

Hybrid Control and Switched Systems


Lecture #17 Hybrid Systems Modeling of Communication Networks

João P. Hespanha

University of California
at Santa Barbara




Motivation



Why model network traffic?

- to validate designs through simulation (scalability, performance)
- to analyze and design protocols (throughput, fairness, security, etc.)
- to tune network parameters (queue sizes, bandwidths, etc.)

Types of models




Packet-level modeling


- tracks **individual data packets**, as they travel across the network
- ignores the data content of individual packets
- **sub-millisecond** time accuracy
- computationally very **intensive**

Fluid-based modeling

- tracks **time/ensemble-average packet rates** across the network
- does not explicitly model individual events (acknowledgments, drops, queues becoming empty, etc.)
- time accuracy of a few **seconds** for time-average
- only suitable to model **many similar flows** for ensemble-average
- computationally very **efficient** (at least for first order statistics)



Types of models



captures fast dynamics
even for a small number of flow

Hybrid modeling


- keeps track of packet **rates** for each **flow averaged** over small time scales
- explicitly models some **discrete events** (drops, queues becoming empty, etc.)
- time accuracy of a few **milliseconds** (round-trip time)
- computationally **efficient**

provide information about both
average, peak, and "instantaneous"
resource utilization
(queues, bandwidth, etc.)

Summary

- Modeling 1st pass: Dumbbell topology & simplified TCP
- Modeling 2nd pass: General topology, TCP and UDP models
- Validation
- Simulation complexity

1st pass – Dumbbell topology



Several flows follow the **same path** and compete for bandwidth in a single **bottleneck link**


Prototypical network to study congestion control

routing is trivial

single queue

B is unknown to the data sources and possibly time-varying

Queue dynamics




When $\sum_f r_f$ exceeds B the queue fills and data is lost (drops)


$$\dot{q} = \begin{cases} \sum_f r_f - B & 0 \leq q \leq q_{\max} \\ 0 & \text{otherwise} \end{cases}$$

$$q = q_{\max}, \sum_f r_f > B \Rightarrow \text{drop} \quad (\text{discrete event - relevant for congestion control})$$

Queue dynamics



Hybrid automaton representation:



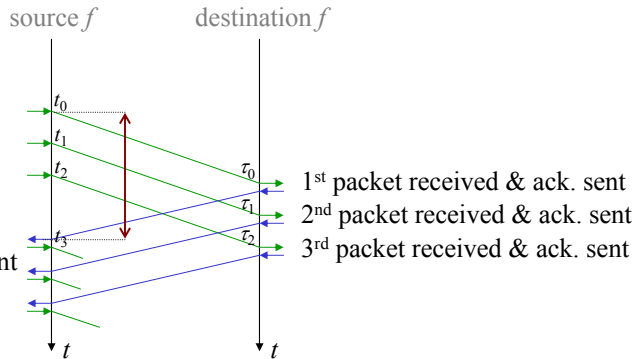
Window-based rate adjustment

w_f (window size) \equiv number of packets that can remain unacknowledged for by the destination

e.g., $w_f = 3$

1st packet sent
2nd packet sent
3rd packet sent

1st ack. received \Rightarrow
4th packet can be sent



w_f effectively determines the sending rate r_f :

$$r_f(t) = \frac{w_f(t)}{RTT(t)}$$

round-trip time


Window-based rate adjustment

w_f (window size) \equiv number of packets that can remain unacknowledged for by the destination

$r_f(t) = \frac{w_f(t)}{RTT(t)} \equiv$ sending rate

$$RTT(t) = T_p + \frac{1}{B}q(t)$$

total round-trip time \rightarrow $RTT(t)$ \leftarrow per-packet transmission time
 propagation delay \leftarrow T_p \leftarrow time in queue until transmission \leftarrow $\frac{1}{B}q(t)$

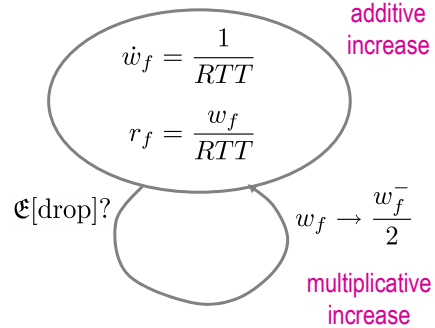


This mechanism is still not sufficient to prevent a catastrophic collapse of the network if the sources set the w_f too large

TCP congestion avoidance

1. While there are no drops, increase w_f by 1 on each RTT (additive increase)
 2. When a drop occurs, divide w_f by 2 (multiplicative decrease)
- (congestion controller constantly probes the network for more bandwidth)

TCP congestion avoidance

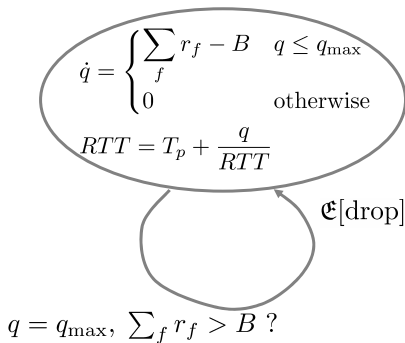


disclaimer: this is a very simplified version of TCP Reno, better models later...


TCP congestion avoidance

1. While there are no drops, increase w_f by 1 on each RTT (additive increase)
 2. When a drop occurs, divide w_f by 2 (multiplicative decrease)
- (congestion controller constantly probes the network for more bandwidth)

Queuing model



TCP congestion avoidance




disclaimer: this is a very simplified version of TCP Reno, better models later...

TCP congestion avoidance

1. While there are no drops, increase w_f by 1 on each RTT (additive increase)
2. When a drop occurs, divide w_f by 2 (multiplicative decrease)
(congestion controller constantly probes the network for more bandwidth)

TCP + Queuing model



disclaimer: this is a very simplified version of TCP Reno, better models later...


Linearization of the TCP model

Time normalization \equiv define a new “time” variable τ by

$$\frac{d\tau}{dt} = \frac{1}{RTT} = \frac{B}{BT_p + q}, \quad \tau(0) = 0$$


1 unit of $\tau \equiv$ 1 round-trip time

TCP + Queuing model




In normalized time, the continuous dynamics become **linear**


Switching-by-switching analysis



$x_k := x(t_k) \equiv$ continuous state before the k^{th} multiplicative decrease



Switching-by-switching analysis



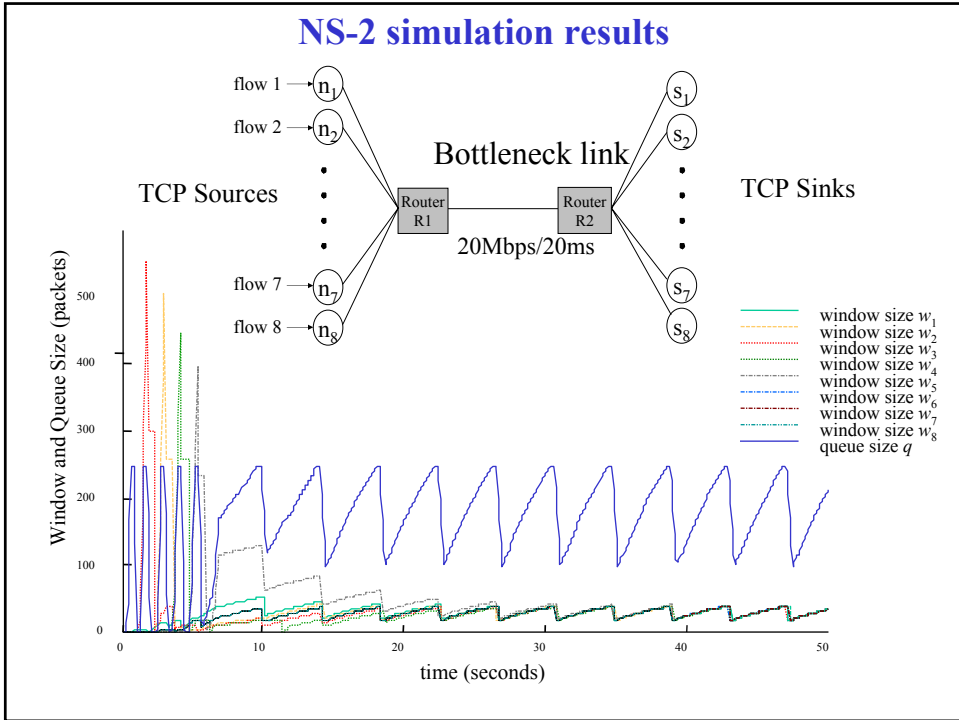
$x_k := x(t_k) \equiv$ continuous state before the k^{th} multiplicative decrease

Theorem. The function T is a contraction. In particular,

$$\|T(a) - T(b)\|_* \leq \frac{1}{2} \|a - b\|_*, \quad \forall a, b$$

Therefore

- $x_k \rightarrow x_\infty$ as $k \rightarrow \infty$
 - $x(t) \rightarrow x_\infty(t)$ as $t \rightarrow \infty$
- $x_\infty \equiv$ constant
 $x_\infty(t) \equiv$ periodic limit cycle



Results

(valid for $n \ll \frac{q_{\max} + BT_p}{2}$)

The diagram shows a timeline from t_0 to t_3 . The phases are: additive increase (from t_0 to t_1), multiplicative decrease (from t_1 to t_2), additive increase (from t_2 to t_3), and multiplicative decrease (from t_3 to the end).

Window synchronization:

$$w_f(t_k) \rightsquigarrow \frac{q_{\max} + BT_p}{n}, \quad \text{as } k \rightarrow \infty \quad \forall f$$

convergence is exponential, as fast as $.5^k$

Steady-state formulas:

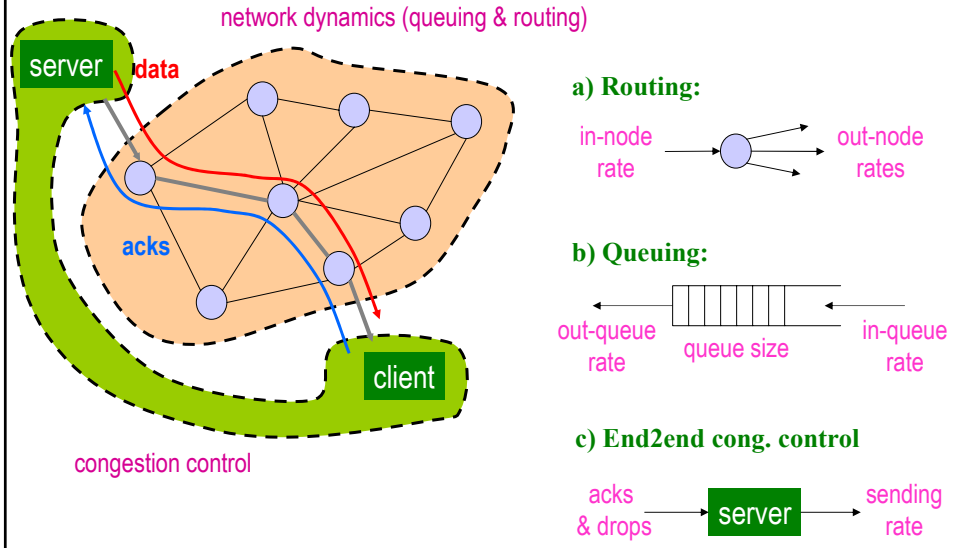
average drop rate $p \approx 2.67 \left(\frac{n}{q_{\max} + BT_p} \right)^2$

average RTT $\overline{RTT} \approx .78 \left(T_p + \frac{q_{\max}}{B} \right)$

average throughput $T \approx \frac{1.27}{\overline{RTT} \sqrt{p}}$ (well known TCP-friendly formula)


2nd pass – general topology

A communication network can be viewed as the interconnection of several blocks with specific dynamics



Routing

determines the sequence of links followed by each flow



Conservation of flows:

$$s_f^\ell = \begin{cases} r_f & \text{flow } f \text{ starts in this node} \\ r_f^{\ell'} & \text{flow } f \text{ comes from link } \ell' \end{cases}$$

in-queue rate of flow f → s_f^ℓ

end2end sending rate of flow f ← r_f


upstream out-queue rate of flow f ← $r_f^{\ell'}$

indexes ℓ and ℓ' determined by routing tables

Routing


determines the sequence of links followed by each flow

Multicast




$$s_f^{\ell_1} = s_f^{\ell_2} = r_f^{\ell'}$$

Multi-path routing



$$s_f^{\ell} = r_f^{\ell'} + r_f^{\ell''}$$

Queue dynamics



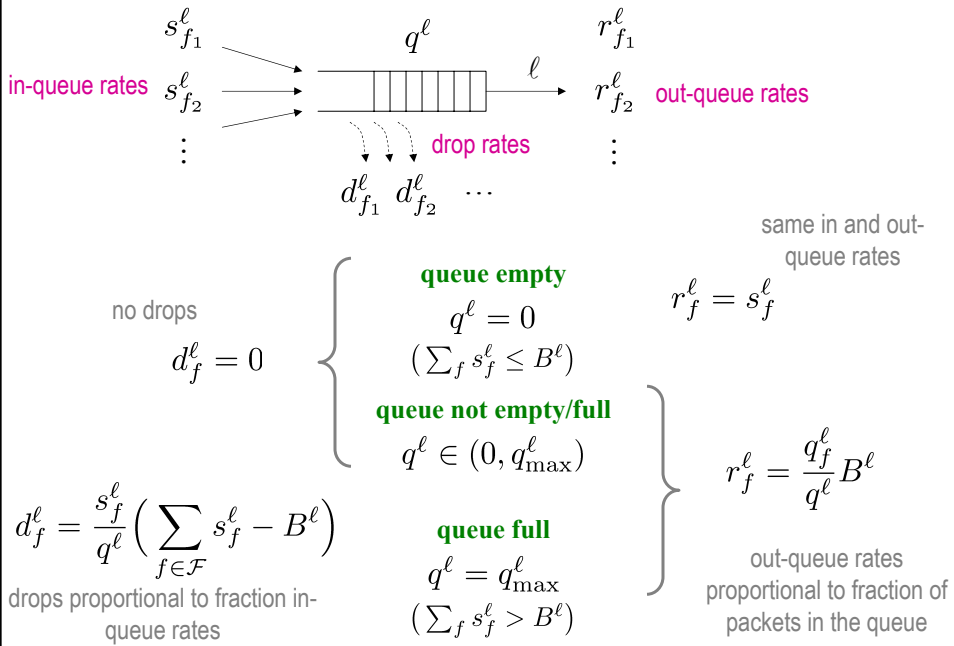
total queue size $\rightarrow q^{\ell} = \sum_{f \in \mathcal{F}} q_f^{\ell} \leftarrow$ queue size due to flow f

the packets of each flow are assumed uniformly distributed in the queue

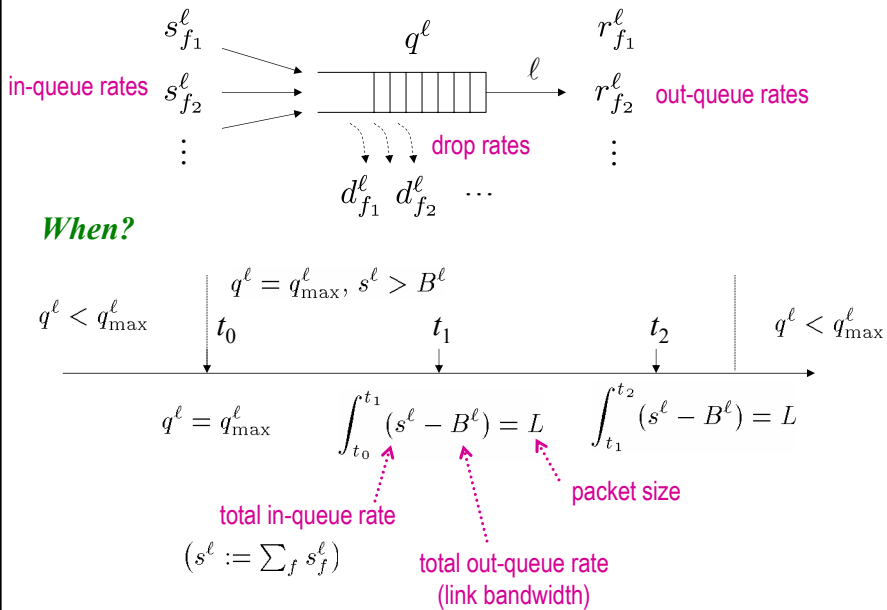
Queue dynamics:

$$\dot{q}_f^{\ell} = s_f^{\ell} - d_f^{\ell} - r_f^{\ell}$$


Queue dynamics




Drops events



Drops events



When?



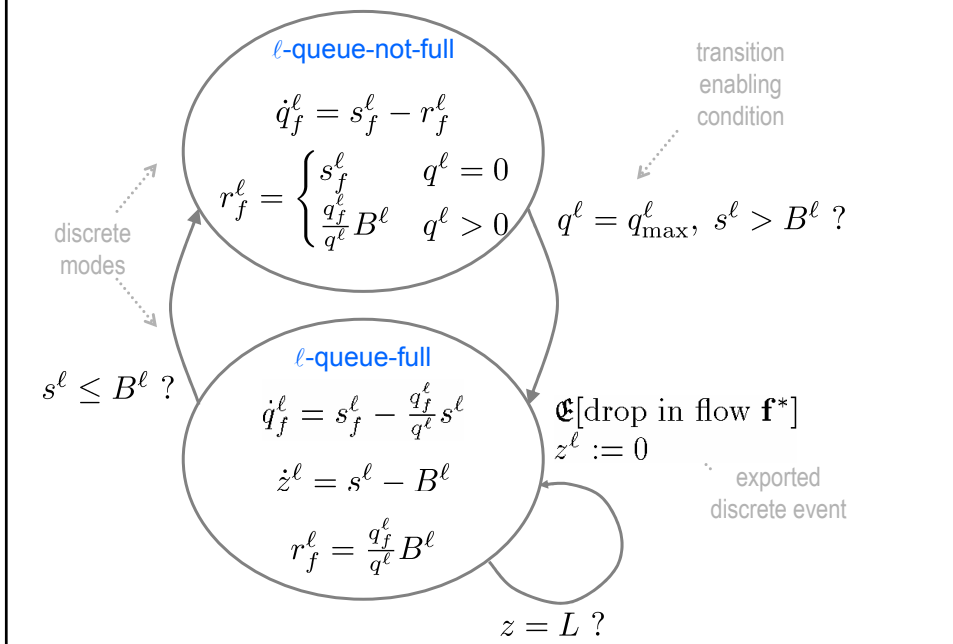
Which flows?

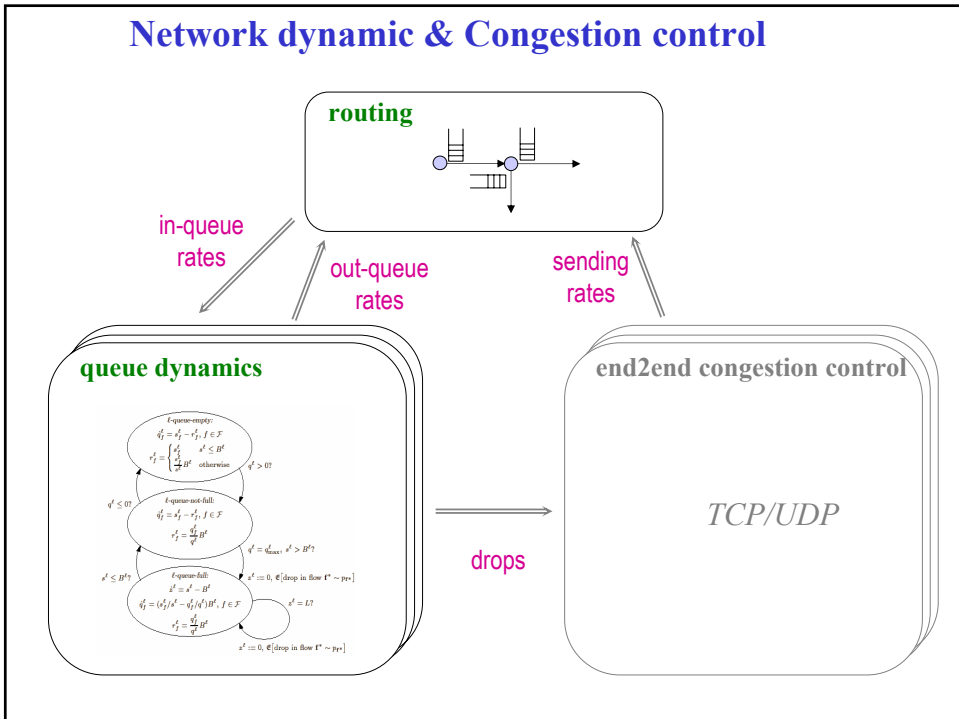
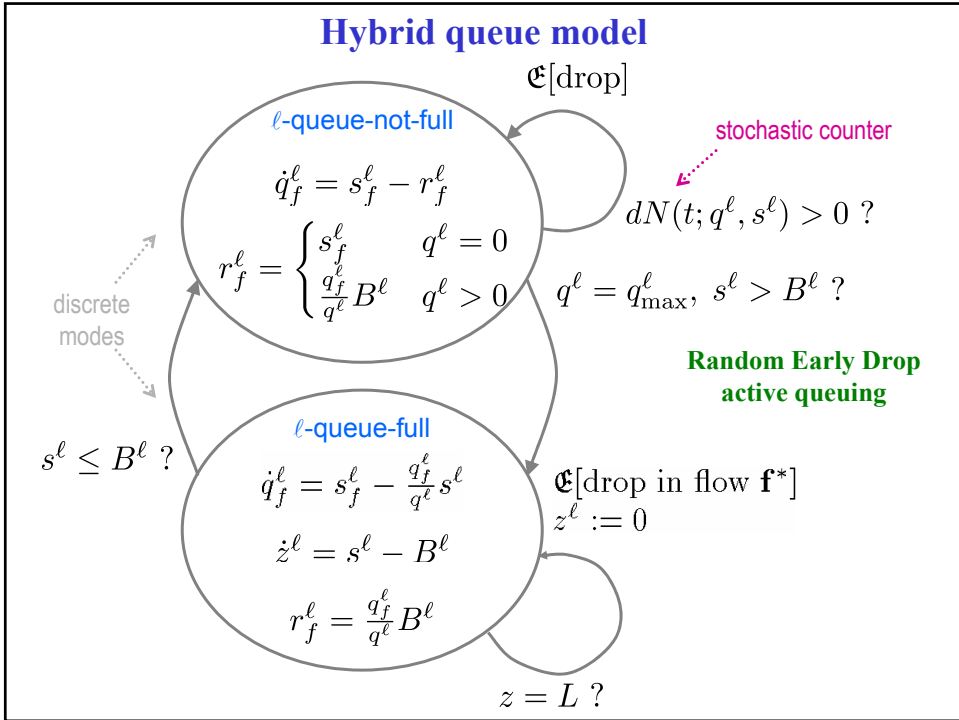
$$P(\mathbf{f}^*(t_k) = f) = \frac{s_f^l}{s^l}, \quad f \in \mathcal{F}$$

(drop tail dropping)

flow that suffers drop at time t_k

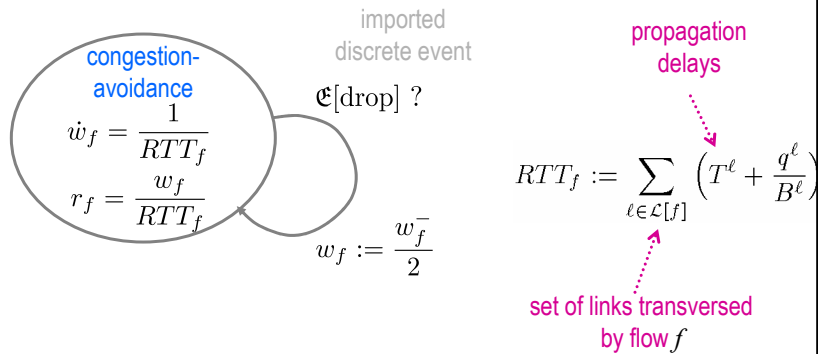
Hybrid queue model





Additive Increase/Multiplicative Decrease

1. While there are no drops, increase w_f by 1 on each RTT (additive increase)
2. When a drop occurs, divide w_f by 2 (multiplicative decrease)
(congestion controller constantly probe the network for more bandwidth)




TCP-Reno is based on AIMD but uses other discrete modes to improve performance

Slow start

In the beginning, pure AIMD takes a long time to reach an adequate window size

3. Until a drop occurs (or a threshold $ssth_f$ is reached), double w_f on each RTT
4. When a drop occurs, divide w_f and the threshold $ssth_f$ by 2

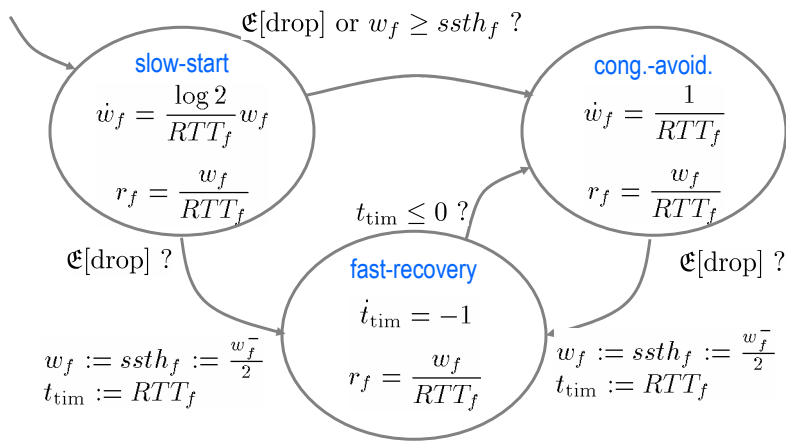


especially important for short-lived flows...

Fast recovery

After a drop is detected, new data should be sent while the dropped one is retransmitted

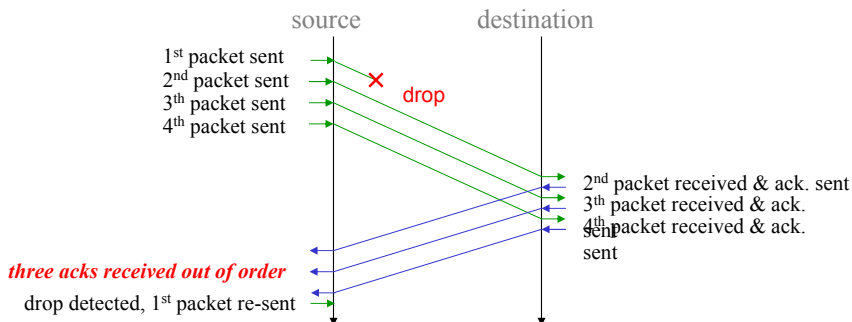
5. During retransmission, data is sent at a rate consistent with a window size of $w_f/2$



(consistent with TCP-SACK for multiple consecutive drops)

Timeouts

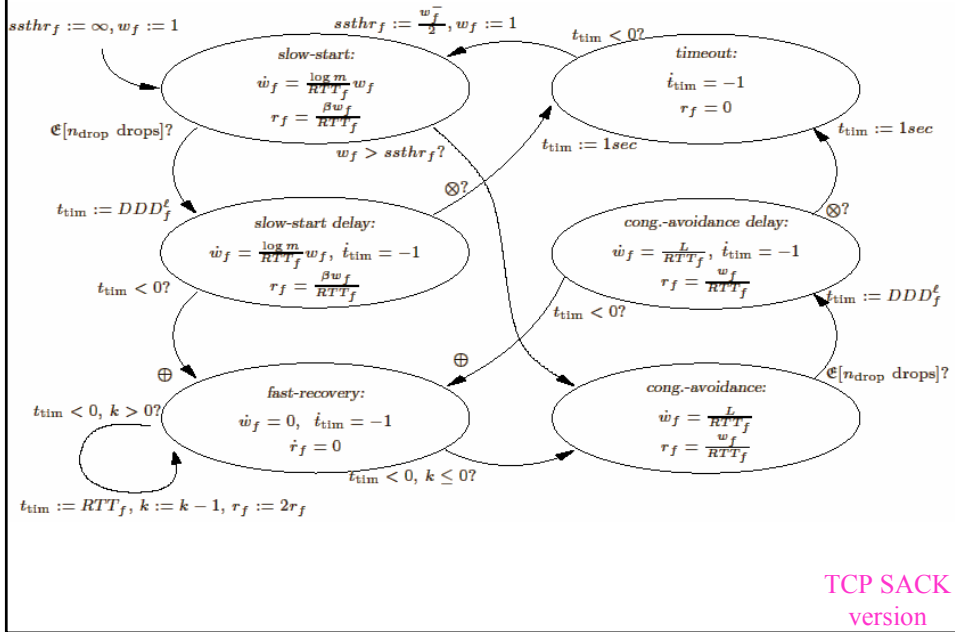
Typically, drops are detected because one acknowledgment in the sequence is missing.



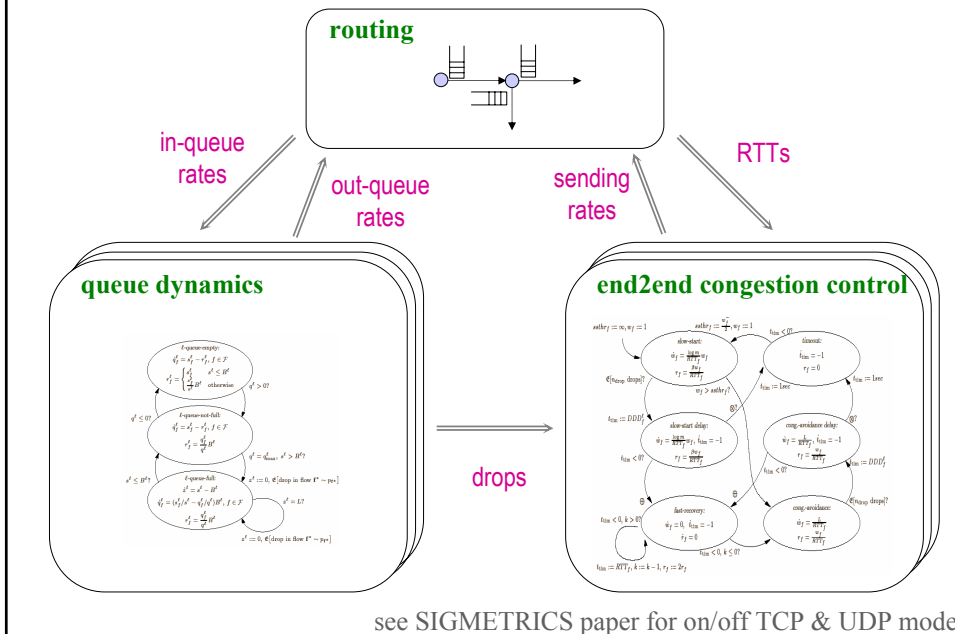
When the window size becomes smaller than 4, this mechanism fails and drops must be detected through acknowledgement timeout.

6. When a drop is detected through timeout:
 - a. the slow-start threshold $ssth_f$ is set equal to half the window size,
 - b. the window size is reduced to one,
 - c. the controller transitions to slow-start

Fast recovery, timeouts, drop-detection delay...



Network dynamic & Congestion control




Validation methodology


Compared simulation results from

- ns-2 packet-level simulator
- hybrid models implemented in Modelica

Plots in the following slides refer to two test topologies




- 10ms propagation delay
- drop-tail queuing
- 5-500Mbps bottleneck throughput
- 0-10% UDP on/off background traffic




- 45,90,135,180ms propagation delays
- drop-tail queuing
- 5-500Mbps bottleneck throughput
- 0-10% UDP on/off background traffic

Simulation traces

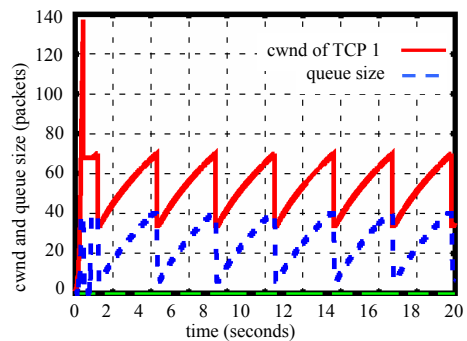


- single TCP flow
- 5Mbps bottleneck throughput
- no background traffic

hybrid model




ns-2



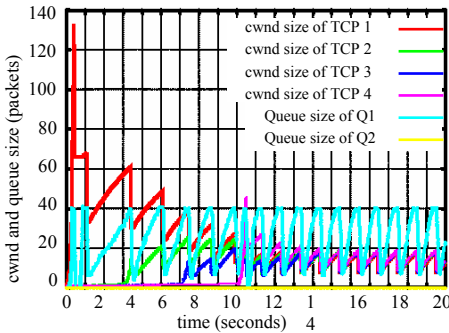
slow-start, fast recovery, and congestion avoidance accurately captured

Simulation traces




- four competing TCP flow (starting at different times)
- 5Mbps bottleneck throughput
- no background traffic

hybrid model




ns-2



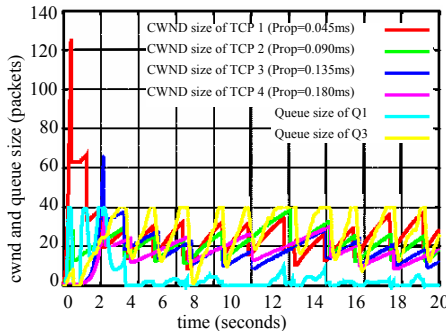
the hybrid model accurately captures flow synchronization

Simulation traces




- four competing TCP flow (different propagation delays)
- 5Mbps bottleneck throughput
- 10% UDP background traffic (exp. distributed on-off times)


hybrid model



ns-2



Average throughput and RTTs

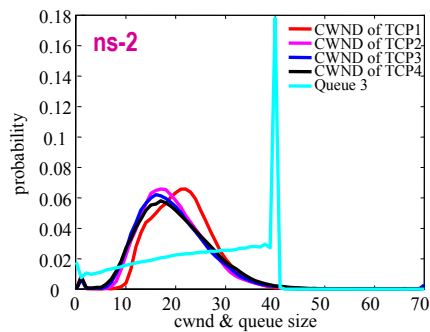



- four competing TCP flow (different propagation delays)
- 5Mbps bottleneck throughput
- 10% UDP background traffic (exp. distributed on-off times)

	Thru. 1	Thru. 2	Thru. 3	Thru. 4	RTT1	RTT2	RTT3	RTT4
ns-2	1.873	1.184	.836	.673	.0969	.141	.184	.227
hybrid model	1.824	1.091	.823	.669	.0879	.132	.180	.223
relative error	2.6%	7.9%	1.5%	.7%	9.3%	5.9%	3.6%	2.1%


the hybrid model accurately captures TCP unfairness for different propagation delays

Empirical distributions



the hybrid model captures the whole distribution of congestion windows and queue size

Execution time



- ns-2 complexity approximately scales with $o(nT)$ (# packets)

- hybrid simulator complexity approximately scales with $o\left(\frac{1}{RTT^2} \frac{n}{T}\right)$

number of flows

per-flow throughput

hybrid models are particularly suitable for large, high-bandwidth simulations (satellite, fiber optics, backbone)