Issues related to the modeling of the TCP protocol

Chadi BARAKAT

INRIA Sophia Antipolis, France
PLANETE group

Email: Chadi.Barakat@sophia.inria.fr
WEB: http://www.inria.fr/planete/chadi

Outline

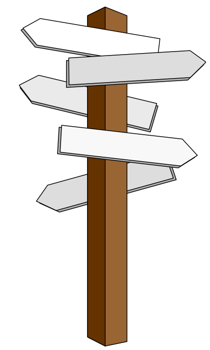

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

Chadi BARAKAT © 2005

TCP is widely used !

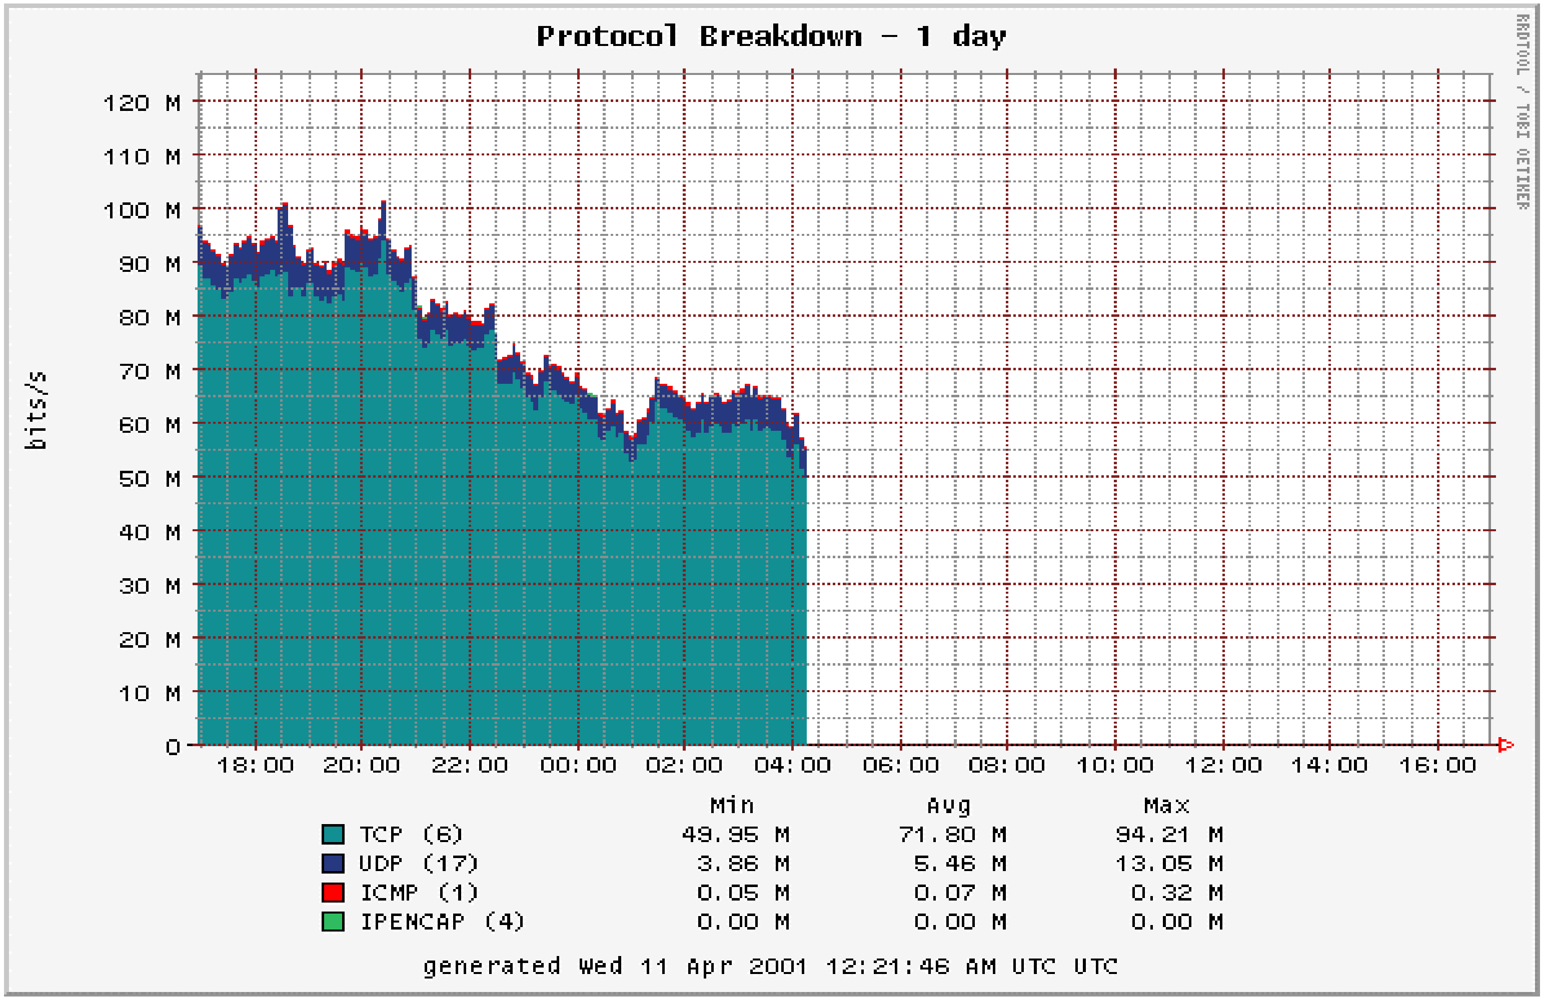

A trace collected on a OC-12 link of Sprint IP backbone in April 2001

TCP modeling: Objectives

- Objectives of a model for TCP: Express the performance of a TCP transfer as a function of some parameters that have physical meaning:
  - Parameters of a model for TCP: Loss rate of TCP packets, round-trip time of the TCP connection, the receiver advertised window, the slow start threshold, initial window size, window increase rate, etc.
- TCP performance measures: Throughput, latency, fairness index, etc.
- TCP models serve to compute (and hence to improve) network, application and user performance.

Refresh: TCP congestion control

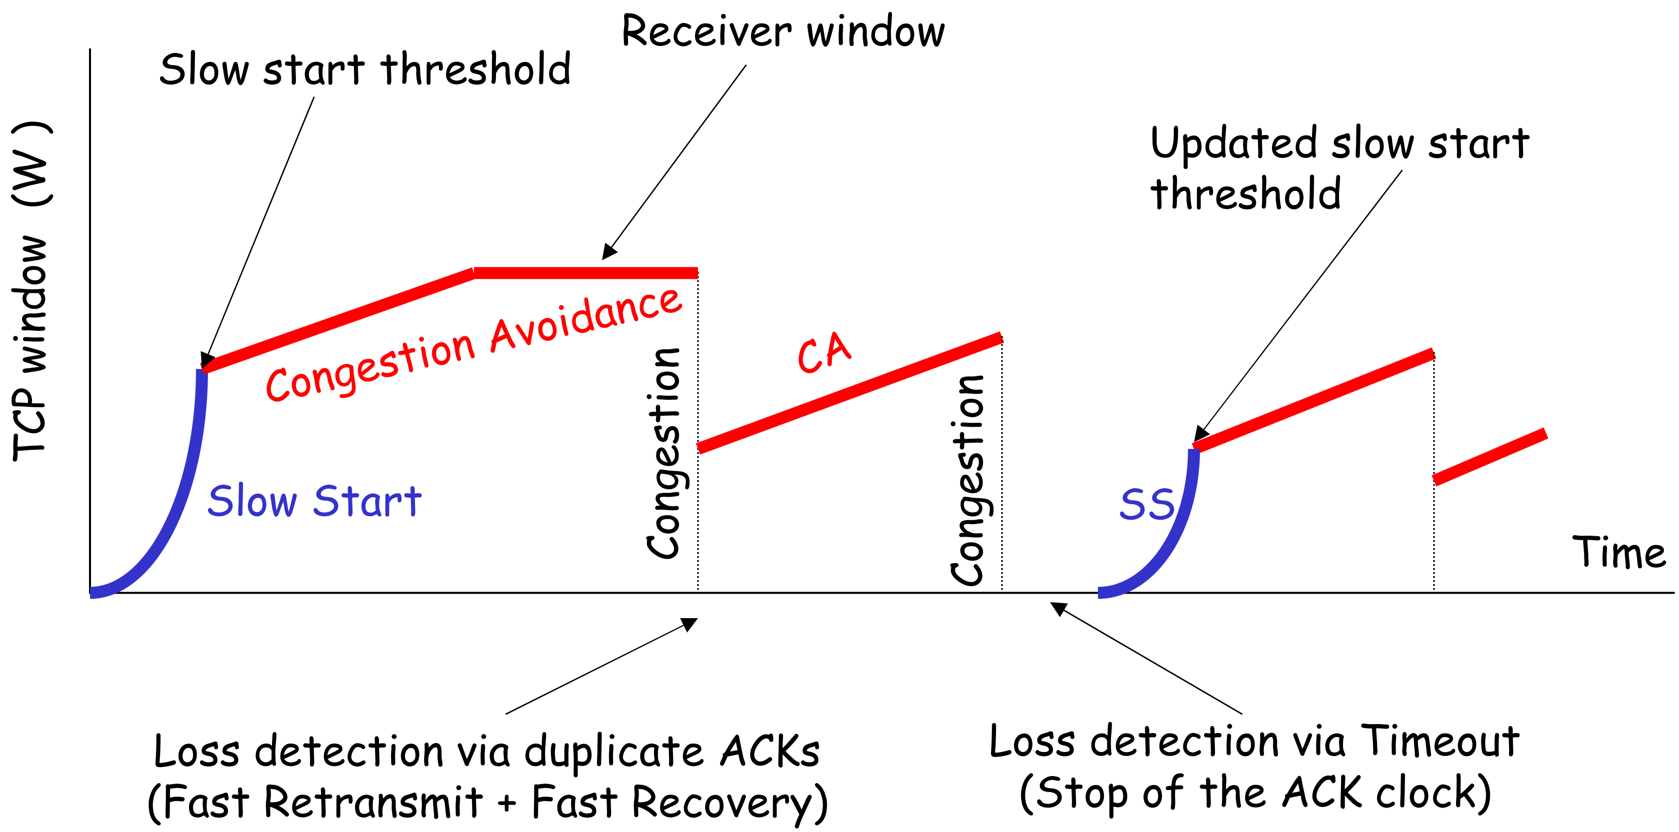

Basics for modeling TCP

- This requires a model for TCP dynamics
  - At the packet level, window level, transmission rate level, etc.
- And it requires a model for the network
  - How does the network drop TCP packets?
  - And by how much does it delay them?
- A simple network model often used: Bernoulli loss model.
  - TCP packets are lost in the network with constant probability p.
  - The round-trip time is constant equal to RTT.
  - As we will see, p and RTT can be computed using network topology, link characteristics, and concurrent traffic.
- Other models are also used in the literature:
  - Processor Sharing models, Deterministic models, bursty loss models, stochastic differential equations, stochastic difference equations, etc.
- Things will become clear later …

Outline

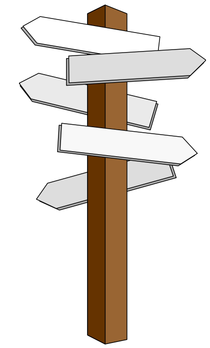

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

A simple model for TCP throughput

[FLO91,MSMO97]

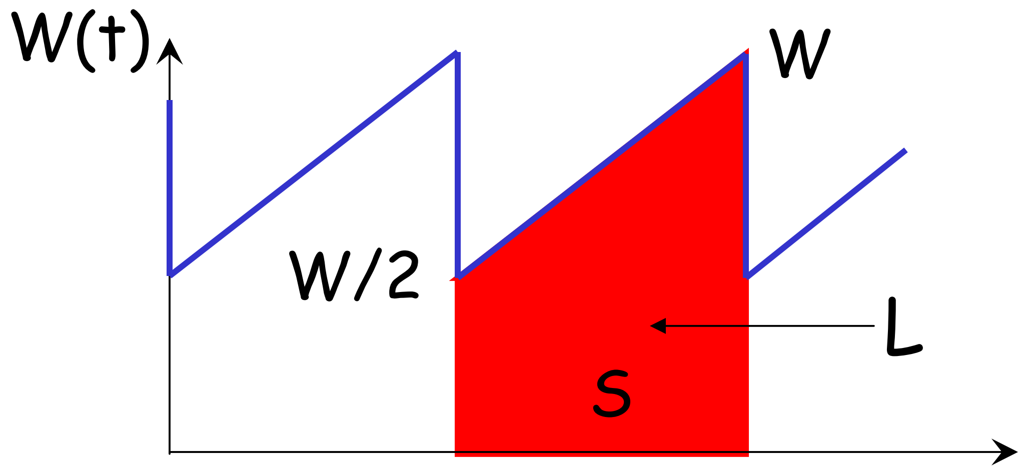

- Assumptions
  - Infinitely long TCP connection.
  - No Timeouts (no slow start).
  - Infinite receiver window.
  - Periodic evolution of the congestion window (time S constant).
  - Congestion window fluid and in packets.
- Throughput $X = \frac{\text{Number of packets transmitted per cycle } L}{\text{Cycle duration } S}$

A simple model for TCP throughput

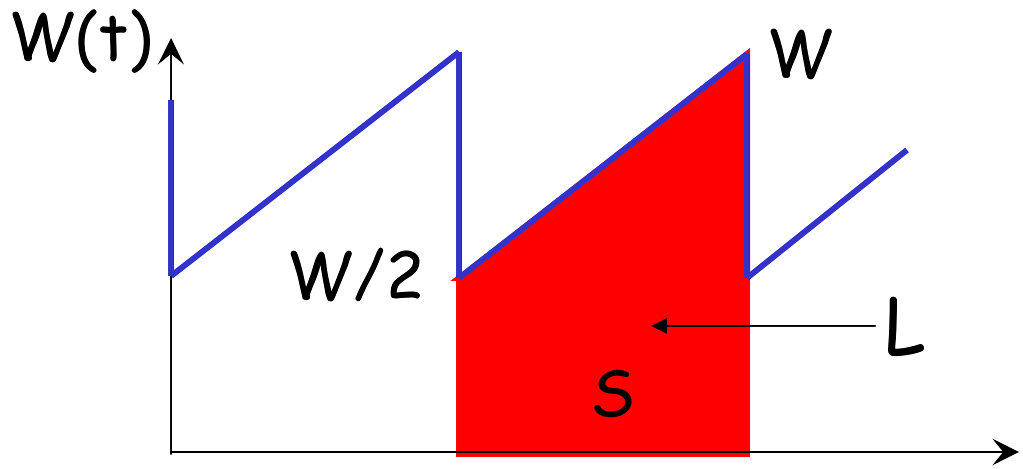

- Each cycle: Congestion window to be increased by (W - W/2).
- In congestion avoidance phase, window increases by 1 packet every round trip time.
- Window size of packets are transmitted per round-trip time.
- Thus, the total number of packets sent per cycle is:
  - $L = W/2 + (W/2 + 1) + (W/2 + 2) + \ldots + (W/2 + W/2) \approx (3/8)W^2$
- And cycle duration S = RTT.W/2
- But, we have $L = 1/p = (3/8)\ W^2$
- Hence,

$$X = \frac{L}{S} = \frac{1}{RTT}\sqrt{\frac{3}{2p}} \text{ packets/s}$$

Square Root Formula !

Comments on the SQRT formula

- The throughput of TCP is inversely proportional to the square root of p, the packet loss probability.
  - Bad performance of TCP over wireless links where packets are lost for other reasons than congestion, e.g. transmission errors.
- The throughput of TCP is inversely proportional to the round-trip time.
  - Bad performance of TCP over satellite links where the round-trip time is large, e.g. 500 ms over GEO satellite links.
  - TCP is unfair against connections with large round-trip times.
- The throughput of TCP is proportional to the packet size MSS.
  - A bulk-data TCP connection has interest to use large packets.

Problems with the SQRT formula

- Does not work when p is high:
  - The fluid assumption does not hold, i.e. the approximation $L \approx (3/8)W^2$.
  - Timeouts become frequent.
  - More than one packet can be lost in the same RTT, resulting in only one division of TCP window. Hence, p overestimates the congestion signal rate, and the SQRT formula underestimates the throughput.
- And it does not work when p is low:
  - At low p, the receiver window is often reached, which limits the throughput.
  - And at low p, the window is large, so it is very probable that the TCP protocol stops being linear increase with time ! (to be clarified later)
- To all that, one has to add the problem with the assumption on the constancy of times between congestion events.

A simple simulation with ns-2

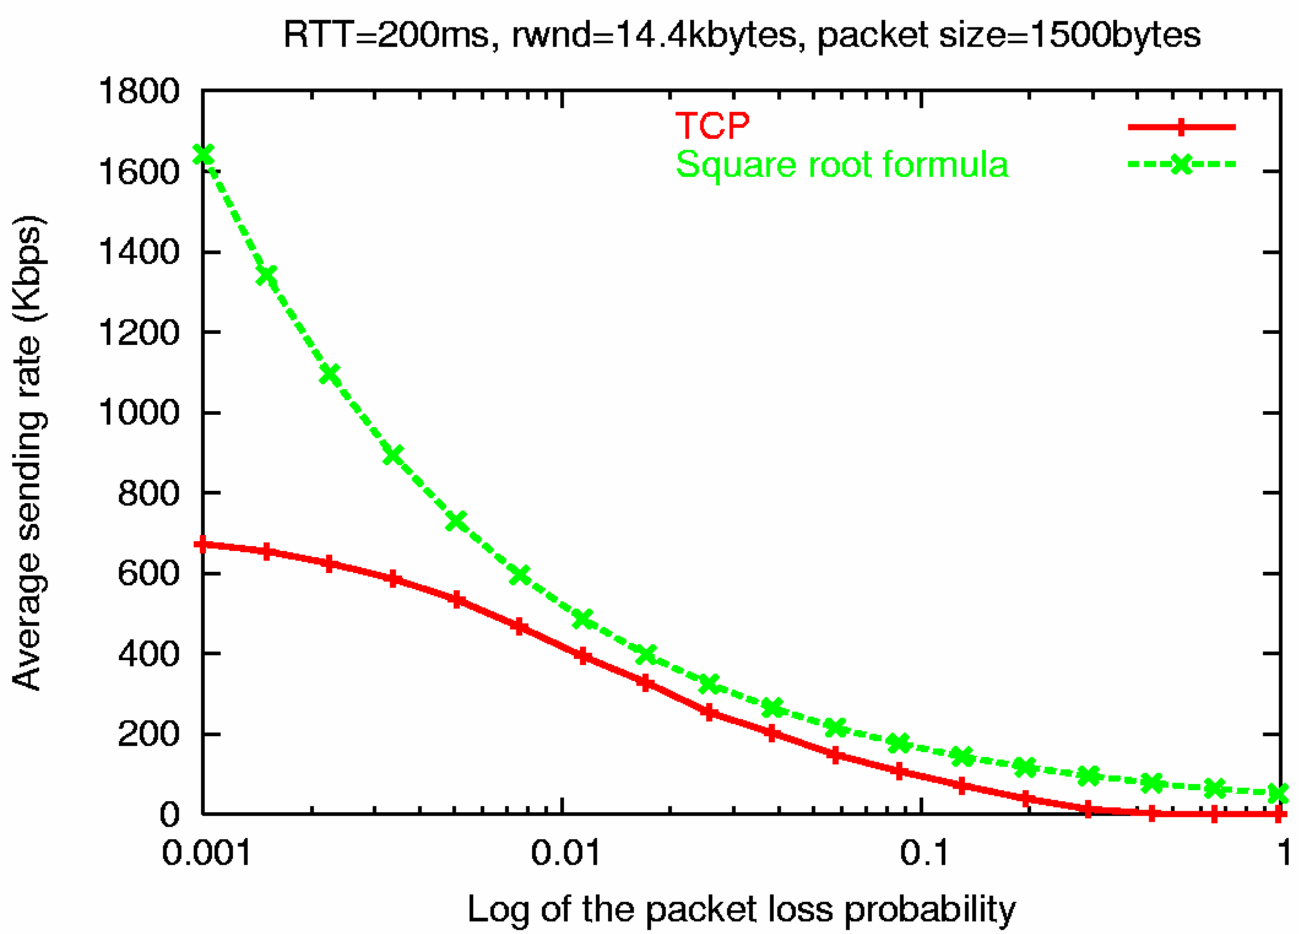

Look at the following document for a description on how to simulate such a scenario:

http://www.inria.fr/planete/chadi/NSCourse-2.pdf

Outline

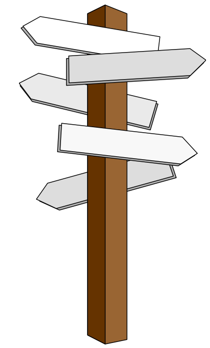

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

Heuristics to enhance the model

- Enhance for the packet nature of TCP (Nagle algorithm).
- Enhance for timeouts.
  - Silence periods incurred by Timeouts, without modeling slow start.
- Enhance for receiver advertised window.

Padhye J., Firoiu V., Towsley D., and Kurose J., "Modeling TCP Throughput: a Simple Model and its Empirical Validation", in Proceedings of ACM SIGCOMM, August 1998.

Chadi Barakat, " TCP modeling and validation", IEEE Network, vol. 15, no. 3, pp. 38-47, May 2001.

Packet nature of TCP: Throughput

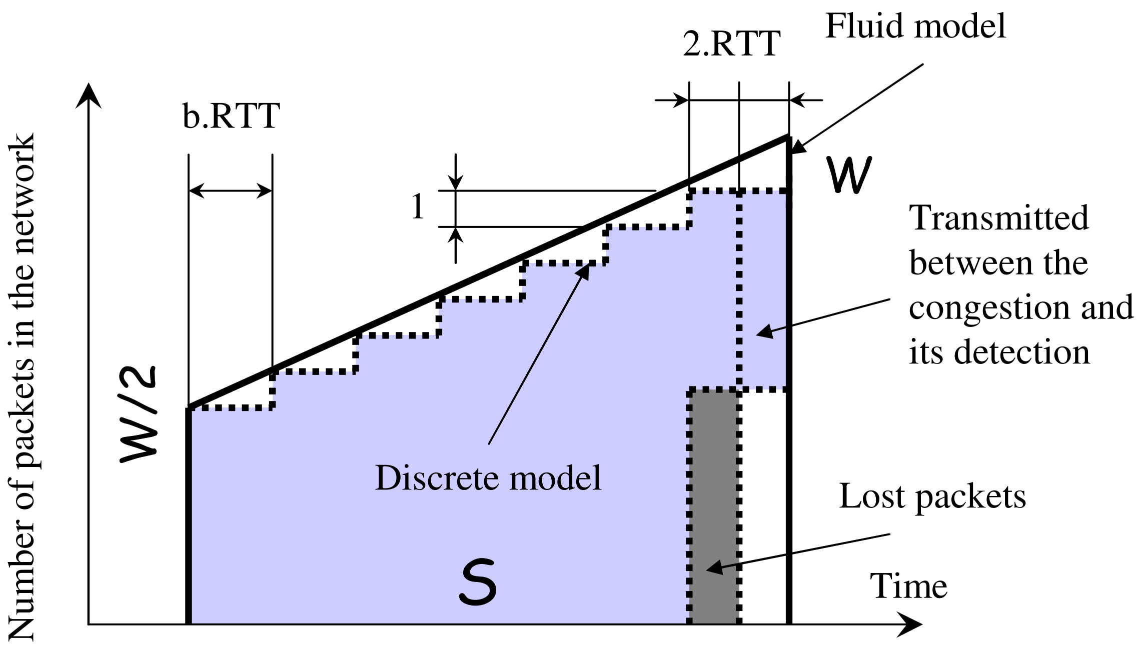

- b = 1: All packets are acknowledged.
- b = 2: One packet over two is acknowledged.
- Half of the window is assumed to be lost upon congestion [PFTK98].

$$X_{discrete} = X_{fluide} - \frac{0.5}{RTT} - \frac{W}{S} = \frac{L}{S} - \frac{0.5}{RTT} - \frac{2}{b.RTT} = \frac{(3/8).W^2}{(W/2).b.RTT} - \frac{0.5}{RTT} - \frac{2}{b.RTT}$$

$$\text{avec } L = b.W/2 + b.(W/2+1) + \ldots + b.W = 1/p + W/2$$

An approximation of W is [PFTK98]: $W = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + (\frac{2+b}{3b})^2}$

Receiver advertised window

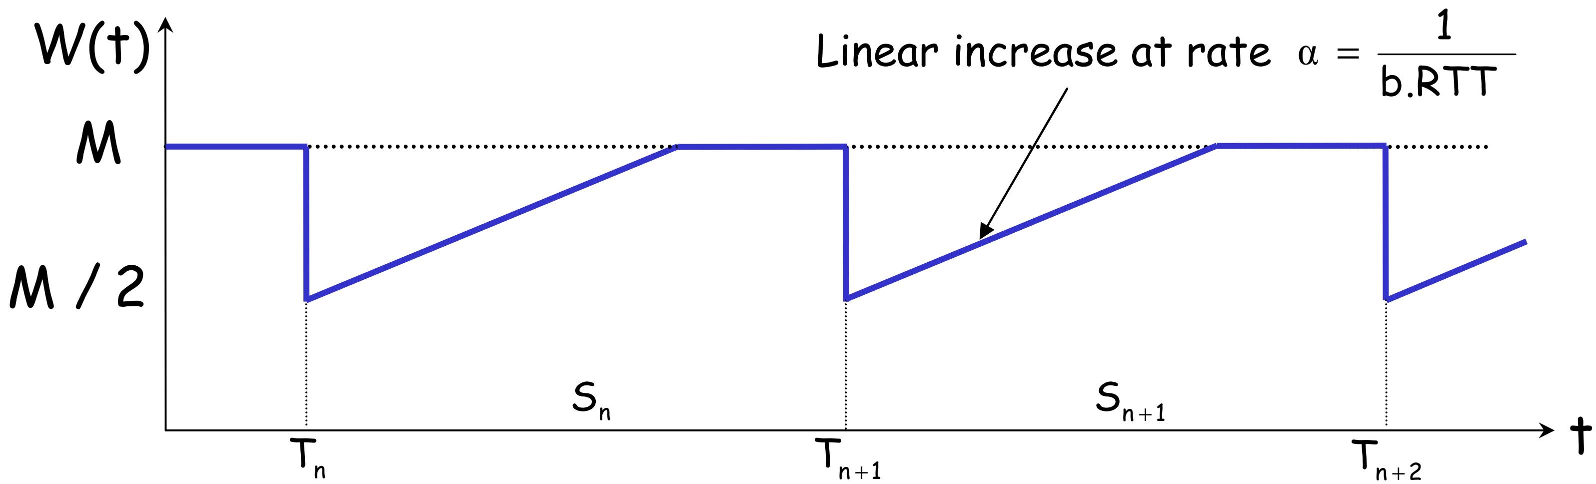

Approximate throughput when the window is reached: $X_M = \frac{1}{RTT}\left(M - \frac{M^2}{8.\alpha.S}\right)$ packets/s

Combined formula:

if $W < M$ then use the throughput obtained with infinite M, $X = (3/8).W^2 / S$

else use $X_M = \frac{1}{RTT}\left(M - \frac{M^2}{8.\alpha.S}\right)$

Account for the packet nature of TCP in the same way as before.

Modeling Timeouts

- TCP may stay idle for a long time before the detection of a loss.

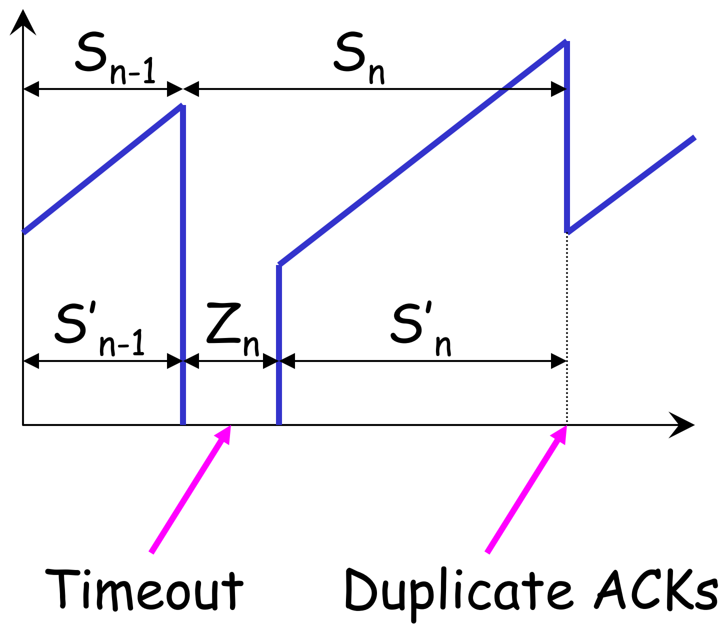

- X = Throughput excluding TO intervals.
- Throughput including TO intervals:

$$X_{TO} = \frac{X}{1 + p.X.Q.Z}$$

$Q = P(Z_n > 0)$

$Z = E[Z_n \mid Z_n > 0]$

One can use for these functions the expressions computed in:

Padhye J., Firoiu V., Towsley D., and Kurose J., "Modeling TCP Throughput: a Simple Model and its Empirical Validation", in Proceedings of ACM SIGCOMM, August 1998.

Modeling Timeouts

$$X_{TO} = \frac{E[L]}{E[S_n]} = \frac{1/p}{E[S'_n] + E[Z_n]} = \frac{1/p}{E[S'_n] + P(Z_n > 0).E[Z_n \mid Z_n > 0]}$$

$$= \frac{1/p}{E[S'_n] + Q.Z} = \frac{1/p}{E[S'_n].(1 + Q.Z / E[S'_n])}$$

Since $X = \frac{1/p}{E[S'_n]}$ **[BAR01]**

We get $X_{TO} = \frac{X}{1 + p.X.Q.Z}$

INRIA SOPHIA ANTIPOLIS

Q(p) and Z(p): Examples

$$Q = \min\left(1, \frac{(1-(1-p)^3)(1+(1-p)^3(1-(1-p)^{W-3}))}{1-(1-p)^W}\right),$$

$$Z = T_0 \frac{1+p+2p^2+4p^3+8p^4+16p^5+32p^6}{1-p},$$

with

$$W = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2},$$

and T_0 often estimated by 4RTT.

A complete formula: Sending rate

- Unconstrained window size: $$W = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2}$$

- If $W < M$

$$X = MSS.\frac{\frac{1-p}{p} + W + Q(W)\frac{1}{1-p}}{RTT(\frac{b}{2}.W + 1) + Q(W)T_0\frac{f(p)}{1-p}}$$

- Otherwise

$$X = MSS.\frac{\frac{1-p}{p} + M + Q(M)\frac{1}{1-p}}{RTT(\frac{b}{8}M + \frac{1-p}{pM} + 2) + Q(M)T_0\frac{f(p)}{1-p}}$$

A complete formula: Throughput

- Unconstrained window size: $$W = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2}$$

- If $W < M$

$$X = MSS.\frac{\frac{1-p}{p} + \frac{W}{2}}{RTT\left(\frac{b}{2}.W + 1\right) + Q(W)T_0\frac{f(p)}{1-p}}$$

- Otherwise

$$X = MSS.\frac{\frac{1-p}{p} + \frac{M}{2}}{RTT\left(\frac{b}{8}M + \frac{1-p}{pM} + 2\right) + Q(M)T_0\frac{f(p)}{1-p}}$$

A simple simulation with ns-2

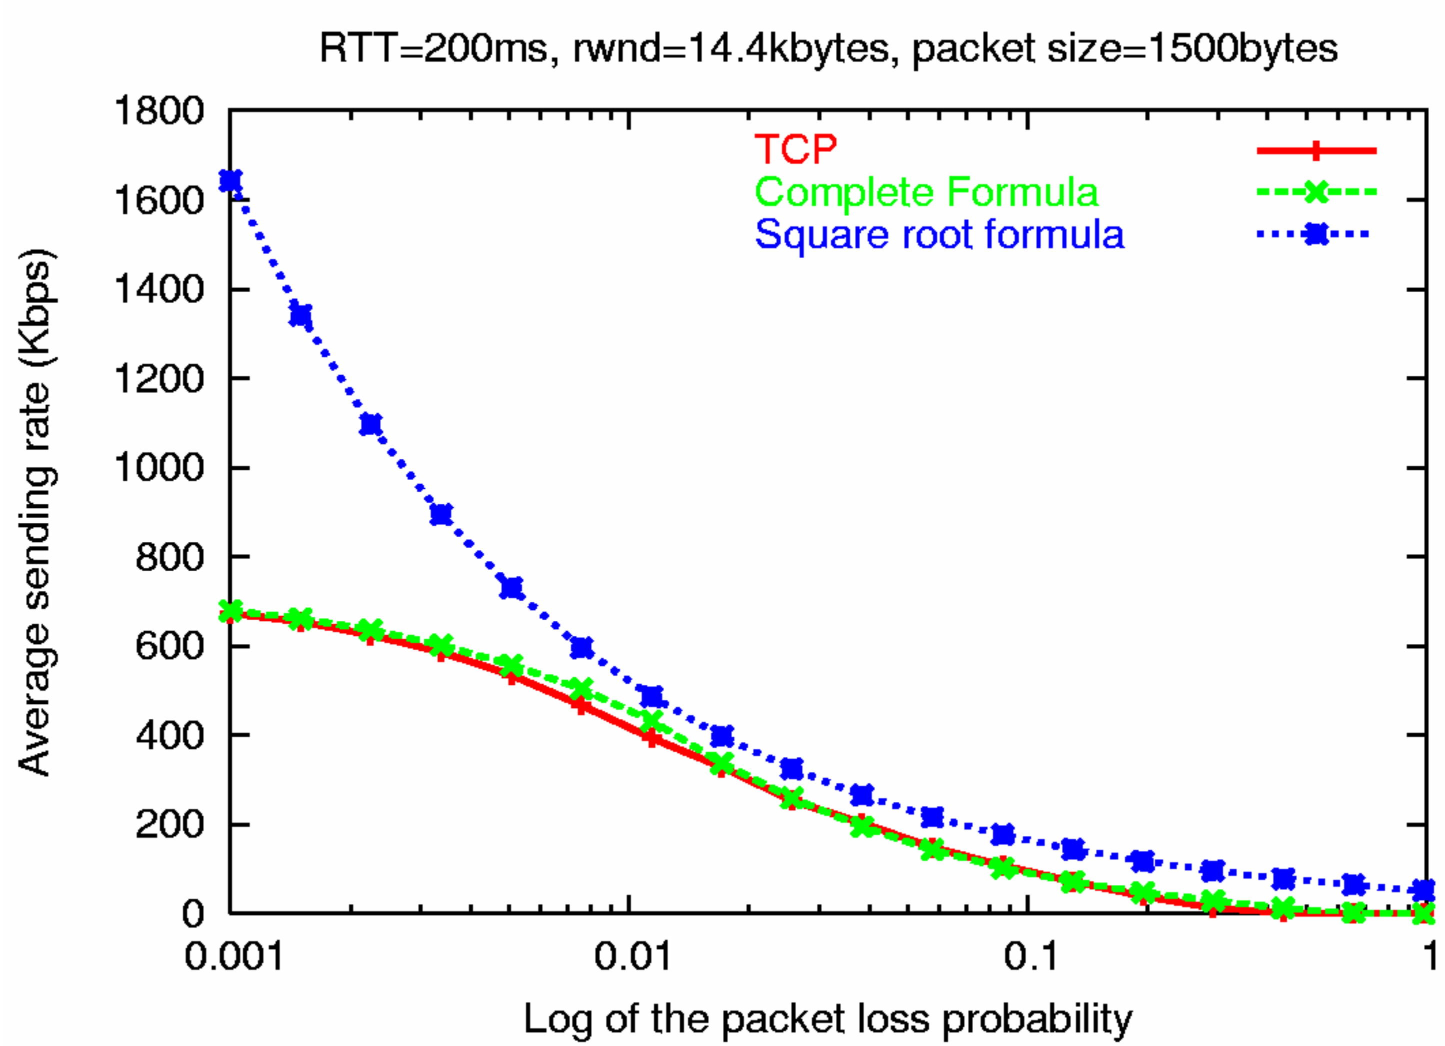

An approximate complete formula

TCP average sending rate:

$$X \approx \min\left(\frac{W_{max}}{RTT}, \frac{MSS}{RTT\sqrt{\frac{2bp}{3}} + T_0 \min\left(1, 3\sqrt{\frac{3bp}{8}}\right)p(1+32p^2)}\right)$$

Multiply X by (1-p) to get the throughput.

- W_{max} - maximum congestion window size (receiver window).
- T_0 - Length of timeout usually approximated by 4 RTT.

Padhye J., Firoiu V., Towsley D., and Kurose J., "Modeling TCP Throughput: a Simple Model and its Empirical Validation", in Proceedings of ACM SIGCOMM, August 1998.

Still some problems

- Finite size of TCP transfers:
  - TCP transfers are not infinite, most of them are of small size (average 10 KB).
  - Transitory phase (slow start) is important !
- Times between congestion events are not constant, they may take any distribution:
  - Exponential, normal, etc.
  - Independent or correlated.
- The round-trip time is not constant, it can be variable and correlated with the window size:
  - Correlation between RTT and the window size appears when the TCP connection has an important share of the total bandwidth. The queuing time in bottleneck routers increases with the window size (at large windows).

Outline

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

INRIA SOPHIA ANTIPOLIS

Outline

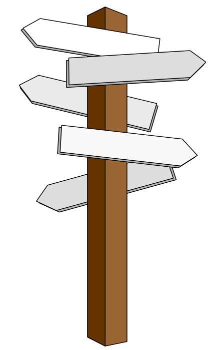

- Advanced modeling of TCP.
  - Complex distributions of inter-loss times.
  - The ACK clock approach
  - The Markov chain approach
  - Short-lived TCP transfers.

Complex distributions for times between loss events

- Previous models only consider periodic losses.
- But, the loss process in the Internet can have a complex distribution that changes from one path to another. Here are some examples:

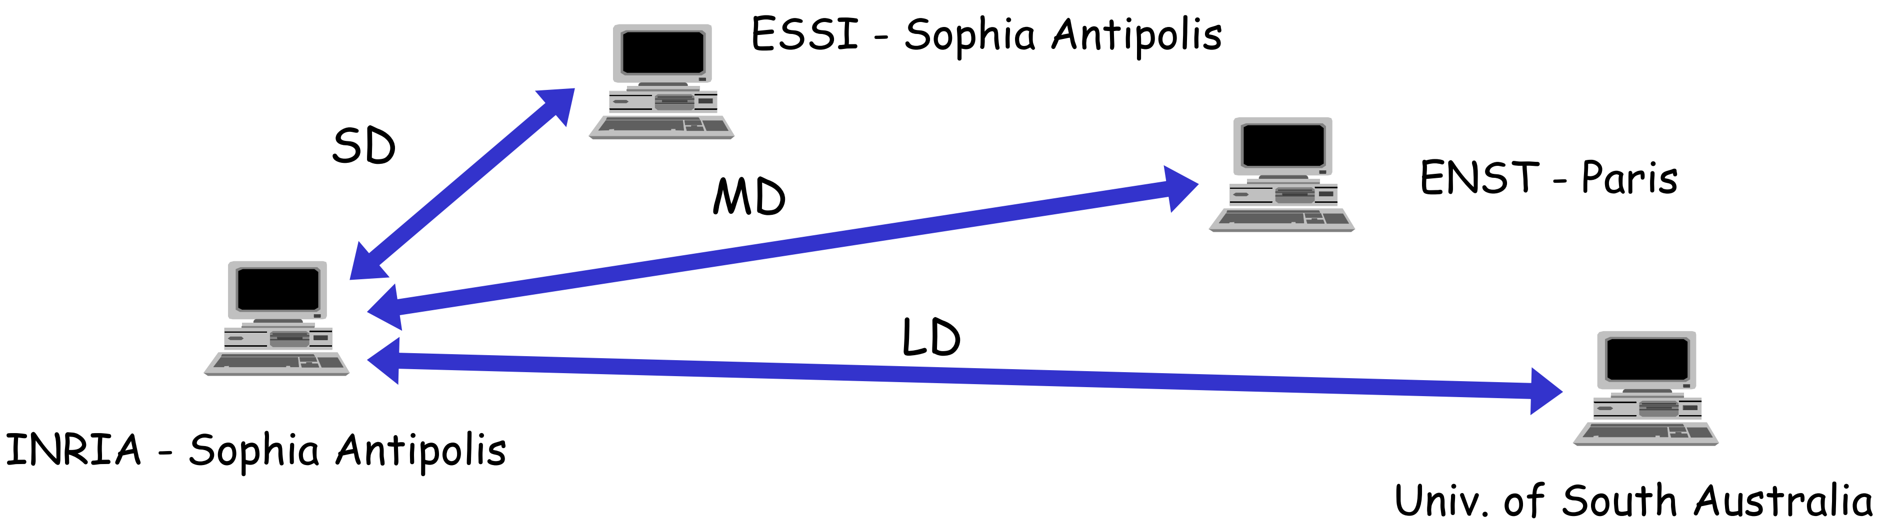

- Three long-lived unlimited-data TCP transfers (New Reno version).
- A tool is developed and run at INRIA to detect loss events.
- Traces are stored in separate files at approximately: 20, 40, 60min for SD, MD, LD resp.

Some inter-loss time distributions

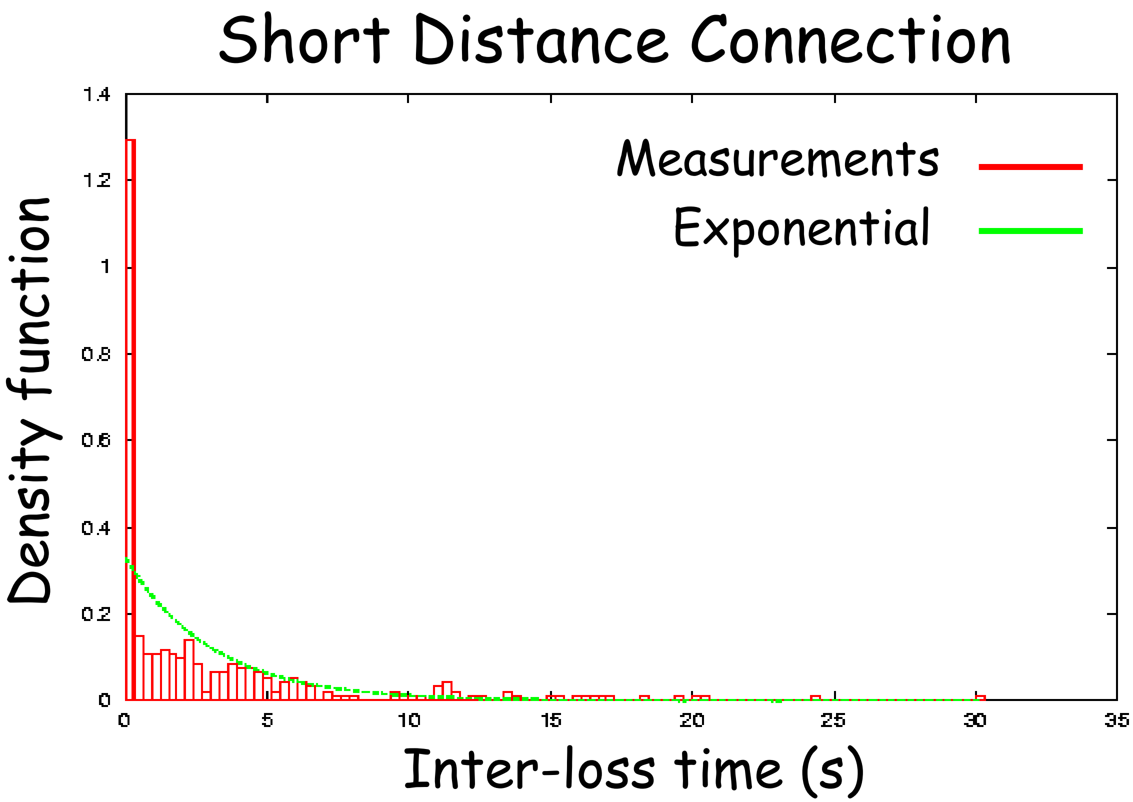

Highly bursty

INRIA SOPHIA ANTIPOLIS

Some inter-loss time distributions

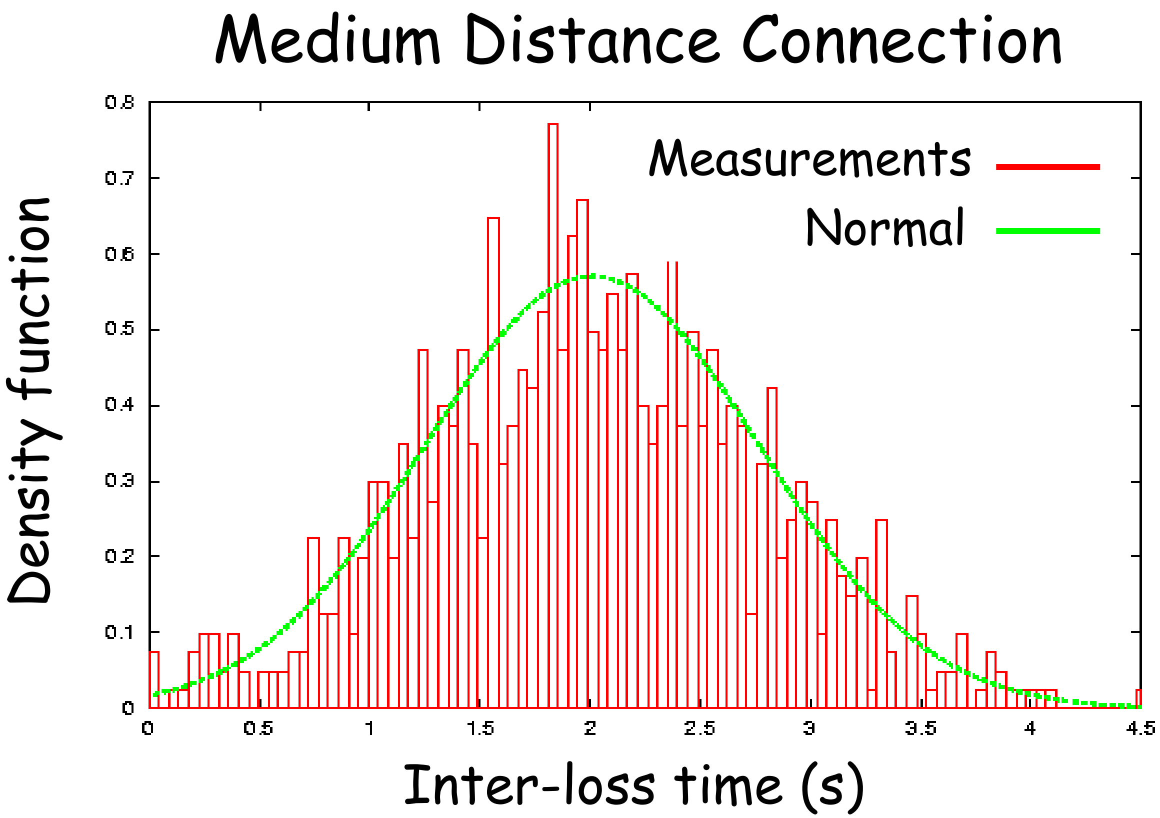

Close to Normal

Some inter-loss time distributions

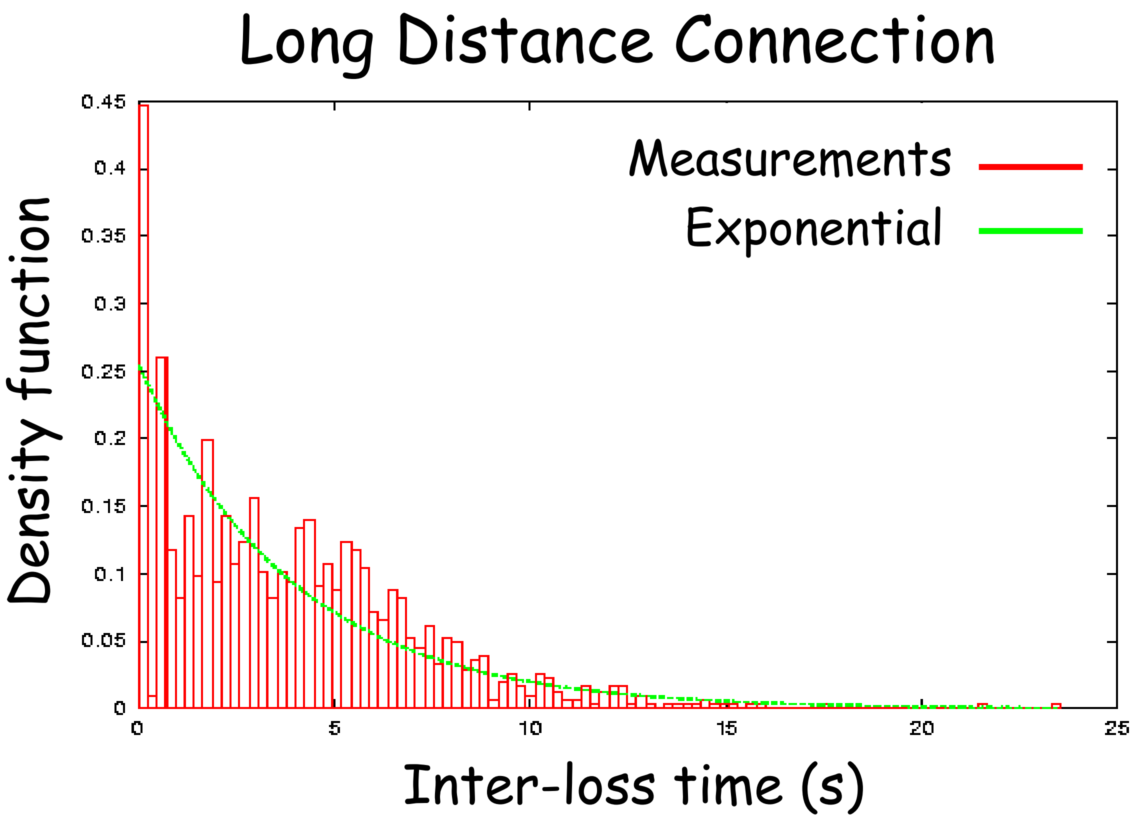

Close to Poisson

How to model the throughput ?

- Same assumptions as those for the square root formula, except that now the loss process is not periodic but GENERAL STATIONARY ERGODIC.

$$X = \frac{E[L]}{E[S_0]} = \lambda . E[L] = \lambda / p$$

λ: rate of loss events.

$$L = \int_{T0}^{T1} \frac{W(t)}{RTT} . dt = \frac{1}{RTT} \left(\frac{1}{2} W_0 S_0 + \frac{1}{2} \alpha S_0^2 \right)$$

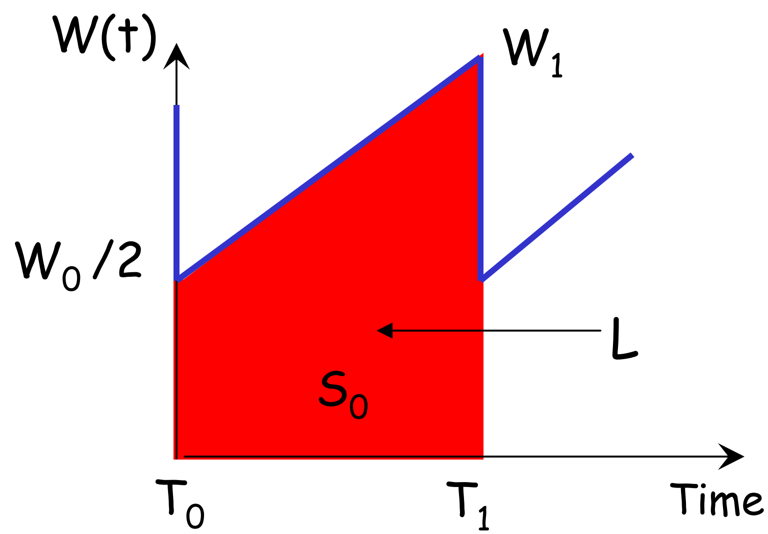

Stationary regime

L: number of packets transmitted in a cycle.

- The difficulty is in the computation of the expectation of $W_0 S_0$, since both random variables are dependent on each other. Both depend on the past values of inter-loss times.

INRIA SOPHIA ANTIPOLIS

Stochastic Difference Equation

- W_n = Window size in packets just before the n-th loss event.
- The dynamics of TCP window is governed by the following **Stochastic Difference Equation:**

$$W_{n+1} = \tfrac{1}{2} W_n + \alpha\, S_n$$

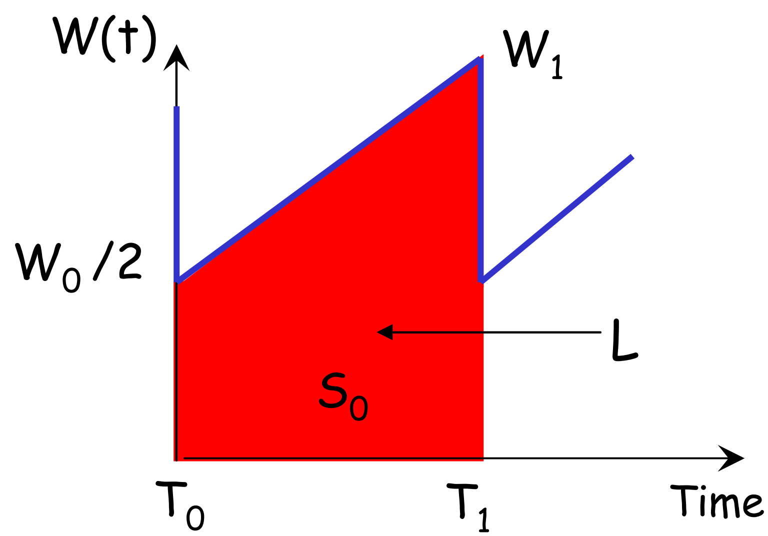

- **Theorem:** Assuming that the process of congestion events is stationary ergodic, the rate of TCP converges to the same stationary regime for any initial state.
  - For T_0 in the stationary regime:

$$W_0 = \alpha \sum_{j=0}^{\infty} \left(\tfrac{1}{2}\right)^j . S_{-j-1}$$

- We can then compute the expectation of $W_0 S_0$ and hence the throughput of TCP as a function of the second order moments of the process $\{S_n\}$.

General model

- Generalization of the square root formula (found under the condition that losses are periodic) for all stationary ergodic loss processes.

Previous result

$$X = \frac{1}{RTT}\sqrt{\frac{3}{2bp}}$$

General result

$$\hat{V} = \frac{E[S_0^{\,2}] - E^2[S_0]}{E^2[S_0]}$$

$$\hat{C}(k) = \frac{E[S_0 S_{-k}] - E^2[S_0]}{E^2[S_0]}$$

$$X = \frac{1}{RTT\sqrt{bp}}\sqrt{\underbrace{\frac{3}{2}}_{\text{Due to average rate of congestion events}} + \underbrace{\frac{1}{2}\hat{V}}_{\text{Due to variance}} + \underbrace{\sum_{k=1}^{\infty}\left(\frac{1}{2}\right)^k .\hat{C}(k)}_{\text{Due to correlation}}}$$

Due to average rate of congestion events

Due to variance

Due to correlation

Eitan Altman, Kostia Avrachenkov, Chadi Barakat, "A stochastic model of TCP/IP with stationary random losses", ACM SIGCOMM, August 2000.

TCP throughput and the loss process

- Fluid model with the different assumptions on the loss process:

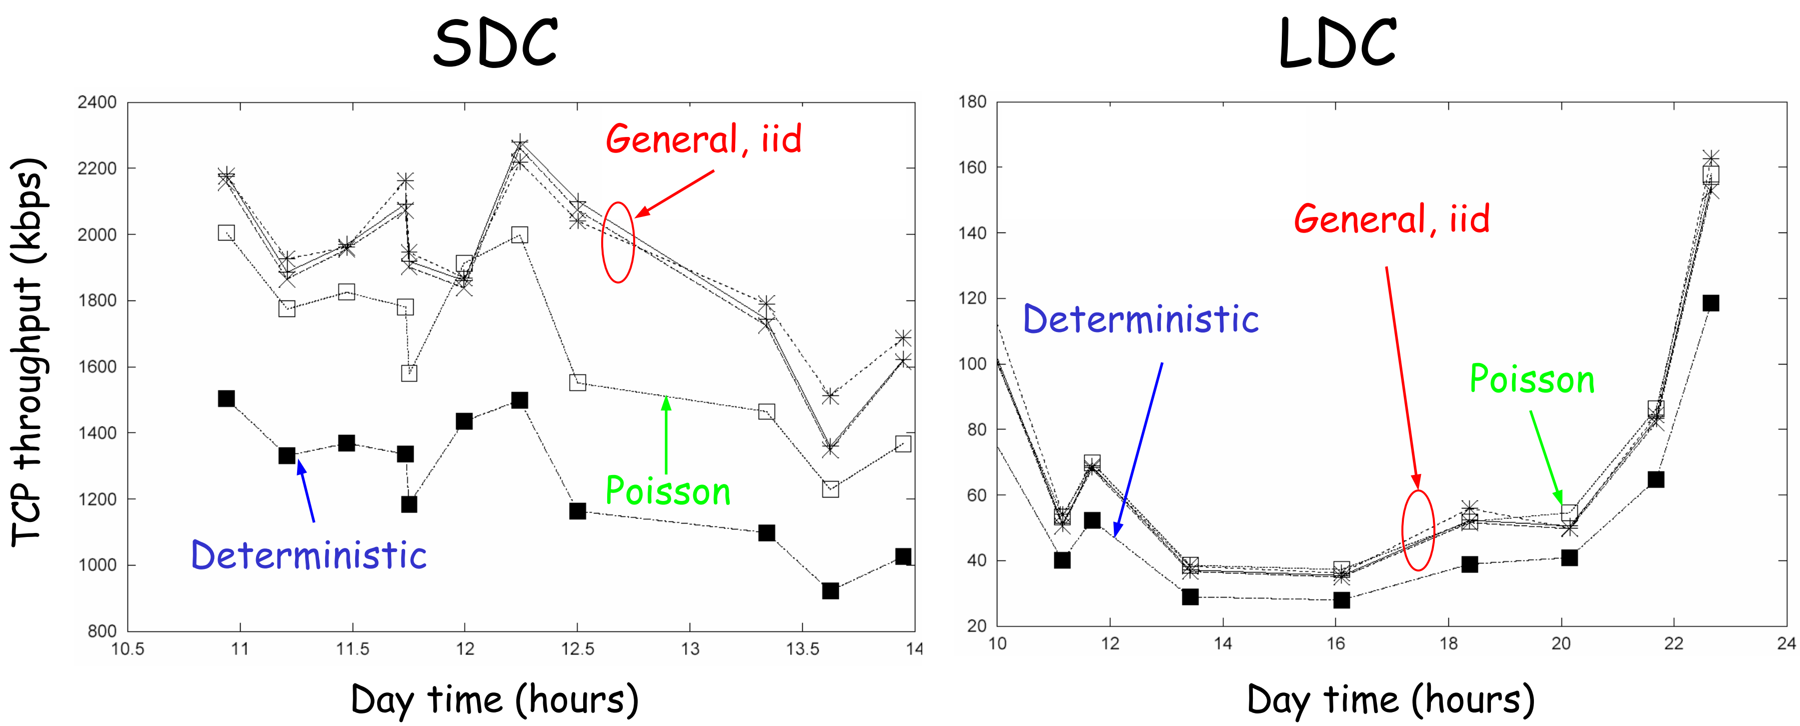

- Periodic losses lead to the smallest TCP throughput among the set of loss processes having the same intensity.
- The throughput of TCP increases with the variability of inter-loss times.

INRIA SOPHIA ANTIPOLIS

Outline

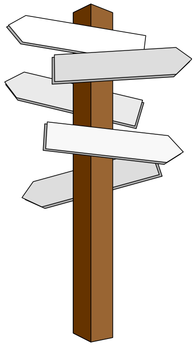

- Advanced modeling of TCP.
  - Complex distributions of inter-loss times.
  - The ACK clock approach
  - The Markov chain approach
  - Short-lived TCP transfers.

The ACK clock approach
From ACK clock to real time

- Instead of looking at the window as a function of time, look at the window when ACKs are received.
  - W_n window size when the n-th ACK is received.
- When an ACK is received, the window either grows or drops if the ACK carries a loss signal.
  - $W_{n+1} = W_n + 1/W_n$ if no loss (congestion avoidance).
  - $W_{n+1} = W_n/2$ if loss.
- Characterize the window in the ACK clock time
  - For example using a Markov chain or the Ott's approach (see next slide).
  - Calculate moments and distribution.
- Then switch back to the real time to characterize the transmission rate and calculate the throughput.

Example: Ott's work [OKM96]

- Take a Bernoulli packet loss process.
- Rescale the ACK clock by p and look at it as a homogenous Poisson process of intensity 1.
- Take a fluid model for the window size in the ACK clock time:

$$\frac{dW}{dACK} = \frac{1}{W} \text{ when no loss and } \frac{dW}{dACK} = -\frac{W}{2} \text{ when there is a loss}$$

$$\text{Let } Z = \frac{1}{2}.p.W^2 \Rightarrow$$

$$\frac{dZ}{dack} = 1 \quad \text{ when no loss and } \frac{dZ}{dack} = -\frac{3.Z}{4} \text{ when there is a loss}$$

- Z is a process that increases linearly and decreases multiplicatively at losses (not function of p!!). Can be studied using standard techniques.
  - In particular one can easily compute the moments and distribution of Z.

Example: Ott's work [OKM96]

- We know everything on Z, thus we know everything on W_n
- But what we need is W(t) and not W_n !
- Palm calculus !
- Take the throughput as example
  - t_n is the time at which the n-th ACK arrives
  - Suppose ACKs uniformly spread over the RTT, hence $t_{n+1} - t_n = RTT/W_n$
  - Assume the ACK process stationary erogodic of intensity λ_{ACK}

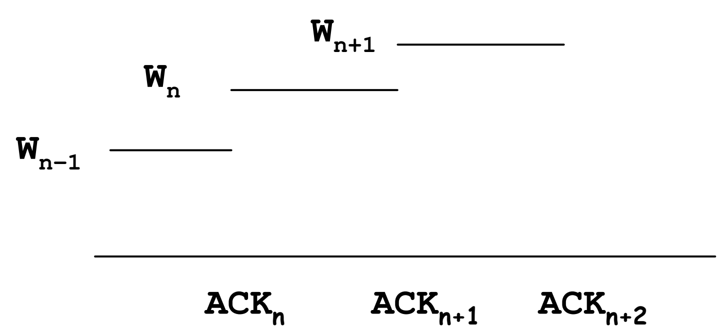

$$X = \frac{E(W(t))}{RTT} = \frac{1}{RTT}.\lambda_{ACK}.E[\int_{t_n}^{t_{n+1}} W_n.\frac{RTT}{W_n}.d\tau] = \lambda_{ACK} = \frac{1}{RTT.E\left[\frac{1}{W_n}\right]}$$

  - Substituting $E[1/W_n] = \sqrt{p}/\sqrt{2}.E[1/\sqrt{Z}]$, one can find

$$X = \frac{1.309}{RTT\sqrt{p}}.$$ Again !! the constant was "1.22" for SQRT

Outline

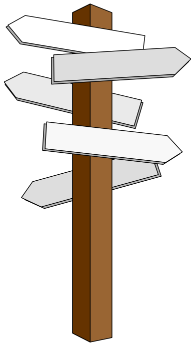

- Advanced modeling of TCP.
  - Complex distributions of inter-loss times.
  - The ACK clock approach
  - The Markov chain approach
  - Short-lived TCP transfers.

The Markov Chain approach

- Take as random variable the window at the beginning of the n-th RTT. Let it W_n.
- Suppose we know the probability that a loss occurs in an RTT and that this probability is not function of the past.
  - For a Poisson loss process λ, Prob{loss in RTT_n} = $1-e^{-\lambda RTT}$
  - For a Bernoulli loss process p, Prob{loss in RTT_n} = $1 - (1 - p)^{W_n}$
    Note how the probability is function of W_n in this case !
- W_n forms then a discrete time Markov chain.
- Study it !
- Then compute throughput and other metrics.
  For example, the throughput $X = E[W_n] / RTT$.
- Think about the case of variable round-trip time ! Answer in [ABR04]

Outline

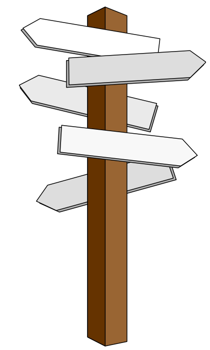

- Advanced modeling of TCP.
  - Complex distributions of inter-loss times.
  - The ACK clock approach
  - The Markov chain approach
  - Short-lived TCP transfers.

Finite size of TCP transfers

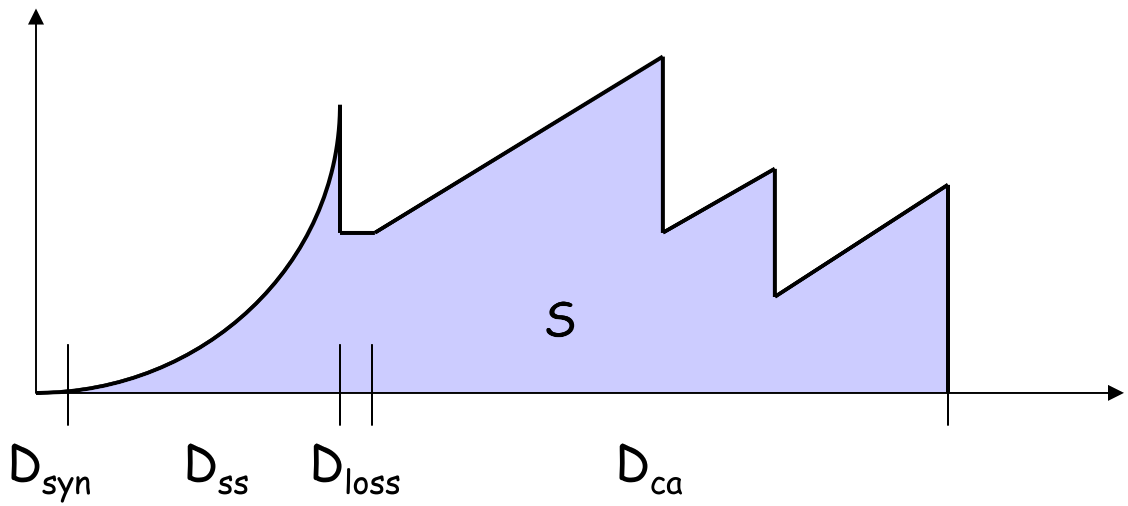

- Average latency: $D = D_{syn} + D_{ss} + D_{loss} + D_{ca}$, and throughput = S/D.
- Compute first D_{ss} as the average time until the first loss.
- Compute the average number of packets transmitted during slow start S_{ss}.
- Subtract S_{ss} from S to compute the number of packets to transmit in Congestion Avoidance, define $S_{ca} = S - S_{ss}$.
- Compute D_{ca} using the throughput of long-lived TCP: $D_{ca} = S_{ca} / X$

Duration of Slow Start

- Let $\gamma = 1 + 1/b$. The window size is multiplied by γ every RTT during SS.
- The average number of packets transmitted in SS:

$$S_{ss} = \left(\sum_{k=0}^{S-1} k.(1-p)^k.p\right) + S.(1-p)^S = \frac{(1-(1-p)^S)(1-p)}{p} + 1$$

- The number of packets transmitted until round-trip time " i " in SS is (assuming that the slow start threshold and receiver window are infinite, w_0 is the initial window size):

$$S_{ss}(i) = w_0 + w_0.\gamma + w_0.\gamma^2 + \ldots + w_0.\gamma^{i-1} = w_0.\frac{\gamma^i - 1}{\gamma - 1} \quad \Rightarrow \quad i = \log_\gamma\left(\frac{S_{ss}(i).(\gamma-1)}{w_0} + 1\right)$$

- The average duration of slow start can then be approximated by:

$$D_{ss} = RTT.\log_\gamma\left(\frac{S_{ss}.(\gamma-1)}{w_0} + 1\right)$$

Finite size of TCP transfers

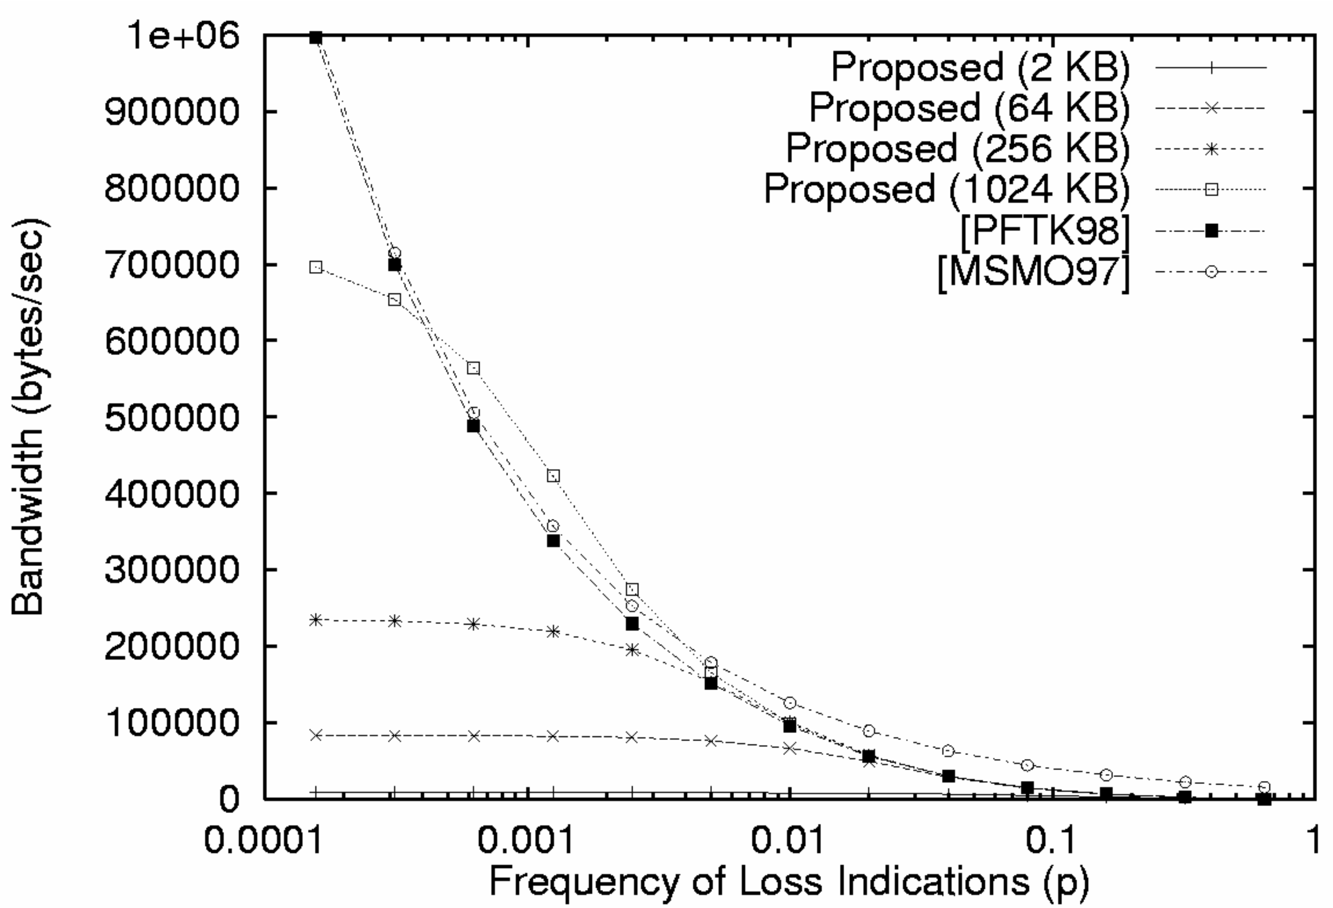

- Proposed = The model proposed in the below paper.
  - Accounts for the transfer size.
- PFTK = The previous "more complete" model.
- MSMO = The simple square root formula.

RTT = 100 ms, MSS = 1460 bytes, receiver window = 10Mbytes

N. Cardwell, S. Savage and T. Anderson, "Modeling TCP Latency", IEEE INFOCOM, March 2000.

Outline

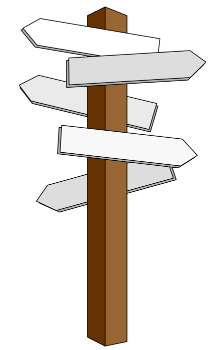

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

Inferring p and RTT from network characteristics: Example

Case of a network crossed by long-lived TCP connections and implementing RED buffers in its routers.

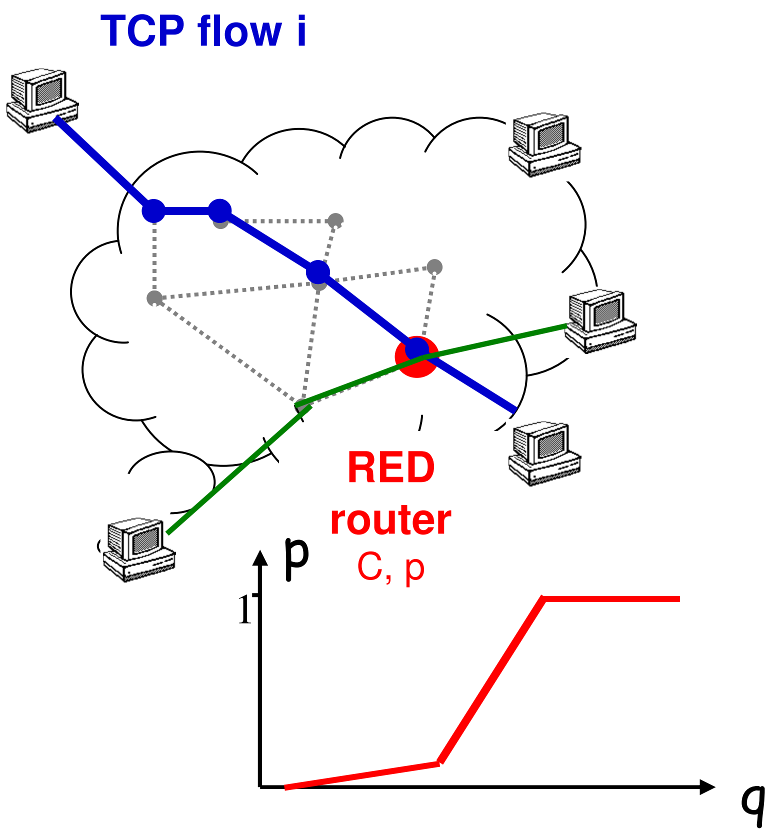

- N TCP flows
  - Throughputs $\{X_i(p_i,RTT_i)\}$.
- V bottlenecked RED routers
  - Capacities $\{C_v\}$.
  - Average queue lengths $\{q_v\}$.
  - Drop probability $\{p_v(q_v)\}$ (relation between p_v and q_v given by the diagram of RED).
- Write TCP throughputs as only a function of $\{q_v\}$: $X_i(\{q_v\})$.
- Write for bottleneck routers:

  $\Sigma_i X_i(\{q_v\}) = C_v$, $v = 1,\dots,V$
- V equations, V unknowns !

Throughput of TCP as a function of network characteristics

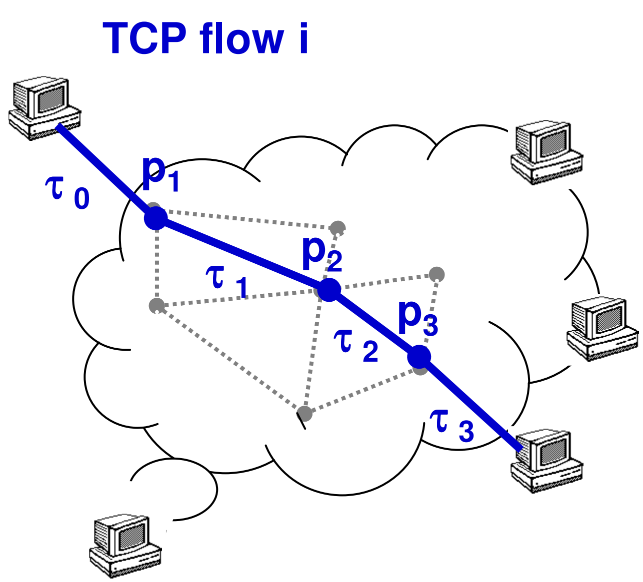

- Assume independence of packet losses in routers:
  $p = 1 - (1-p_1)(1-p_2)(1-p_3)$.
  When losses are rare, we can write:
  $p = p_1 + p_2 + p_3$.
- $RTT_i = 2(\tau_0 + \tau_1 + \tau_2 + \tau_3) + q_1/C_1 + q_2/C_2 + q_3/C_3$.
- The throughput X_i (p_i, RTT_i) can be written as a function of $\{q_v\}$ by using $p_v = p\ (q_v)$.
- TCP rate at 3 : $X_i\ /\ (1-p_1)$.
- TCP rate at 2: $X_i\ /\ (1-p_3)\ /\ (1-p_2)$.
- TCP rate at 1, or TCP sending rate:
  $X_i\ /\ (1-p_3)\ /\ (1-p_2)\ /\ (1-p_1)$.

Some results: tandem network

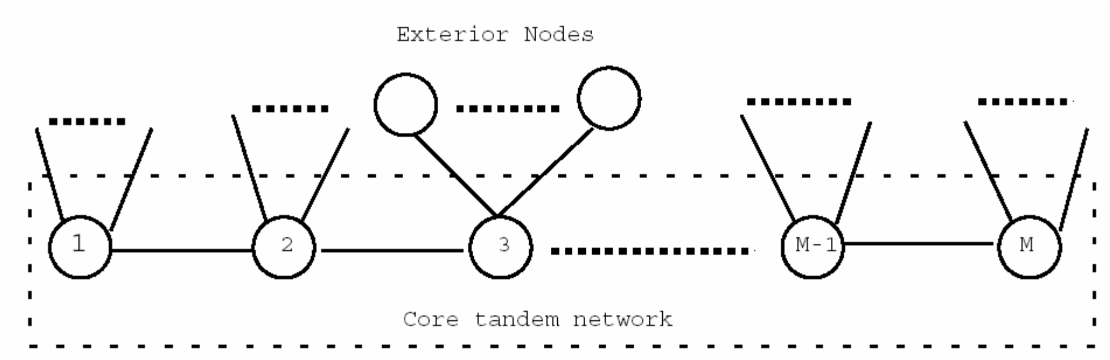

- 16 simulation scenarios:
  - Different configurations of RED parameters.
  - Number of routers in [5 - 10].
- For each scenario, the two way propagation delay takes different values in [20 - 120] ms.
- And the bandwidth takes different values in [2 - 6] Mbps.

Tian Bu and Don Towsley, "Fixed Point Approximation for TCP behavior in an AQM Network", Proceedings of ACM SIGMETRICS 2001.

Some results: tandem network

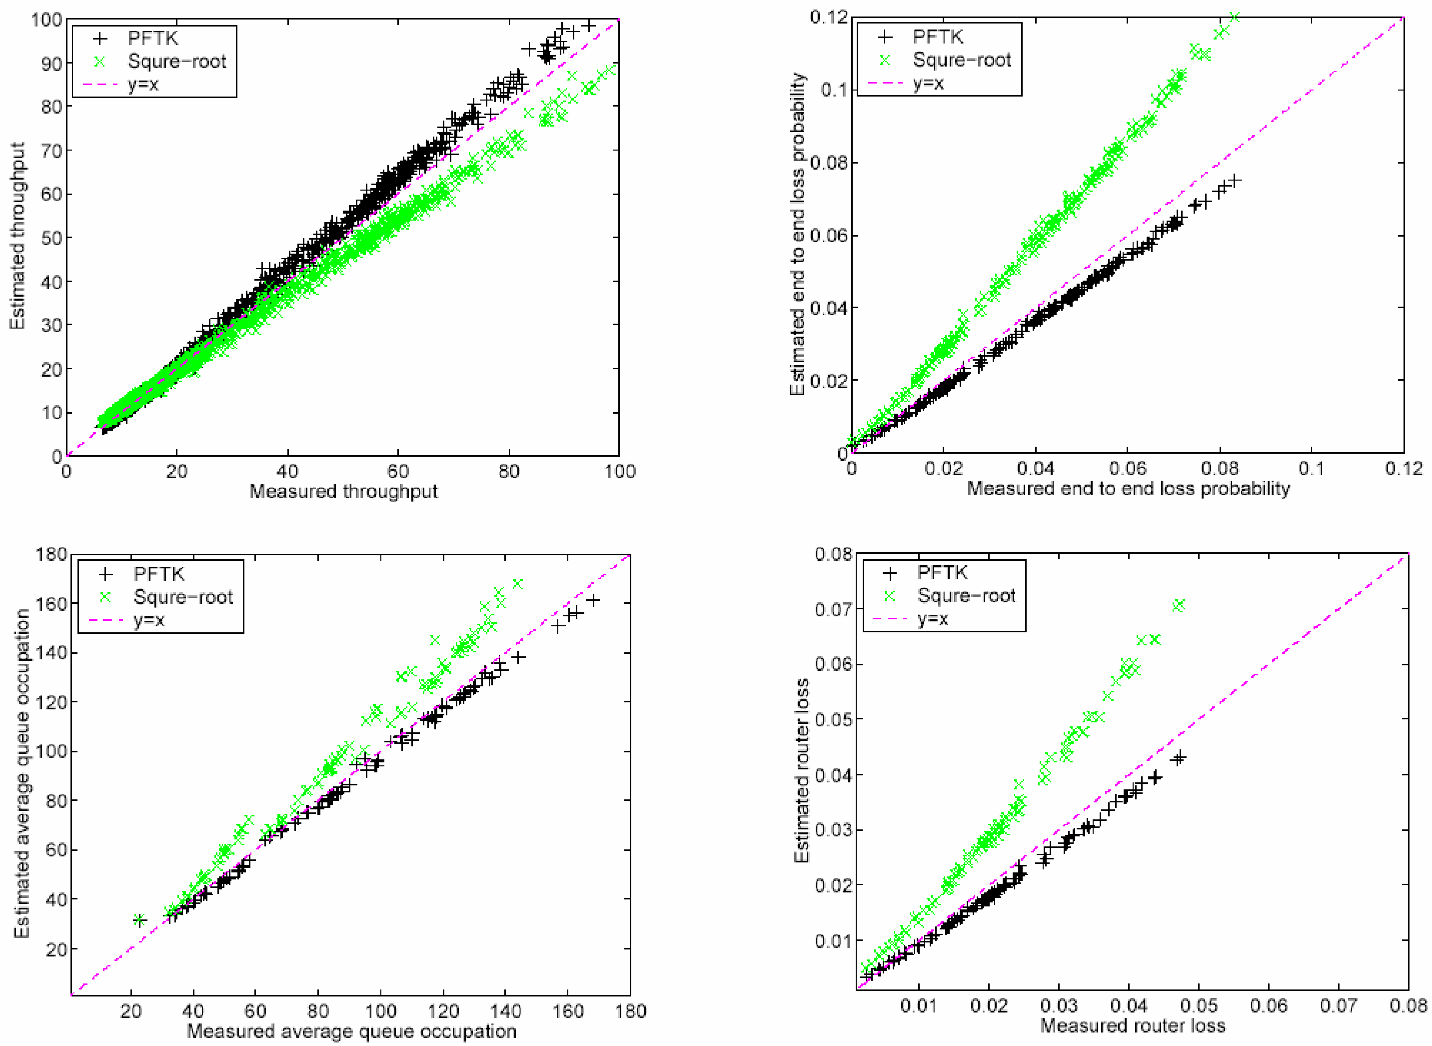

General model for a bottleneck router

- Model for the router:
  - λ be the arrival rate of packets to the router.
  - **$p(\lambda)$** be the packet loss rate.
  - **$q(\lambda)$** be the average queue length.
- Model for TCP traffic:
  - Short TCP sessions of arrival rate λ_0 and of size **$\{S_n\}$**.
  - **N** long-lived TCP sessions of average sending rates **$X(p,RTT_i)$**.
- Combining both models together (Fixed-point approach):
  - **$RTT_i = \tau_i + q(\lambda)/C$**.
  - Write one equation in λ,

$$\lambda = \frac{\lambda_0 E[S_n]}{1-p} + \sum_{i=1}^{N} X(p(\lambda), \tau_i + q(\lambda)/C)$$

  - Solve this equation for λ, then for the loss rate, the average queue length, and the performance measures at the TCP level.
- The difficulty is in the computation of **$p(\lambda)$** and **$q(\lambda)$** !
  - We don't know how the packet arrival process looks like.

Modeling the router as M/M/1/K

- Packets arrive according to a Poisson process, and have exponentially distributed service times.
- $$p(\lambda) = \frac{\rho^K(1-\rho)}{1-\rho^{K+1}}, \text{ with } \rho = \lambda / C.$$
  $$q(\lambda) = \frac{\rho}{1-\rho} - \frac{(K+1)\rho^{K+1}}{1-\rho^{K+1}}$$
- A unique solution to the fixed point method exists for all $\lambda_0 E[S_n] < C$.

Urtzi Ayesta, Kostya Avranchenkov, Eitan Altman, Chadi Barakat, Parijat Dube, "Multilevel Approach for Modeling TCP/IP", in proceedings of ITC-18, Berlin, September 2003.

- The accuracy of the Poisson assumption on arrivals increases with the degree of multiplexing of sessions.

Outline

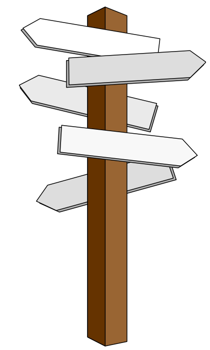

- Introduction.
- A simple model for TCP throughput: Square Root Formula.
- Enhancing the square root formula to account for the packet nature of TCP, timeouts, and receiver window.
- Advanced modeling of TCP.
- Inferring the parameter of a model for TCP.
- Conclusions

Conclusions

- Overview of the main techniques for fine-grained modeling of TCP.
- Still many other issues to be well understood:
  - The sub-linearity of TCP window growth on paths where the window and the RTT are correlated.
  - An accurate modeling of the impact of the receiver window size.
  - An accurate modeling of the TCP recovery phase (I would say this is the weakest component in actual models).
- But all this is useless if there is no accurate model for the loss process in the Internet.
  - Is it Bernoulli ? is it Poisson ? is it Markov Modulated Poisson ?
  - Effort can be wasted on modeling non realistic or non general scenarios.
- And may be finally simple models are the most useful …

References

- [AAB00] Eitan Altman, Kostia Avrachenkov, and Chadi Barakat, "A stochastic model of tcp/ip with stationary random losses", ACM SIGCOMM'2000.
- [ABR04] Eitan Altman, Chadi Barakat, Victor M. Ramos R., "Analysis of AIMD protocols over paths with variable delay", in proceedings of IEEE INFOCOM, Hong Kong, March 2004.
- Urtzi Ayesta, Kostya Avranchenkov, Eitan Altman, Chadi Barakat, Parijat Dube, "Multilevel Approach for Modeling TCP/IP", in proceedings of ITC-18, Berlin, September 2003.
- [BAR01] C. Barakat, "TCP modeling and validation", IEEE Network, vol. 15, no. 3, pp. 38-47, May 2001.
- [BT01] Tian Bu and Don Towsley, "Fixed Point Approximation for TCP behavior in an AQM Network", Proceedings of ACM SIGMETRICS 2001.
- [CSA00] N. Cardwell, S. Savage and T. Anderson, "Modeling TCP Latency", IEEE INFOCOM, March 2000.
- [FLO91] Floyd, S., Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic. Computer Communication Review, Vol.21, No.5, October 1991.
- [MSMO97] M. Mathis, J. Semke, J. Mahdavi, T. Ott, ``The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm'', Computer Communication Review, vol. 27, no. 3, pp. 67-82, Jul. 1997.
- [OKM96] T. Ott, J. Kemperman, M. Mathis, "The stationary behavior of ideal TCP congestion avoidance", manuscript, Aug. 1996.
- [PFTK98] Padhye J., Firoiu V., Towsley D., and Kurose J., "Modeling TCP Throughput: a Simple Model and its Empirical Validation", in Proceedings of ACM SIGCOMM, August 1998.