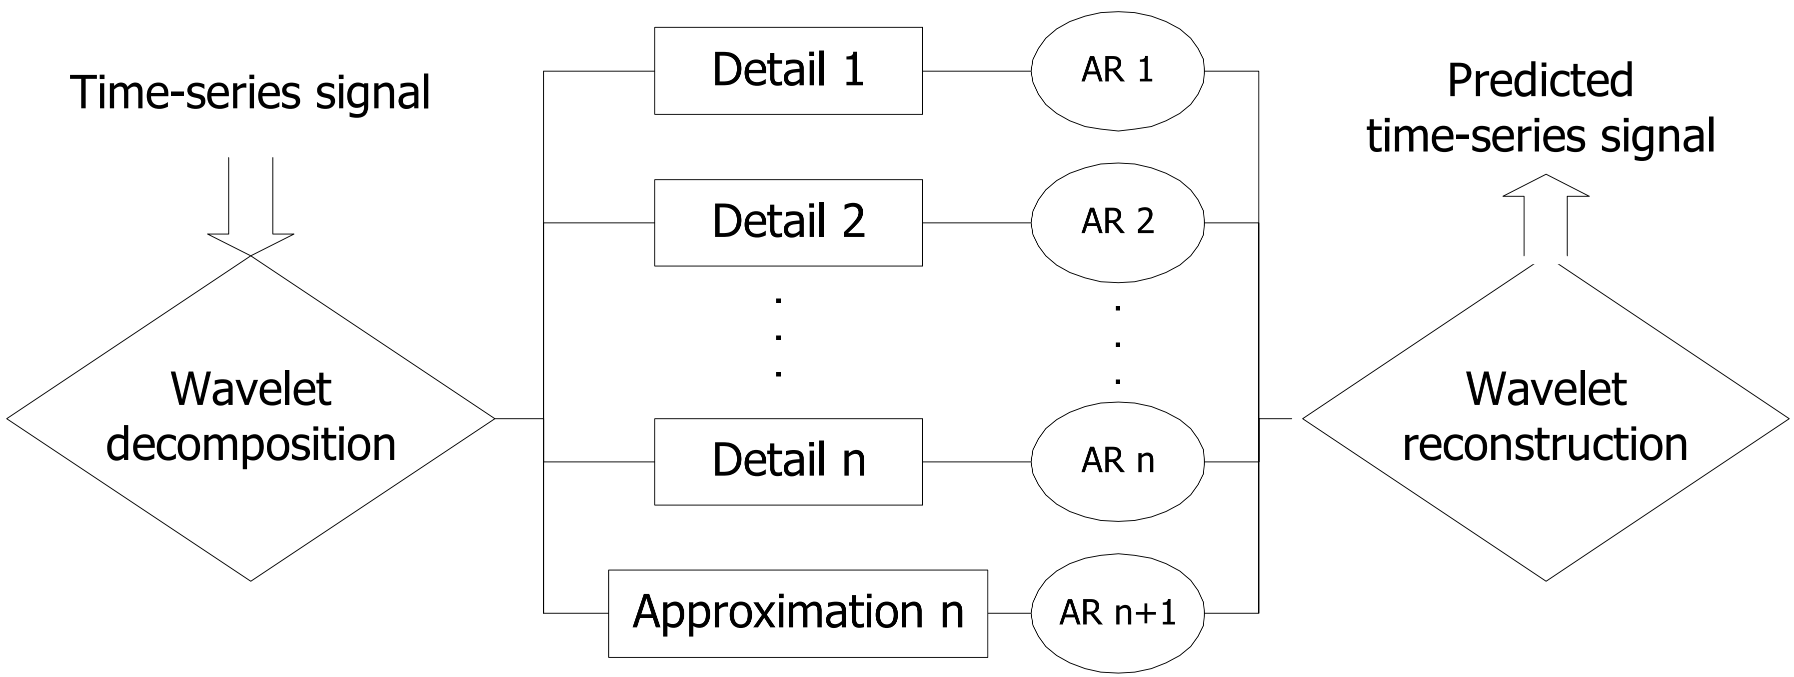

Combined wavelet method and AR model for short-term TCP connection prediction

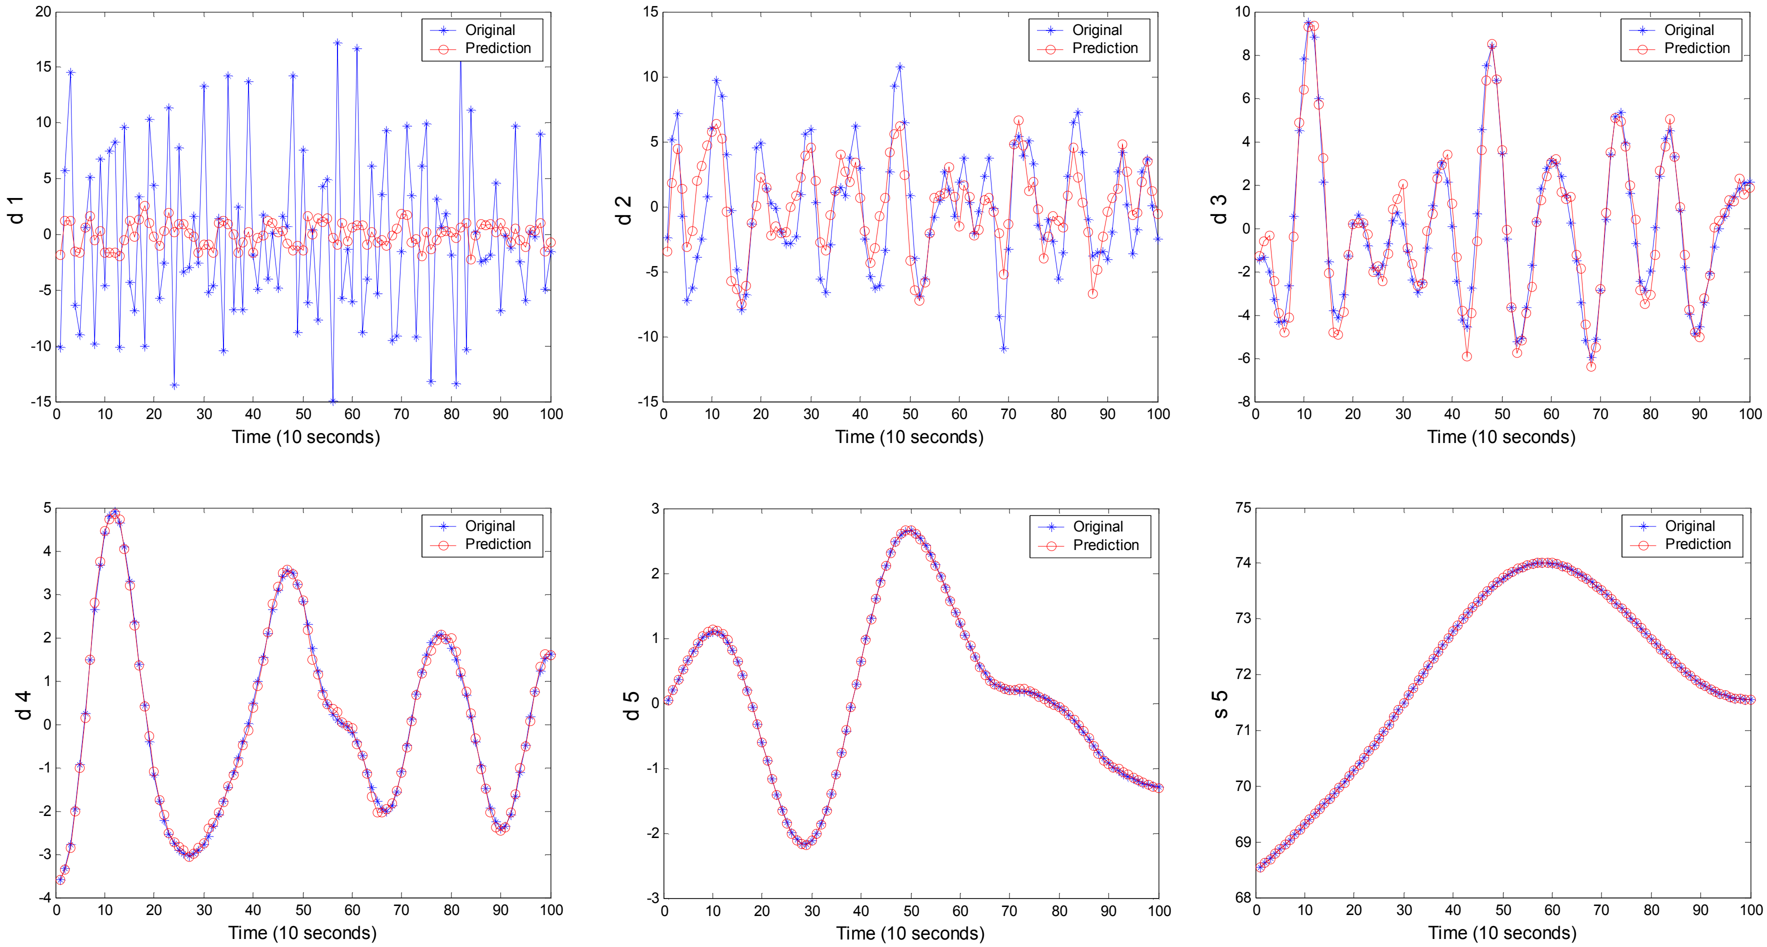

Combined wavelet method and AR model for short-term TCP connection prediction

NMSE

Predictor	d1	d2	d3	d4	d5	d5	signal
AR+atrous	0.9854	0.3711	0.0380	0.0019	1.43e-004	5.3e-006	0.4261
AR	-	-	-	-	-	-	1.2654

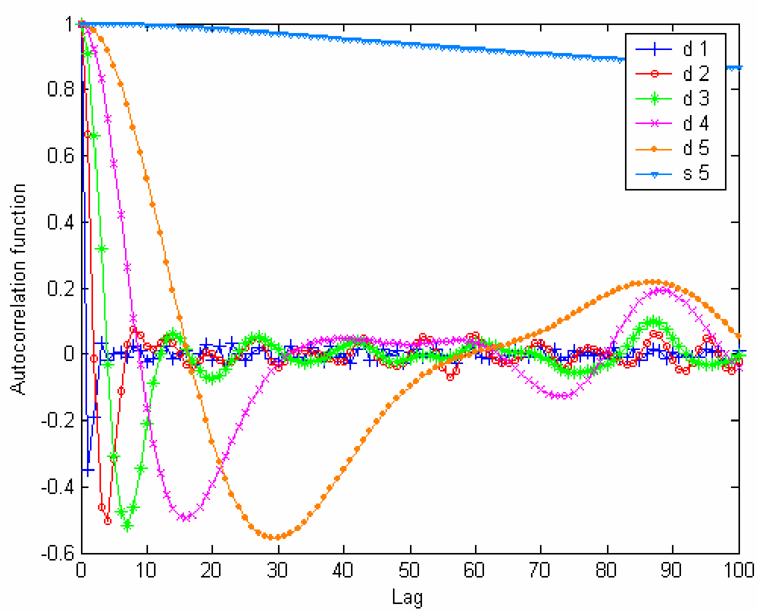

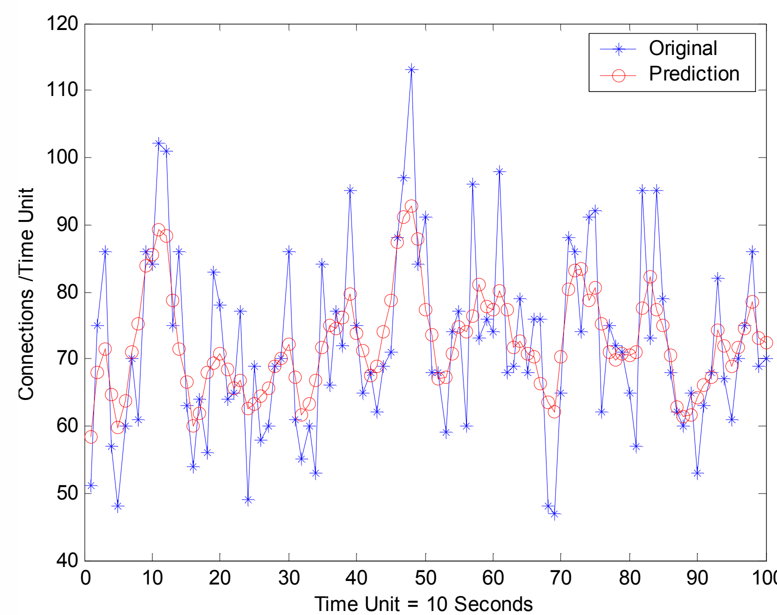

Performance of the predictor versus the length of the predictors

- L is the prediction length:

$$x(n+k) = f(x(n-1), x(n-1), \ldots.., (n-L+1))$$

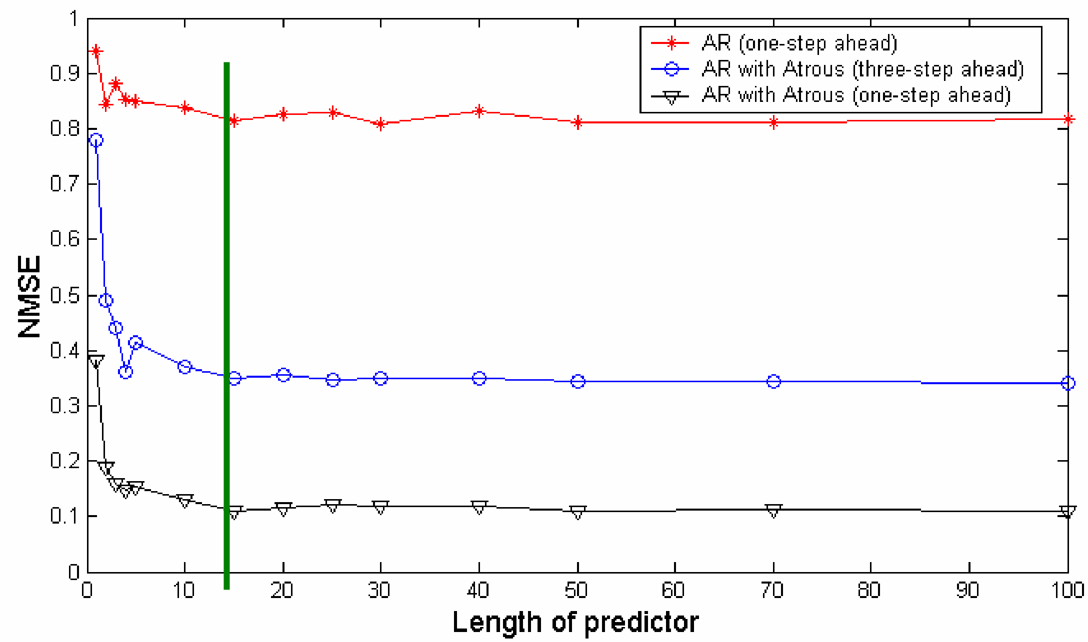

- Beyond a certain point, including more historical traffic data does not necessarily improve the predictability

Conclusions

- Traffic collection
- Traffic analysis and modeling
- Traffic prediction

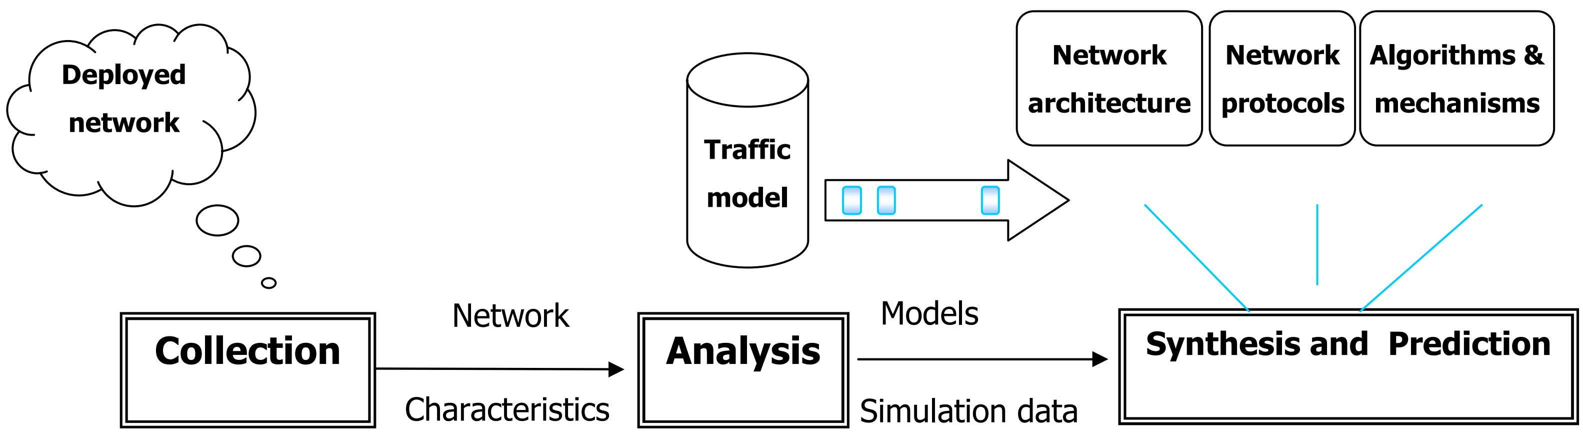

Conclusions

- Analysis of collected traffic data:
  - Web application and TCP protocol dominate the collected traffic
  - packet size distribution is bimodal, most bytes are transferred in big packets
  - few web servers account for majority of the traffic
  - TCP connection distribution among the hosts follow the discrete lognormal distribution
  - Hurst estimators perform differently, wavelet and Whittle estimators may provide more accurate estimate

Conclusions

- TCP modeling:
  - Weibull: inter-arrival time
  - Lognormal: downloaded bytes per TCP connection
- Traffic prediction
  - downloaded traffic is more difficult to predict compared to predict the uploaded traffic
  - combining atrous wavelet and AR method can improve the short-term predictability
  - beyond a certain point, including more historical traffic data does not necessarily improve the predictability

References

- W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic (extended version)", *IEEE/ACM Transactions on Networking*, vol. 2, No. 1, pp. 1-15, February 1994.
- M. E. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: evidence and possible causes", in *IEEE/ACM Transactions on Networking,* no.6, pp. 835-846, December, 1997.
- M. S. Taqqu and V. Teverovsky, "On estimating the intensity of long-range dependence in finite and infinite variance time series," in *A Practical Guide to Heavy Tails: Statistical Techniques and Applications.* Boston, MA, Birkhauser, pp. 177-217, 1998.
- P. Abry and D. Veitch, "Wavelet analysis of long-range dependence traffic", *IEEE Transactions on Information Theory*, vol.44, No. 1, pp. 2-15, January 1998.
- V. Paxson, "Empirical derived analytic models of wide-area TCP connections". *IEEE/ACM Trans. Networking*, vol. 2, no. 4, pp.316-336, 1994.
- A. Feldmann, "Characteristics of TCP Connection Arrivals", in *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, Eds. New York: Wiley, pp. 367-399, 2000.
- Zhiqiang Bi, Christos Faloutsos and Flip Korn, The "DGX" Distribution for Mining Massive, Skewed Data, in *Proc. of ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, pp.17-26, August 2001.
- V. Teverovsky *Estimation Methodologies* http: //math.bu.edu/people/murad/methods/

References

- P. Barforf and M. Crovella "Generating representative Web workloads for network and server performance evaluation" in *Proc of ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, p.151-160, June 22-26, Madison, Wisconsin, US. 1998.
- G. Box, and G. Jenkins, *Time Series Analysis: Forecasting and control* 2nd ed. Oakland, CA: Holden-day, 1976, pp. 208-329.
- P. Barford, A. Bestavros, A. Bradley, M. Crovella, "Changes in web client access patterns: characteristics and caching implications in world wide web", *Special Issue on Characterization and Performance Evaluation*, vol. 2, pp. 15-28, 1999.
- D. Veitch and P. Abry, "A statistical test for the time constancy of scaling exponents," *IEEE Trans. on Signal Processing*, vol. 49, no. 10, pp. 2325-2334, October 2001.
- S. Ostring and H. Sirisena, "The influence of long-range dependence on traffic prediction," *IEEE Int. Conf. On Communication*. vol. 4, pp. 1000-1005, 2001
- A. Aussem and F. Murtagh, "Web traffic demand forecasting using wavelet-based multiscale decomposition," in *International Journal of Intelligent Systems*, vol. 16, pp. 215—236, 2001.
- S. Ostring and H. Sirisena, "The influence of long-range dependence on traffic prediction," in *Proc. Int. Conf. Communication,* vol. 4, pp. 1000-1005, June, 2001.
- D. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot, "Long-Term Forecasting of Internet Backbone Traffic: Observations and Initial Models," in *Proc. IEEE INFOCOM 2003,* San Francisco, CA, 2003, pp. 1178-1188..

Thanks!