



Overview: Traffic Modeling & Performance Analysis

Ian Akyildiz, Georgia Tech
John Baras, University of Maryland
William Cleveland, Purdue University
Jennifer Neville, Purdue University
Ness Shroff, Ohio State University
Mark Ward, Purdue University



Wide Range of Network Traffic Types and Performance Issues



- Packet-level traffic
 - performance for disparate applications: bulk data transfer and realtime
 - scheduling, queuing, load, and throughput
- Connection-level traffic
 - information assurance
 - monitoring and control for attacks and intrusions
- Battlefield spatial and temporal traffic among personnel
 - temporal traffic characterization at traffic nodes
 - assess loads on communications systems

2

Packet-Level Modeling and Performance



- **The Network Setting:** packets arriving at a network device for further network operations such as
 - queuing and network layer forwarding and transmission over a link
 - scheduling for link layer forwarding and protocol conversion
- The traffic is the arrival times and sizes of the packet sequence

- Models for traffic
 - validity: must account for observed traffic phenomena such as nonlinearity and long-range dependence
 - mathematical study of models to develop simple explanations of phenomena and to devise approximations for specific tasks to enable tractable solutions

3

Packet-Level Modeling and Performance



- Characterize properties of queueing delay and jitter
 - run simulations with validated traffic models generating arrivals and sizes
 - delay and jitter are "data" whose statistical properties we seek to characterize
 - use approximations of traffic models to derive queueing properties and then check if derived results are good approximations to the jitter and delay properties

- Performance
 - realtime: measured by packet delay and jitter
 - file transfer: measured by delay
- Control
 - provisioning: determining amount of traffic
 - algorithms for different network layers

4

Connection-Level Modeling and Performance



- **The Network Setting:** connections occurring for hosts on an “inside” network, each involving one inside host and with an inside or outside host
- The traffic is at the connection level
 - each connection summarized by ≈ 15 variables derived from packet headers and timestamps
 - connection traffic: sequences through time of connection variables
- Control
 - surveillance, alarms, automated defensive actions, and rapid forensics
 - requires the processing of evolving connection traffic: billions of connections at each moment

5

Connection-Level Modeling and Performance



- Models for the time sequences of the traffic variables
 - must take account of the network topology through characterization of locality
 - must take into account the operation of protocols, chiefly applications, that create relationships among the connections
 - must take into account legitimate human network usage and attacker-intruder patterns that create relationships among connections
- This whole area is in its infancy
 - where packet-level traffic study for performance and control was in 1990
 - must establish statistical and mathematical foundations
 - foundations must be based on empirical study
 - some packet-level statistical properties like heavy-tailed distributions are critical

6

Tactical Network Modeling and Performance



- **The Network Setting:** soldiers on the move, central command, and events
- Traffic
 - heterogeneous
 - can be highly correlated at both time and space domains
 - can be modeled jointly by spatial correlation functions, event source models, and mobility models
- Research challenges
 - investigate the effect of spatially co-located point sources on the statistical properties of the traffic.
 - investigate the characteristics of traffic generated from a large area modeled by the field source.
 - develop a theoretical framework to find the relationship between the long range dependent (LRD) characteristics of the traffic and the location they are generated from.
 - apply the theoretical techniques from military networks to commercial networks.
 - realtime performance issues

7

Data Collection for Empirical Study



- Packet trace data collection
 - all packets
 - network and transport layer headers plus timestamps
 - collection and analysis challenging
 - need permissions and cooperation of local network administrators
 - privacy a major matter
 - databases get large very fast
- Connection summary collection
 - process all headers and timestamps
 - keep only connection summaries
- Our Analytic Computational Environment
 - linux cluster
 - R interactive system for data analysis: based on S language (1999 ACM Software System Award)
 - Snow: within-R distributed computing
 - Hadoop: parallel R processes running same analyses on data subsets
 - distributed file system reproduces data and distributes across cluster
 - distributed computational engine

8

Data Collection: Operations



- Global Crossing VoIP (completed)
 - partnership with Global Crossing (GBLX) and Cariden, Inc.
 - 100 mbps Ethernet link on the GBLX IP network
 - voice gateway
 - edge router
 - two days of traces
 - times stamps plus UDP, RTP, and IP headers
 - both directions of 328,631 calls
 - 1/2 terabyte of data
- Purdue Statistics subnet (ongoing)
 - wireline and wireless traffic
 - software collection using pcap library at a router span port
 - packet traces
 - 262×10^6 packets/day
 - 19.6 gigabytes/day
 - connection summaries
 - 96×10^3 connections/day
 - expect 15 summary statistics/day
 - 10^5 -row 15-column array/day

9

Data Collection: Operations



- Purdue 1-gig campus link (implementing)
 - partnership with CAIDA (Cooperative Association for Internet Data Analysis)
 - wireline and wireless
 - hardware collection with Endace 1-gig card
 - monitor supplied by CAIDA
 - much higher traffic rate than Statistics subnet
 - many more inside hosts than Statistics subnet
 - awaiting link choice and installation
- Purdue bracketed wireless collection (planning)
 - bracketed collection provides very powerful study of performance at a wireless interface
 - simultaneous collection at two points with synchronized clocks
 - point 1: a single-user host running test traffic over a wireless access point
 - point 2: the first router for the access point
 - can implement some control procedures on the single-user host and other selected hosts

10

Research Project: Modeling for Performance: LRD — Causes, Effects from the Network, Models



- Causes of LRD traffic in mobile wireless networks: What controls the LRD characteristics?
 - dependence on technologies used
 - dependence on protocols used
 - dependence on types of traffic demand and user behaviors
 - measurements-experiments-testbeds-validation
- Selecting models for LRD phenomena local-global
 - Beyond graphs — time varying multi-fractal time varying conditional Markov Random Fields (MRF)
 - statistical physics methods and analogies - phase transitions - critical parameters
 - complex dynamical systems — chaos
- Loss models with LRD traffic — Fixed Point models — validation
- Effects of topology and connectivities on LRD characteristics

11

Research Project: Modeling for Performance: Fractional Sum-Difference Models for Packet Arrivals



- Simple models that account for nonlinearity and long-range dependence
 - The key discovery: nonlinear monotone transformation of interarrivals results in sum of white noise plus a near-fbm process
 - Extensive validation shows excellent fit for traffic rates down to TCP footprint limit (about 3-5 mbps)
- Outcome
- simplicity allows mathematical study
 - explicit formulas for traffic properties previously only described empirically
 - much simpler intuition for traffic phenomena such as scaling properties

12

Research Project: Parametric and Semi-Empirical Modeling of VoIP Traffic



- “Calls” are a mixture of voice, fax, modem logins, credit card authentications, etc.
- Call = connection
- Statistical properties
 - call arrival process
 - call durations
 - silence suppression
 - packet-level traffic
- Two forms of modeling
 - semi-empirical
 - full mathematical statistical modeling

13

Research Project: Modeling of Queueing Delay and Jitter Stochastic Processes



- Simulation study of delay and jitter using semi-empirical or mathematical statistical models for traffic inputs to the queue
- Mathematical study using approximations of the traffic models to enable tractable results for delay and jitter

15

Research Project: Delay Analysis for Wireless Networks



- Goal: Analyze the expected delay
 - multi-hop wireless networks
 - taking interference constraints into account
 - valid for any traffic load (not just heavy traffic load with high utilization)
- Approach: Build on method of Lyapunov drifts
 - derive tight system-wide lower bound
 - derive upper bounds based on scheduling policy
 - find throughput-optimal algorithms whose performance is close to lower bound or
 - near-delay-optimal
 - obtain insights on designing low-complexity sub-optimal schedulers with provably efficient delay bounds.

14

Research Project: Provisioning Links



- Determine amount of traffic that can be put on a link (or series of links) and maintain high performance
- Realtime scenarios: performance is measured by packet delay and jitter
 - realtime traffic on link by itself
 - realtime traffic with priority queueing on link with best-effort traffic
 - realtime traffic with no priority queueing on link with best-effort traffic
- File transfer: performance is measured by delay
- Seek to characterize and clarify the nature of “multiplexing gains”

16

Research Project: Modeling and Analysis of Tactical Networks



- Heterogeneous structure and hierarchical architecture of tactical networks result in correlation in multiple scales
- Investigating the impact of the correlation characteristics on the network traffic can help to predict and control network traffic

Time-Space Modeling

- 1 Select an event source model
Select a mobility model for the soldier
- 2 Build ON-OFF communications traffic model for each soldier
- 3 Evaluate the distributions of ON-periods & OFF-periods
- 4 Combine the traffic at the commander based on spatial correlation
- 5 Analyze and model the commanders traffic

17

Research Project: Modeling Connection Traffic



- Extensive study of properties provides valid models
- Initial thrust for enterprise networks (civilian and government)
- Center modeling on inside hosts

18



Multifractal and Gaussian Fractional Sum-Difference Models for Best-Effort Packet Arrivals

William S. Cleveland, Purdue University



Modeling Multiplexed Best-Effort Packet-Arrival Traffic



- All early work and most work since: time-aggregated arrivals
 - e.g., packet arrival counts in fixed time intervals such as 1 ms
 - W. Leland and M. Taqqu and W. Willinger and D. Wilson (1994), On the Self-Similar Nature of Ethernet Traffic
 - Vern Paxson and Sally Floyd (1994), Wide-Area Traffic: The Failure of Poisson Modeling
- First studies of arrivals without aggregation
 - interarrival time sequence
 - Rudolf H. Riedi and Matthew S. Crouse and Vinay J. Ribeiro and Richard G. Baraniuk (1999), A Multifractal Wavelet Model with Application to Network Traffic
 - J.B. Gao and I. Rubin (2001), Multiplicative Multifractal Modeling of Long-Range-Dependent Network Traffic
 - J. Cao and W. Cleveland and D. Lin and D. Sun (2001), On the Nonstationarity of Internet Traffic

2

Time Aggregation



- Aggregation to make modeling easier
 - arrival process is long-range dependent and nonlinear
 - arrival measurements give rise to very large datasets
 - aggregation offers an “opportunity” to deal only in Gaussian models
 - aggregation can reduce dataset size by orders of magnitude
 - these are not good reasons for aggregates
- Aggregation to make performance analysis feasible
 - mathematics for performance analysis becomes tractable
 - can carry out large-scale simulation for performance analysis
 - these are good reasons for aggregates

Solution: aggregate after we have the models

3

Modeling the Interarrival Process — 1



- The arrival point process is
 - the fundamental mathematical process
 - what the routers see
- A simple, valid model for the interarrivals could provide
 - a mechanism for exploring mathematically how well an aggregation approximation represents the actual arrival point process
 - investigation of other approximations

4

Modeling the Interarrival Process — 2



■ Properties of Internet traffic have been described empirically by nonparametric statistics that estimate population quantities and not mathematically (i.e., formulas with which we can work)

- power spectrum
- autocorrelations
- process variance

■ A simple, valid model for the interarrivals could provide

- mathematical formulas in place of just numeric values, for use in performance studies

5

Modeling the Interarrival Process — 3



■ t_u = interarrival sequence and $a_r[m] = m$ -aggregate for $m = 1, 2, \dots$

$m = 1$: $| t_1 | t_2 | t_3 | t_4 | t_5 | t_6 | t_7 | t_8 | \dots$

$m = 2$: $| t_1 + t_2 | t_3 + t_4 | t_5 + t_6 | t_7 + t_8 | \dots$

$m = 3$: $| t_1 + t_2 + t_3 + t_4 | t_5 + t_6 + t_7 + t_8 \dots$

■ Basic intuition about traffic has been driven to a large extent by time-scaling properties:

- statistics of $a_r[m]$ as a function of m
- some statistics very complex
 - moments in multifractal analysis
 - multifractal wavelets
 - hard to relate back to the traffic

■ A simple, valid model for the interarrivals could provide

- a simpler intuition for traffic properties and ultimately performance issues, in place of complex statistics

6

Fractional Sum-Difference (FSD) Models



Introduced casually in J. Cao and W. Cleveland and D. Lin and D. Sun (2001), On the Nonstationarity of Internet Traffic

Very simple models with three free parameters

Vehicle to show change in traffic statistics with traffic load

Not fully validated

Properties not fully explored

Mathematical implications for traffic phenomena and performance not explored

7

Current Project: Validation and Mathematical Properties of FSD Models



Multifractal, Multiplicative, and Gaussian FSD Models

Completed extensive validation study

- multifractal FSD model recreates interarrival traffic statistics extremely well

Discovered that an even simpler multiplicative, multifractal FSD model is a very good approximation of the best-fitting multifractal FSD model

Using the multiplicative FSD model, deriving interpretable formulas and developing a new and more fundamental intuition that relates to the traffic more effectively than

8

A Simplifying Monotone Transformation



Multifractal Domain

t_u = interarrival sequence

$$t_u = \Omega(z_u)$$

Multifractal FSD

$$z_u = s_u + n_u$$

s_u and n_u are independent

s_u is a moving sum of length 2 of the Hosking discrete analog of fBm

n_u is white noise

Gaussian Domain

z_u = Gaussian image of t_u

$$z_u = \Omega^{-1}(t_u)$$

Gaussian FSD

T = cdf of t_u

Z = cdf of normal with mean 0 and variance 1

$$\Omega = Z^{-1}(T)$$

For Internet packet interarrivals, T is Weibull

9

The Multiplicative FSD Approximation



In place of Ω , we use \log , so $\Omega^{-1} = \exp$

This means that

$$t_u \approx e^{s_u + n_u}$$


t_u is approximately the product of the exponent of a near-fBm process and the exponent of Gaussian white noise

10

Multiplicative FSD Reproduces Multifractal Properties



$M_{a[m]}(q) = q$ th moment of m -aggregate for $q = -10, -5, -2, 2, 5, 10$ vs. function of m



11

The Variance-Time Plot



A canonical Internet traffic analysis method

$\log(\text{Variance}(a_r[m]/m))$ vs. $\log(m)$

For self-similar sequence (process), linear with slope $2d - 1$ where $d = (\text{Hurst parameter} - 0.5)$, and $0 \leq d < 0.5$

In the first studies (early 90s) of Internet traffic the plot was linear so “self-similar” was used to describe the traffic

By the late 90s

- the plot was showing up convex
- slope = -1 for small m
- slope = $2d - 1$ for large m
- so “asymptotically self-similar” was used to describe the traffic
- there was a tendency to ascribe this to changing Internet technology

12

Heuristics for Variance-Time Plot for the Gaussian FSD



$z_u = s_u + n_u$ where $s_u = \text{near-fBm}$ and $n_u = \text{white noise}$

Aggregation is a low-pass filter whose pass-band moves toward the origin as m increases, reflecting at first the rapid attenuation of n_u , whose variance goes down like m^{-1} , and then settling in on s_u , whose variance goes down like m^{2d-1}

This intuition can be extended to the multifractal domain via the multiplicative model

13

FSD Model for Increased Traffic Loads



$z_u = s_u + n_u$ where $s_u = \text{near-fBm}$ and $n_u = \text{white noise}$

As the load increases (more multiplexing), the fraction of the variance of z_u due to n_u increases toward 1

We can relate this increase to the traffic rate

This means the multifractal FSD is tending toward Poisson

14