Global Modeling of Backbone Network Traffic

Stilian Stoev
Department of Statistics
University of Michigan, Ann Arbor

INFOCOM Mini–Conference
San Diego, 2010

(Joint work with George Michailidis and Joel Vaughan. Partially funded by the NSF grant DMS–0806094.)

Outline

Preliminaries

The Internet2 Backbone Network (fka Abilene)

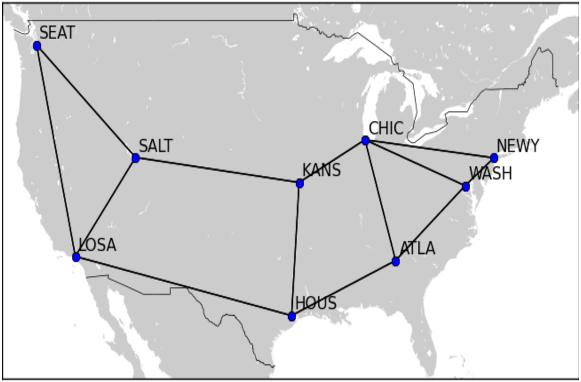

- ◦ 9 nodes, $L = 26$ uni–directional links, $J = 72$ routes (flows).
- ◦ most links are 10 Gbs/s, the rest are 20 Gbs/s.

Routing Matrix

- Consider a network with n nodes and L links.

◦ Traffic is routed between every pair of nodes along predetermined routes.

◦ Resulting in $J = n(n-1)$ flows associated with each (*source*, *dest*) pair.

- The routing matrix:

$$A = (a_{\ell j})_{L \times J}, \quad \text{with } a_{\ell j} = \begin{cases} 1 & , \text{ route } j \text{ involves link } \ell \\ 0 & , \text{ otherwise,} \end{cases}$$

describes the physical paths of the flows.

Our Framework

Notation:

∘ $X_j(t)$ is the traffic load (eg in bytes/sec) during $(t, t+1]$ along flow j, $1 \le j \le J$.

∘ $Y_\ell(t)$ is the traffic load during $(t, t+1]$ on link ℓ, $1 \le \ell \le L$.

∘ But how can one define or even measure flow–level traffic?

Key Assumption: On the time scale of interest traffic propagates instantaneously along the network, eg time units $>$ RTT.

The Routing Equation: With $\vec{X}(t) = (X_j(t))_{j=1}^J$ and $\vec{Y}(t) = (Y_\ell(t))_{\ell=1}^L$:

$$\vec{Y}(t) = A\vec{X}(t), \qquad \text{since } Y_\ell(t) = \sum_{j=1}^{J} a_{\ell j} X_j(t).$$

Objective

Problem: Construct a mechanistic model that captures and explains the *statistical behavior* of the traffic on *all* links and across time in the network.

- This is too ambitious!

◦ We will make some progress.

Guiding Principles:

◦ Include relevant network–specific information.

◦ Incorporate existing successful research on single–link modeling.

◦ Understand model limitations:
"All models are wrong, but some are useful" (George Box)

◦ Validate and apply the model to real data.

Global Traffic Models

Single–flow Models

• Well–established self–similarity and long–range dependence phenomena are present on single–link and single–flow level traffic.

• A number of physically interpretable and accurate models have been developed since the early 1990's. For details, see eg Taqqu, Willinger and Sherman (1997) and the monograph by Park and Willinger (2000), among many others.

On–Off Models: Multiple users share a flow and have On/Off type behavior:

$$X^{(i)}(t) = \begin{cases} 1 & , \quad i\text{–th user is On} \\ 0 & , \quad i\text{–th user is Off} \end{cases}$$

◦ The **durations** of the On/Off periods are independent and heavy–tailed:

$$\mathbb{P}\{U_{\rm on} > x\} \sim c_{\rm on} x^{-\alpha} \text{ and } \mathbb{P}\{U_{\rm off} > x\} \sim c_{\rm off} x^{-\alpha}$$

as $x \to \infty$, with

$$1 < \alpha < 2.$$

• So, the durations have finite means but infinite variances.

Limit Theorems for On–Off Sources

On flow j, the cumulative traffic is:

$$X_j(t) = \sum_{i=1}^{M} X_j^{(i)}(t),$$

where $X_j^{(i)}(t)$'s are the On/Off individual 'user' processes.
• What is the statistical behavior of the cumulative traffic fluctuations:

$$X_j^*(T,M) := \int_0^T X_j(t)dt - T\mathbb{E}X_j(0),$$

over judicious time scales T and number of users M?

Classical results: Taqqu et al (1997), Mikosch et al (2000)
• (fast growth) $M(T)/T^{\alpha-1} \to \infty$, as $T \to \infty$.

$$\mathcal{L} \lim_{T\to\infty} \frac{1}{T^H\sqrt{M(T)}} X_j^*(Tt, M) = B_H(t), \quad (t \geq 0),$$

where $B_H(t)$ is the fractional Brownian motion with self–similarity parameter

$$H = (3-\alpha)/2 \in (1/2, 1).$$

Limit Theorems for On–Off Sources (cont'd)

- (slow growth): $M(T)/T^{\alpha-1} \to 0$, as $T \to \infty$.

$$\mathcal{L} \lim_{T\to\infty} \frac{1}{(TM(T))^{1/\alpha}} X_j^*(Tt, M) = \Lambda_\alpha(t), \quad (t \geq 0)$$

where $\Lambda_\alpha(t)$ is the Lévy stable motion with independent increments and index α.

What is the significance? Traffic time series on large time scales T

$$X_j(t) = X_j^*(Tt) - X_j^*(T(t-1)), \quad t = 1, 2, \cdots$$

are well–modeled by:

(a) fractional Gaussian noise (fGn) $G(t) = B_H(t) - B_H(t-1)$ if many users share the flow.

(b) Lévy stable noise $S(t) = \Lambda_\alpha(t) - \Lambda_\alpha(t-1)$ if a few users are present.

- The long–range dependent fGn–type regime is more robust theoretically and in practice.

Illustration

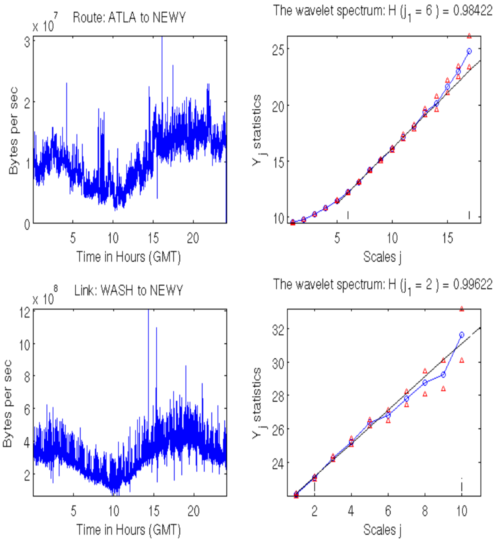

Network–Wide Models

Simplifying Assumptions:

1. Traffic propagates instantaneously.
2. Multiple independent users share each flow.
3. The traffic flows along different routes are independent

The routing equation $\vec{Y}(t) = A\vec{X}(t)$, then yields:
∘ **(fast regime)** If $M(t)/T^{\alpha-1} \to \infty$, $T \to \infty$, then

$$\mathcal{L}\lim_{T\to\infty} \frac{1}{T^H\sqrt{M(T)}} \int_0^{Tt} (Y_\ell(\tau) - \mu_{Y_\ell})d\tau = B(tf_\ell),$$

where $B(\cdot)$ is a certain functional fBm process, and the f_ℓ's are related to the routing matrix.
∘ **(slow regime)** If $M(t)/T^{\alpha-1} \to 0$, $T \to \infty$, then

$$\mathcal{L}\lim_{T\to\infty} \frac{1}{(TM(T))^{1/\alpha}} \int_0^{Tt} (Y_\ell(\tau) - \mu_{Y_\ell})d\tau = \Lambda(tf_\ell),$$

where Λ is a functional Lévy stable motion.

Intuition, Scope of Validity and Interpretation

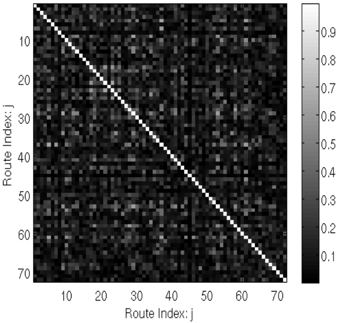

In reality:

• The flows are nearly uncorrelated.

○ True for non–congested networks.

○ Under heavy traffic the forward/reverse dependence should be modeled.

• We developed global models of the traffic fluctuations for all links and across time. Are they applicable or is this too good to be true?

○ The models work well! **Caveats:** time scales, non–congested network, multiple users, user homogeneity.

○ The functional fBm i.e. fast growth regime often prevails in practice.

○ Knowledge of the routing is key!

Network Kriging

An Application to Network Prediction

Fast Growth Regime

Over judicious time scales e.g. $> 100ms$ or $>$ RTT, the traffic loads on all L links

$$\vec{Y}(t) = (Y_\ell(t))_{1 \le \ell \le L}, \quad t = 1, 2, 3, \cdots ,$$

can be modeled as a stationary Gaussian vector time series with:

$$\mathbb{E}\vec{Y}(t) = \vec{\mu}_Y \ \text{ and } \ \mathrm{Cov}(\vec{Y}(t)) = \Sigma_Y = A\Sigma_X A^t,$$

where $\Sigma_X = \mathrm{diag}(\sigma^2_{X_j}, \ 1 \le j \le J)$ and A is the routing matrix.

Problem I: (Network Kriging) Observed are $\mathcal{O} \subset \{1, \cdots, L\}$ links at time t. Predict the unobserved links.

Problem II: (Temporal Prediction) Observed are $\mathcal{O} \subset \{1, \cdots, L\}$ links at times $t \le t_0$. Predict the unobserved and observed links at time $t_0 + h$.

Instantaneous Prediction or Kriging

Partition $\vec{Y}(t)$ into observed and unobserved components:

$$\vec{Y}(t) = \begin{pmatrix} \vec{Y}_u(t) \\ \vec{Y}_o(t) \end{pmatrix} \quad \text{and} \quad A = \begin{pmatrix} A_u \\ A_o \end{pmatrix}.$$

- The best m.s.e. predictor of $\vec{Y}_u(t)$ via $\vec{Y}_o(t)$ is:

$$\widehat{Y}_u(t) = A_u\mu_X + A_u\Sigma_X {A_o}^t(A_o\Sigma_X {A_o}^t)^{-1}(\vec{Y}_o(t) - A_o\mu_X).$$

◦ and the mean squared error matrix is:

$$\mathbb{E}\Big((\widehat{Y}_u(t) - \vec{Y}_u(t))(\widehat{Y}_u(t) - \vec{Y}_u(t))^t | Y_o(t)\Big) \\ = A_u\Sigma_X A_u^t - A_u\Sigma_X A_o^t(A_o\Sigma_X A_o^t)^{-1}A_o\Sigma_X A_u^t$$

Bonus: Confidence sets for $\vec{Y}_u(t)$ via $\widehat{Y}_u(t)$ and the m.s.e.

Illustration: The Good

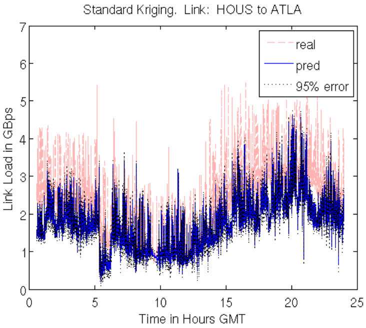

Illustration: The Bad

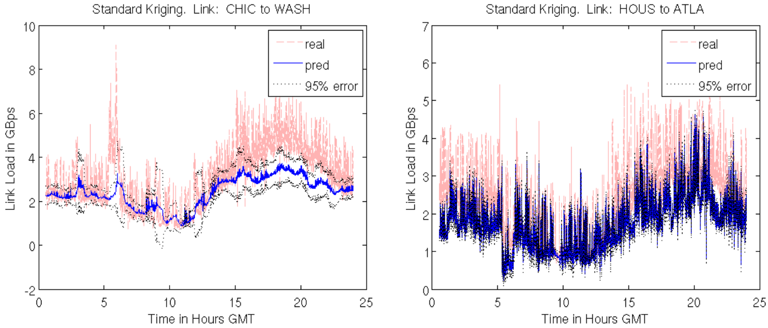

- The degree of predictability depends on the routing!

Illustration: The Ugly

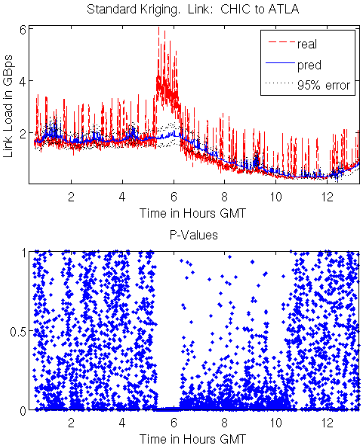

- The P–values clearly indicate the anomaly.
- When all links are observed, kriging provides tools for anomaly detection.

Summary & Future Work

- Constructed a global network traffic model applicable to:
  ◦ moderate an large scales (eg $>$ RTT).
  ◦ non–congested, backbone networks.
  ◦ monitoring, statistical anomaly detection, and network prediction.
  ◦ understanding the intricate network–wide traffic dependence.
- Limitations & future work
  ◦ forward/reverse flows are dependent under congestion
  ◦ the limit regimes are sometimes impure
  ◦ no knowledge of the traffic means – addressed in Vaughan, Stoev and Michailidis (2010)

Some References

D.B. Chua, E.D. Kolaczyk, and M. Crovella (2006) Network kriging. *IEEE Journal on Selected Areas in Communications*, 24(12):2263–2272, Dec. 2006.

M.E. Crovella and A. Bestavros (1996) Self-similarity in World Wide Web traffic: evidence and possible causes. In *Proceedings of the 1996 ACM SIGMETRICS. International Conference on Measurement and Modeling of Complex Systems*, pages 160–169, May 1996.

B.D'Auria and G. Samorodnitsky (2005) Limit behavior of fluid queues and networks. *Oper. Res.*, 53(6):933–945, 2005.

A. Erramilli, P. Pruthi, and W. Willinger (1995) Self-similarity in high-speed network traffic measurements: Fact or artifact? In *Proceedings of the 12th Nordic Teletraffic Seminar NTS12, Espoo, Finland*, pages 299–310, 1995.

Internet2: http://www.internet2.edu/observatory/

A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E.D. Kolaczyk, and N. Taft (2004) Structural analysis of network trafic flows. *Proceedings of Sigmetrics*, 2004.

Leland, Taqqu, Willinger and Wilson (1993) On the self-similar nature of Ethernet traffic. Proceedings of the ACM/SIGCOMM'93, San Francisco, September 1993.

T. Mikosch, S. Resnick, H. Rootzén, and A. Stegeman (2002) Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *The Annals of Applied Probability*, 12(1):23–68, 2002.

K. Park and W. Willinger (2000) *Self-Similar Network Traffic and Performance Evaluation*. J. Wiley & Sons, Inc., New York, 2000.

Y. Paschalidis, and G. Smaragdakis (2009) Spatio-temporal network anomaly detection by assessing deviations of empirical measures. *IEEE/ACM Trans. on Networking*, 17: 685–697, 2009.

S. Stoev, G. Michailidis, and J. Vaughan (2009) On Global Modeling of Network Traffic.
http://www.stat.lsa.umich.edu/~ sstoev/global_tr.pdf

M. S. Taqqu, W. Willinger, and R. Sherman (1997) Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review*, 27(2):5–23, 1997.

J. Vaughan, S. Stoev, and G. Michailidis (2010) Network–wide statistical modeling and prediction of computer traffic. Working paper,