A Study of the Coexistence of Heterogeneous Flows in Data Network

Adrian Sai-wah Tam

swtam3@ie.cuhk.edu.hk

Department of Information Engineering

The Chinese University of Hong Kong

About

- *What is friendly to TCP actually?*
  - We are going to redefine “friendly”
- *Is TCP-friendly the only friendly way of transport?*
  - We will show something is also friendly, under a new definition

Data Networks

- Telephone network: Circuit switching
  - One circuit for one user, with bandwidth guarantee
- Computer network: Packet switching
  - One channel shared by many users, no bandwidth guarantee

Data Networks

- Telephone network: Circuit switching
  - One circuit for one user, with bandwidth guarantee
- Computer network: Packet switching
  - One channel shared by many users, no bandwidth guarantee

Do we have applications in data networks that prefer circuit switching-like services?

Data Flows

Dichotomy: Elastic vs Inelastic

Data Flows

Dichotomy: Elastic vs Inelastic

- Elastic flow can adapt to network conditions
  - It still functions if the network is slow, low bandwidth, high delay, . . .
  - Example: HTTP, FTP

Data Flows

Dichotomy: Elastic vs Inelastic

- Elastic flow can adapt to network conditions
  - It still functions if the network is slow, low bandwidth, high delay, . . .
  - Example: HTTP, FTP
- Inelastic flow cannot adapt
  - If bandwidth/delay is below the desired level, it is nearly useless
  - Example: VoIP, streaming

Problem Statement

- Elastic flows are adaptive to the available bandwidth
- Inelastic flows do not react to congestion, with constrain on min. data rate and delay

Problem Statement

- Elastic flows are adaptive to the available bandwidth
- Inelastic flows do not react to congestion, with constrain on min. data rate and delay

How should the elastic and inelastic flows coexist in the Internet?

Solution A: No control

- Use UDP for multimedia use
- Use RTP on top of UDP to keep track of the packet arrival time
- Problem: fairness with elastic flows is not guaranteed
  - A fear of congestion collapse is on the rise

Solution B: TCP Friendly

- IETF is working on this solution
- Requires inelastic flows to adapt, but allows them to adapt smoothly
- Inelastic flows *need* to be fair when using the network

Solution C: Admission Control

- Similar to circuit switching approach
- Multimedia stay inelastic
  - Do not insist equal sharing of bandwidth
- Before you use, make sure the network can support you!

Which one is better?

Compare the merit of different controls.

- Evaluation 1: Utility based
  - compare for superiority among controls
- Evaluation 2: Stochastic Differential Equations
  - for blocking probability formula

Evaluation 1: Utility

- The network is serving many flows
- Each flow has some utility function
- Different controls $\Rightarrow$ Different bw. allocation
- The network's utility = Sum of the flows' utility
- Add up the utility of different flows — the better traffic control should yield higher total utility

Credit: Scott Shenker (1995)

Utility

Utility

- Elastic: $U(x) = \log(x)$
  - Following Frank Kelly (proportional fairness, paper in 1997)
  - A concave function and monotonically increasing

Utility

- Elastic: $U(x) = \log(x)$
  - Following Frank Kelly (proportional fairness, paper in 1997)
  - A concave function and monotonically increasing
- Inelastic: $U(x) = \sin^k(x)$
  - Steep decay in utility if the allocation is lower than desired rate
  - Over-allocation yields no value
  - This is known as a sigmoidal function

Utility

$$U_E(x) = \log(1+x); \quad U_I(x) = \sin^{50}(\frac{\pi}{2}\frac{\min(x,\alpha)}{\alpha})$$

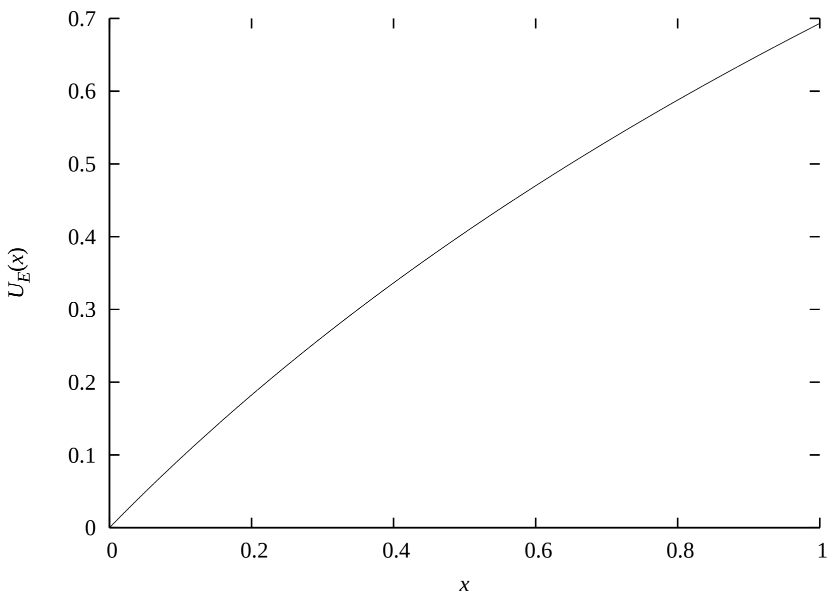

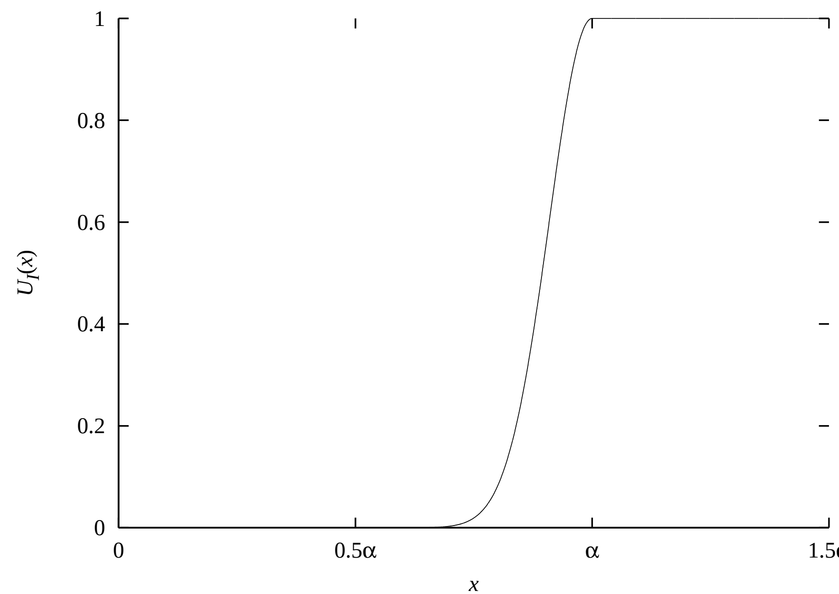

Model for Evaluation

Approximation by fluid model

- Network conditions are sensed by the flows instantly and the controls take effect immediately
- Single bottleneck link network

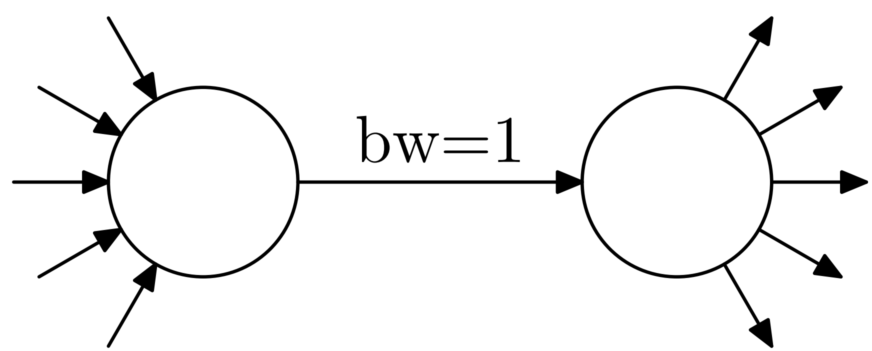

Markov Chain Model

- Applied with the fluid assumption
- Network as a stochastic models of flows
- State space: no. of elastic and inelastic flows, (n, m)
- Stochastic arrival, but the service rate depends on the flow controls

Flow Controls for Inelastic

1. No Control — multimedia over UDP
2. Congestion Control — TCP-friendly
3. Admission control in an “aggressive” way
4. Admission control in a “conservative” way

NC: No Control

Each inelastic flow uses α of bandwidth

- If there are n elastic and m inelastic flows,

	No.	Each	Total
Inelastic	m	α	$m\alpha$
Elastic	n	$\frac{1-m\alpha}{n}$	$1-m\alpha$
Total			1

- If $m\alpha > 1$, elastic flows get nothing and each inelastic flow has α/m

CC: Fair Share Congestion Control

- If there are n elastic and m inelastic flows,

	No.	Each	Total
Inelastic	m	$\frac{1}{m+n}$	$\frac{m}{m+n}$
Elastic	n	$\frac{1}{m+n}$	$\frac{n}{m+n}$
Total			1

- If $\frac{1}{m+n} > \alpha$, each inelastic flow will use only α. Then each elastic flow will have

$$\frac{1 - m\alpha}{n} > \frac{1}{m+n}$$

AC-A: Aggressive Admission Ctrl

- Assume an inelastic flow always take α of bandwidth
- Guarantee each elastic flow gets ϵ or more when admitting inelastic flows

	No.	Each	Total
Inelastic	m	α	$m\alpha$
Elastic	n	$\frac{1-m\alpha}{n}$	$1-m\alpha$
Total			1

- Admission only if $n\epsilon + (m+1)\alpha \leq 1$
- Typically $0 < \epsilon \ll \alpha$

AC-C: Conservative Admission Ctrl

- $\epsilon = \alpha$
- Admission only if $(n + m + 1)\alpha \leq 1$
- We call this the “TCP-friendly admission control”

Markov Chain: NC

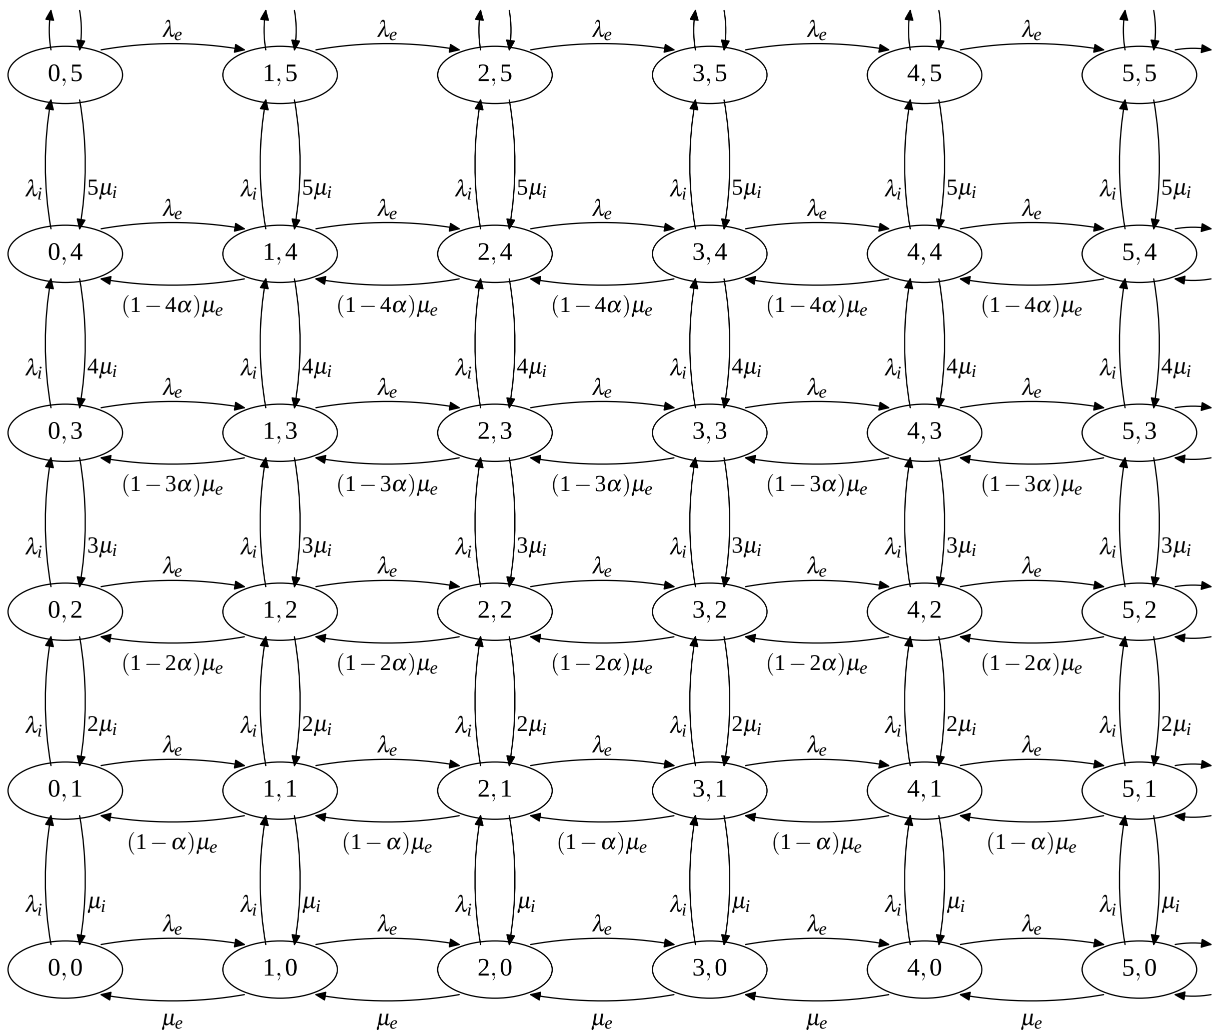

Markov Chain: CC

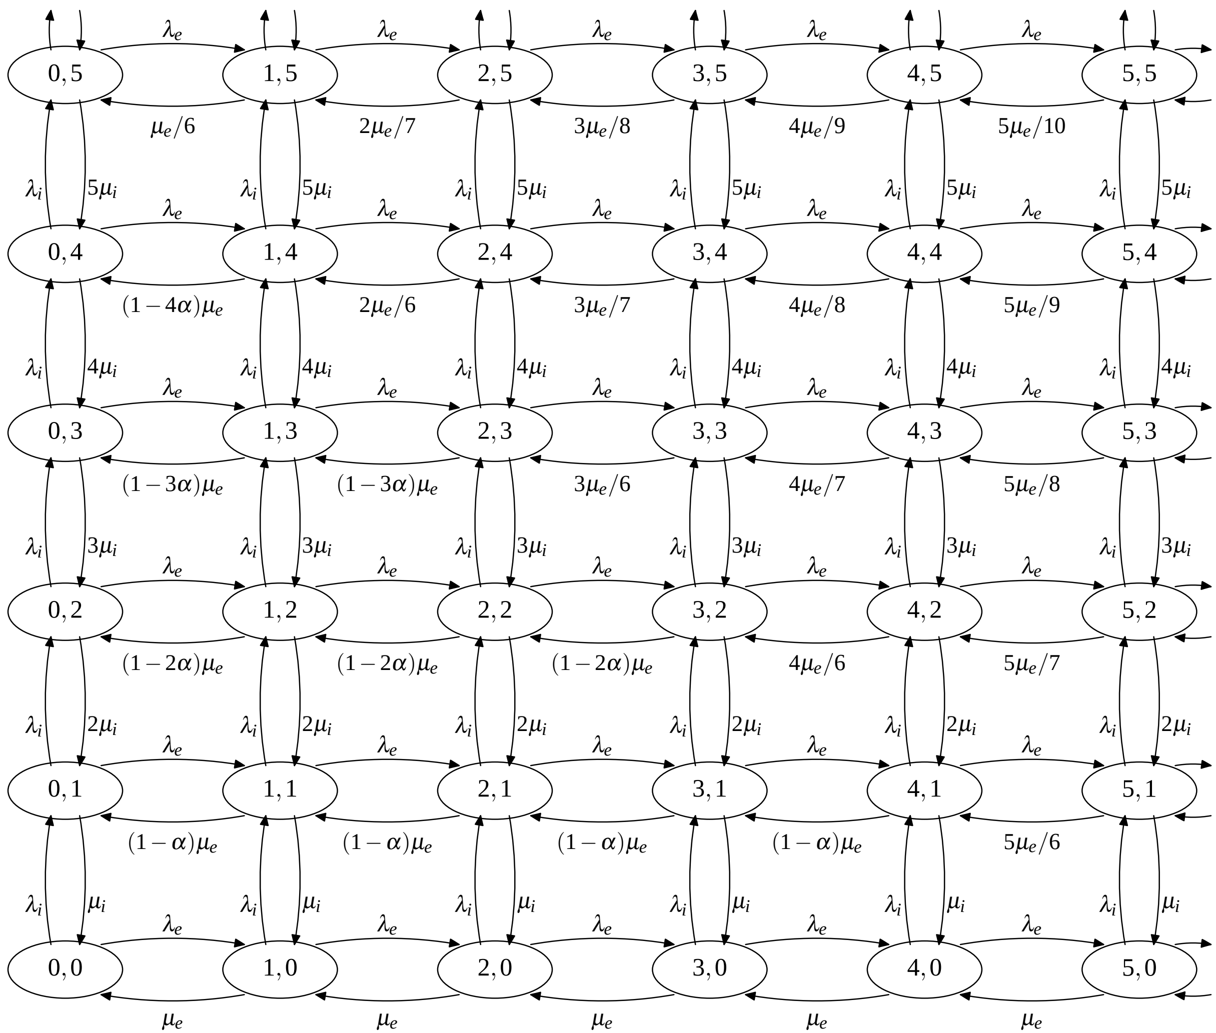

Markov Chain: AC-A

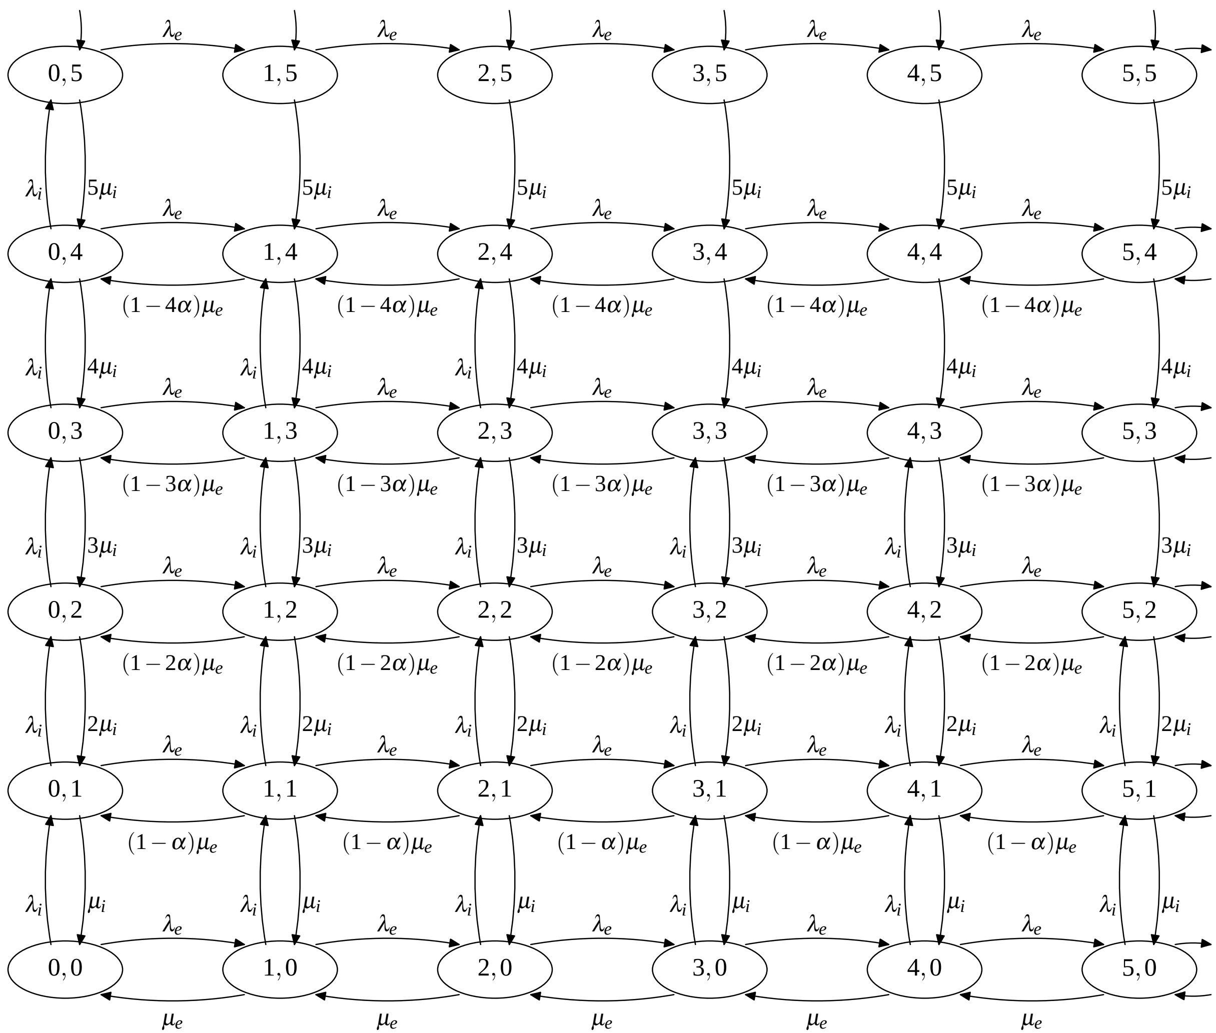

Markov Chain: AC-C

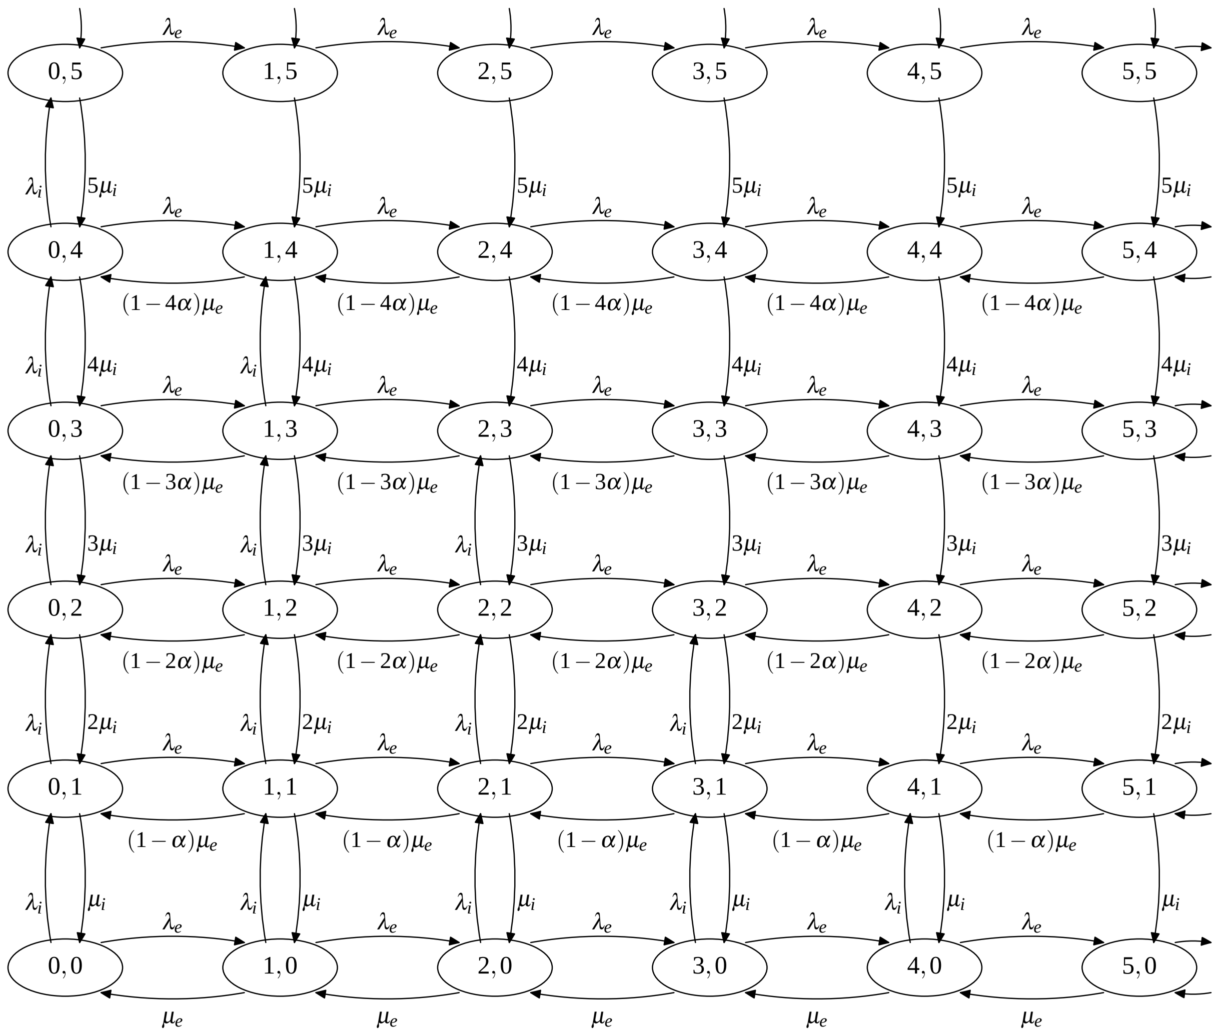

Markov Chain: Summary

Transition rates of Markov Chain:

		$(n,m) \to (n,m+1)$	$(n,m) \to (n+1,m)$	$(n,m) \to (n,m-1)$	$(n,m) \to (n-1,m)$
NC	$m\alpha \leq 1$	λ_i	λ_e	$m\mu_i$	$(1-m\alpha)\mu_e$
	$m\alpha > 1$	λ_i	λ_e	$m\mu_i$	0
CC	$(n+m)\alpha \leq 1$	λ_i	λ_e	$m\mu_i$	$(1-m\alpha)\mu_e$
	$(n+m)\alpha > 1$	λ_i	λ_e	$m\mu_i$	$\frac{n}{n+m}\mu_e$
AC-A	$n\epsilon + (m+1)\alpha \leq 1$	λ_i	λ_e	$m\mu_i$	$(1-m\alpha)\mu_e$
	$n\epsilon + (m+1)\alpha > 1$	0	λ_e	$m\mu_i$	$\max(0,(1-m\alpha)\mu_e)$
AC-C	$(n+m+1)\alpha \leq 1$	λ_i	λ_e	$m\mu_i$	$(1-m\alpha)\mu_e$
	$(n+m+1)\alpha > 1$	0	λ_e	$m\mu_i$	$\max(0,(1-m\alpha)\mu_e)$

Define: $\rho = \rho_e + \alpha\rho_i; \quad \rho_e = \lambda_e/\mu_e; \quad \rho_i = \lambda_i/\mu_i$

Simulation

- Simulating the Markov chain
- Result: AC-C > AC-A, CC > NC

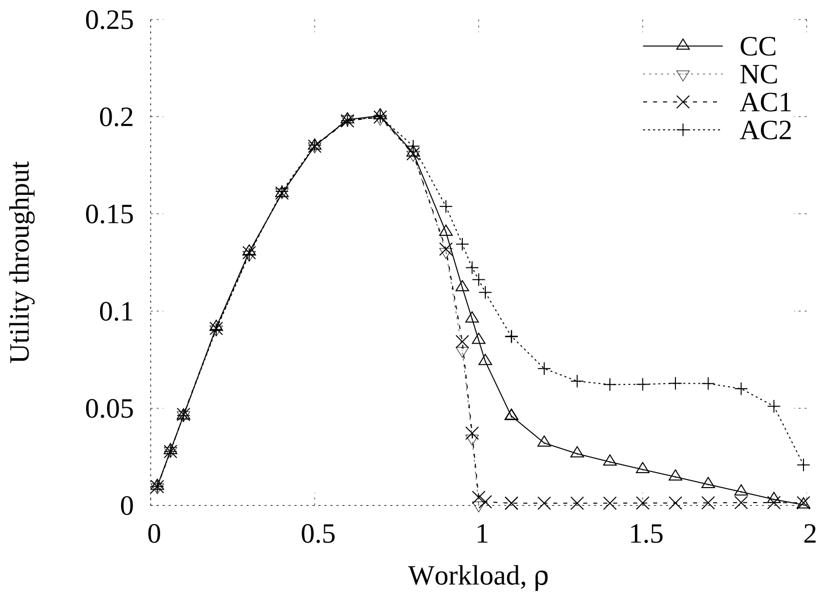

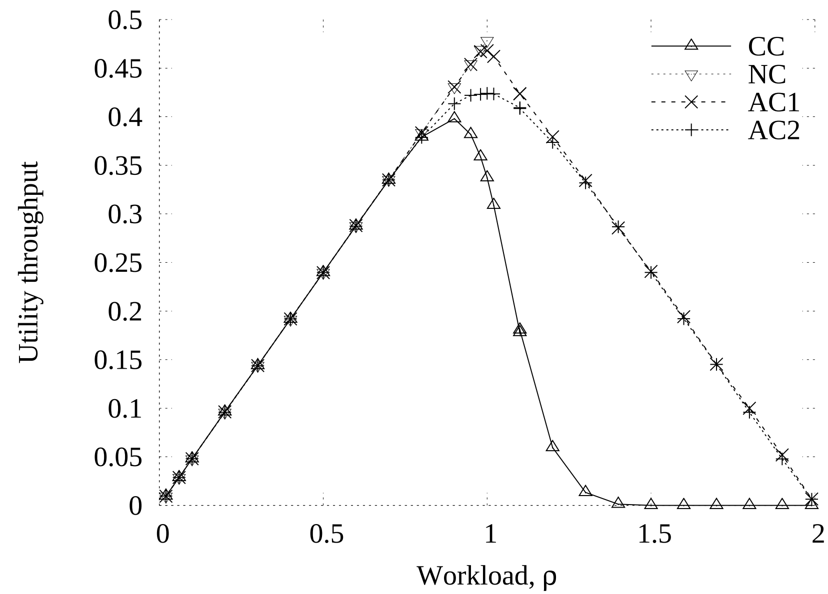

*The above is just one of the many cases, showing equal offered load from elastic and inelastic flows

Different evaluations

- Evaluation 1: Utility based
  - compare for superiority among controls
- Evaluation 2: Stochastic Differential Equations
  - for blocking probability formula

Evaluation 2: Blocking Probability

- We have shown that using admission control (esp. the conservative type) can make both elastic and inelastic flows happier
- Comparing different admission controls do not need utility functions
- The performance of admission control is determined solely by the blocking probability

Evaluation 2: Blocking Probability

- Consider only the admission control models
- Make use of Poisson Counter Driven Stochastic Differential Equation
- Defining
  - τ to be the total number of bytes yet to be transferred by all the existing flows, and
  - N_i, N_e to be Poisson counters marking the arrival of inelastic and elastic flows

Evaluation 2: Blocking Probability

Equation:

$$d\tau = -\mathbf{1}(\tau > 0)dt + S_e dN_e + I(n,m)S_i dN_i$$

evaluates to:

$$1 - P_{\text{block}} = \frac{\Pr[\tau > 0] - \rho_e}{\alpha \rho_i}$$

Evaluation 2: Blocking Probability

$$1 - P_{\text{block}} = \frac{\Pr[\tau > 0] - \rho_e}{\alpha \rho_i}$$

- $\Pr[\tau > 0]$ is the probability that the network is not idle
- Intuitively, we can approximate by:

$$\Pr[\tau > 0] \approx \min(\rho, 1)$$
$$\rho = \rho_e + \alpha \rho_i$$

Evaluation 2: Blocking Probability

$$1 - P_{\text{block}} = \frac{\Pr[\tau > 0] - \rho_e}{\alpha \rho_i}$$

- $\Pr[\tau > 0]$ is the probability that the network is not idle
- Intuitively, we can approximate by:

$$\Pr[\tau > 0] \approx \min(\rho, 1)$$
$$\rho = \rho_e + \alpha \rho_i$$
$$\therefore \quad 1 - P_{\text{block}} \approx \frac{\min(\rho, 1) - \rho_e}{\alpha \rho_i}$$

Selfish is not good

$$1 - P_{\text{block}} \approx \frac{\min(\rho, 1) - \rho_e}{\alpha \rho_i}$$

- No ϵ in the equation!
- Whichever AC models, the same P_{block}
  - Being aggressive and selfish does not improve the performance
  - In terms of social welfare, AC-C should be chosen instead of AC-A

Selfish is not good

$$1 - P_{\text{block}} \approx \frac{\min(\rho, 1) - \rho_e}{\alpha \rho_i}$$

- No ϵ in the equation!
- Whichever AC models, the same P_{block}
  - Being aggressive and selfish does not improve the performance
  - In terms of social welfare, AC-C should be chosen instead of AC-A

(pseudo-Nash equilibrium)

Selfish is not good

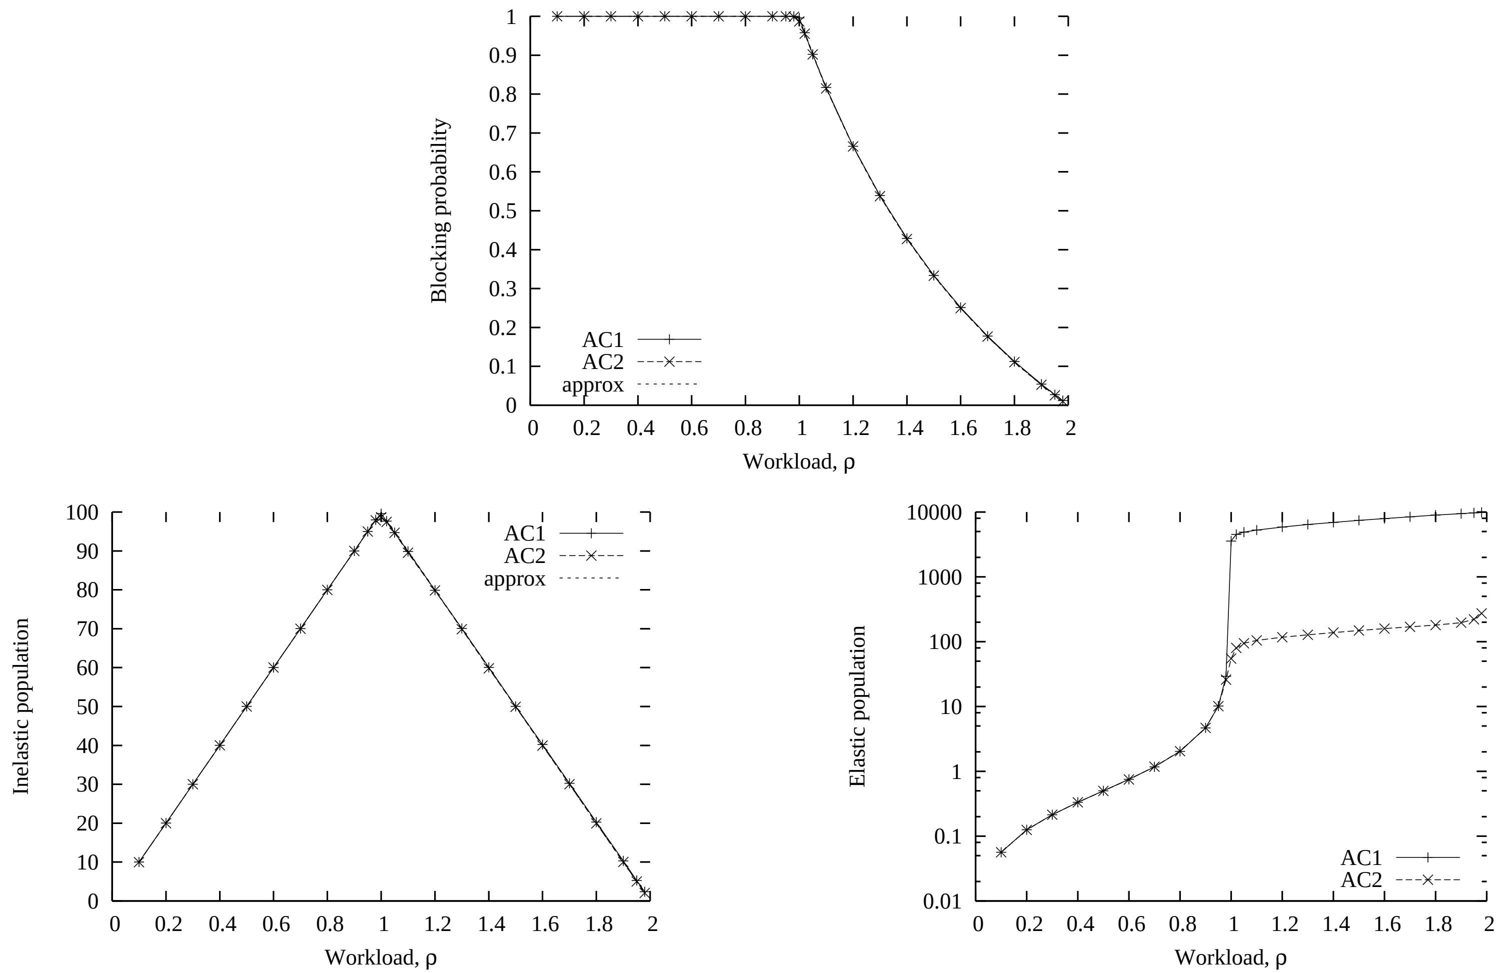

Conclusion

Conclusion

- We argue for multimedia flows it is better to use admission control then TCP-friendly congestion control

Conclusion

- We argue for multimedia flows it is better to use admission control then TCP-friendly congestion control
- To make admission control TCP-friendly is easy:
  - Work as if you are normal TCP first
  - If (attained the rate you want)
    continue with your desired rate
    otherwise
    quit

Conclusion

- It does not pay to be too aggressive! You won’t get any advantage in the long run

References

- *Network fairness for heterogeneous applications*
  Dah Ming Chiu and Adrian Tam
  In Proc. SIGCOMM Asia Workshop 2005
- *A case for TCP-friendly admission control*
  Adrian Tam, Dah Ming Chiu, John Lui, Y. C. Tay
  Submitted to INFOCOM 2006
- *Redefining fairness in the study of TCP-friendly traffic controls*
  Dah Ming Chiu and Adrian Tam
  Submitted to IEEE Trans. on Networking