

Network architectures evolution and teletraffic theory: general principles and open issues

Michele Pagano

e-mail:m.pagano@iet.unipi.it

Dipartimento di Ingegneria dell'Informazione
Università di Pisa




ACMPT 2017
Moscow, October 25th

Outline

- 1 Telephone networks and Classical Teletraffic Theory
 - Public Switched Telephone Network
 - Teletraffic theory
- 2 Network architecture of the Internet
 - Packet Switching Principles
 - Congestion control
 - TCP Linux
- 3 Internet traffic features
 - POTS vs. Internet
 - Why fractals?
 - Queueing performance
 - Modelling LRD through MAPs
- 4 Conclusions

Plain Old Telephone Service (POTS)



- **Circuit switching**
 - Resources reserved for call
 - link bandwidth
 - switch capacity
 - Dedicated resources: no sharing
 - Call set-up required
- **Teletraffic theory and POTS:** one of the most successful applications of mathematics in industry

POTS and teletraffic theory

- **Highly static nature of PSTN**

- *Homogeneous systems*
- *Limited variability*

⇒ Existence of universal laws

- **Agner Krarup Erlang (1878-1929)** – Danish mathematician, the father of queueing theory, has worked for the Copenhagen Telephone Company (KTAS in Danish) for almost 20 years
 - The **Erlang B formula** gives the (steady-state) probability that a trunk is not available as a function of the load and the number of trunks in a loss system
 - The **Erlang C formula** gives the (steady-state) probability that an arrival must wait before beginning service in a delay system

General features of POTS traffic

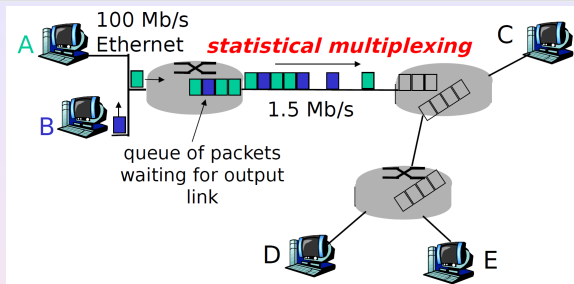
- **Poisson nature of call arrivals at links where traffic is highly aggregated**
 - *Parsimonious traffic model*

Palm–Khintchine theorem

The superposition of a large number of independent equilibrium renewal processes, each with a small intensity, behaves asymptotically like a Poisson process


- **Call holding times follow more or less an exponential distribution**
 - *Insensitivity property of the Erlang B formula (Kosten, 1948)*
- Mathematically tractable models could be used to predict accurately many performance measures of interest

Packet Switching



- Each end-end data stream is divided into *packets*
 - Users share network resources: **statistical multiplexing**
 - Each packet uses full link bandwidth
 - Resources are used *as needed*
- **Resource contention**
 - Aggregate resource demand can exceed the amount available
 - Need for end-to-end **congestion control** mechanisms

Node delay




- Processing delay d_{proc}
- Queueing delay d_{queue}
- Transmission delay d_{trans} — the delay between the times that the first and the last bits of the packet are transmitted

$$d_{\text{trans}} = L/R$$


where L is the packet length and R is the transmission rate

- Propagation delay d_{prop}

Packet length distribution




Packet length distribution



Transmission Control Protocol (TCP)


- Connection-oriented transport protocol that provides a reliable byte-stream data transfer service between pairs of processes
- Key features:
 - Connection-oriented
 - Multiplexing/Demultiplexing
 - Reliability
 - Flow Control
 - Congestion Control — TCP sensitive to network conditions



TCP Congestion Control

- TCP seeks to
 - achieve high utilization
 - *control* congestion
 - share bandwidth
- TCP Congestion Control is **window-based**
 - **cwnd** – state variable that limits how much data TCP is allowed to have in transit
 - A TCP source calculates **cwnd** according to the level of congestion *it perceives to exist* in the network
- TCP assumes **packet losses are caused by congestion**
- Behaviour of **cwnd**
 - no losses \Rightarrow more bandwidth is available \Rightarrow **cwnd** \nearrow
 - loss of a packet \Rightarrow network congestion \Rightarrow **cwnd** \searrow
- **Differentiation between major and minor congestion events**
 - Introduction of **Fast Recovery** mechanism

Classical TCP Congestion Control (TCP Reno)




Main TCP Linux Variants

Cubic and Reno (NewReno)

- Loaded into the kernel, via standard kernel module mechanism
- Information available in `/proc/sys/net/ipv4`
 - `tcp_allowed_congestion_control` — cubic reno
 - `tcp_congestion_control` — **cubic**

Key features of TCP CUBIC

- Cubic growth of cwnd
- Default since 2.6.19 Linux kernel



```
ls -a /lib/modules/$(uname -r)/kernel/net/ipv4/tcp*
```





kernel 2.6.35-30

- tcp_bic.ko
- tcp_highspeed.ko
- tcp_htcp.ko
- tcp_hybla.ko
- tcp_illinois.ko
- tcp_lp.ko
- tcp_probe.ko
- tcp_scalable.ko
- tcp_vegas.ko
- tcp_veno.ko
- tcp_westwood.ko
- tcp_yeah.ko

kernel 4.4.0-97

- tcp_bic.ko
- tcp_cdg.ko
- tcp_dctcp.ko
- tcp_diag.ko
- tcp_highspeed.ko
- tcp_htcp.ko
- tcp_hybla.ko
- tcp_illinois.ko
- tcp_lp.ko
- tcp_probe.ko
- tcp_scalable.ko
- tcp_vegas.ko
- tcp_veno.ko
- tcp_westwood.ko
- tcp_yeah.ko

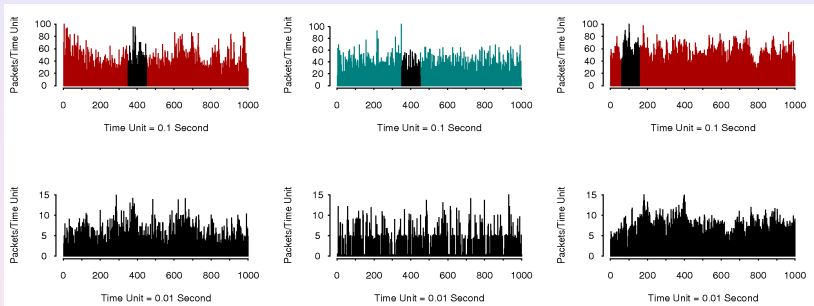
Behaviour of some Linux TCP Variants



Statistical features of data traffic

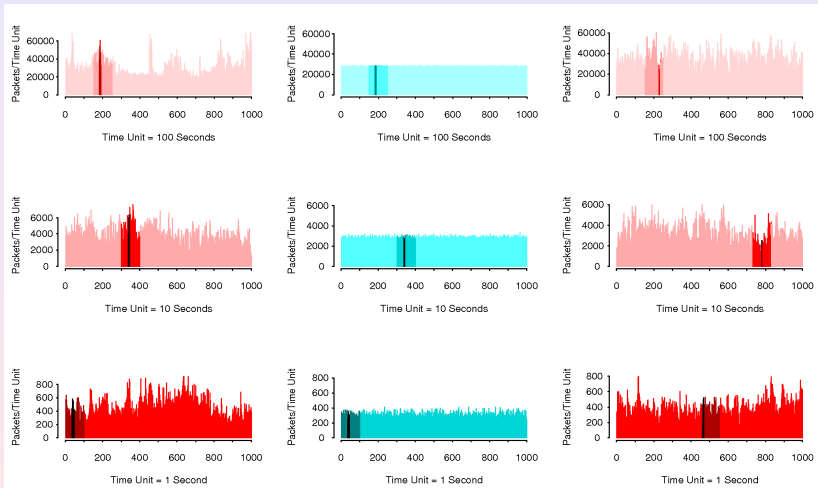
- **End-to-end congestion control**
 - Traffic is shaped by the conditions each connections has encountered in its past
- **Data traffic is much more variable than voice traffic**
 - Mice and elephant data flows
 - Highly different rates
 - Heterogeneous applications
 - High burstiness
- **Mathematics of high or extreme variabilities**
 - Temporal high variability \Rightarrow Long Range Dependence
 - Spatial high variability \Rightarrow Heavy-tailed distributions
 - Lack of a dominant *time scale* \Rightarrow Fractals

Traffic Self-similarity



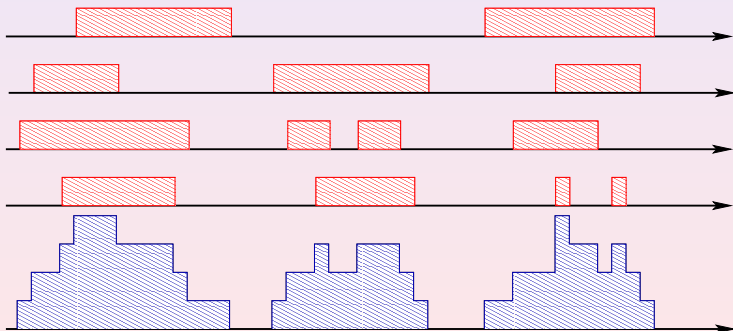
M. S. Taqqu, W. Willinger, R. Sherman *Proof of a fundamental result in self-similar traffic modeling*, Computer communication review, 1997

Traffic Self-similarity



Limit theorems for aggregated traffic

- Let us consider a network with a **high number of hosts** communicating with each other
- Each source is modeled according to a **binary On-Off alternating renewal process**



Flow characterization

Fluid on-off model

$$I(t) = \begin{cases} 0 & t \in \text{Off interval} \\ 1 & t \in \text{On interval} \end{cases}$$

- Denote by μ_{on} and μ_{off} the mean sojourn times in **On** and **Off** states
- Assume that

$$\bar{F}_{\text{on}} \simeq \ell_{\text{on}} x^{-\alpha_{\text{on}}} L_{\text{on}}(x) \quad (1 < \alpha_{\text{on}} < 2)$$

$$\text{or} \quad \sigma_{\text{on}}^2 < \infty \Rightarrow \alpha_{\text{on}} \stackrel{\Delta}{=} 2$$

and

$$\bar{F}_{\text{off}} \simeq \ell_{\text{off}} x^{-\alpha_{\text{off}}} L_{\text{off}}(x) \quad (1 < \alpha_{\text{off}} < 2)$$

$$\text{or} \quad \sigma_{\text{off}}^2 < \infty \Rightarrow \alpha_{\text{off}} \stackrel{\Delta}{=} 2$$

Superposition of On–Off sources: Key result

- Aggregate cumulative packet counts for N IID sources in the interval $[0, tT]$

$$A_N(tT) = \int_0^{tT} \left(\sum_{k=1}^N I_k(u) \right) du$$

- Convergence to fractional Brownian motion (fBm) $Z_H(t)$ for high values of N and T


$$\lim_{T \rightarrow \infty} \lim_{N \rightarrow \infty} \frac{1}{T^H \sqrt{L(T)} \sqrt{N}} \left(A_N(tT) - TN \frac{\mu_{\text{on}}}{\mu_{\text{on}} + \mu_{\text{off}}} t \right) \stackrel{(d)}{=} \sigma_{\text{lim}} Z_H(t)$$

where

$$H = \frac{3 - \alpha_{\text{min}}}{2} \in (1/2, 1) \quad \text{and} \quad \alpha_{\text{min}} = \min(\alpha_{\text{on}}, \alpha_{\text{off}})$$

- Noah Effect at source level (heavy tails) produces aggregate network traffic that exhibits the Joseph Effect (self-similarity)

Single server queue with infinite buffer



- $A(t) = mt + X(t)$ – Gaussian traffic model
- Constant service rate $C > 0$ and $r \triangleq C - m > 0$

Logarithmic large buffer asymptotic (LDT result)

$$\mathbb{P}(Q > b) \asymp \sup_{t \in \mathbb{R}} \exp\left(-\frac{(b + rt)^2}{2v(t)}\right) = \exp\left(-\inf_{t \in \mathbb{R}} \frac{(b + rt)^2}{2v(t)}\right)$$

where $v(t) = \mathbb{D}X(t)$

Large buffer asymptotic ($b \rightarrow \infty$)

Exact asymptotic for fBm

$$\mathbb{P}(Q > b) \sim \frac{\alpha(H)}{\sqrt{2\pi\beta(H)}} \cdot b^\gamma \exp(-\Theta b^{2-2H})$$

where $\gamma = 2H - 3 + \frac{1}{H}$


$$\alpha(H) \triangleq \frac{\mathcal{H}_{2H}\sqrt{\pi}}{2^{(1-H)/2H}\sqrt{H}} \left(\frac{H}{r(1-H)}\right)^{H-1} \left(\frac{1}{1-H}\right)^{(2-H)/H}$$

$$\beta(H) \triangleq \left(\frac{r(1-H)}{H}\right)^H \frac{1}{1-H}$$

Pickands constants

$$\mathcal{H}_\alpha \triangleq \lim_{T \rightarrow \infty} \frac{1}{T} \cdot \mathbb{E} \exp \left(\sup_{t \in [0, T]} \left(\sqrt{2} B_{\alpha/2}(t) - t^\alpha \right) \right)$$

where $\alpha \in (0, 2]$ and $B_{\alpha/2}$ – fBm with Hurst parameter $\alpha/2$

Asymptotic ($b \rightarrow \infty$) for fBm

Approximating LRD

Motivations

- Difficulty in handling the complex mathematical structure of self-similar processes
- Correlation beyond a *certain threshold* does not influence the queueing performance

Superposition of MMPP(2)

- Superposition of d two-state MMPPs
- 2^d state MMPP

$$\mathbf{D}_0 = \bigoplus_{i=1}^d \mathbf{D}_0^{(i)} \quad \mathbf{D}_1 = \bigoplus_{i=1}^d \mathbf{D}_1^{(i)}$$

MMPP fitting - Bellcore Data

$$d = 4, n = 5, \bar{\lambda} = 6.3, H = 0.85 \text{ and } \rho = 0.15$$

	$\lambda^{(i)}$	$\delta_1^{(i)}$	$\delta_2^{(i)}$
IPP(1)	2.062	$4 \cdot 10^{-1}$	$4 \cdot 10^{-1}$
IPP(2)	1.469	$8.618 \cdot 10^{-3}$	$8.618 \cdot 10^{-3}$
IPP(3)	0.427	$1.857 \cdot 10^{-4}$	$1.857 \cdot 10^{-4}$
IPP(4)	0.602	$4 \cdot 10^{-6}$	$4 \cdot 10^{-6}$
Poisson	$\lambda_P = 4.020$		

MMPP vs. IPP

It is always possible to interpret the **superposition of i SPPs** as a **superposition of i IPPs** and a **Poisson process**

Conclusions

- IETF (Internet Engineering Task Force) motto
 - We reject kings, presidents and voting.*
 - We believe in: rough consensus and running code*
- Differences between POTS and Internet
 - Circuit vs. packet switching
 - Homogeneous vs. heterogeneous systems
 - Limited variability vs. burstiness
 - Poisson vs. LRD
 - Exponential distribution vs. heavy tails
- Relevance of LRD in terms of network performances
 - Buffers are not the solution!
- Big challenges for mathematicians and statisticians
 - Parsimonious modelling
 - Parameter estimations
 - Queueing performance

