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

Data Network

- Telephone network: Circuit switching
  - At any instant of time, the circuit is used by only one user, with bandwidth guarantee
- Computer network: Packet switching
  - At any instant of time, data network is shared by a number of users, but no guarantee on bandwidth

Traffic

Dichotomy: Elastic vs Inelastic

- Elastic traffic can adapt to network conditions
  - It still functions if the network is slow, low bandwidth, high delay, . . .
- Inelastic traffic cannot adapt
  - If bandwidth/delay is below the desired level, it is nearly useless

Traffic: Examples

- Traditional TCP applications are elastic: HTTP, FTP, etc.
- Multimedia application are generally inelastic
  - e.g. VoIP, streaming, etc.

Problem Statement

How should the elastic and inelastic traffic coexist in data networks?

Why is there a Problem?

- Elastic traffic adapts via congestion control algorithms
  - e.g. TCP flows increase/decrease their rates depending on congesion feedback
  - That’s why they are *adaptive*
- Inelastic traffic cannot do congestion control
  - They need *specific* data rate & delay
  - example: UDP flows
- Will such inelastic traffic cause unfair bandwidth allocation?

Existing Solution: 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 traffic is not guaranteed

Proposal 1: TCP Friendly

- IETF is working on this problem
- The current solution requires inelastic traffic to adapt
- Inelastic flows *need* to be fair when they use the network

Proposal 2: Admission Control

- Similar to circuit switching approach
- Multimedia stay inelastic
- Before you use, make sure the network can support you!
- Frank Kelly, Laurent Massoulié, Peter Key, Alan Bain, James Roberts, Thomas Bonald, Gunnar Karlsson, . . .

Which one is the best?

- My research topic
- Compare different traffic controls based on modeling and analysis

Different approaches

- Evaluation 1: Utility based
- Evaluation 2: Stochastic Differential Equations

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 at different controls, the best control should give a higher value

Credit: Scott Shenker (invited paper in 1995)

Utility

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

Utility

$$U_E(x) = \log(1+x); \quad U_I(x) = \begin{cases} \sin^{50}\left(\frac{\pi}{2}\frac{x}{\alpha}\right) & \text{if } 0 \le x \le \alpha \\ 1 & \text{if } x > \alpha \end{cases}$$

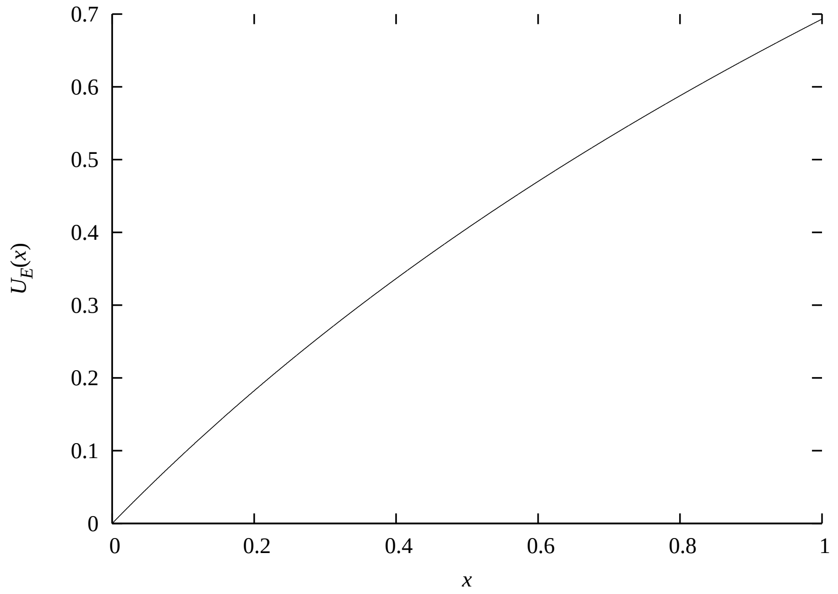

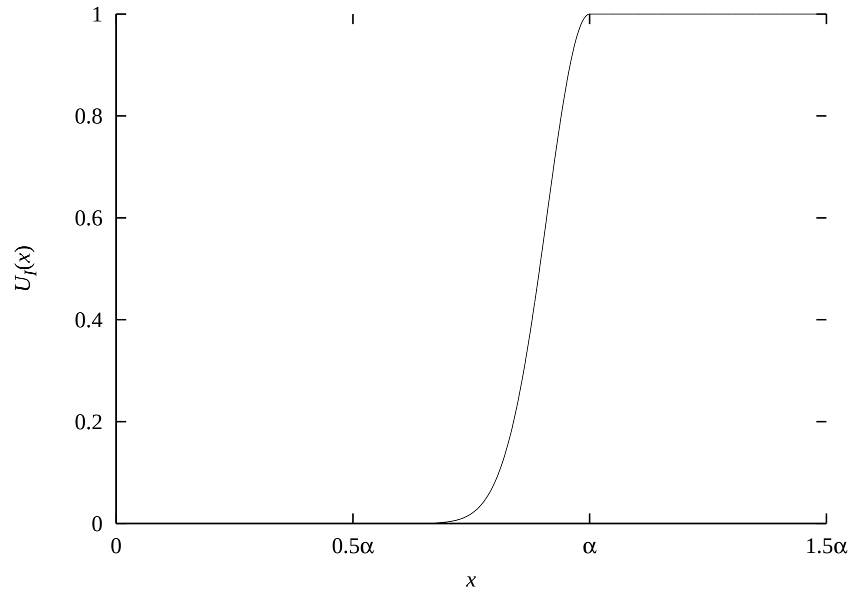

Model for Evaluation

Approximation by fluid model

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

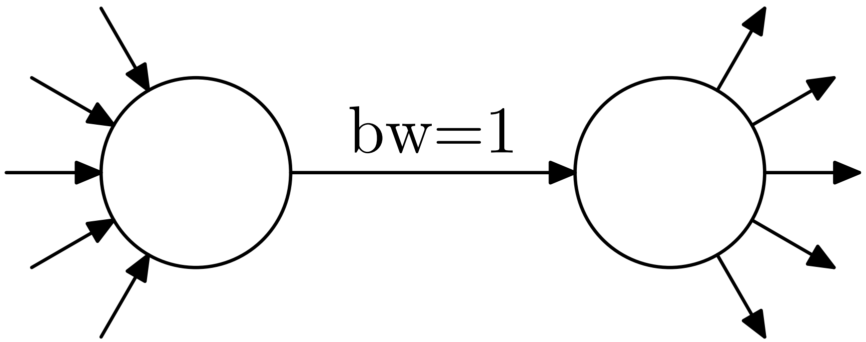

Markov Chain Model

- Network as a stochastic models of flows
- State space: no. of elastic and inelastic flows
- Stochastic arrival, but the service rate depends on the traffic controls

Traffic 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 traffic 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

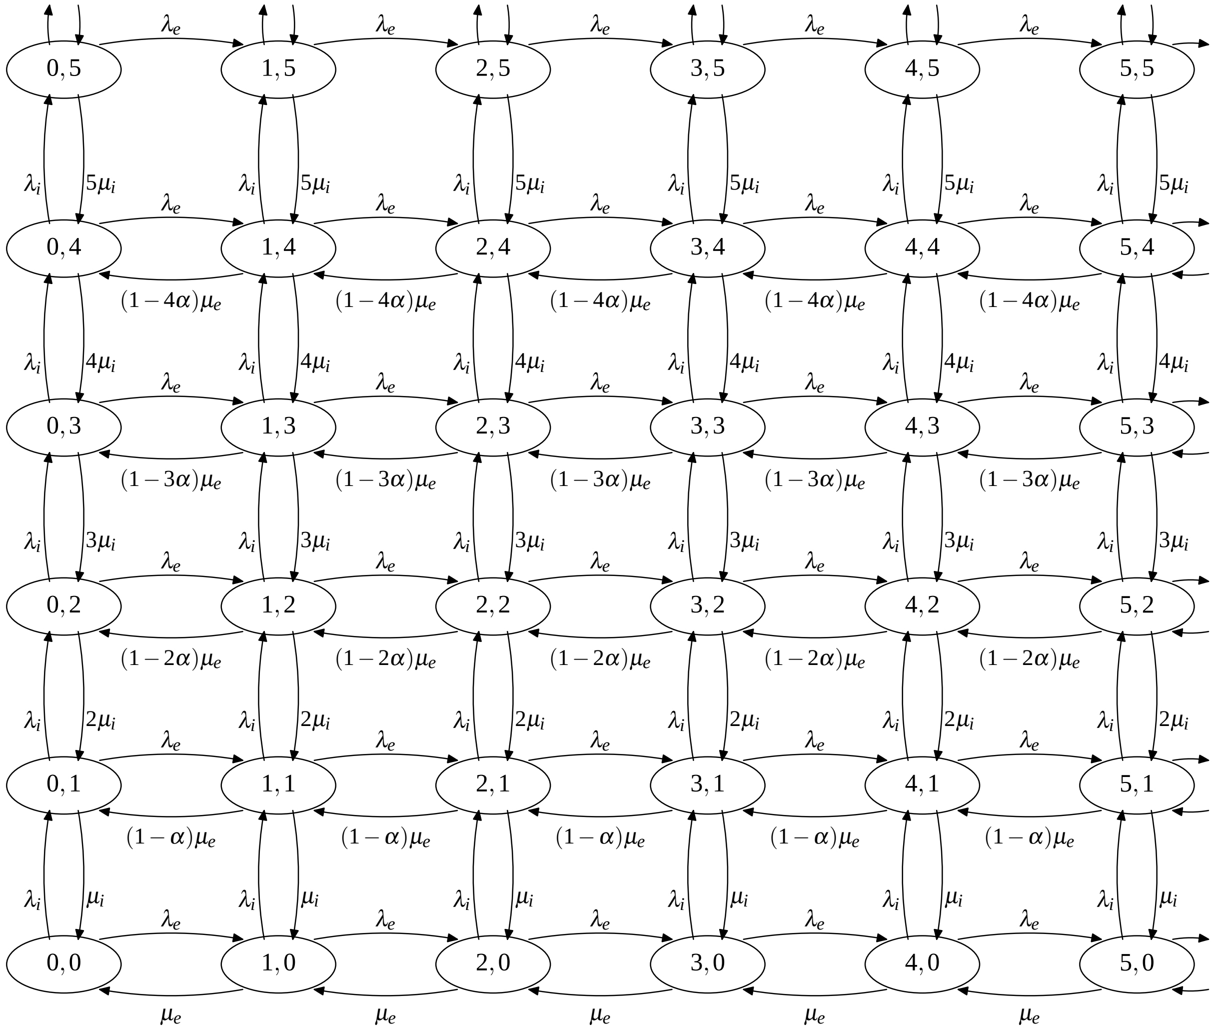

Markov Chain: CC

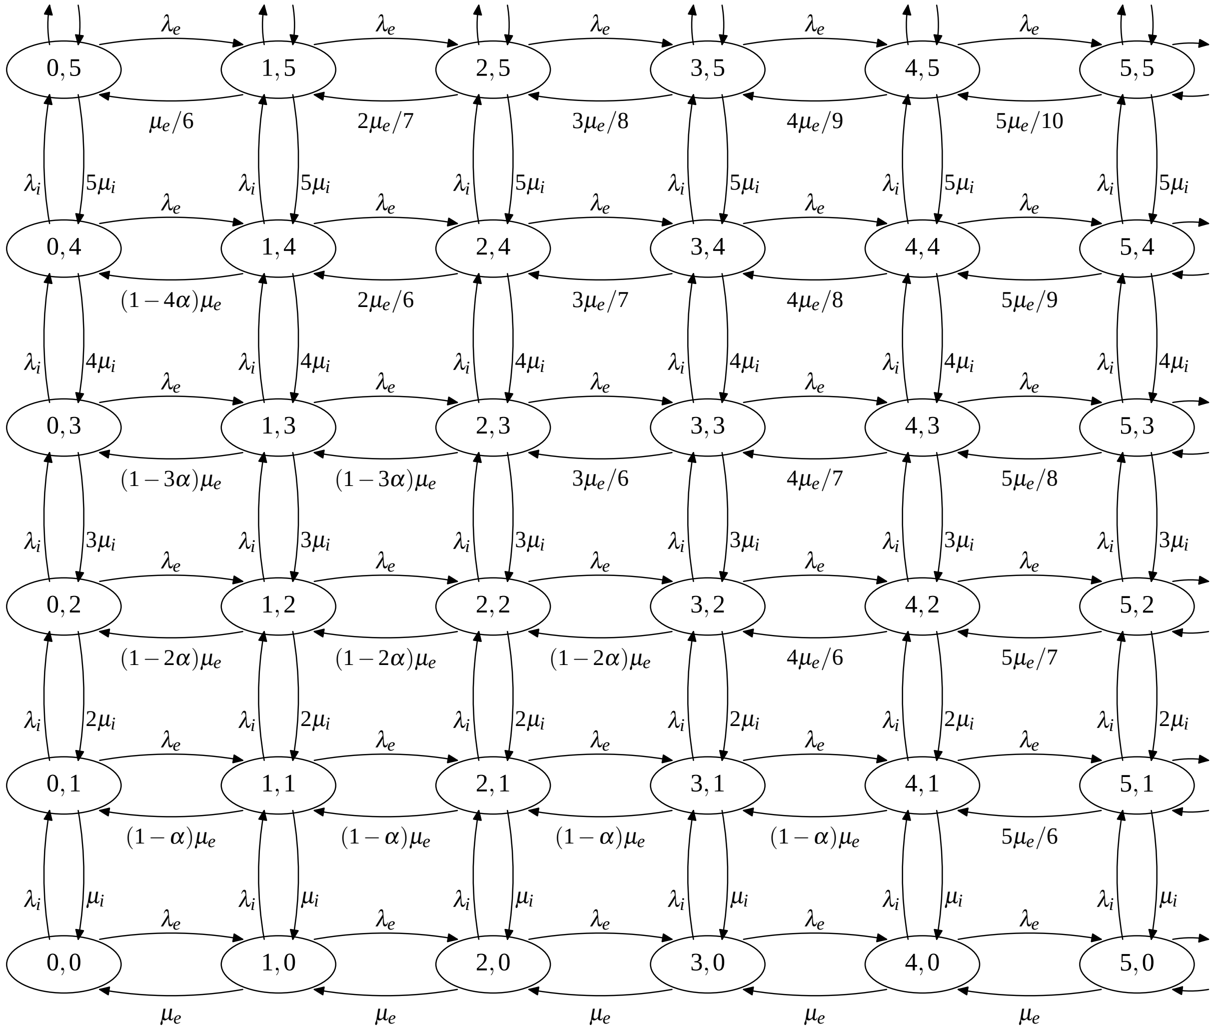

Markov Chain: AC-A

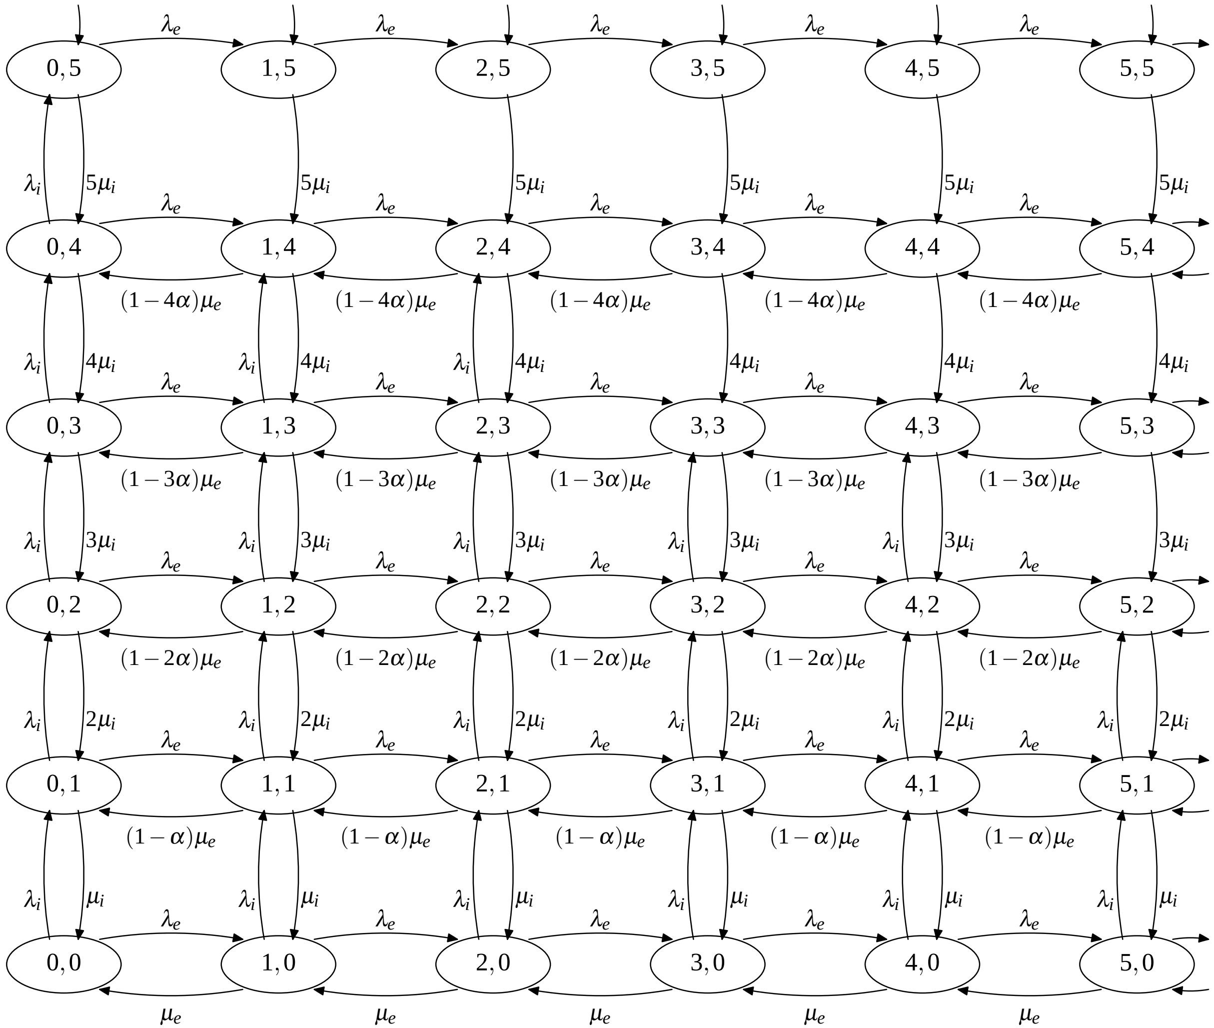

Markov Chain: AC-C

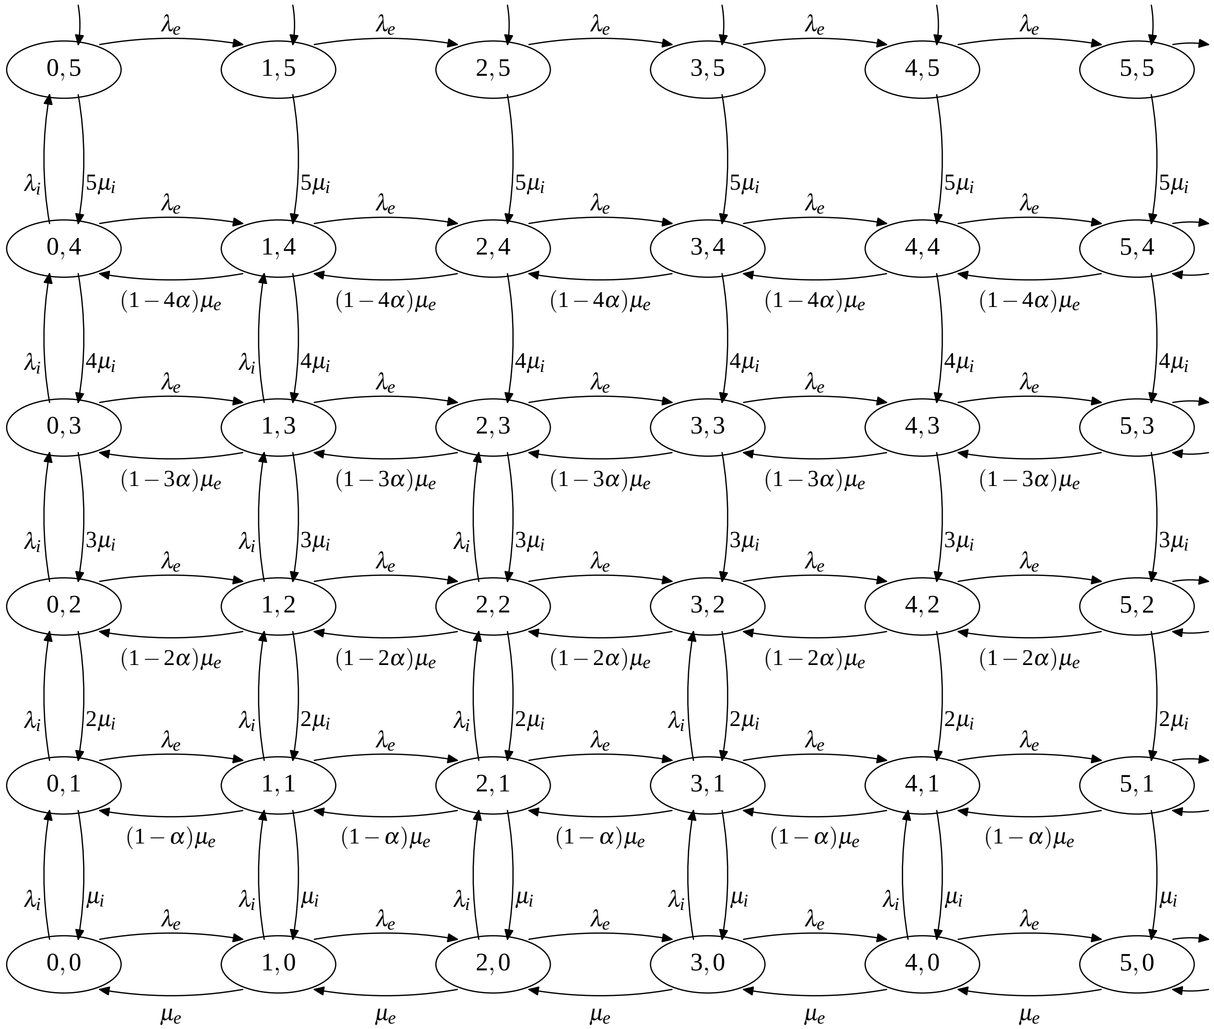

Markov Chain: Summary

		$(n,m) \to (n,m+1)$	$(n,m) \to (n+1,m)$	$(n,m) \to (n,m-1)$	$(n,m) \to (n-1,m)$
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	$m\mu_i$	$\max(0,(1-m\alpha)\mu_e)$

Simulation

- We can solve numerically
- Easier: simulation, for different workloads
- Result: AC-C > AC-A, CC > NC

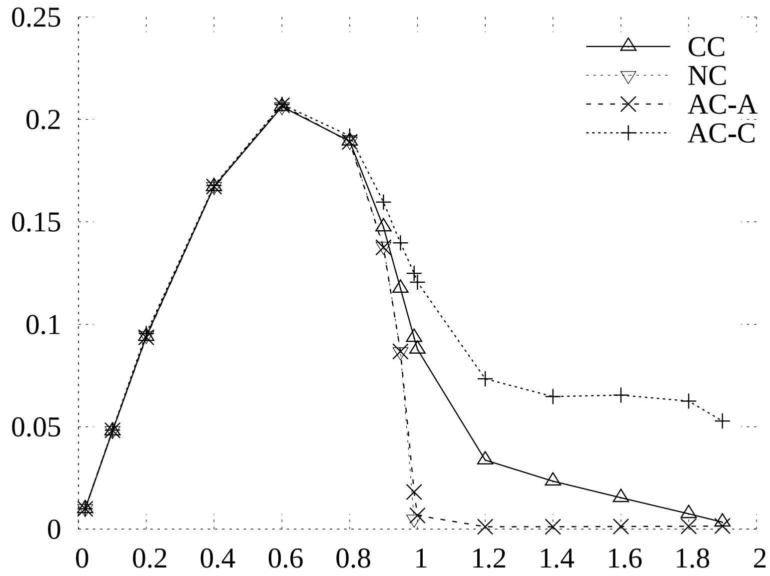

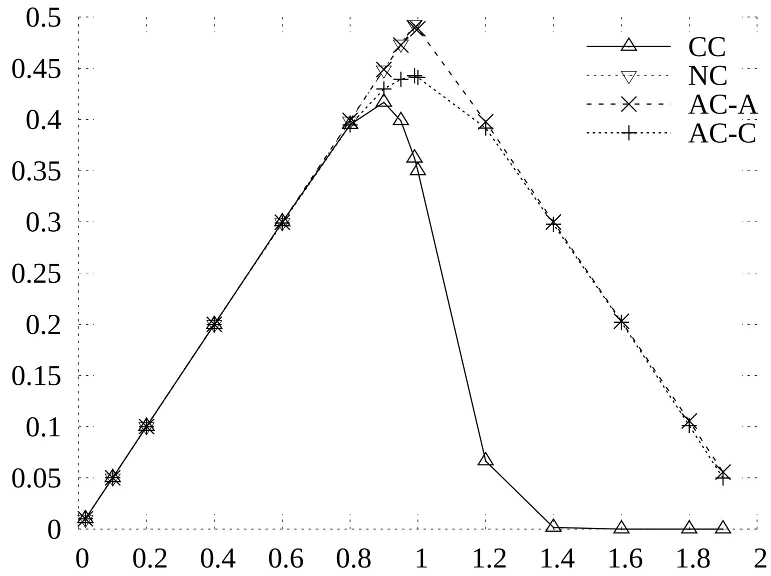

The above is just one of the many cases, showing equal offered load from elastic and inelastic traffic

Different approaches

- Evaluation 1: Utility based
- Evaluation 2: Stochastic Differential Equations

Evaluation 2: Blocking Probability

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

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

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

$$\begin{aligned} dE[\tau] &= E[-\mathbf{1}(\tau > 0)]dt + E[S_e dN_e + I(n,m)S_i dN_i] \\ &= -\Pr[\tau > 0]dt \\ &\quad + E[S_e]E[dN_e] + \Pr[I(n,m) = 1]E[S_i]E[dN_i] \end{aligned}$$

$$\frac{dE[\tau]}{dt} = -\Pr[\tau > 0]dt + \rho_e + \Pr[I(n,m) = 1]\alpha\rho_i$$

Evaluation 2: Blocking Probability

Setting $\frac{dE[\tau]}{dt} = 0$,

$$\Pr[I(n,m) = 1] = \frac{\Pr[\tau > 0] - \rho_e}{\alpha \rho_i}$$

this is the admission probability, i.e. $1 - P_{\text{block}}$

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

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

- We do not have ϵ in the equation!
- Whichever AC models we use, the resulting P_{block} is the same
  - Being aggressive and selfish does not improve the performance
  - In terms of social welfare, AC-C should be chosen in place of AC-A

Selfish is not good

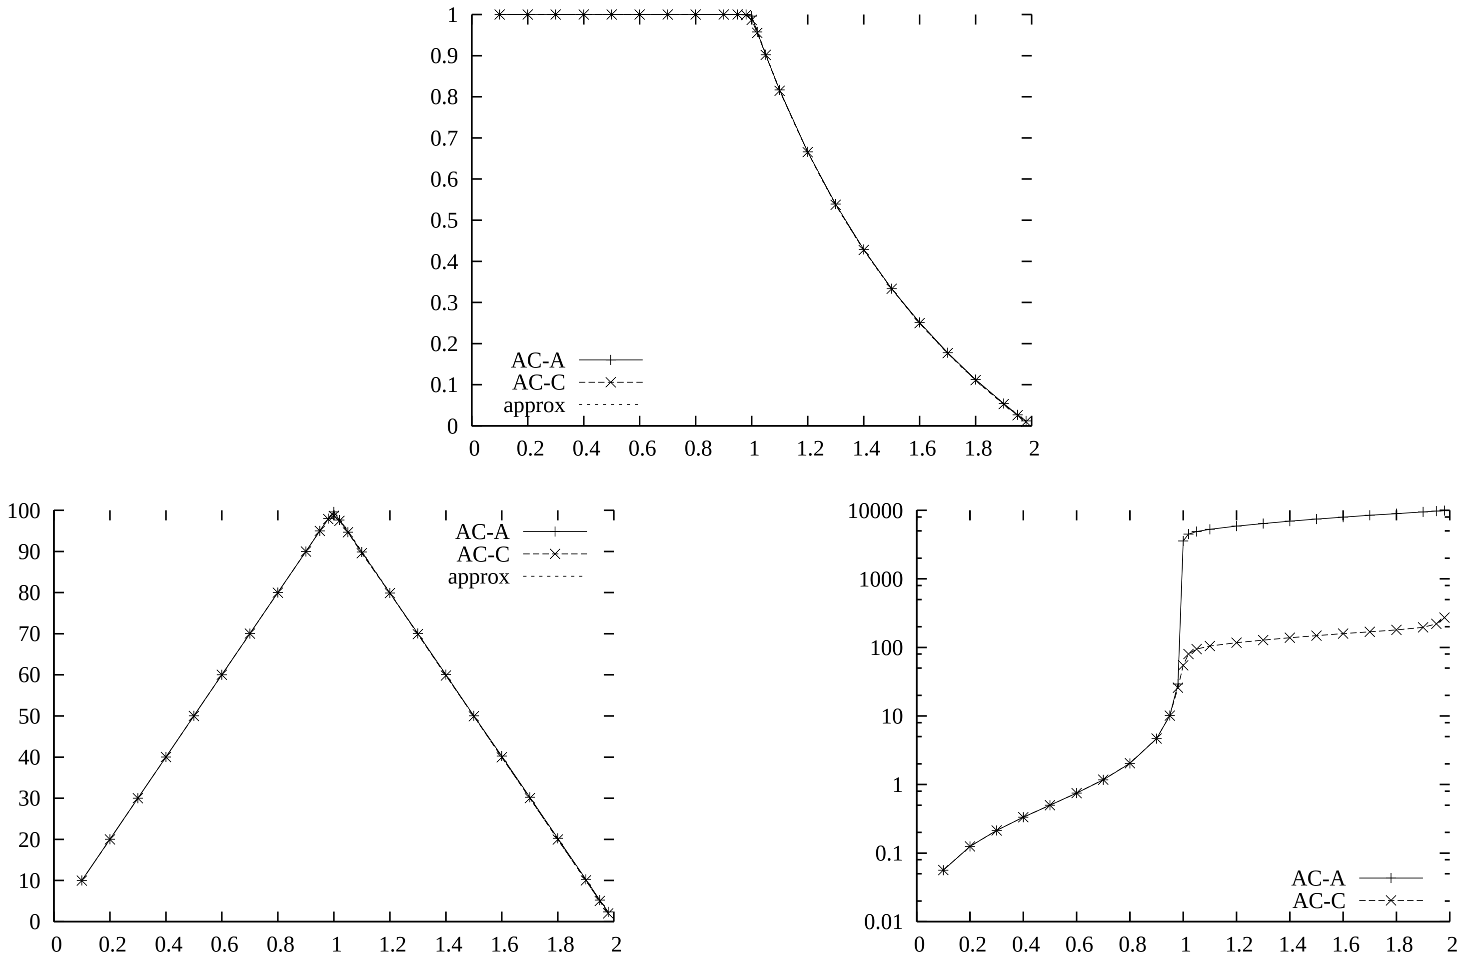

Conclusion

- We argue for multimedia traffic, it is better to use admission control then TCP-friendly congestion control
- To make admission control TCP-friendly is easy: Think you are normal TCP first, and see if you will still get what you want
  - If not, then quit, otherwise continue
  - Remember: It is not worth 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*
  A journal paper in progress