

EP2210 – Performance evaluation of communication networks

Viktor Fodor

Laboratory for Communication Networks

School of Electrical Engineering, KTH

Course homepage

http://www.kth.se/student/kurser/kurs/EP2210?l=en_UK

<https://www.kth.se/social/course/EP2210/>

Fall 2012 (HT 2012)

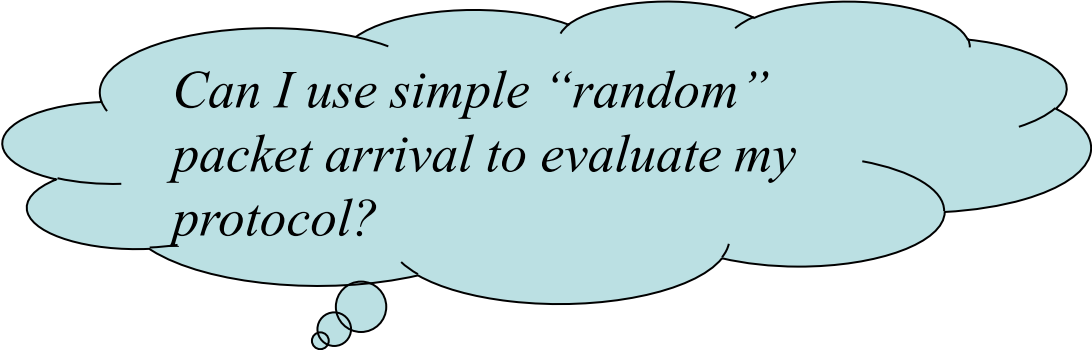
EP2210 – Performance evaluation of communication networks

Course objectives:

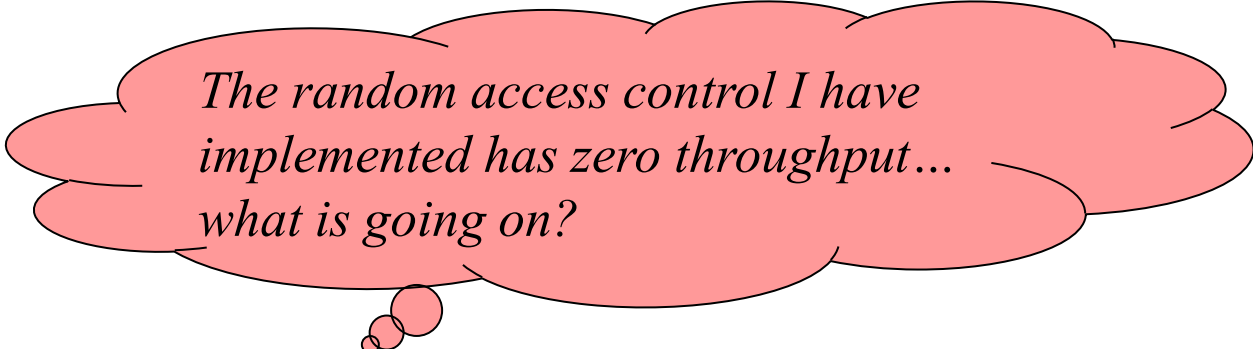
- Advanced networking course
- Discuss mathematical modeling in some main areas of networking
 - Learn techniques to address performance related questions
 - Discuss some of the significant results – and read the original papers
 - Improve our “paper reading” (and writing) skills

Topics

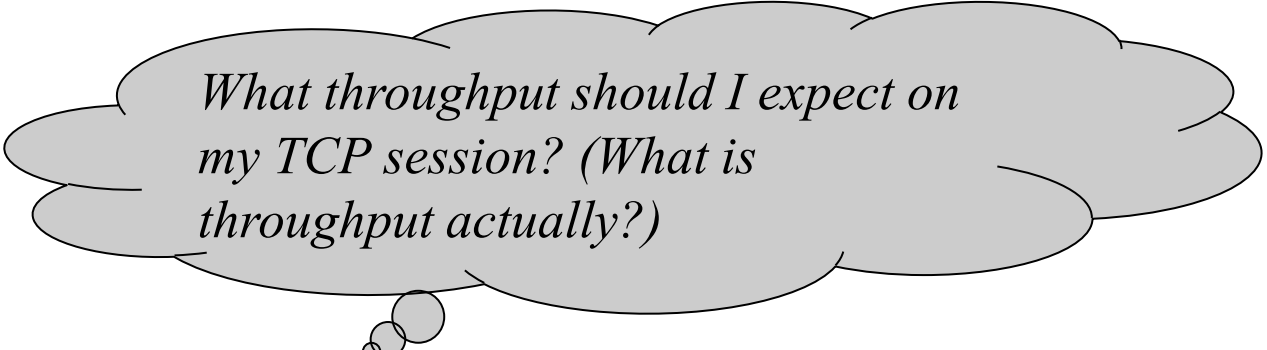
1. Traffic modeling
2. Multiple access protocols
3. Congestion control



Can I use simple “random” packet arrival to evaluate my protocol?



The random access control I have implemented has zero throughput... what is going on?



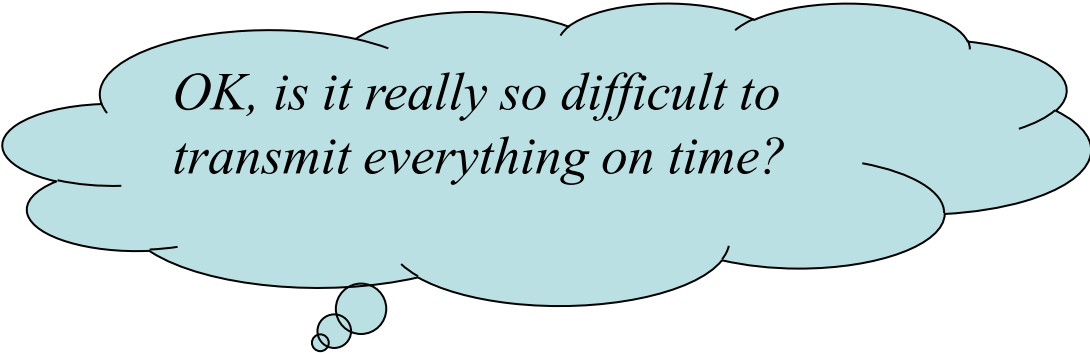
What throughput should I expect on my TCP session? (What is throughput actually?)

Topics

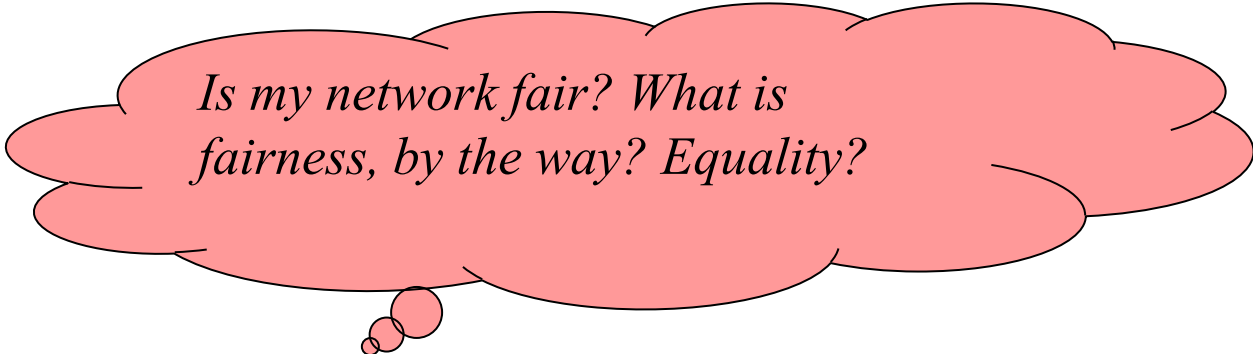
3. Scheduling

4. Fairness

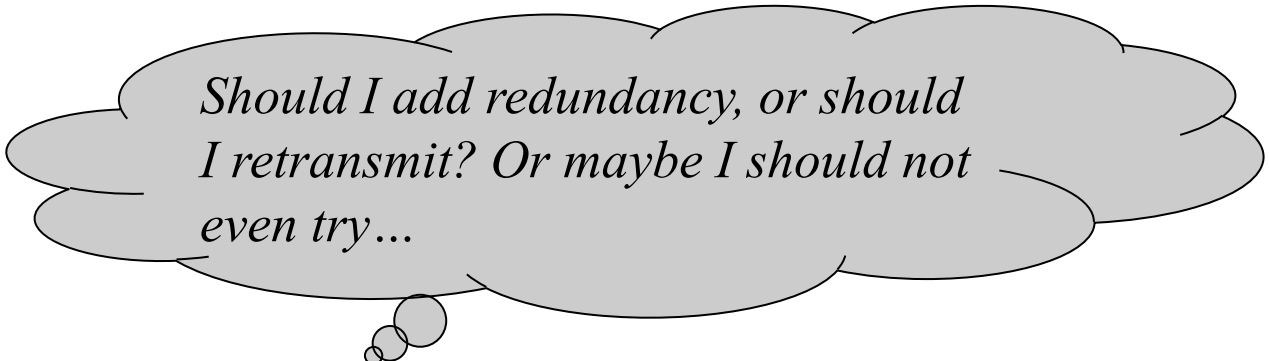
5. Multimedia
communication



*OK, is it really so difficult to
transmit everything on time?*



*Is my network fair? What is
fairness, by the way? Equality?*



*Should I add redundancy, or should
I retransmit? Or maybe I should not
even try...*

Course setup

- Scheduled activities:
 - 12 lectures of 2 hours
 - make-up test opportunity (Oct 9)
 - project presentations (Oct 11)
- 2 lectures per subject
 - first lecture – introduction and simple models
 - second lecture – advanced models, discussion of papers, phd student presentations
- Midterm tests (5 tests altogether, lectures 3, 5, 7, 9, 12)
- Home assignments (3 home assignments altogether, submitted at lectures 5, 9, 12)
- Project

Requirements

- Read all the papers
 - covering the lecture and for home reading
- Home assignments
 - questions to answer
 - numerical examples (e.g., matlab)
 - independent solutions, submit one paper copy at the lecture or at STEX
 - tell me in advance if you can not submit on time (minus points)
- Tests
 - ca. 20 minutes
 - questions on the lecture material and about the papers (open book)
 - make-up test after the course (missed or weak results)

Requirements

- Project
 - in groups of ca. 2 students
 - subject selected from subject list or on your own (discuss with the instructor)
 - comparative review of 4-5 papers in the area
 - written report of 4-5 pages
 - presentation of the project
 - schedule: defined later
 - groups: defined later by the instructor based on the study results in the course
 - good reports from earlier years are available on the web

Grading

- Tests: 50%
- Home assignments: 30%
- Project 20% (same for all project members)
 - detailed on the web-page under Projects
- Grading guidelines (approx):
 - 90%:-A, 80%-B, 70%-C, 60%-D, 50%-E, 45%-Fx

Requirements – graduate students

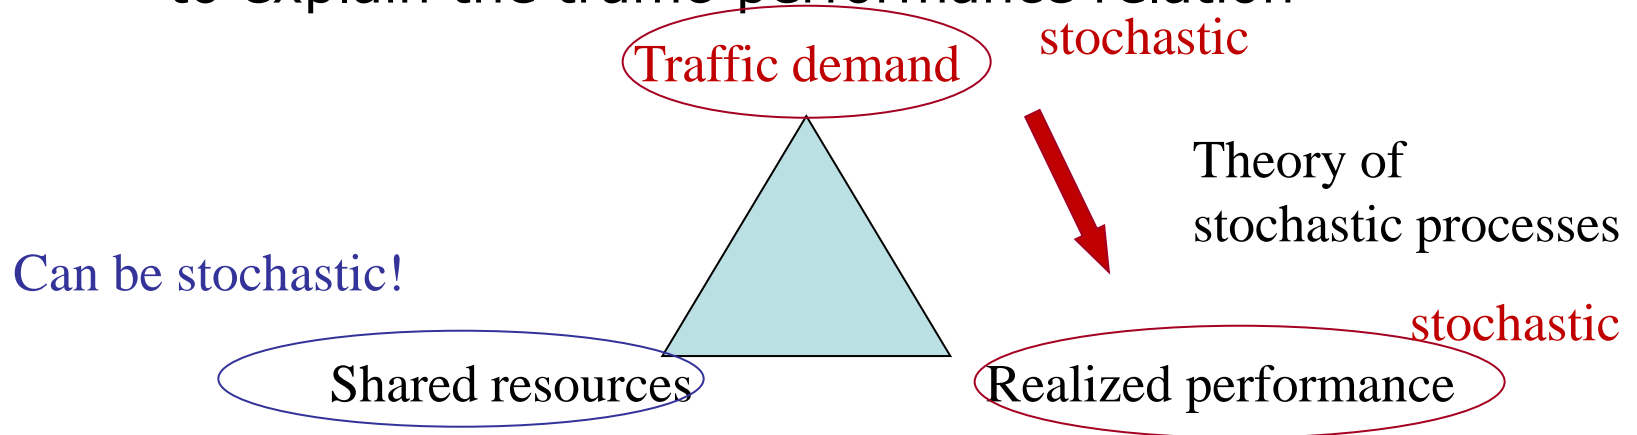
- Paper presentation (for 9ECTS)
 - select a lecture topic before lecture 3 (shared presentations)
 - ca. 20 minutes presentation on one of the lectures
 - short meeting right after today's lecture about the details
- Small project – during or after the course (for +3ECTS)
 - select a lecture topic
 - prepare a small simulator to support a mathematical model or problem definition, the simulator could be used for demonstration
 - see examples under Course Material

Traffic theory - Traffic models

- Topics:
 - Traffic modeling – traffic objects
 - Markov processes recall
 - Traffic models: markovian and non-markovian models
- Lecture material:
 - A. Adas, "Traffic models in broadband networks," IEEE Communications Magazine, July 1997.
 - J. Roberts, "Traffic theory and the Internet," IEEE Communications Magazine, January 2001.
 - V. Frost, B. Melamed, "Traffic modeling for telecommunications networks", IEEE Communications Magazine, March 1994.
 - I. Kaj, „Stochastic modeling“, 5.2.2-5.3.1.

Teletraffic theory




- Teletraffic theory:
 - to model dynamic resource sharing systems
 - to explain the traffic-performance relation



- Traffic: arrival intensity, holding time, packet length (distribution or moments)
- Resources: link bandwidth, router buffer, server capacity
- Performance: utilization, loss, delay, delay variation, perceptual quality

Traffic modeling

- To describe the network traffic demand
- Statistical characterization
- Traffic objects



- **Flow** (one instance of communication, TCP or UDP session) — Skype call
- **Burst** (Active/passive periods) — Talk/listen
- Sequence of **packets** — IP packets
- Multiplexed packets — IP packets at a router

Traffic modeling


- Packet level:
 - Characteristics of the sequence of packets
 - packet arrival process
 - according to some stochastic/deterministic arrival process (e.g. Poisson arrival at a router...)
 - saturated source model: there is always packet to send at the source
 - packet size distribution
- Flow level (burst level is similar too, but rarely used):
 - flow arrival process
 - e.g., flows from all the laptops in a WLAN are generated according to a Poisson process
 - flow duration distribution
 - flow characteristics – how traffic is generated within a flow

Flow characteristics

- Models that describe the distribution of the sequence of packets for a flow level model


- packet scale model**

- arrival process and packet size distribution
- queuing theory
- used typically in this course
- may lead to very complex models on flow level



- fluid models**

- transmission as a continuous stream
- parameter: flow rate $r(t)$
- system of differential equations
- often more tractable on the flow level



Flow types - Terminology

- **Flow** - one instance of an application
 - Reasonable to classify according to application types
- **Elastic flow**
 - The application requires the transmission of a given amount of information, some delay is acceptable – that is, transmission is elastic in time
 - e.g., file transfer over TCP
 - packet transmissions determined by the transport protocol (e.g., TCP) and the background traffic
- **Streaming flow**
 - The application has strict delay limitations, late information is dropped
 - e.g., VoIP over UDP
 - packet transmissions determined by source characteristics (e.g., coding)

Traffic modeling


- Which traffic object should we use for modeling?
- Select the object to which the considered traffic control is applied
 - buffer dimensioning – sequence of packets
 - error control – loss of individual packets
 - PACKET LEVEL MODELS
 - video rate control
 - routing
 - FLOW LEVEL MODELS

Group work

- Should we use packet or flow level models in the following problems? In the case of flow level models, what kind of flow characterization is necessary?
1. What is the probability that a packet collides and therefore needs to be retransmitted when using CSMA/CA protocol?
 2. Several Skype calls are using the same communication link. What is the utilization of the link ($utilization = \{average\ rate\ of\ traffic\} / \{link\ transmission\ rate\}$)
 3. Several flows are multiplexed at a router with limited buffer. What is the probability that consecutive packets of a flow are dropped due to buffer overflow?

Recall – Markov chains

- Basic tools of queuing theory
- Stochastic process
 - Discrete state space
 - Discrete or continuous time (change of state)
 - Markovian property: the future of the process does not depend on the past, only on the present




- Discrete time Markov chains
 - State transition probability matrix $\mathbf{P} = \{p_{ij}\}$
 - $\underline{p}_{n+1} = \underline{p}_n \mathbf{P}$
 - If steady state exists, the stationary state probability is given by $\underline{p} = \underline{p} \mathbf{P}$
 - Holding time of a state is geometric with parameter $1 - p_{ii}$
 - $\Pr(\tau_i = k) = p_{ii}^{k-1} (1 - p_{ii})$, $\bar{\tau}_i = 1/(1 - p_{ii})$
 - Geometric distribution is memoryless

Recall – Discrete time Markov chains

- E.g., to model the packet loss process at a receiver
 - States: packet received or lost (0,1)
 - Captures the burstiness of the loss process (see Gilbert model later in the course)
 - If a packet is lost (state 1), the next one is lost with probability $p_{11}=0.8$
 - If a packet is received (state 0), the next one is received with probability $p_{00}=0.9$
 - Packets lost in a row $\sim Geo(1 - p_{11})$, in average $1 / (1 - p_{11}) = 5$
 - Steady state probability of receiving or losing a packet:

$$\{p_0, p_1\} = \{p_0, p_1\} \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix}$$

$$\{p_0, p_1\} = \{p_0, p_1\} \begin{bmatrix} 0.9 & 0.1 \\ 0.2 & 0.8 \end{bmatrix}$$



Q: What is the probability that an arbitrary packet gets lost?

Recall – Continuous time Markov chains


- Continuous time Markov chains

- State transition is possible at any time
- State transition intensity matrix $\mathbf{Q} = \{q_{ij}\}$, $q_{ii} = -\sum q_{ij}$
- $\underline{p}'(t) = \underline{p}(t)\mathbf{Q}$
- If steady state exists, the stationary state probability is given by $\underline{0} = \underline{p}\mathbf{Q}$
- Holding time of a state is Exponential with parameter $-q_{ii}$
 - $P(\tau_i < t) = 1 - e^{q_{ii}t}$, $\bar{\tau}_i = \frac{1}{-q_{ii}}$
 - The exponential distribution is memoryless
- E.g., good (0) or bad (1) state of a wireless channel
- Steady state probabilities:

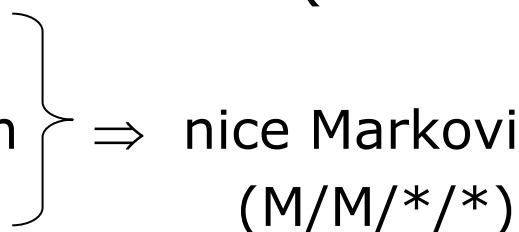
$$\{0,0\} = \{p_0, p_1\} \begin{bmatrix} -3 & 3 \\ 2 & -2 \end{bmatrix}$$

Q: What is the probability that the system is in state 1?

Q: What is the average time spent in state 1?



Traffic modeling

- Mathematical tools to describe traffic on packet and flow level
- Traditional telephone networks (from Erlang):
 - Poisson call arrival
 - exponential call duration
 - constant rate

\Rightarrow nice Markovian models
(M/M/*/*)
- Similar models are possible for data networks:
 - Poisson flow/packet arrival process
 - Exponential flow size (e.g., file length), packet size

Markovian traffic models

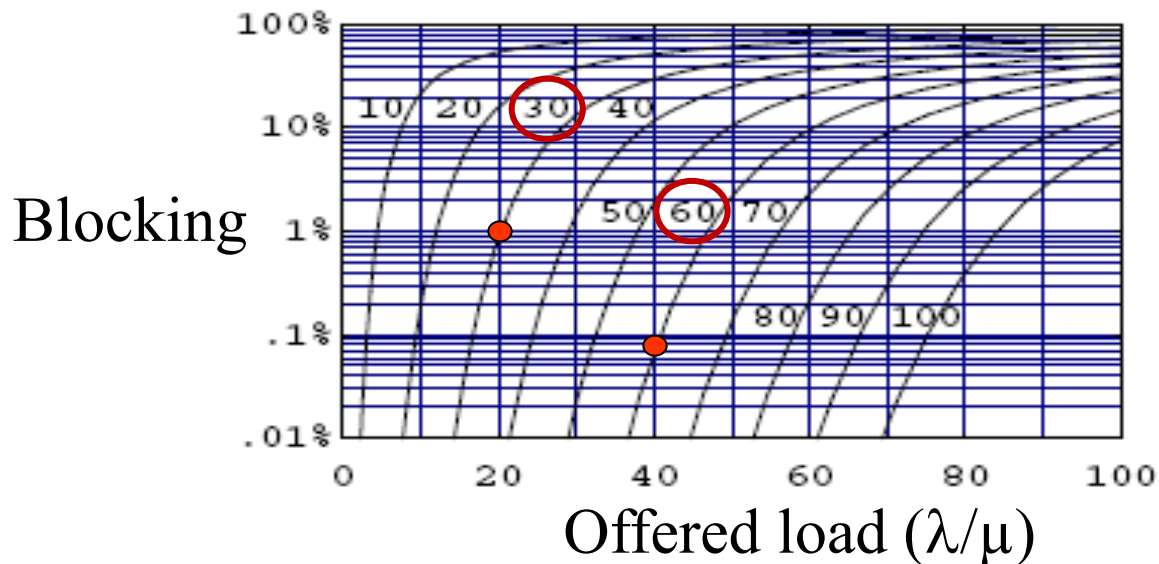
- Poisson process: $P\{N(t)=n\}=e^{-\lambda t}(\lambda t)^n/n!$
- Exponential distribution: $P(t \leq T)=1-e^{-\lambda T}$, $f(t)=\lambda e^{-\lambda t}$
- Recall – some basic results
- Exponentially distributed interarrival and service times
 - Poisson arrival: exponential interarrival time
 - Exponential distribution is memoryless – simple modeling
 - Tail function $P(t > T)=e^{-\lambda T}$ – **exponential decay in t**
 - e.g., the probability that a packet size is larger than T decreases exponentially in T .
 - Consecutive values (interarrival time, service time) are independent, therefore **auto-covariance is zero**
$$\text{Cov}(k)=E[(X_i-E[X])(X_{i+k}-E[X])]=E[X_i X_{i+k}]-E[X]^2=0$$

Markovian traffic models

- Exponential interarrival and service times in queues (M/M/*/*)
- **Buffering** is efficient, does not cause large delays:
- E.g, distribution of the number of users in an M/M/1 queue:
 - $p(n) = (1 - \rho)\rho^n$, $\rho = \lambda x$
 - $P(n \geq N) = \rho^N$ – the probability that the queue length is at least N decays exponentially fast.

Markovian traffic models

- **Multiplexing** is efficient, decreases the blocking probability
- E.g, M/M/m/m
 - Multiplexing: higher aggregate arrival intensity \rightarrow higher offered load
- Blocking given by the Erlang-B curves



$B(\text{load}, \text{servers})$

$B(20, 30) \approx 1\%$

$B(40, 60) \approx 0.1\%$

Markov modulated models

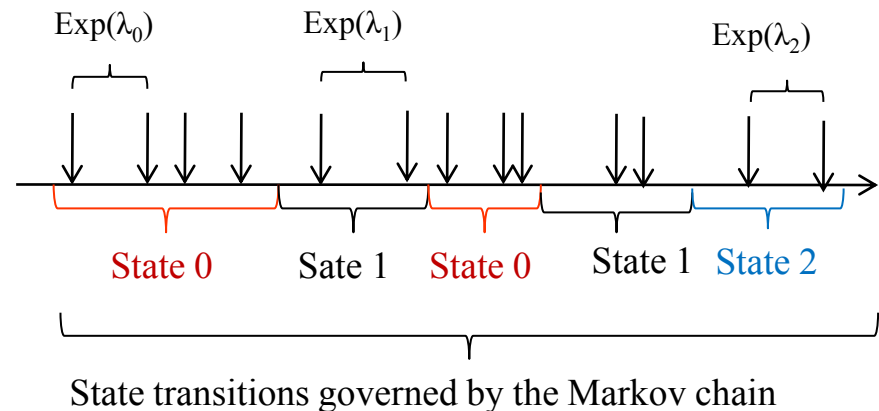

- However, we know that packet arrivals are not Poissonian
 - the arrival rate changes with time (traffic control, coding)
- First step towards modeling traffic sources:
- **Markov-modulated traffic models**
 - to capture “burstiness” (changing arrival rate)
 - while keeping the simplicity of modeling

Markov modulated models

- Packet level models


- Markov-modulated Poisson Process (MMPP)

- A Markov chain is given that describes the state of the source
 - The packet generation process is Poisson in each state, but with different intensity (state $i \rightarrow \lambda_i$)
 - Burstiness is captured by the state transitions in the Markov chain




Markov modulated models

- Packet level models with two states




– Interrupted Poisson Process (IPP)

- Most simple MMPP
- two states $\lambda_0=0$, $\lambda_1=\lambda$



– ON-OFF model


- two states, no arrivals in state 0 and fixed (d) packet interarrival times in state 1 (deterministic arrival process)



Markov modulated models

- Fluid models

- When individual units (e.g., packets) have little impact
- Traffic as a continuous stream with a parameterized flow rate (state $i \rightarrow r_i$)
- Flow rate changes described by a Markov-chain



Markov modulated fluid


MMPP traffic models

Example

- A traffic source is modeled as follows:
 - 2 state MMPP
 - Transition intensity $q_{12}=0.5$, $q_{21}=0.5$ (transitions per sec)
 - Transmission rates: 100 packets/s and 400 packets/s
 - Packet size: 500Bytes
1. What is the mean time in states 1 and 2 respectively?
 2. What is the probability that the source is in state 1 (state 2) at an arbitrary point of time?
 3. What is the average packet interarrival time in state 1?
 4. What are the transmission rates in the two states in bit per sec?
 5. What is the average transmission rate?
 6. If 5 such sources are multiplexed, what is the probability that the instantaneous rate is 8Mbps or larger?

Markovian traffic models

- Auto-covariance and auto correlation function decays exponentially
 - Auto-covariance: $\text{Cov}(t) = E[(X_i - E[X])(X_{i+t} - E[X])]$
 - Auto-correlation: $r(t) = \text{Cov}(t) / V[X]$
- Simplest example: on-off fluid model
 - The auto-correlation of the state of the system (on or off)
 - $r(t) = e^{-(\alpha + \beta)t}$
- What does it mean: the system has some memory about the past, but only for a short time – (we introduce the concept of short range dependence later)



Markovian traffic models

Modeling voice traffic

- Compare the average delay at a multiplexer, if
 - Real voice source packets are multiplexed in a simulator
 - Poisson arrival is assumed with the same average rate
 - 2 state MMPP model is used
 - Some advanced technique is used
- Results:
 - Poisson arrival approximation underestimates delays (queue lengths)
 - MMPP seems to fit well at high load regime as well




Fig. 2. Expected delay for a packetized voice multiplexer.

Traffic modeling

- Are Markovian traffic models enough to model network traffic sources?
- Or do we need other models?

Modeling Internet traffic

- Read: J. Roberts, "Traffic theory and the Internet,"
"As a first approximation, it is not unreasonable to assume that **individual flows also occur as a Poisson process**. To ignore the correlation between flow arrivals within the same session is not necessarily significant when the number of sessions is large. It is also true that results derived under the simple Poisson assumption are also often true under more general assumptions.
The **size of elastic flows (i.e., the size of the documents transferred) is extremely variable and has a so-called heavy-tailed distribution**: most documents are small (a few kilobytes) but the number which are very long tend to contribute the majority of traffic. The precise nature of the size distribution is important in certain circumstances, such as describing the **resulting self-similar packet arrival process**, and can have a significant impact on performance in some multiplexing schemes.
The **duration of streaming flows also generally has a heavy-tailed distribution**. Furthermore, the **packet arrival process within a variable rate streaming flow is often self-similar**."

Modeling Internet traffic


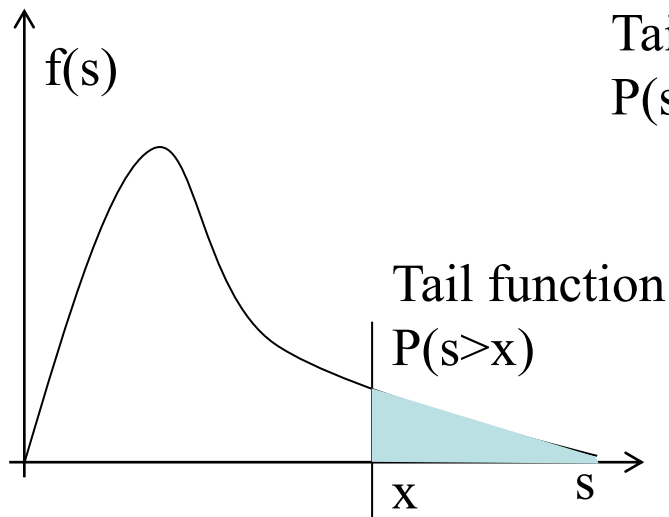
- Elastic flows - controlled by TCP
 - e.g., file transfer
 - arrival of flows: independent activity of a large number of users → *Poisson*
 - size: *heavy tail*
 - traffic characteristics: extreme variability introduced by TCP and heavy tailed flows
 - *self-similar* packet arrival process
- Streaming flows - determined by the source coding
 - arrival of flows: *Poisson*
 - duration: extreme variability, *heavy tail*
 - traffic characteristics (rate): often *self-similar* due to coding
- Conclusion:
 - Simple Markovian or Markov Modulated source models may not work

Non-markovian models

- We have to address the followings
 - what is heavy-tail distribution
 - what is self-similarity (and related: what is long range dependence)
 - when is it possible to apply Markovian models





Heavy-tail distributions, self-similarity, and long-range dependence

- What are the limitations of Markovian models?
- Example 1
 - Telephone call holding time measurements
 - Exponential assumption: $P(s > x) = e^{-\mu x}$
 - Statistics (for large s): $P(s > x) \sim x^{-\alpha}$, $\alpha > 0$
 - Decay is slower than exponential: **heavy-tail distribution**



Heavy-tail distributions, self-similarity, and long-range dependence

- Example 2
 - Packet arrivals in 40 hours Ethernet traffic (Bellcore '89)
 - Number of packet arrivals in increasing time intervals:




Ethernet measurement

Poisson

- Self-similar nature of packet arrival process

Heavy-tail distributions, self-similarity, and long-range dependence

- Example 3 – the effect of (long range dependent) self similarity
- LRD-SS source characteristics changes the network performance significantly.
 - E.g., mean queue-length at routers/multiplexers
 - Blocking and loss probabilities
- Therefore
 - we have to take it into account at the performance evaluation
 - have to understand how it emerges and whether it is possible to avoid it




Short overview

- Heavy-tail distributions
- Long-range dependence
- Self-similarity



Heavy-tail distributions

- Exponential distribution: $P(s > x) = e^{-\mu x}$
- Heavy-tail distribution:
 - $P(s > x) \sim x^{-a}$, $x \rightarrow \infty$, $a > 0$
 - the asymptotic shape is hyperbolic
- Pareto distribution: often used heavy tail distribution (e.g., for file size length):
 - $f(x) = ab^a/(x^{a+1})$,
 $a > 0$ (shape),
 b is the minimum possible value (base)
 - $P(s > x) = 1 - F(x) = (b/x)^a$
 - $E[x] = ab/(a-1)$ for $a > 1$
otherwise the mean is not finite



Heavy-tail distributions

- Why heavy-tail is a problem?
- Can be proved:
 - Superposition of ON-OFF processes where the distribution of the ON periods is heavy tailed (e.g., Pareto) gives **long-range dependent self-similar** process



Long-range dependence

- Consider X_i stochastic process, $i=1,2,3$ (discrete time)
 - Discrete process, samples from a continuous time process or integral over the interval
- $E[X]$, $V[X]$ finite
- Auto-covariance: $\text{Cov}(k) = E[(X_i - E[X])(X_{i+k} - E[X])]$
- Auto-correlation: $r(k) = \text{Cov}(k) / V[X]$
- **Short-range dependent:**
 - $\sum_{k=1}^{\infty} r(k) < \infty$: the consecutive samples are correlated, but the correlation decreases fast with k
- **Long-range dependent:**
 - $\sum_{k=1}^{\infty} r(k) = \infty$: the consecutive samples are correlated, the correlation is preserved for long period.
- Note: long-range dependence is an asymptotic definition for large lags (k).


Long-range dependence

- Short-range dependent:
 - $\sum_{k=1}^{\infty} r(k) < \infty$: the consecutive samples are correlated, but the correlation disappears fast
- Long-range dependent:
 - $\sum_{k=1}^{\infty} r(k) = \infty$: the consecutive samples are correlated, the correlation can be preserved for long period.

- MMPP is short range dependent. E.g., on-off fluid:

$$r(k) = e^{-(\alpha+\beta)k} \quad (\text{geometric serie})$$

$$\sum_{k=1}^{\infty} r(k) = \sum_{k=1}^{\infty} e^{-(\alpha+\beta)k} < \infty \quad \text{since } e^{-(\alpha+\beta)} < 1$$



- What is the relationship between long range dependence and self-similarity?

Self-similarity

- Consider X_i stochastic process, $i=1,2,3$ (discrete time)
- $E[X]$, $V[X]$ finite



time unit: 0.01s $\rightarrow X_1, X_2, \dots$

X_1, X_2, \dots

$$X_1^{(2)} = \frac{X_1 + X_2}{2}, X_2^{(2)} = \frac{X_3 + X_4}{2}$$

$$X_1^{(m)} = \frac{X_1 + \dots + X_m}{m}, X_2^{(m)} = \frac{X_{m+1} + \dots + X_{2m}}{m}$$

$$r^{(m)}(k) = \frac{Cov^m(k)}{V[X^m]}$$



time unit: 10s $\rightarrow X_1^{(1000)}, X_2^{(1000)}, \dots$

- (Second order) Self-similar: auto-correlation $r^m(k)=r(k)$, for all m and k
- Asymptotically self-similar: if above true for large m and k

Self-similarity and long-range dependence (LRD)

- Second-order self-similar: $r^m(k) = r(k)$, for all m and k
- $r(k)$ has specific form (can be proved):

$$r(k) \sim H(2H-1) \frac{1}{k^{2(1-H)}}, \quad 0 < H < 1, H \neq 0.5$$

- H : Hurst parameter, the parameter of a self-similar process
- Self-similarity and LRD

$$r(k) = H(2H-1) \sum_{k=1}^{\infty} \frac{1}{k^{2(1-H)}} \quad \text{is a harmonic serie}$$

$$\sum_{k=1}^{\infty} r(k) = \infty \quad \text{if } 2(1-H) < 1 \rightarrow H > 0.5$$

→ A self-similar process is LRD if $0.5 < H < 1$. This is the interval when SS makes trouble.

- Often the terms self-similarity and long-range dependence are used for the same thing.

Markovian vs. SS/LRD models

- How should we choose traffic model for performance evaluation?

Examples

A cause of self-similarity in TCP traffic

[Fukuda, K.](#), [Takayasu, M.](#), [Takayasu, H.](#) **Source:**

International Journal of Communication Systems, v 18, n 6, Aug. 2005, p 603-17

Abstract: We analyze the statistical properties of aggregated traffic flows generated by TCP, in order to clarify a possible cause of self-similarity in Internet traffic. Using ns-2 simulation... The main contribution of our work is to show that TCP itself contains a mechanism for generating self-similarity...

Examples

Overflow and losses in a network queue with a self-similar input

[Tsybakov, B.](#) , [Georganas, N.D.](#) **Source:** *Queueing Systems, Theory and Applications*, 2000.

Abstract: ...considers a discrete time queuing system that models a communication network multiplexer which is fed by a self-similar packet traffic. The model has ... an input traffic which is an aggregation of independent source-active periods having Pareto-distributed lengths and arriving as Poisson batches. **The new asymptotic upper and lower bounds to the buffer-overflow and packet-loss probabilities are obtained.**

Examples

Internet traffic modeling: Markovian approach to self-similar traffic and prediction of loss probability for finite queues

[Kasahara, S.](#) **Source:** *IEICE Transactions on Communications*, 2001.

Abstract: It has been reported that IP packet traffic exhibits self-similar nature and causes the degradation of network performance... However recent studies have revealed that the loss probability of finite queuing systems can be well approximated by the Markovian input models. This paper studies the time-scale impact on the loss probability of MMPP/D/1/K system where the MMPP is generated so as to match the variance of the self-similar process over specified time-scales.

Markovian vs. SS/LRD models

- How should we choose traffic model for performance evaluation?
 - SS/LRD
 - complex models, possible to use for simulation but mathematical models are not that tractable
 - LRD captures asymptotic behavior but not short time characteristics
 - Markovian models
 - can capture correlations on arbitrary – finite – time scale
 - easier to use in mathematical models
- We have to choose models according to the dominant time scale we consider.

Summary (1/2)

- Network traffic modeling
 - Flows, bursts and packets
 - Elastic and streaming flows
 - Packet scale and fluid models for flow characterization
- Markovian traffic models
- Markov modulated traffic models
 - The rate is modulated by a Markov chain to capture burstiness
 - Can describe short term correlation

Summary (2/2)

- Heavy-tail, long-range dependence, and self-similarity
 - Asymptotic characteristics (infinite time...)
 - Heavy-tail: the tail function of the distribution has only hyperbolic decrease: $P(s > x) \sim x^{-\alpha}$, $x \rightarrow \infty$,
multiplexing heavy-tail flows leads to self-similarity
 - Long-range dependence: correlation is preserved over long timescales: $\sum_{k=1}^{\infty} r(k) = \infty$
 - Self-similarity: the correlation is preserved irrespective of time aggregation: $r^m(k) = r(k)$
 - Self similarity is characterized by H , the Hurst parameter, and the SS process is LRD if $0.5 < H < 1$
 - LRD-SS flows lead to inefficient multiplexing and long queues

Home reading

- Home reading for Thursday: W.E. Leland et al, "On the self-similar nature of Ethernet Traffic," ACM Sigcomm, 1993.
 - Sections 1,2,3, and 6. You are allowed to read the other sections as well!
- Home reading for Tuesday next week: A. Adas, "Traffic Models in Broadband Networks", IEEE Communications Magazine, July 1997
 - Markov and Embedded Markov models in detail
 - including the MMPP example for video coding
 - Regression models are not part of the course material, but are interesting reading
 - Long-range dependent traffic models, not including fractional ARIMA and fractional Brownian Motion