Fluid Modeling of TCP in Wireless Networks

Majid Ghaderi

Department of Computer Science
University of Calgary
Email: mghaderi@cs.ucalgary.ca
Web: http://www.ucalgary.ca/~mghaderi

A typical cellular network

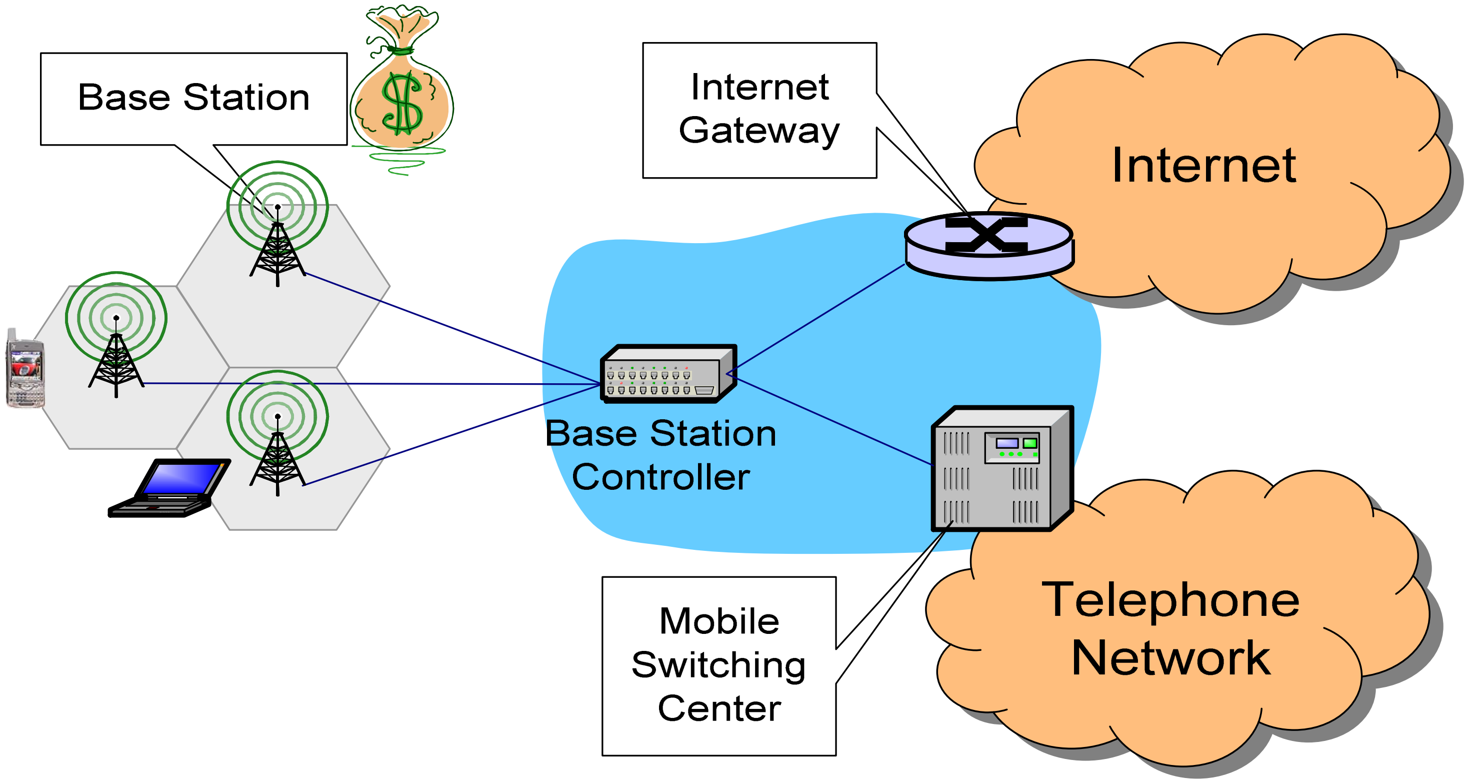

- Wireless bandwidth is limited and expensive

How expensive is spectrum?

Rogers	$	999,367,000
Telus	$	879,889,000
Bell	$	740,928,000
Quebecor	$	554,549,000
Globalive	$	442,099,000
Data A/V	$	243,159,000
Shaw	$	189,519,000
SaskTel	$	65,690,000
MTS	$	40,773,750
Bragg	$	25,620,000

For 20 MHz nation-wide

Canadian spectrum auction in July 2008 raised $4.25 billion.

In the US, Auction 73 in Jan 2008 raised $19.592 billion!

TCP in wireless networks

- TCP is the dominant transport protocol
- TCP has poor throughput over wireless channels
  - Interprets channel errors as sign of congestion
- Various solutions to improve throughput
  - Transport layer mechanisms: change TCP
    [Balakrish95, Ludwig00, Chan04]
  - Link layer mechanisms: apply FEC, ARQ and power control
    [Barakat02, Liu02, Baccelli06, Barman04]

What is the problem?

- Modern 3G/4G wireless systems (CDMA2000, EV-DO, WiMax)
  - Low bit-error-rate
  - **On-the-fly rate adaptation**
    - Network dynamically changes channel rate
    - Goal: maximize MAC throughput subject to some target frame-error-rate (FER)
    - TCP cognizant!
    - How to exploit this for TCP benefit?

This talk

- Optimizing rate adaption to maximize TCP throughput
  - Determine a rate adaptation policy (i.e., scheduler) at MAC
  - Determine the set of rates at PHY

Outline

- TCP-aware resource allocation
- Single TCP session
- Multiple TCP sessions

References

[1] M. Ghaderi, A. Sridharan, H. Zang, D. Towsley and R. Cruz,
TCP-aware resource allocation in CDMA networks, in Proc. *ACM Mobicom* 2006.

[2] M. Ghaderi, A. Sridharan, H. Zang, D. Towsley and R. Cruz,
Modeling TCP in a multi-rate multi-user CDMA system, in Proc. *Networking* 2007.

TCP overview

- Reliable end-to-end communication
- Congestion control
  - Increase linearly if no packet loss
  - Decrease to half if packet loss

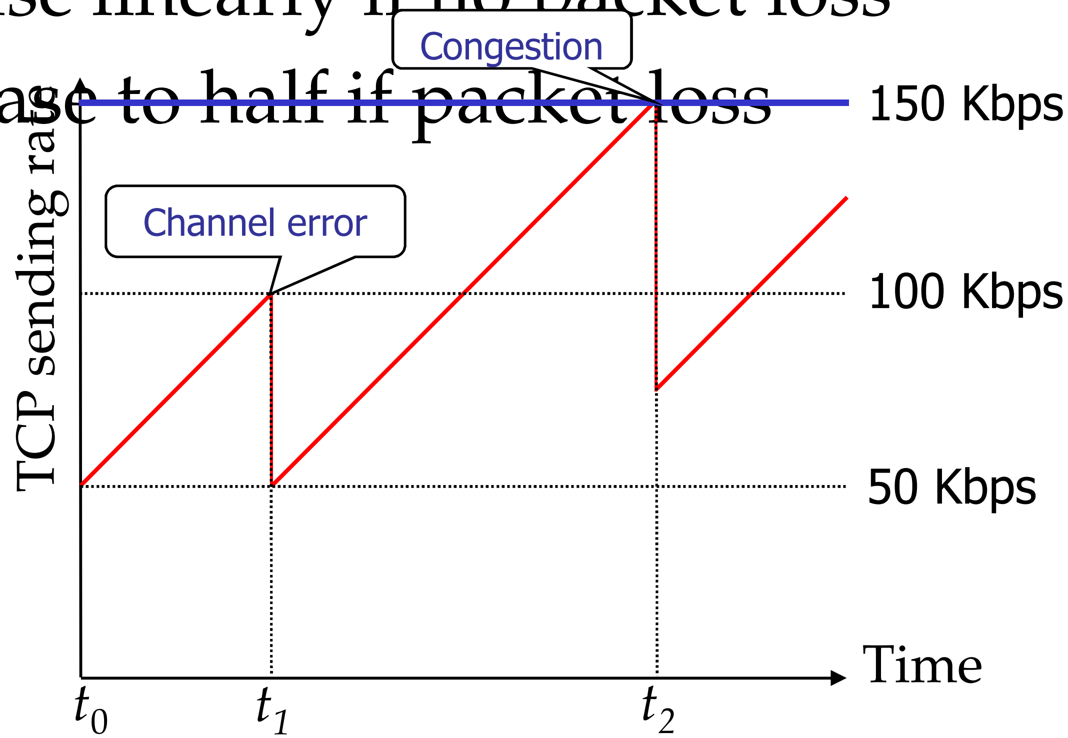

Previous work: static resource allocation

- Optimal fixed operating point
  - A priori optimization of system parameters (e.g., power level, coding rate) that maximize TCP throughput
- Optimizing Forward Error Correction (FEC) coding rate
  - Trade-off between channel rate and packet error probability
  - A single coding rate that maximizes TCP throughput → **Static coding**

Coding trade-off: rate vs. error

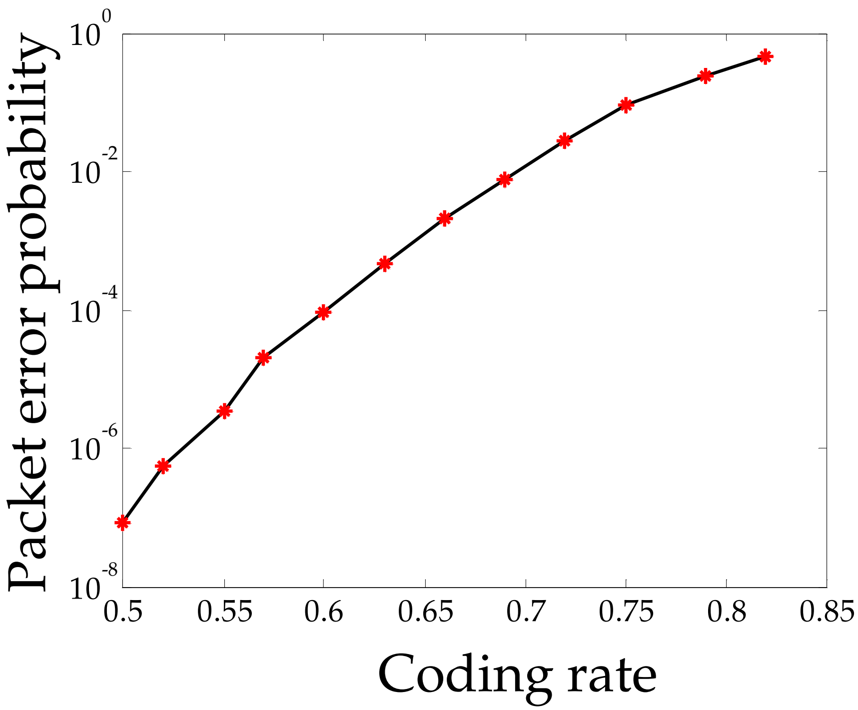

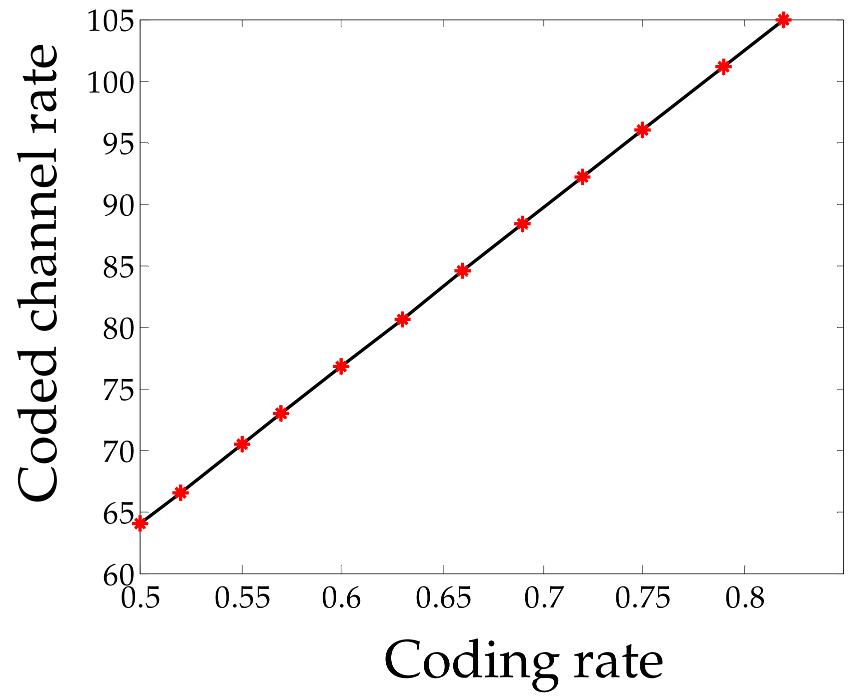

$$\text{Coding rate }(\rho) = \frac{\text{DATA}}{\text{DATA} + \text{FEC}}$$

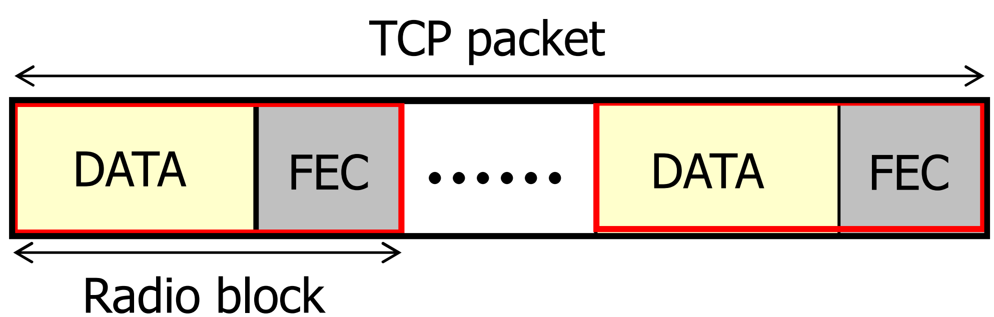

Static coding

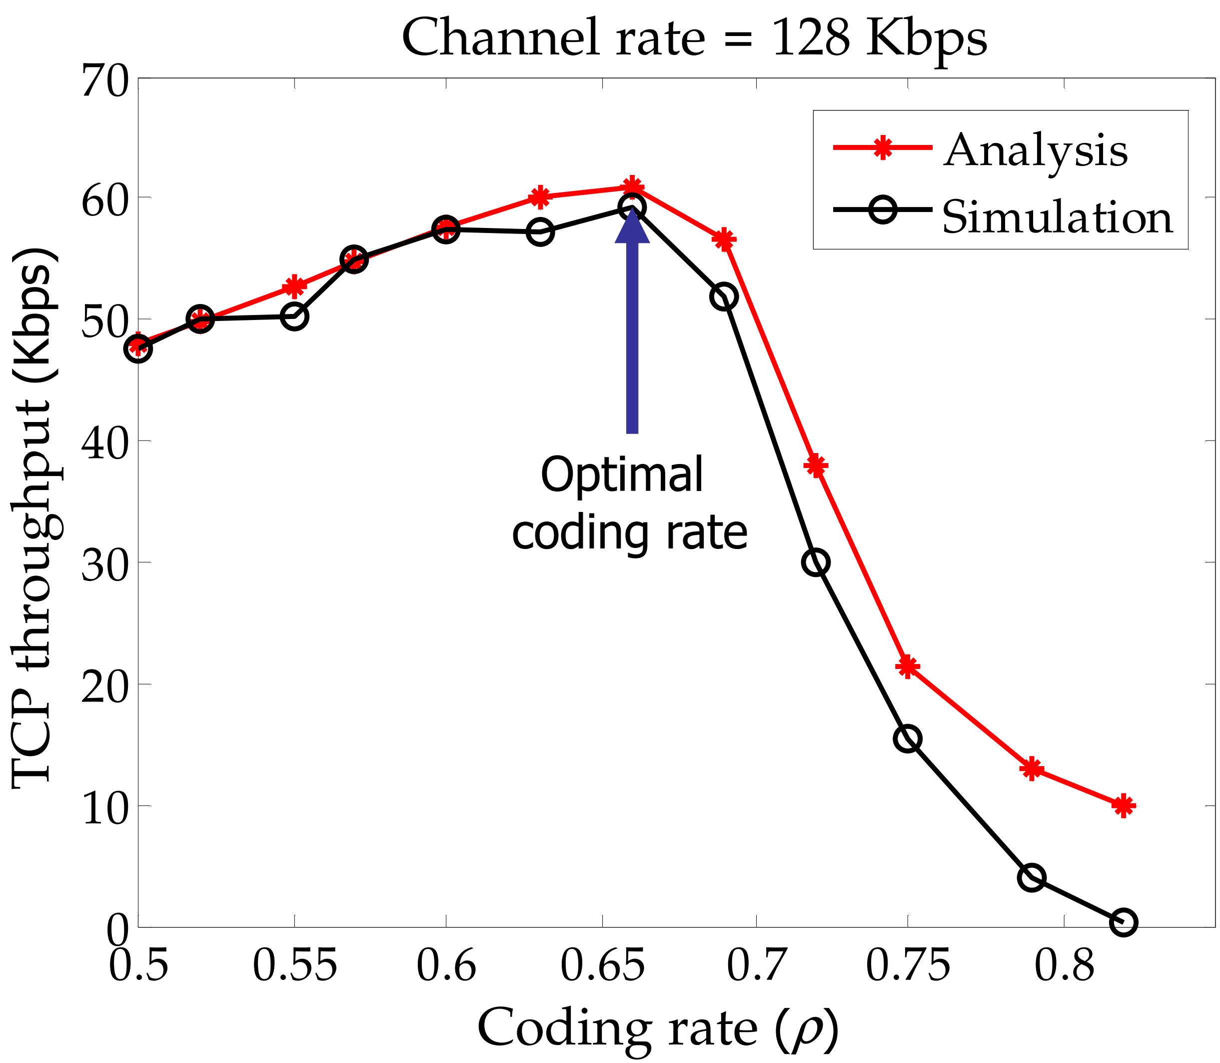

Max = 58.8 Kbps
ρ= 0.66

Adaptive coding

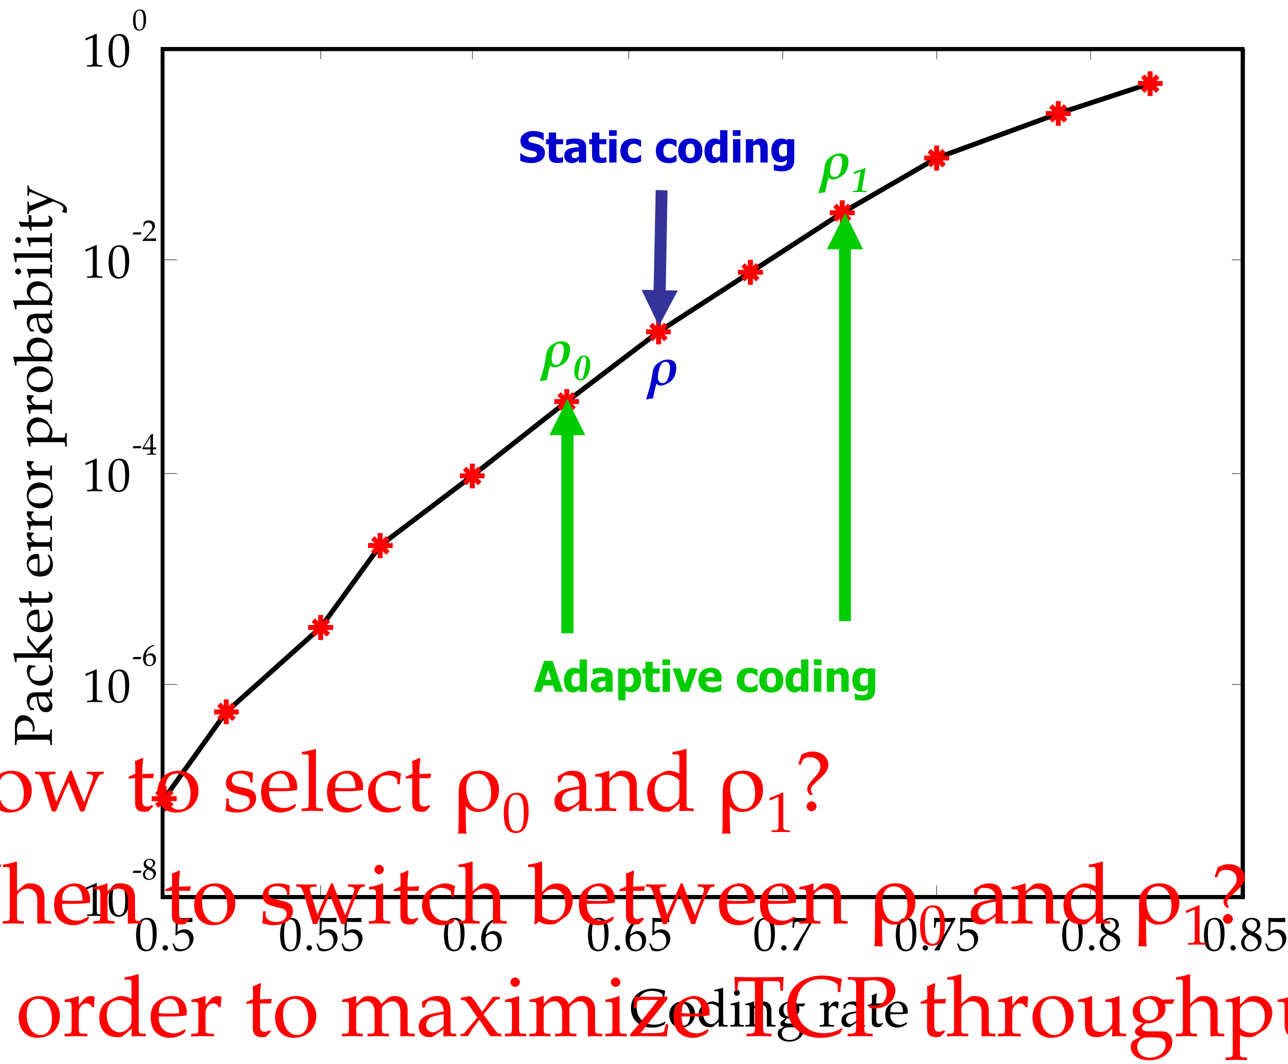

- How to select ρ_0 and ρ_1?
- When to switch between ρ_0 and ρ_1?
- In order to maximize TCP throughput

Outline

- TCP-aware resource allocation
- Single TCP session
- Multiple TCP sessions

Our approach: TCP-aware rate allocation

- Jointly optimize both the MAC and PHY layer parameters with respect to TCP dynamics
  - **Adaptive rate allocation:** Allocate channel rates based on TCP sending rate
  - **Channel rate optimization:** Choose a set of channel rates that jointly maximize TCP throughput across all potential rates at PHY

MAC: adaptive rate allocation

- How to allocate the channel rates?
- Assume two channel rates $C_0, C_1 (C_0 \leq C_1)$
- *C(t)*: allocated *channel rate* at time *t*
- *X(t)*: TCP *sending rate* at time *t*
- Wireless scheduler operates as follows:

$$C(t) = C_0 \text{ if } X(t) \leq C_0$$

$$C(t) = C_1 \text{ if } X(t) > C_0$$

PHY: channel rate optimization

- How to select the channel rates?
- p_i : packet error probability when channel rate is $C_i, i = 0,1$
- R_i : RTT when channel rate is C_i
- A channel is specified by (C_i, p_i, R_i)
- Objective:

$$\arg\max_{\text{all possible } (C_i, p_i, R_i)} \overline{X}\{(C_0, p_0, R_0), (C_1, p_1, R_1)\}$$

Mean TCP Throughput

TCP fluid model for static rate allocation

- p: packet error probability $p << 1$
- R: round-trip-time
- X: TCP throughput
- W: TCP window size

$$X(t) = \frac{W(t)}{R} \Rightarrow \Delta X(t) = \frac{\Delta W(t)}{R}$$

Fluid approximation of TCP throughput

- #packets transmitted in Δt:

  $= X(t)\Delta t$

- Congestion: if at least 1 packet is lost

$$p_c = \Pr\{\text{congestion in } \Delta\text{t}\} = 1 - (1-p)^{X(t)\Delta t}$$

$$\cong pX(t)\Delta t \qquad \text{for } p << 1$$

Fluid approximation of TCP throughput

- AIMD:
  - if no congestion: $\Delta W(t) = \Delta t/R$

    Note: window increases by 1 packet in R if there is no congestion
  - if there is congestion: $\Delta W(t) = - W(t)/2$

$$\Delta W(t) = (1 - p_c)\frac{\Delta t}{R} - p_c \frac{W(t)}{2}$$

Fluid approximation

Fluid approximation of TCP throughput

$$\Rightarrow \Delta X(t) = (1 - p_c)\frac{\Delta t}{R^2} - p_c \frac{X(t)}{2}$$

$$= (1 - pX(t)\Delta t)\frac{\Delta t}{R^2} - pX(t)\Delta t \frac{X(t)}{2}$$

$$\overset{\Delta t \to 0}{=} \frac{\Delta t}{R^2} - p\Delta t \frac{X(t)^2}{2}$$

$$\Rightarrow \lim_{\Delta t \to 0} \frac{\Delta X(t)}{\Delta t} = \frac{1}{R^2} - p\frac{X(t)^2}{2}$$

$$\Rightarrow \frac{d}{dt}X(t) = \frac{1}{R^2} - p\frac{X(t)^2}{2}$$

Stochastic differential equation

Illustration of TCP throughput: static rate allocation

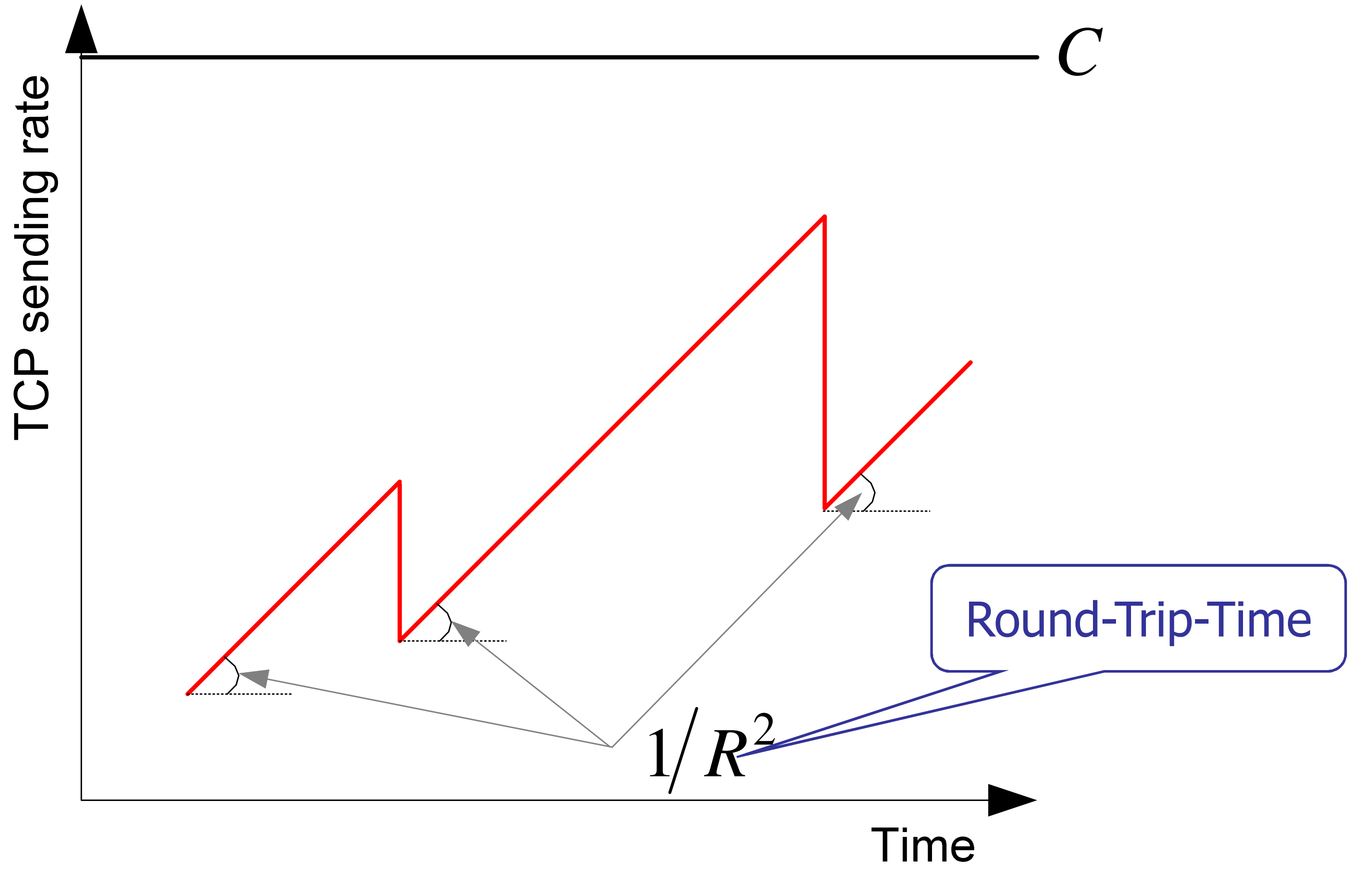

Illustration of TCP throughput: adaptive rate allocation

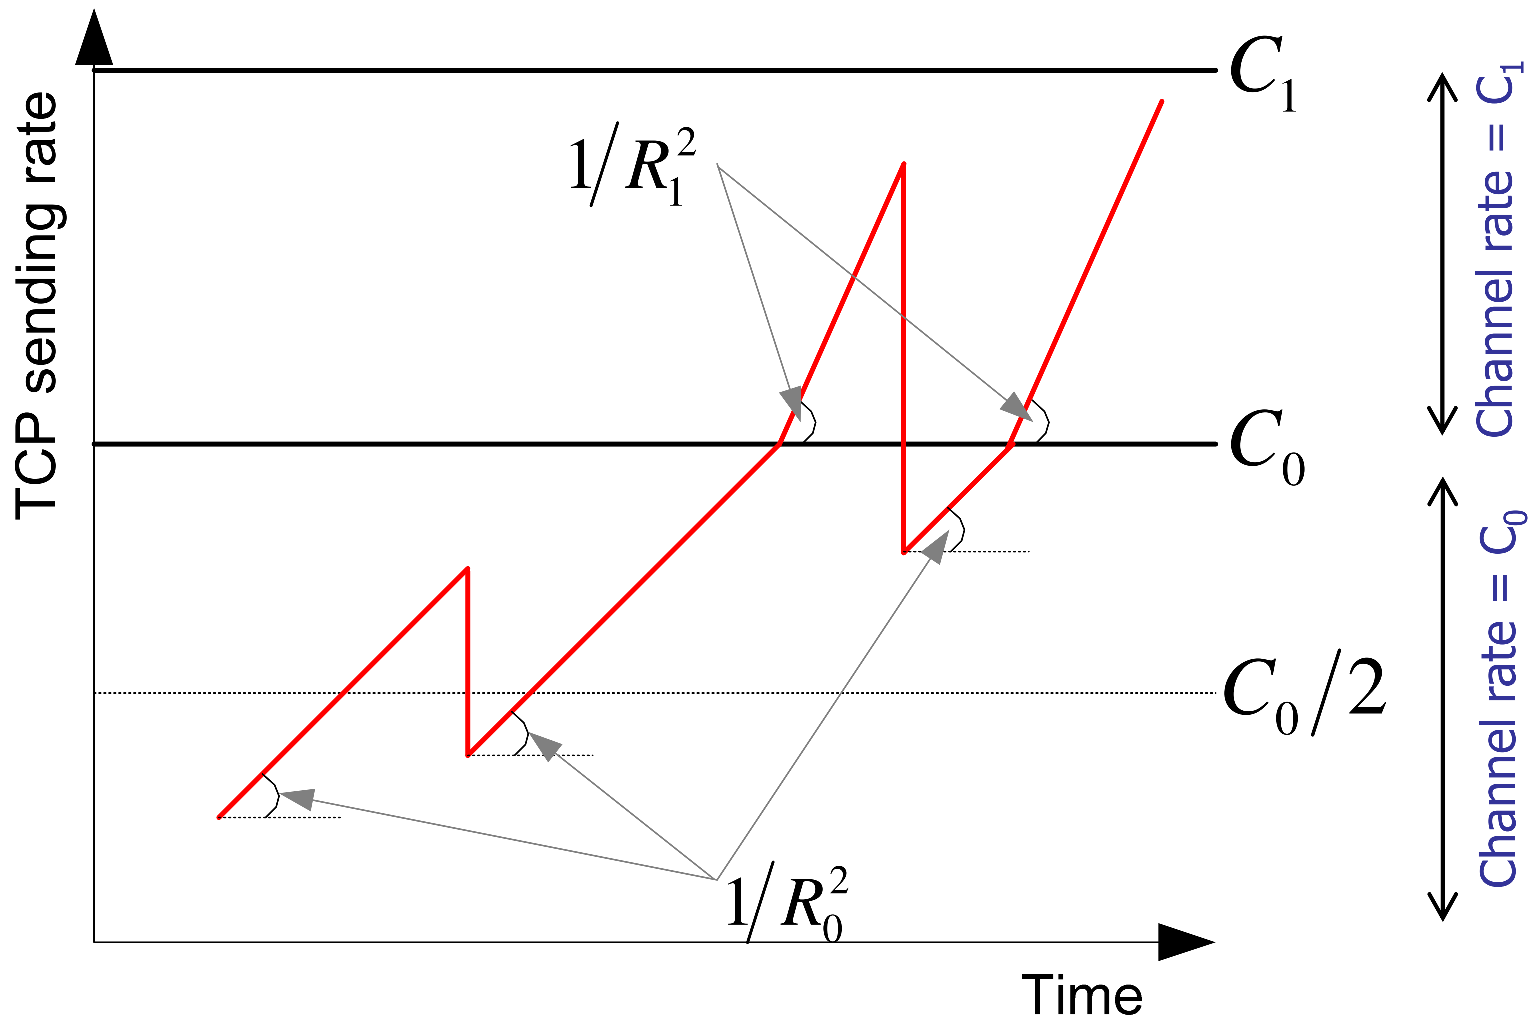

Analytical model

- Fluid model for TCP throughput
  - A set of ODEs for the probability distribution of TCP throughput
  - Explicitly incorporates impact of two different Round-Trip-Times, two different packet error probabilities and two channel rates
  - Explicitly differentiates between losses due to congestion and channel errors

Deriving balance equations

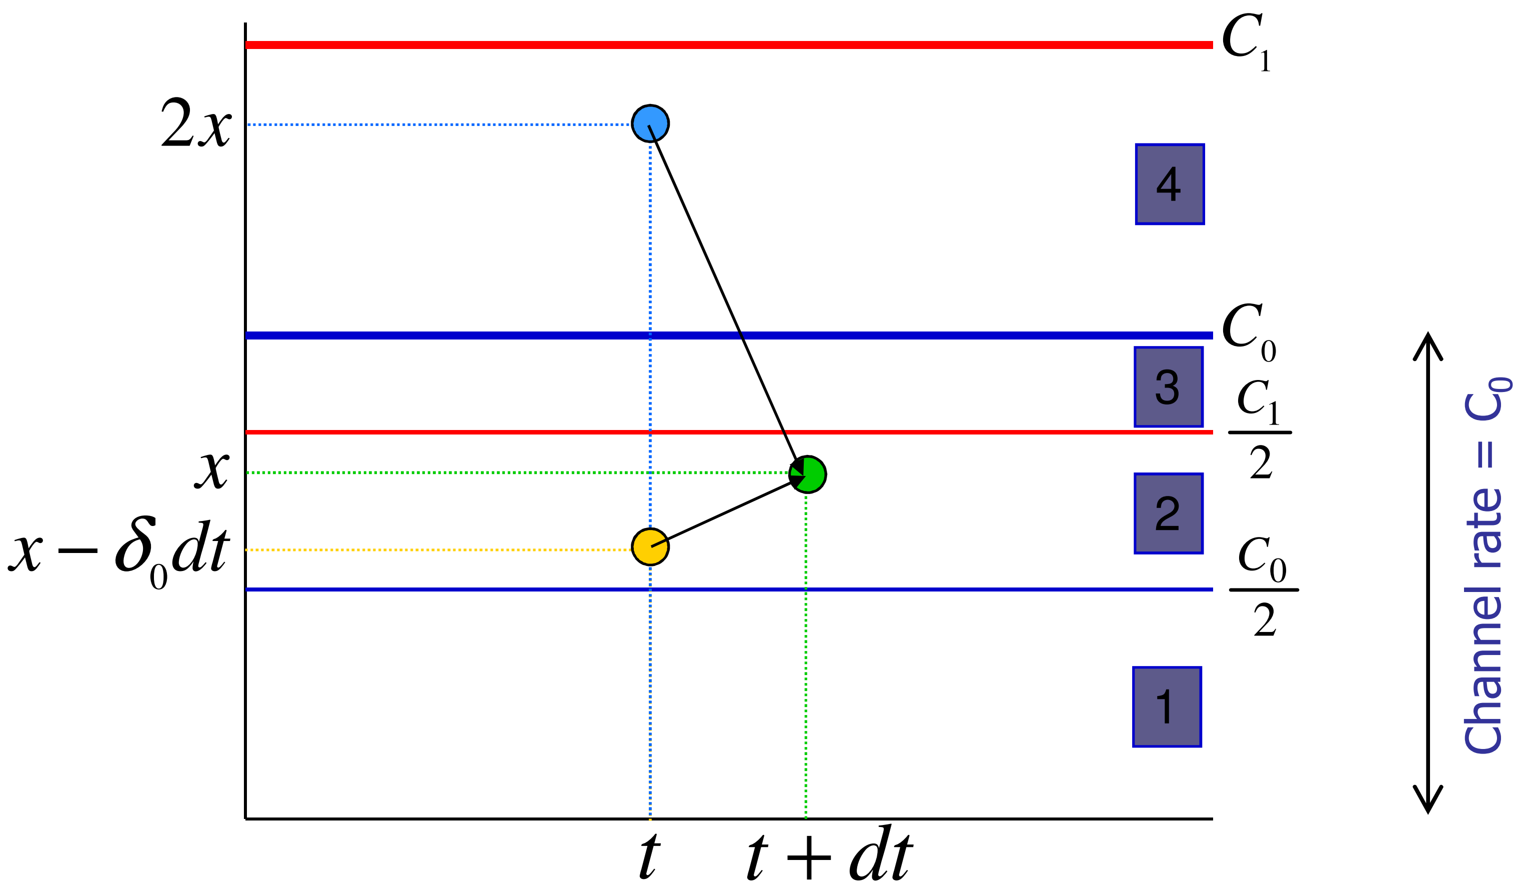

$$P_0(x, t+dt) = P_0(x - \delta_0 dt, t)(1 - \gamma_0 x dt) + P_1(2x, t)(\gamma_1 2x dt)$$

$$P_1(x, t+dt) = 0$$

Diffusion equation

$$\frac{P_0(x,t+dt)-P_0(x,t)}{dt}+\delta_0\frac{P_0(x,t)-P_0(x-\delta_0 dt,t)}{\delta_0 dt}=$$

$$-\gamma_0 xP_0(x-\delta_0 dt,t)+2\gamma_1 xP_1(2x,t)$$

let $dt \to 0$

$$\frac{\partial P_0(x,t)}{\partial t}+\delta_0\frac{\partial P_0(x,t)}{\partial x}=-\gamma_0 xP_0(x,t)+2\gamma_1 xP_1(2x,t)$$

let $t \to \infty$

$$\delta_0\frac{dP_0(x)}{dx}=-\gamma_0 xP_0(x)+2\gamma_1 xP_1(2x)$$

ODEs for throughput evolution

1. $0 < x < C_0 / 2$

$$\begin{cases} \delta_0 \dfrac{d}{dx} P_0(x) = -\gamma_0 x P_0(x) + 2\gamma_0 x P_0(2x) \\ P_1(x) = 0 \end{cases}$$

2. $\mathrm{C}_0 / 2 < x < C_1 / 2$

$$\begin{cases} \delta_0 \dfrac{d}{dx} P_0(x) = -\gamma_0 x P_0(x) + 2\gamma_1 x P_1(2x) \\ P_1(x) = 0 \end{cases}$$

3. $\mathrm{C}_1 / 2 < x < C_0$

$$\begin{cases} \delta_0 \dfrac{d}{dx} P_0(x) = -\gamma_0 x P_0(x) \\ P_1(x) = 0 \end{cases}$$

4. $\mathrm{C}_0 < x < C_1$

$$\begin{cases} P_0(x) = 0 \\ \delta_1 \dfrac{d}{dx} P_1(x) = -\gamma_1 x P_1(x) \end{cases}$$

Mean TCP throughput

- Solving ODEs using Mellin transform

$$\hat{f}(u) = \int_0^\infty f(x) x^{u-1} dx$$

Hint: $\overline{X} = \hat{f}(2)$

- TCP Throughput in the form of power series

$$\overline{X} = \frac{\Delta(2) + \sum_{k \geq 0} (\phi_{00})^k \Pi_k(2) \psi(2 + 2k)}{\Delta(1) + \sum_{k \geq 0} (\phi_{00})^k \Pi_k(1) \psi(1 + 2k)}$$

Model is accurate for low packet error probability

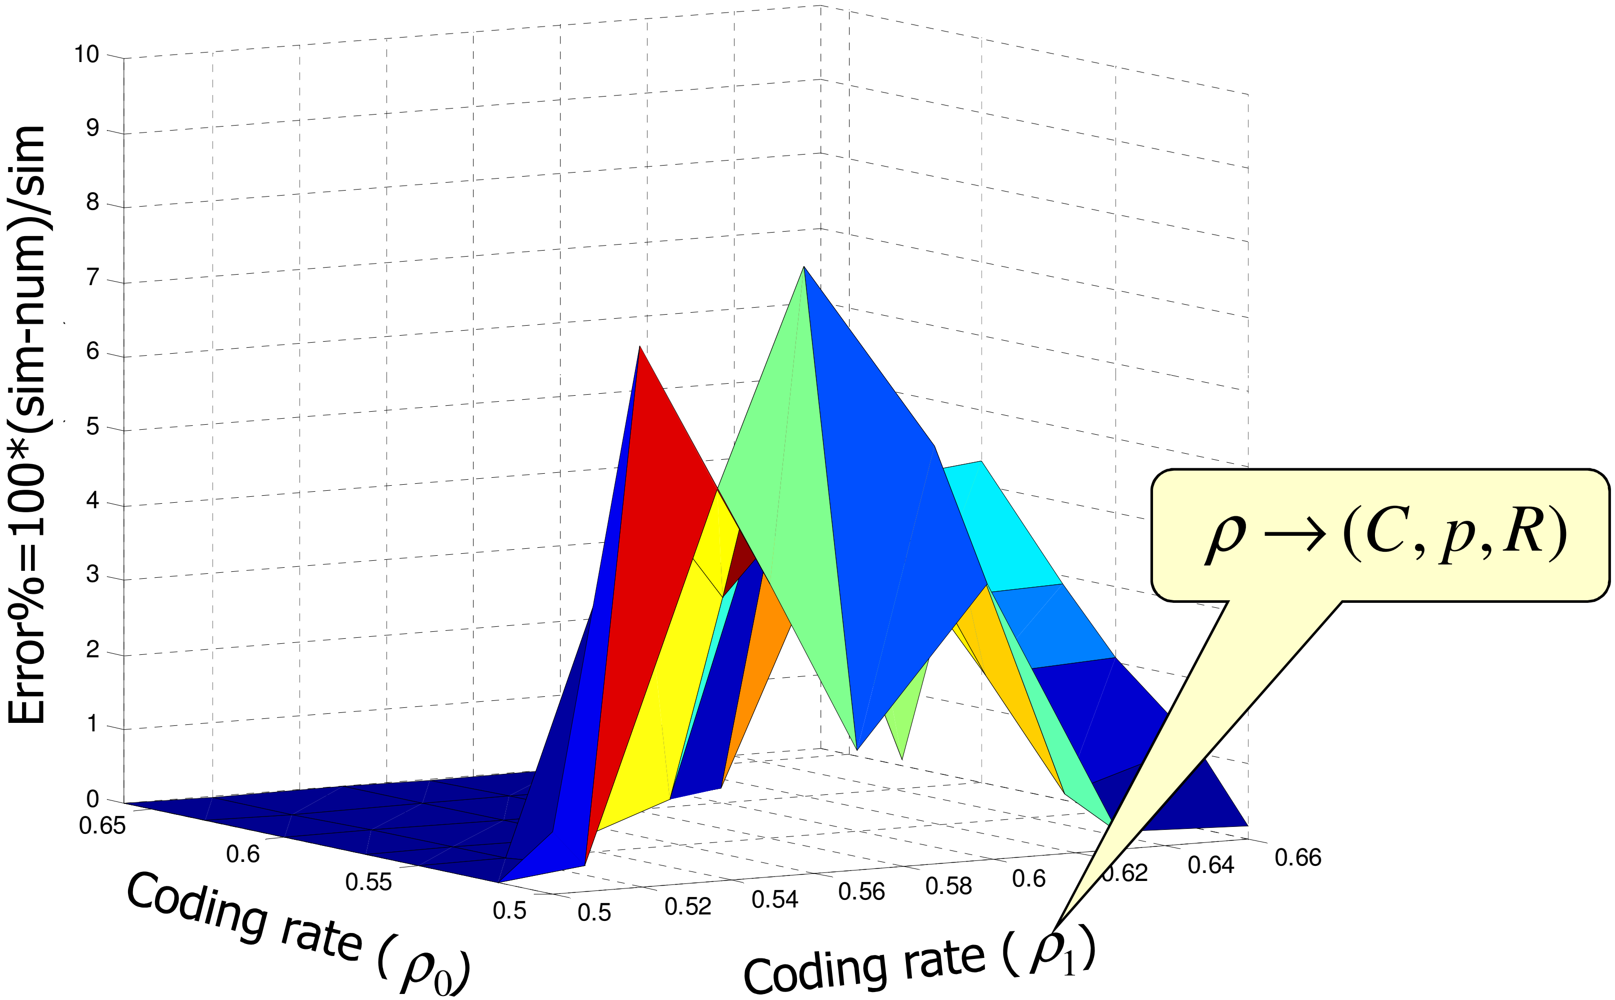

FEC comparison

- Coding trade-off:
  - Channel rate increases by increasing the coding rate
  - Packet error probability decreases by decreasing the coding rate
- Static coding picks one coding rate
- Adaptive coding picks two coding rates

Static versus Adaptive FEC

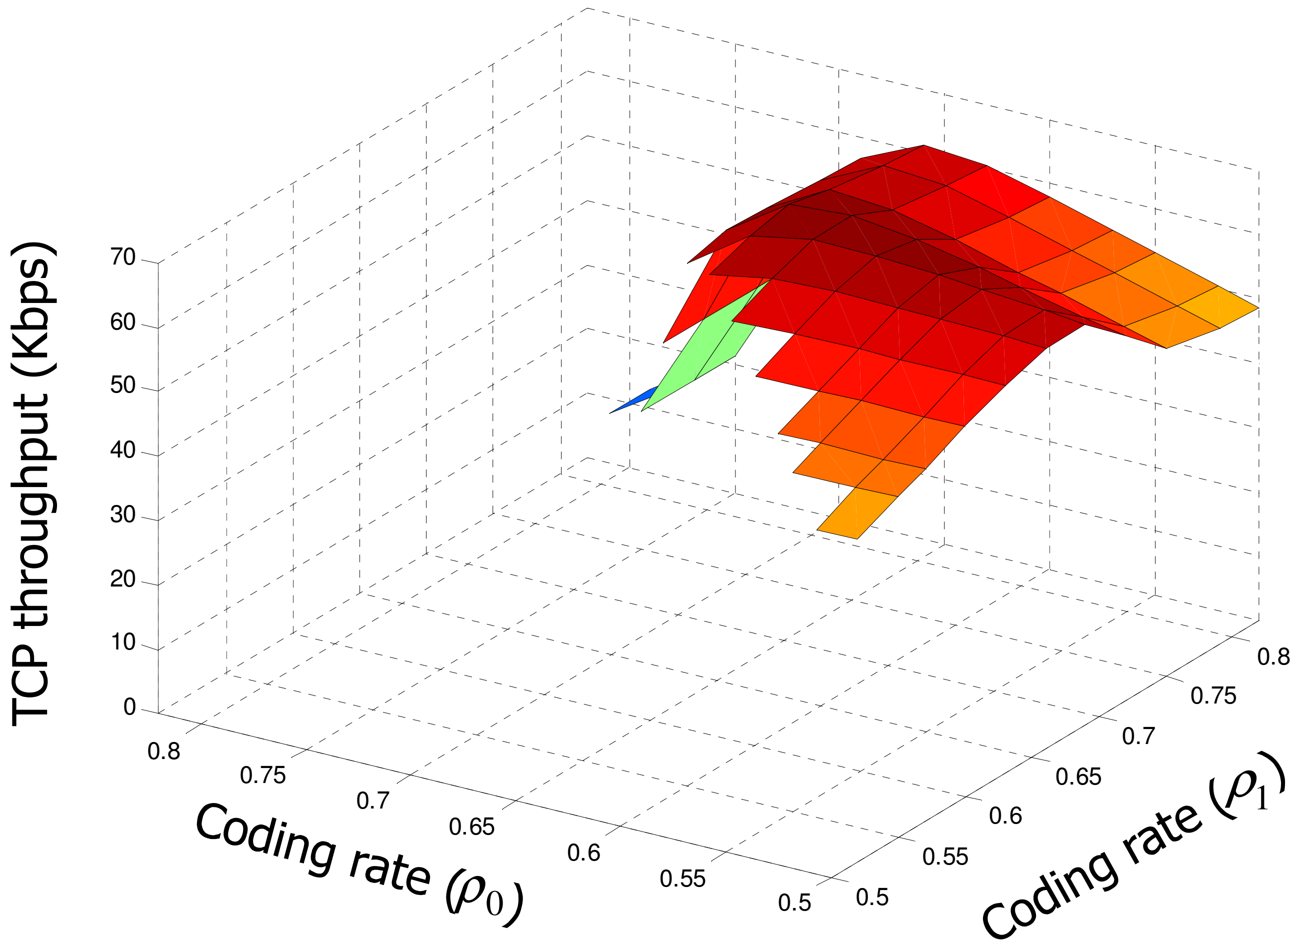

Channel = 128 Kbps

Max = 65.19 Kbps

ρ_0 = 0.63

ρ_1 = 0.72

Gain = 11%

Orthogonal Walsh codes

- Walsh codes are used in CDMA systems
- The channel rate can be increased by decreasing the code length
- This however increases error rate
  - Somewhat mitigated with higher energy

Static versus Adaptive Walsh codes

Energy Profile	Analysis		Simulations	
	Tput Gain	Energy Savings	Tput Gain	Energy Savings
E_1	10.8%	-3.5%	14.95%	-3.5%
E_2	15.8%	-4.8%	20.5%	4.2%

Higher throughput with less energy!

Multiple channel rates

- Simulation-based comparison using *ns2*

$$C(t) = \begin{cases} C_0 & \text{if } X(t) \le C_0 \\ C_1 & \text{if } C_0 < X(t) \le C_1 \\ C_2 & \text{otherwise} \end{cases}$$

BER	Two-rate Tput (Kbps)	Three-rate Tput (Kbps)	Gain
10^{-2}	62.9	64.3	2.2%
10^{-3}	81.4	82.6	1.5%
10^{-4}	91.7	91.7	0%

Outline

- TCP-aware resource allocation
- Single TCP session
- Multiple TCP sessions

Problem statement

- Only a few users can be simultaneously supported on high-rate channels
- Example: CDMA2000 1xRTT
  - Supports **30** users at 9.6 Kbps (called **fundamental channel**)
  - At most **2** users can be simultaneously allocated a 153.6 Kbps channel (called **supplemental channel**)
- How to allocate supplemental channels to competing TCP sessions?

Some notation

- N low-rate fundamental channels
  (i.e., N users in the system)

$$\text{fundamental channel} \rightarrow (C_0, p_0, R_0)$$

- K ($\leq N$) high-rate supplemental channels

$$\text{supplemental channel} \rightarrow (C_1, p_1, R_1)$$

Probabilistic preemptive scheduling

- If a session requests a supplemental channel and less than *K* supplemental channels occupied
  - Always assign a supplemental channel
- If all *K* supplemental channels are occupied
  - Randomly preempt a high-rate session with probability α
  - Deny requesting session with probability $1- \alpha$

Extended single session model

- Single session TCP throughput
  - q : acceptance probability
  - v : preemption rate

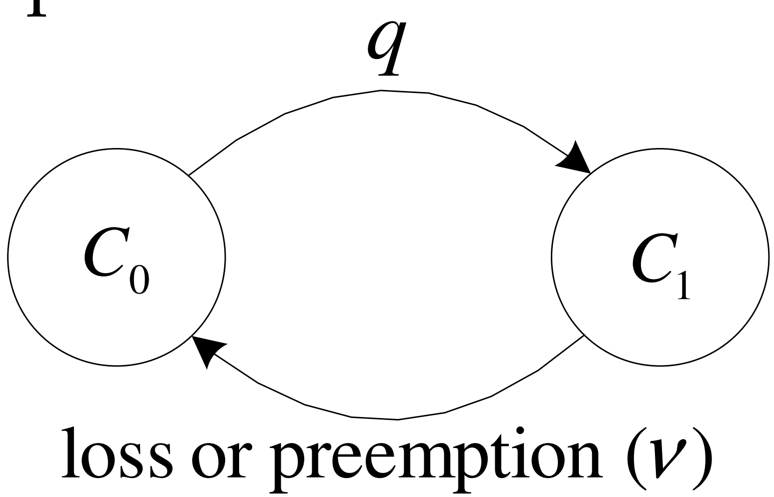

- q and v depend on supplemental channel occupancy

Inter-session interactions

- Supplemental channel occupancy
  - λ : supplemental channel request rate
  - μ : supplemental channel release rate

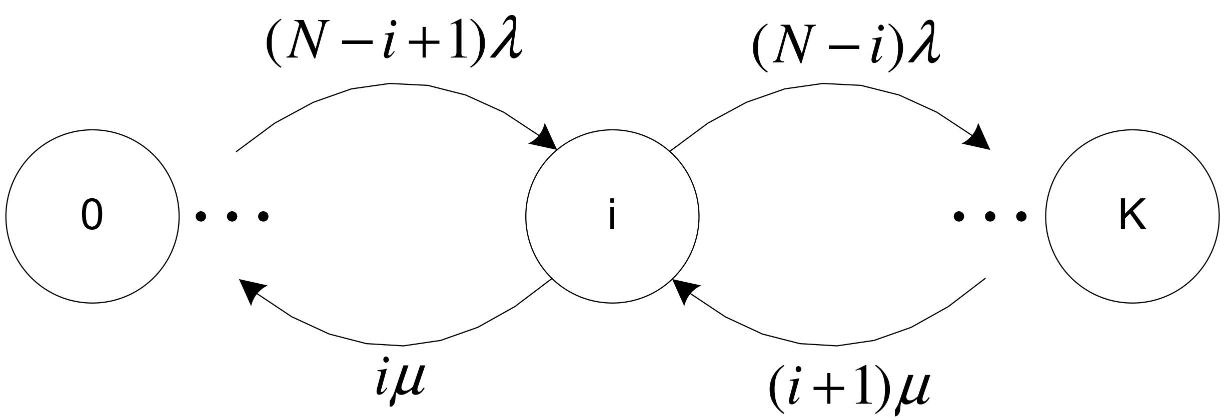

- λ and μ come from the single session model

Fixed point model

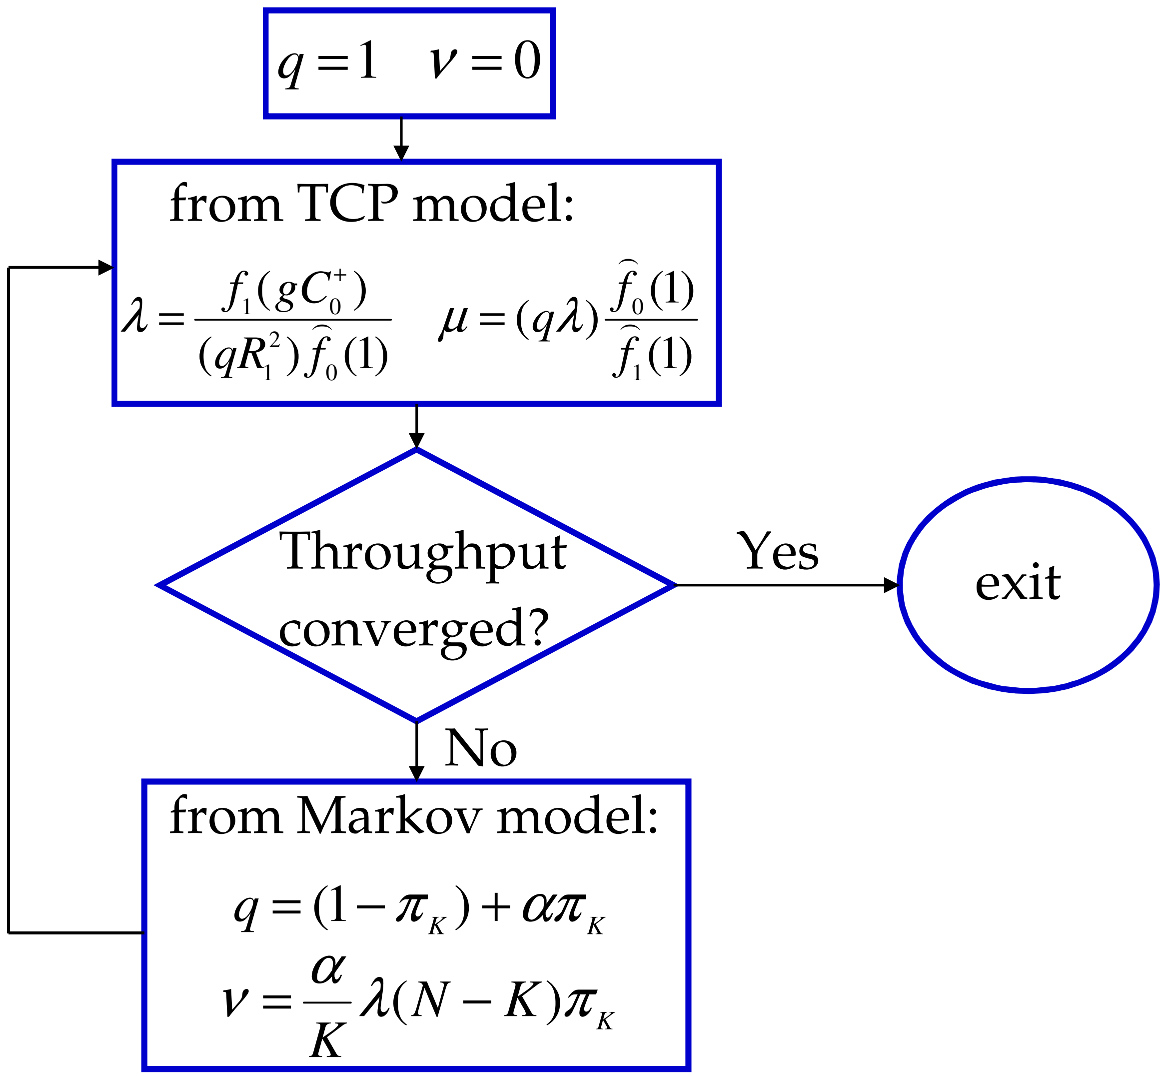

Fixed point model is accurate

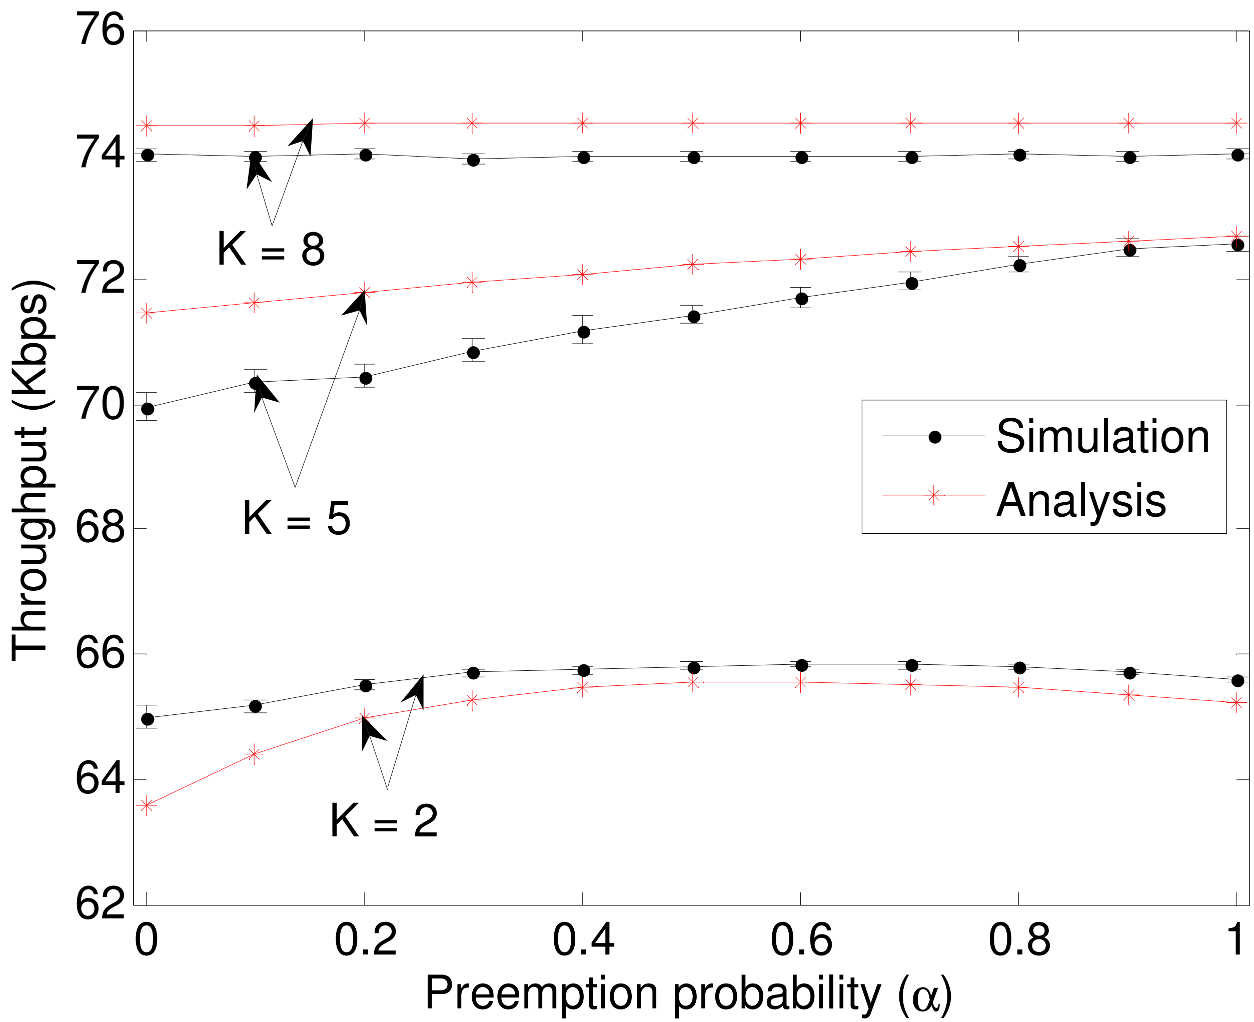

$N = 10$

$C_0 = 76.8\,\text{Kbps}$

$C_1 = 102.4\,\text{Kbps}$

Network dimensioning

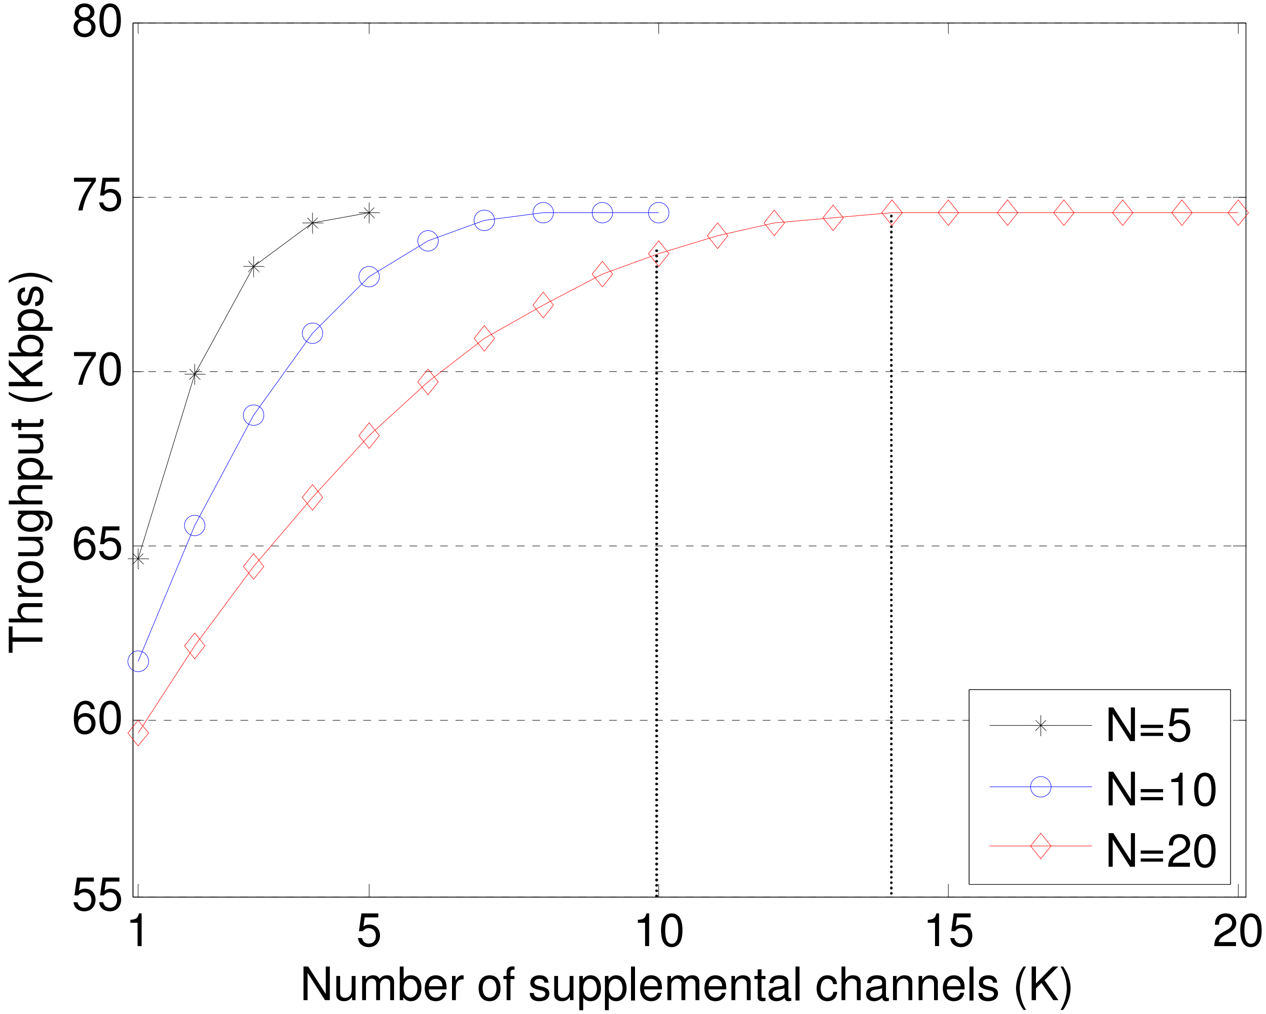

$C_0 = 76.8\,\text{Kbps}$

$C_1 = 102.4\,\text{Kbps}$

Summary

- TCP-aware rate allocation
- Analytical model to capture TCP dynamics with adaptive rate allocation
- Gains from 10% to 20% for a single TCP session compared to optimal static rate allocation
- Extensions to multiple TCP sessions