RESOURCE ALLOCATION IN NETWORKS

Jean Walrand - BANGALORE - January 2012

CONTENTS

1. Introduction
2. Utilities & Choice
3. Network Economics
4. Distributed Algorithms
5. Pricing

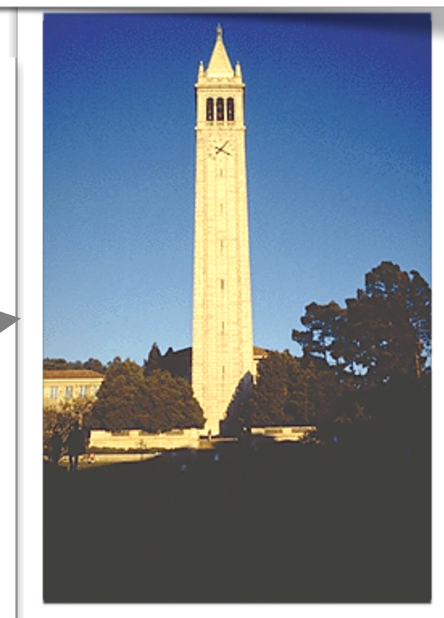

ACKNOWLEDGMENTS

Many thanks to the organizers of the workshop for inviting me.

This material is based on a course that I taught a few times with Dr. Abhay Parekh at Berkeley and on a manuscript we are about to finish [PW12]. Many thanks to Abhay!

The work reported here includes many contributions from colleagues and students and is supported in part by NSF and by MURI Research Grants of AFRL and ONR.

Dr. Abhay Parekh

SOME CONTRIBUTORS

CONTENTS

- **Introduction:** Overview, Resources, Good Allocations.
- **Utilities and Social Choice:** Preferences, Utilities, Efficiency, Fairness, Strategy-Proof Allocations, Distributed Allocations, Cooperative vs. Non-Cooperative.
- **Network Economics:** Game Theory, Tragedy of the Commons, Price of Anarchy in Routing (Tardos-Rougharten), Net Neutrality (Musacchio-Schwartz-W.), Economics of Security (Jiang-Anantharam-W., Schwartz et al., Gueye-W.), Network Upgrades (Duan-Huang-W.)
- **Distributed Algorithms:** Primal/Dual Decomposition (Kelly, Low); Backpressure Protocols (Srikant, Modiano-Neely), Wireless Backpressure (Jiang-W.)
- **Pricing:** Congestion Pricing, Time-Shifting (Jiang-Anantharam-W.), Tokens (Lee-Mo-W.), Paris Metro Pricing, Mixed Spectrum Auction (Silva-Beltran-W.), Network Upgrades (Duan-Huang-W.), Collaboration (Duan et al.)

TABLE OF CONTENTS

TABLE OF CONTENTS

RESOURCE ALLOCATION IN NETWORKS

Topics:

- Spectrum Auctions
- WiFi Association
- Routing
- Congestion Control
- P2P peer selection
- Data Center Scheduling
- Pricing Wireless Data
- Upgrade to 4G
- Sharing Revenue
- Ad Hoc Protocols
- Network Neutrality
- Security Investments

RESOURCE ALLOCATION IN NETWORKS

Challenges:

Efficient
Fair
Distributed
Robust
Strategy-proof
Uncertainty
Incentives

1. INTRODUCTION - 1.1 OVERVIEW

Key Observations:

1) Rapid growth in demand
2) Vast variability in demand across users

Questions:

How to handle this situation?

Traffic shaping?

Traffic discrimination?

User/Source discrimination?

Pricing?

1. INTRODUCTION - 1.1 OVERVIEW

Key Observation 1:

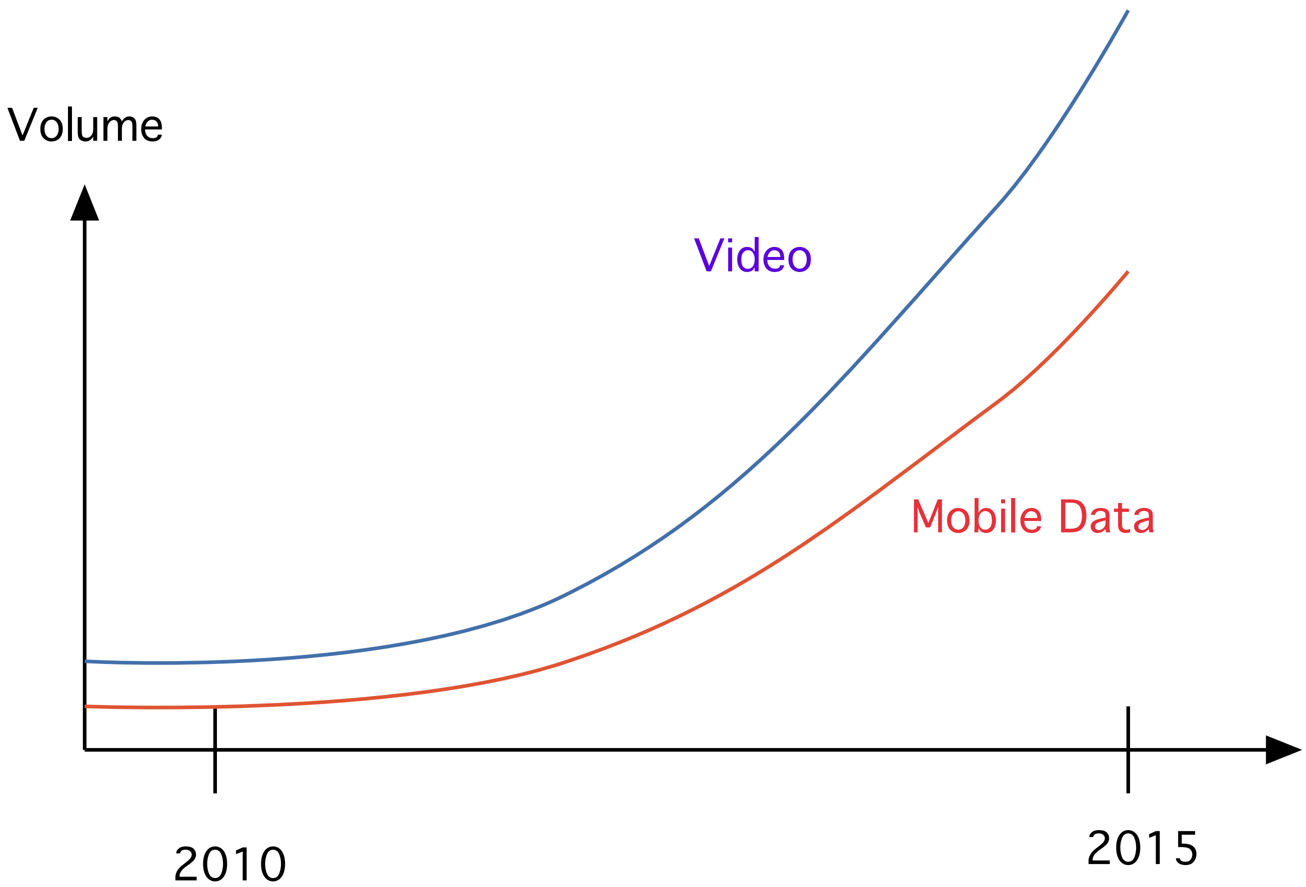

Key Observation 2:

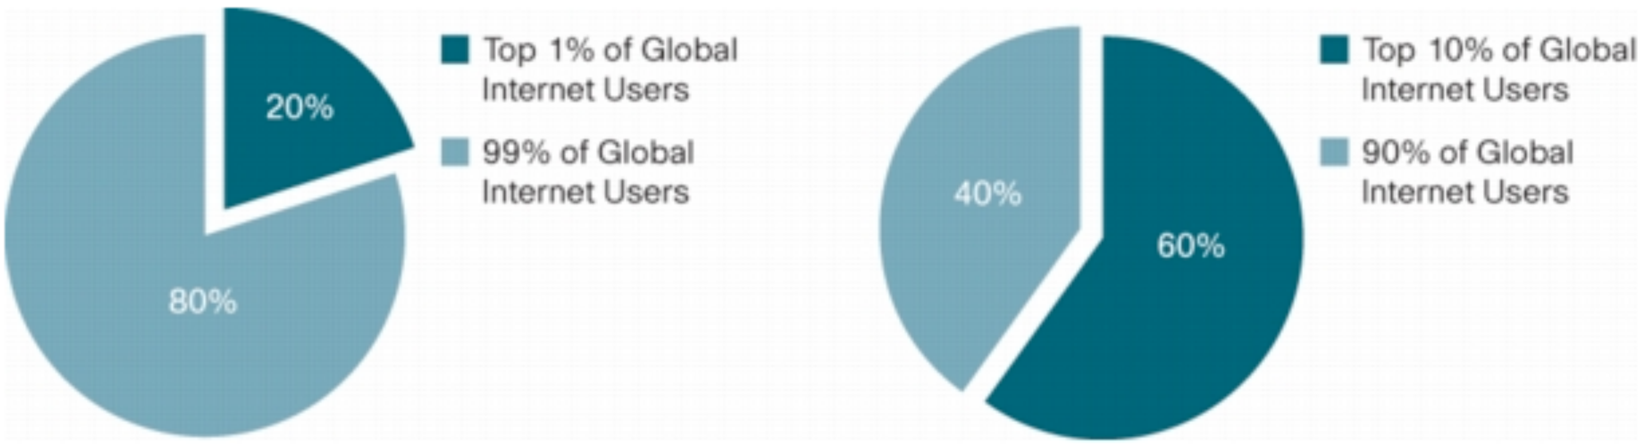

Source: VNI Usage, 2010

1. INTRODUCTION - 1.1 OVERVIEW

EXAMPLES:

Verizon to offer data 'Turbo' API to developers, fees to users

By Dieter Bohn on November 2, 2011 06:47 pm

1. INTRODUCTION - 1.1 OVERVIEW

EXAMPLES:

An Open Technology Initiative Report

Bandwidth Caps for Residential High-Speed Internet in the U.S. and Japan

By Chiehyu Li and James Losey, New America Foundation

August 10, 2009

Lowest Bandwidth Cap									
United States					Japan				
ISP	**Speed Down/Up**	**Monthly Price**	**Monthly Cap Down/Up**	**P2P Limits**	**ISP**	**Speed Down/Up**	**Monthly Price**	**Monthly Cap Up Only**	**P2P Limits**
Cable One	1.5Mbps/150Kbps	$26	1GB total	No	i-revo	100Mbps/100Mbps	$51-130	150GB	No
Cox	768Kbps/256Kbps 1.5Mbps/384Kbps	$19.95 $29.99	3GB/1GB 4GB/1GB	Yes	BB. Excite	47Mbps/5Mbps 100Mbps/100Mbps	$63 $60	420GB	Yes
Time Warner	768Kbps/128Kbps 1.5Mbps/256Kbps 7Mbps/384Kbps 10Mbps/512Kbps	$19.95 $34.95 $49.95 $59.90	40GB total (proposed)	Yes	Internet Initiative Japan	100Mbps/100Mbps	$48-77	450GB	No
AT&T	768Kbps/384Kbps	$19.95	20GB (proposed)	No	SoftBank ODN	100Mbps/100Mbps	$53-72	450GB	No

EXAMPLES:

comcast. Customer Central

Home | Account & Bill | Users & Settings | Alerts

› Overview Billing High-Speed Internet Internet 2go Cable TV Digital Voice

Frequently Asked Questions about Excessive Use

Comcast is committed to providing the best online experience for all its customers. Please review our frequently asked questions regarding excessive use below in order to find out what Comcast is doing for you.

How does Comcast define "excessive use"?

Am I at any risk of reaching the "excessive use" threshold?

Why is this policy in place?

Is Comcast going to offer a pay-per-gigabyte option for customers who go over 250 GB in a month?

How is your data usage threshold evaluated on an ongoing basis? What customer input do you seek?

What will happen if I exceed 250 GB of data usage in a month?

EXAMPLES:

BE AFRAID

AT&T to Cap Broadband Usage Starting Monday

at&t bandwidth caps Broadband

by James Plafke | 2:00 pm, May 1st, 2011

0 159 21

+1 Like Tweet

Starting Monday, **AT&T** will place a **cap** on its **broadband Internet** services. Following Comcast's broadband cap and AT&T being the number two carrier in the U.S., the majority of the U.S. broadband Internet will now be saddled with a cap. AT&T will be placing a 150 GB monthly cap on its DSL users and a 250 GB monthly cap for its UVerse users.

AT&T plans to charge an extra $10 a month if users go over their designated cap, with a recurring $10 charge for each 50 GB over the cap the user goes. Not counting torrent and other nefarious data usage, Wired points out that the 150 GB and 250 GB caps may at first seem like a decent amount, but in this day and age of Netflix and online gaming, those caps could easily be reached. Wired does a bit of the Netflix math, and states that streaming standard content ranges anywhere from .3 GB an hour to 1.0 GB an hour, eas streaming HD content can max out at 2.3 GB an hour. This may not seem like a lot,

EXAMPLES:

Netflix Traffic Overtakes Web Surfing In US

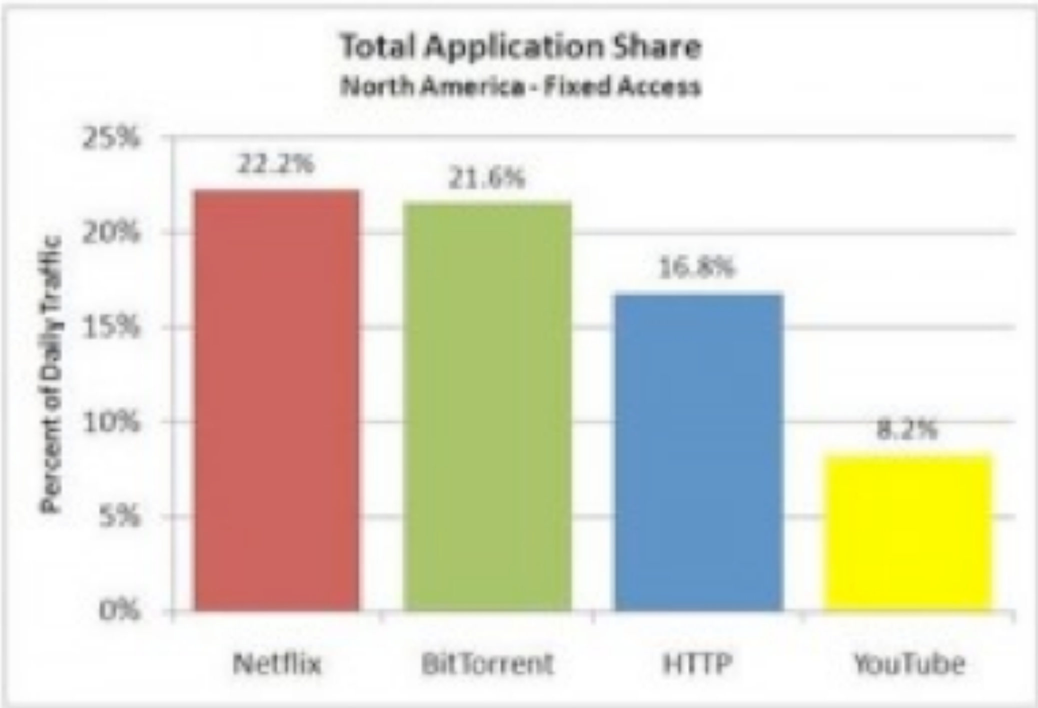

May 2011

1. INTRODUCTION - 1.1 OVERVIEW

EXAMPLES:

Figure 1. Cisco Forecasts 6.3 Exabytes per Month of Mobile Data Traffic by 2015

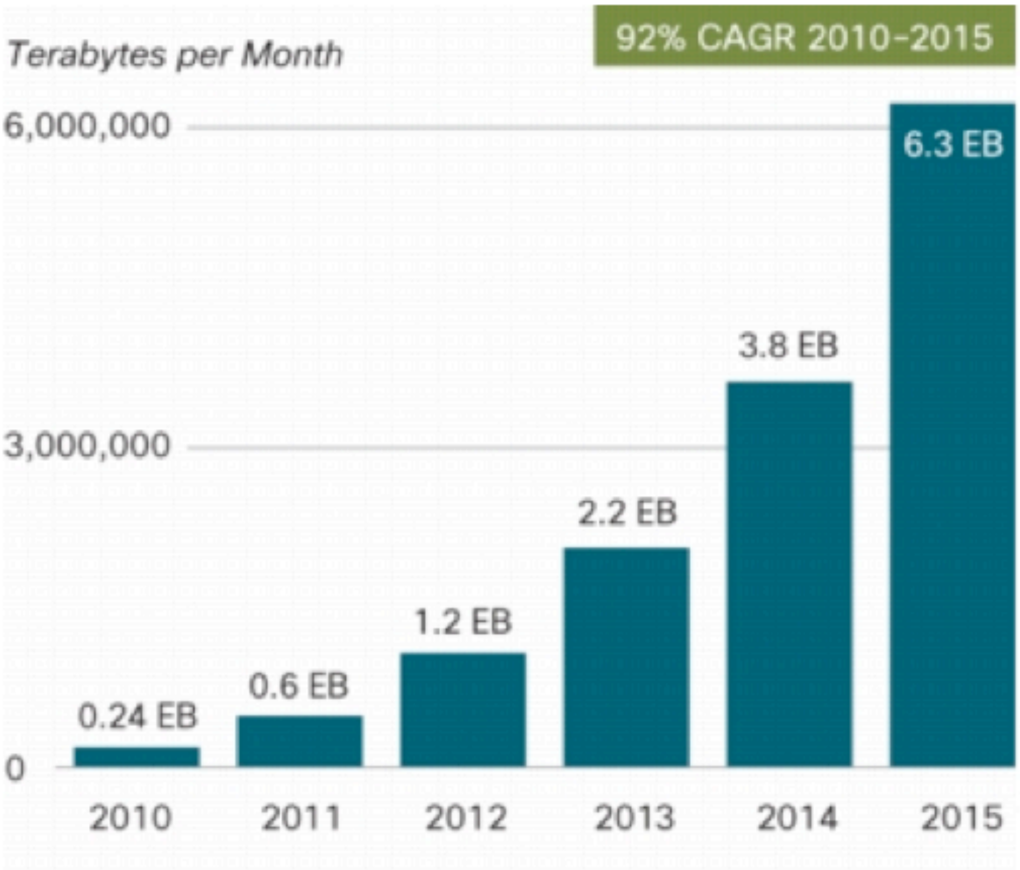

Source: Cisco VNI Mobile, 2011

EXAMPLES:

Figure 4. High-End Devices Can Multiply Traffic

1. INTRODUCTION - 1.1 OVERVIEW

EXAMPLES:

Figure 5. Mobile Video Will Generate 66 Percent of Mobile Data Traffic by 2015

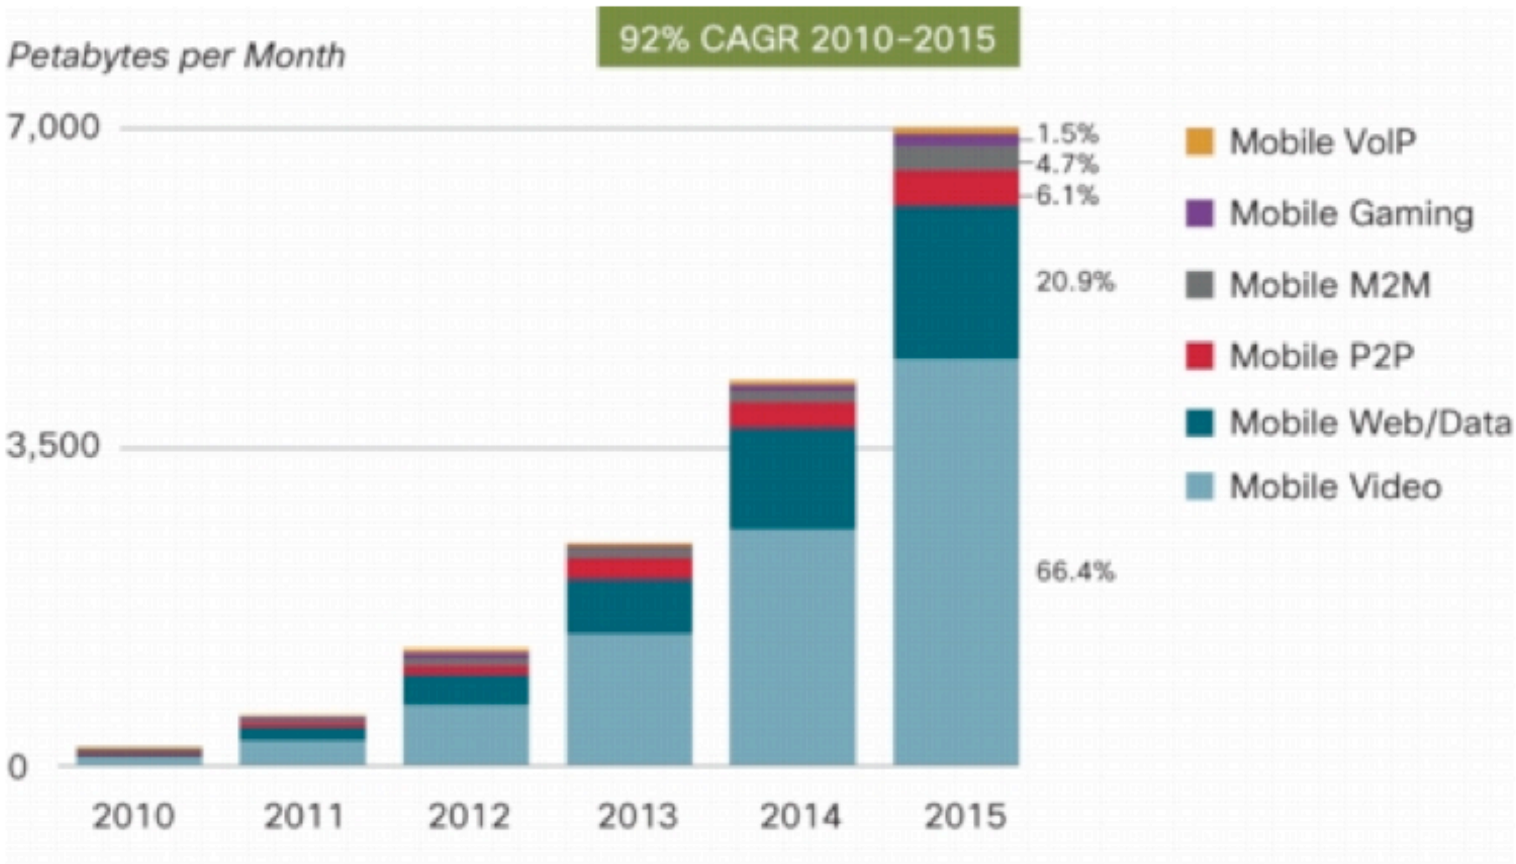

VoIP traffic forecasted to be 0.4% of all mobile data traffic in 2015.

Source: Cisco VNI Mobile, 2011

EXAMPLES:

Table 2. Global IP Traffic, 2010-2015

IP Traffic, 2010-2015							
	2010	**2011**	**2012**	**2013**	**2014**	**2015**	**CAGR 2010-2015**
By Type (PB per Month)							
Fixed Internet	14,955	20,650	27,434	35,879	46,290	59,354	**32%**
Managed IP	4,989	6,839	9,014	11,352	13,189	14,848	**24%**
Mobile data	237	546	1,163	2,198	3,806	6,254	**92%**

NETWORK ECONOMICS?

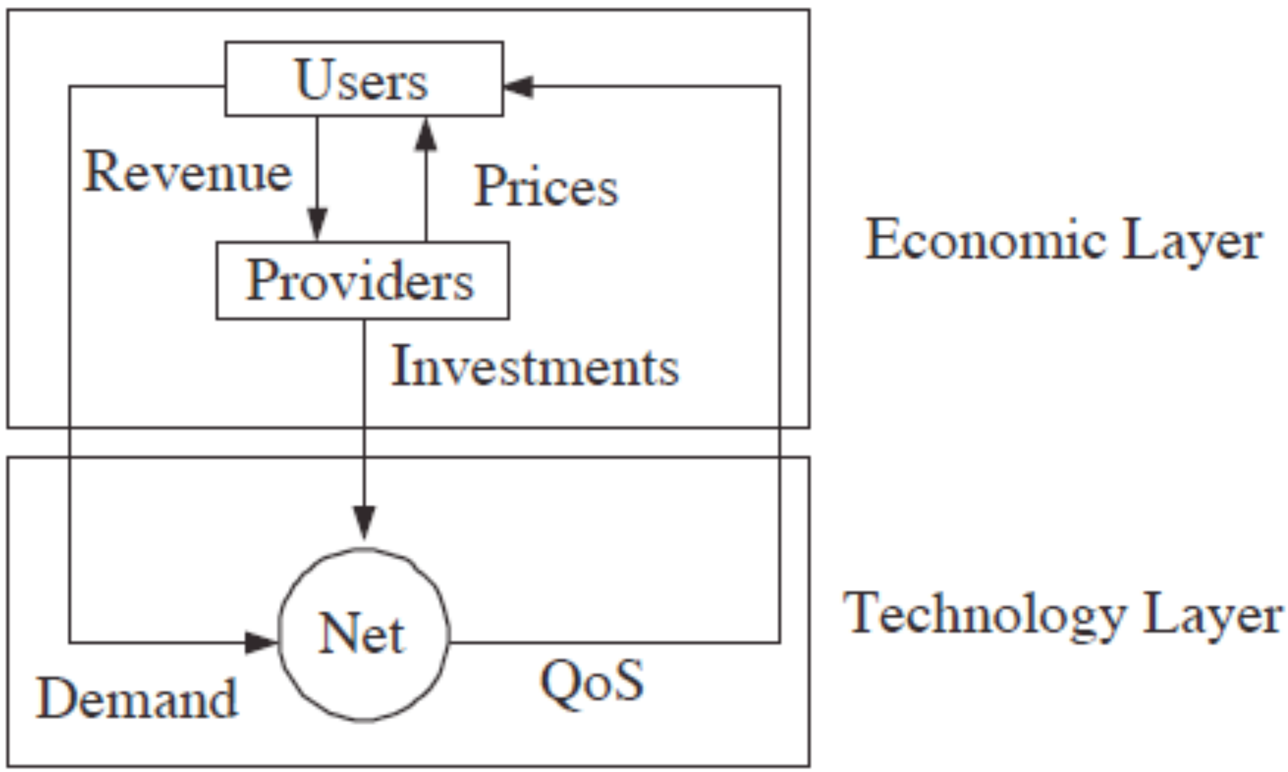

Users and providers respond to economic incentives and affect the network.

[WAL08b]

MAIN POINTS:

- A range of resources to be allocated over very different time scales.
- Examples:
    - New spectrum: Years
    - New TCP connection: Seconds
    - Medium access control: Milliseconds

1. INTRODUCTION - 1.2 RESOURCES

Time Scale	Action	Examples
Years	Acquire	Spectrum, Right-of-Way, Energy Contract, Land, Buildings,...
Months	Deploy	Base Stations, Fibers, Routers, Servers, Access Points, New Technology, ...
Weeks	Configure	SLAs, Channel Allocation, Wavelengths, Tariffs, ...
Seconds	Connect	Associate, Access Control, Circuit Setup, TCP handshake, Load Balancing, New Table Entry, Routing, P2P, Servers...
Milliseconds	Forward, Transmit	Scheduling, Window, CSMA,

KEY NOTIONS:

- Efficiency
- Fairness
- Strategyproofness
- Scalability
- Robustness

1. INTRODUCTION - 1.3 GOOD ALLOCATIONS

Efficiency	Full use of resources: - User welfare (sum of user utilities) Pareto: Cannot improve one without degrading another - Social welfare (sum of all utilities minus cost)
Fairness	Cannot grossly disadvantage some users - How to make this precise?

1. INTRODUCTION - 1.3 GOOD ALLOCATIONS

Goals	Efficiency & Fairness ■ Efficient (Pareto, ...) ■ Fair
Properties	■ Scalable ⇒ Distributed ■ Strategy-Proof (Work when strategic users may cheat) ■ Robust (Work in face of uncertainty & dynamics)

SCALABLE: Share N Sources among 4 users.

Each user specifies 2^N numbers! ← Information

Number of possible allocations: ← Computation

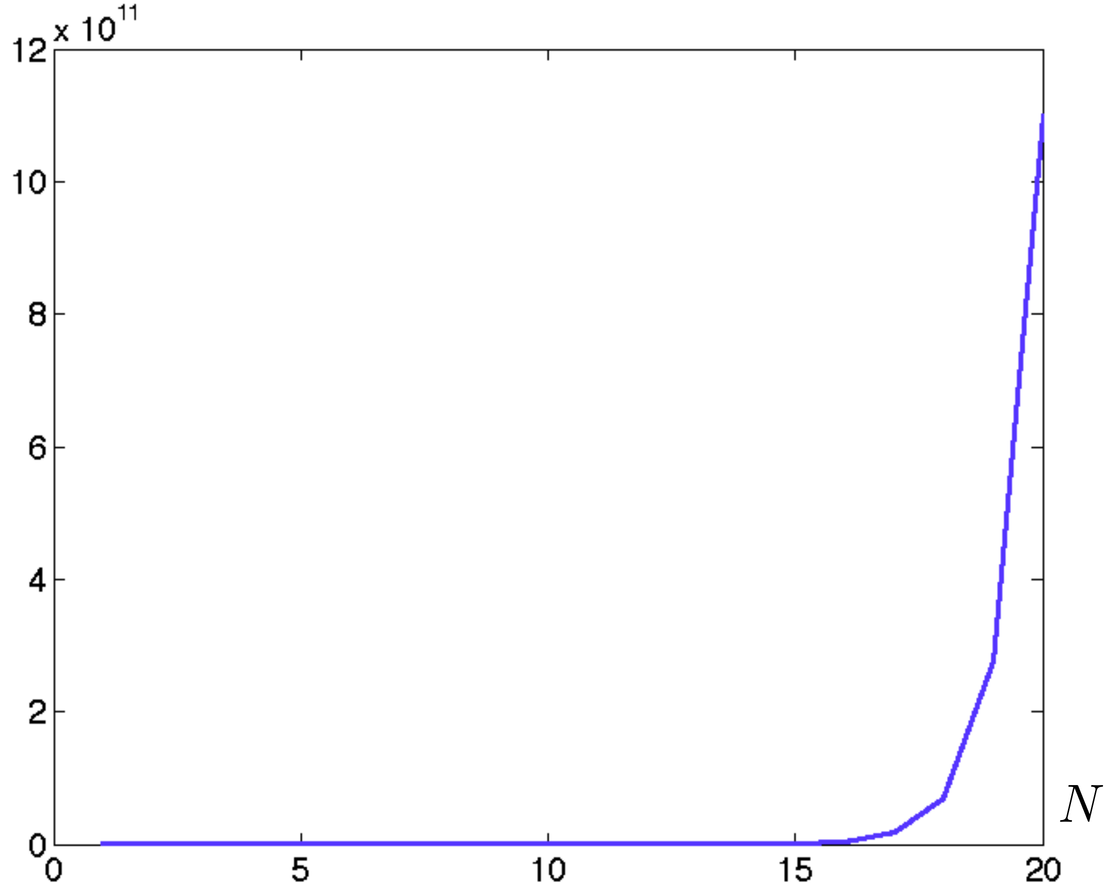

STRATEGY-PROOF:

Say $U_i(x) = \log(x), i = 1, 2.$

Maximize $U_1(x_1) + U_2(x_2)$ subject to $x_1 + x_2 \leq 10.$

$$\Rightarrow x_1 = x_2 = 5.$$

Assume user 1 pretends that her utility is $w.\log(x)$ for some $w > 1.$

$\Rightarrow$ Maximize $w.\log(x_1) + \log(x_2)$ subject to $x_1 + x_2 \leq 10.$

$$\Rightarrow x_1 = 10\frac{w}{w+1} \text{ and } x_2 = 10\frac{1}{w+1}.$$

$\Rightarrow$ strong incentive to cheat: Not strategy-Proof.

ROBUSTNESS

Say $U_i(x) = \log(x), i = 1, 2.$

Maximize $U_1(x_1) + U_2(x_2)$ subject to $x_1 + x_2 \leq 10.$

$\Rightarrow x_1 = x_2 = 5.$

If user 1 goes away, the allocation should change to $(0, 10)$.

If the actual capacity is 9 instead of 10,

the allocation should change to $(4.5, 4.5)$.

The allocation should adjust quickly as conditions change.

ROBUSTNESS

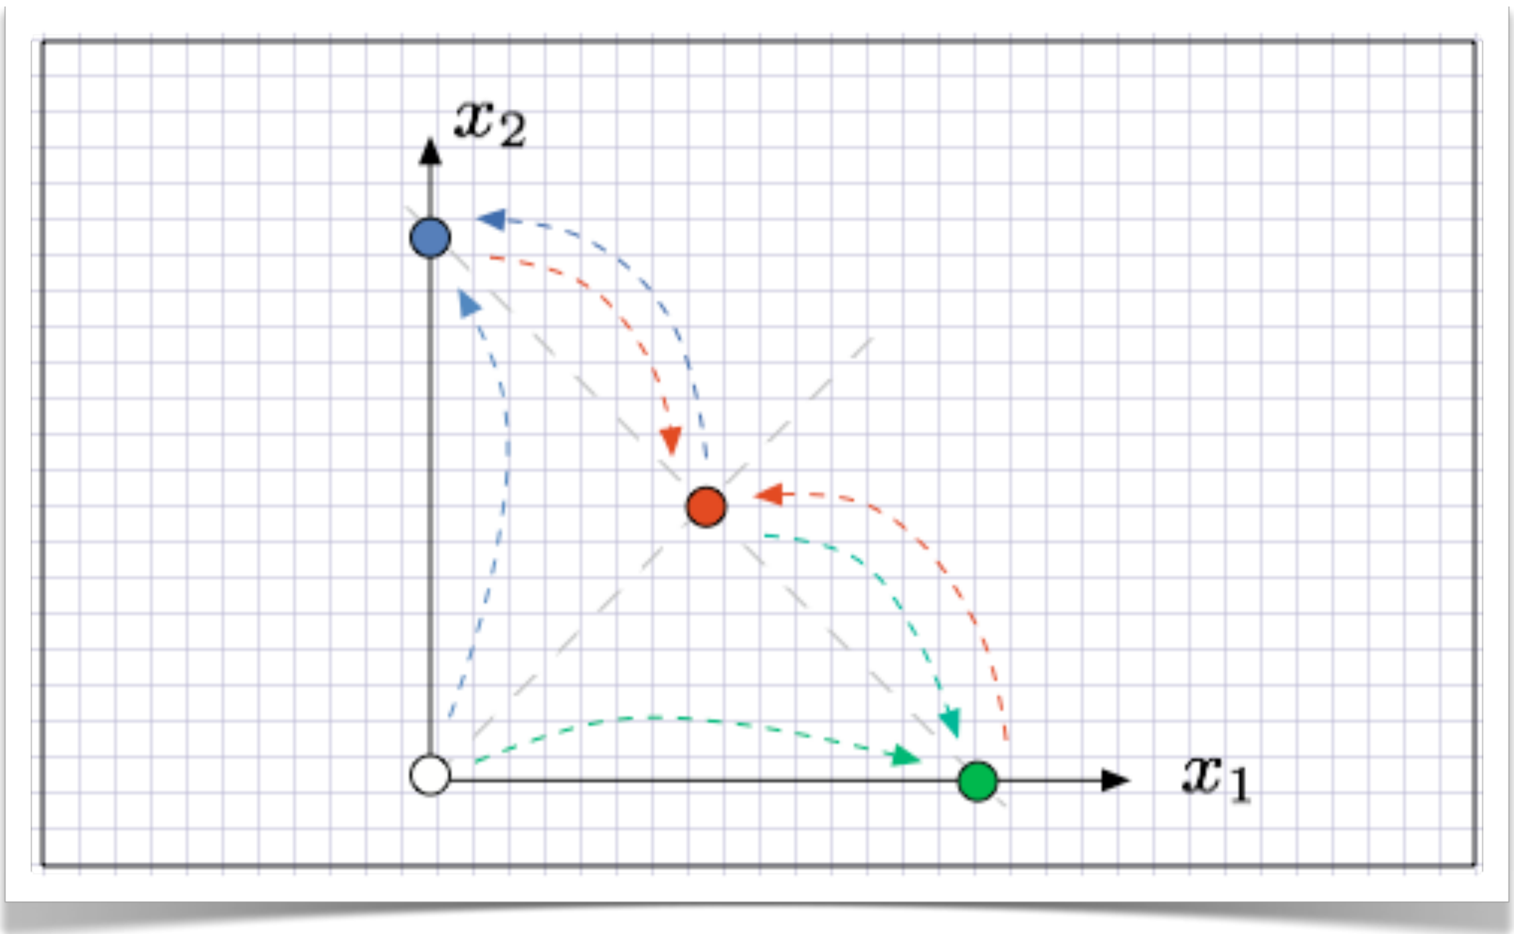

The allocation must adapt quickly to changing conditions.

1. INTRODUCTION - 1.4 SUMMARY

- New Technology, New Problems
  - Ad hoc wireless
  - Distributed control for smart grid
  - Pricing schemes for 4G services
  - Incentive schemes for P2P, security
  - Incentives for more efficient network utilization
- Resource Allocation on Multiple Time Scales
  - Acquire, Deploy, Provision, Connect, Transmit
- Good Allocations must be
  - Efficient, Fair, Scalable, Strategy Proof, Robust

TABLE OF CONTENTS

2. **Utilities & Choice**

What is the value of resources to users?

- Preferences
  - Transitivity?
  - Monetary value?
  - Aggregation?
- Utilities
  - Concavity?
  - Networking Examples
- Knowledge
  - Reveal Utilities?
  - Explicit vs. Implicit?

MAIN POINTS:

- Goal: Decide how to reconcile conflicting desires of different users
- Users' desires are specified by preferences or utilities
- Preferences may be intransitive (typically in case of multiple attributes) --> not much one can do
- When preferences are transitive, one can assign an "ordinal utility"
- Combining ordinal preferences is hopeless
- When users have cardinal utilities, one has a number of options.

PREFERENCES

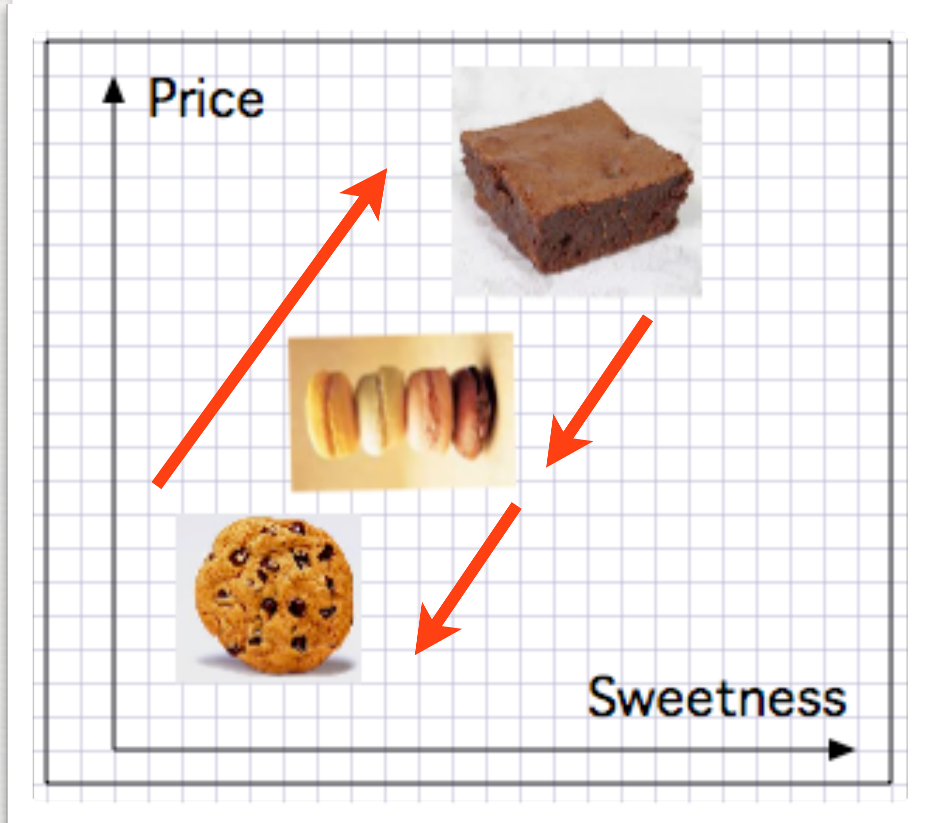

Similarly sweet: Prefer cheaper

Always prefer significantly sweeter

A → B = Prefer B to A

Note: Intransitive (two attributes)

PREFERENCES (Continued)

Not much better: Prefer cheaper

Always prefer much better

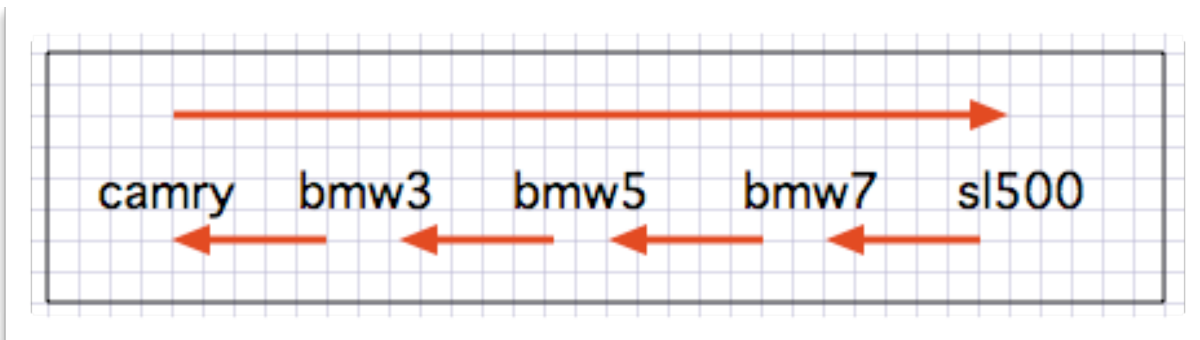

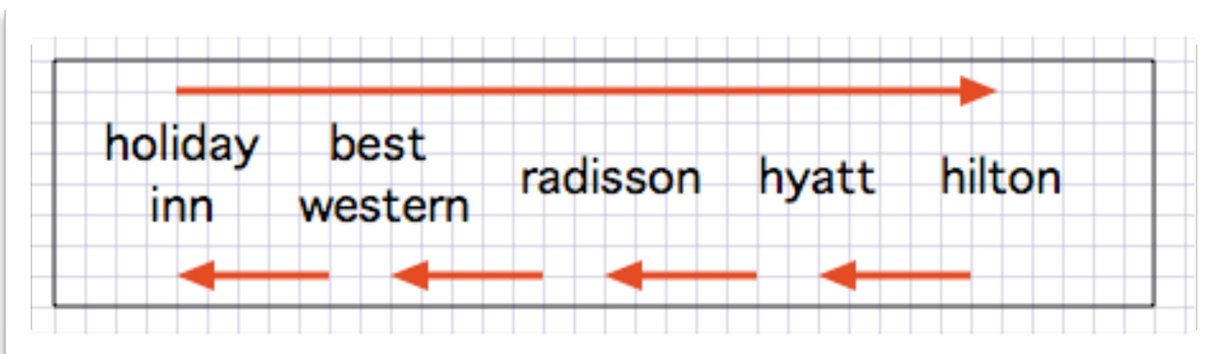

Note: Intransitive (two attributes)

PREFERENCES (Continued)

Often, preferences are transitive

Broadband Services:

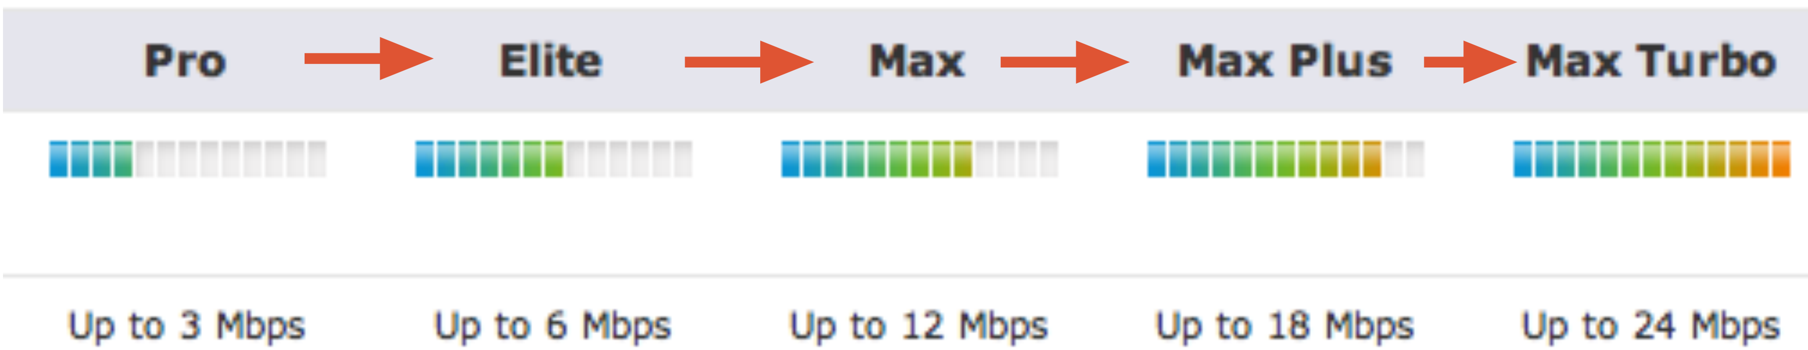

PREFERENCES (Continued)

Often, preferences are transitive

Activities:

Bob:

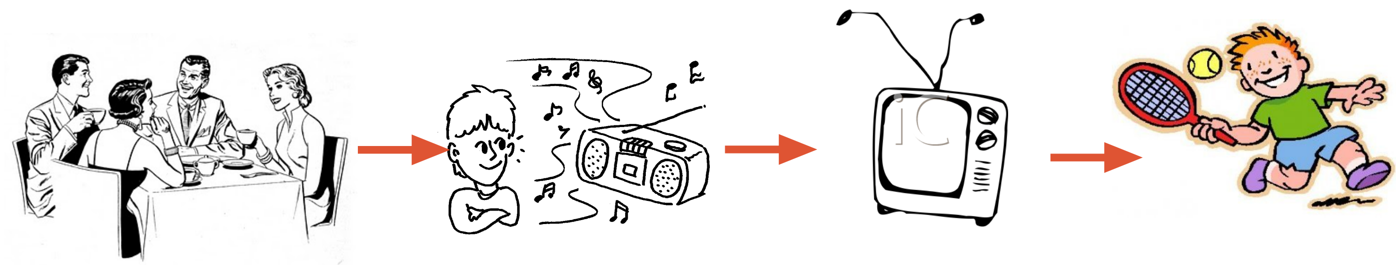

Alice:

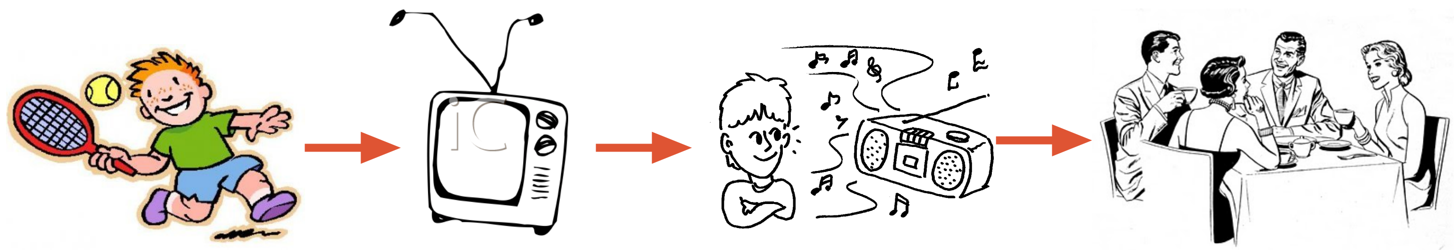

UTILITIES:

Transitive preferences can be mapped into ordinal utilities

Examples:

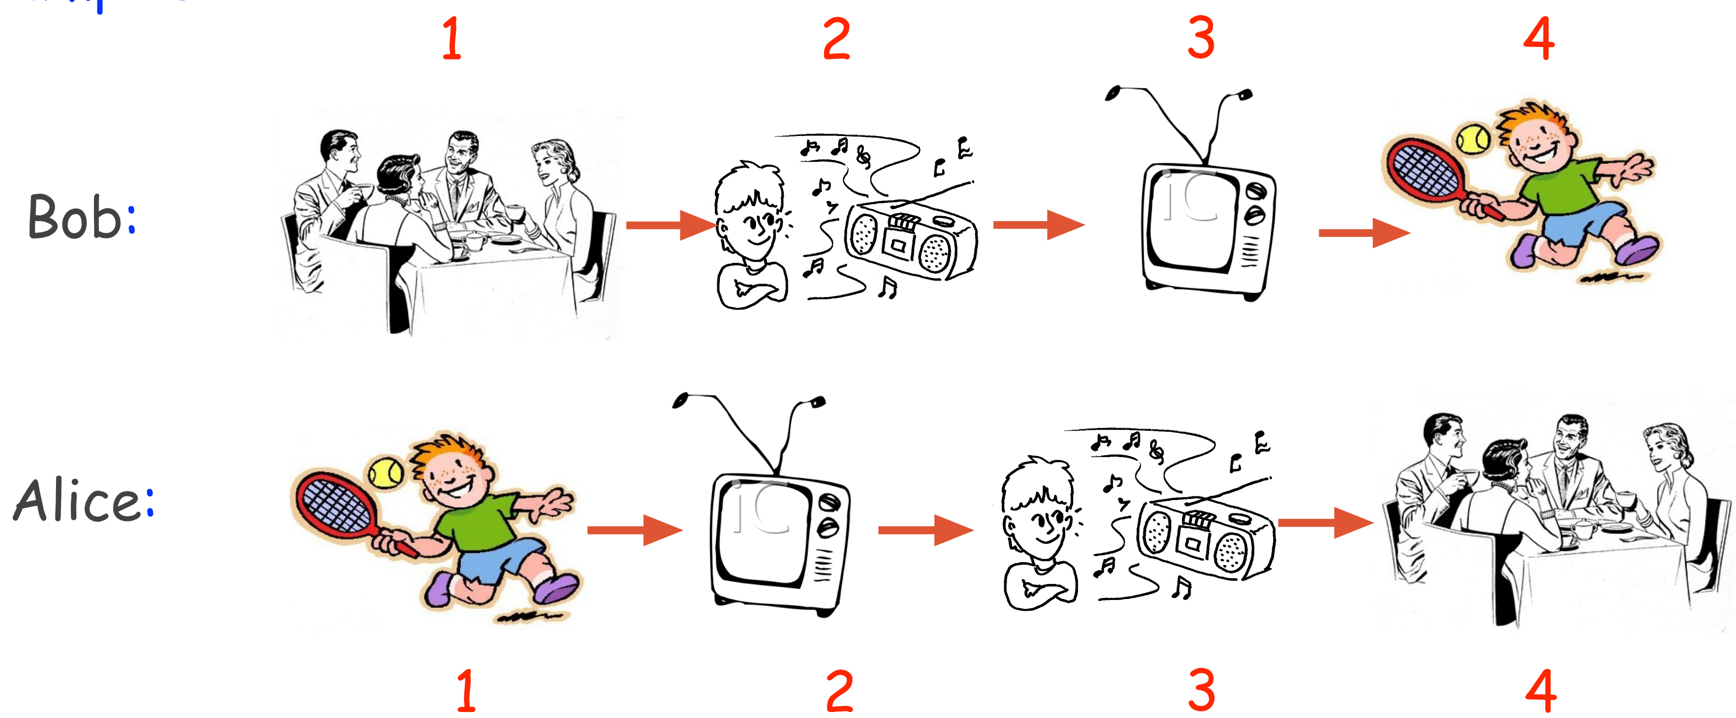

UTILITIES:

One may be able to map transitive preferences into `cardinal utilities.' The meaning is `how much more do you prefer B to A.'

Examples:

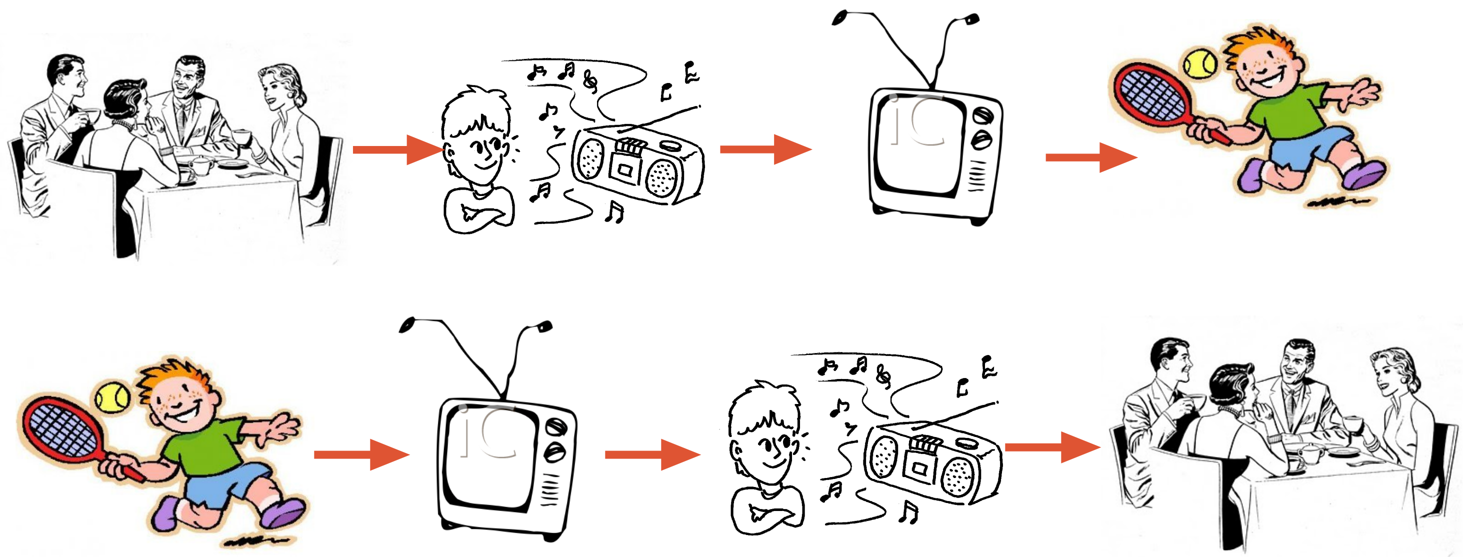

ECONOMICS VIEWPOINT:

Everyone (everything?) has a price.

For Alice:

 =  + 

Thus, all utilities have some monetary value, supposedly exchangeable.

In this case, one has monetary utilities.

COMBINING PREFERENCES

We could try to combine transitive preferences

Bob:

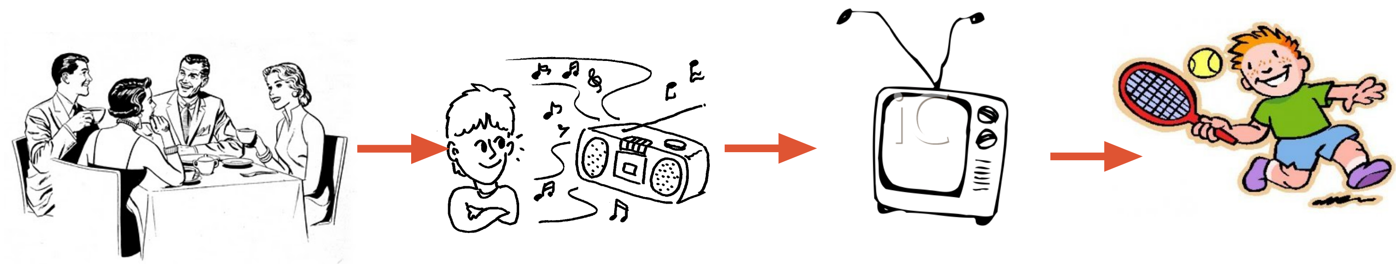

Alice:

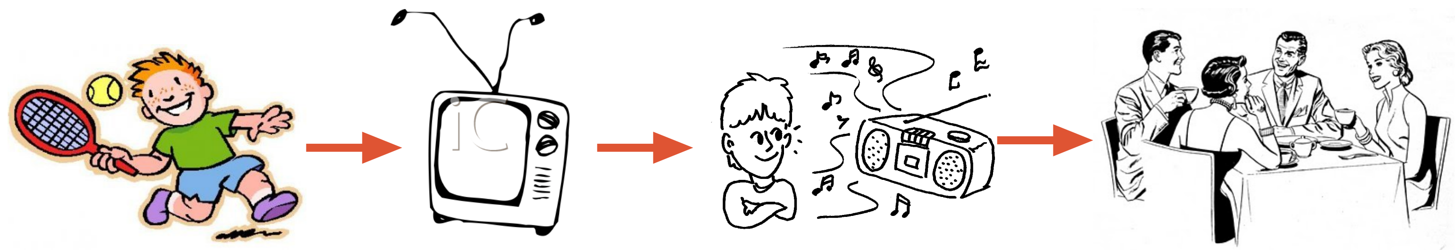

Social Choice SC: Individual Preferences → Global Preferences

Example: Preferences above →

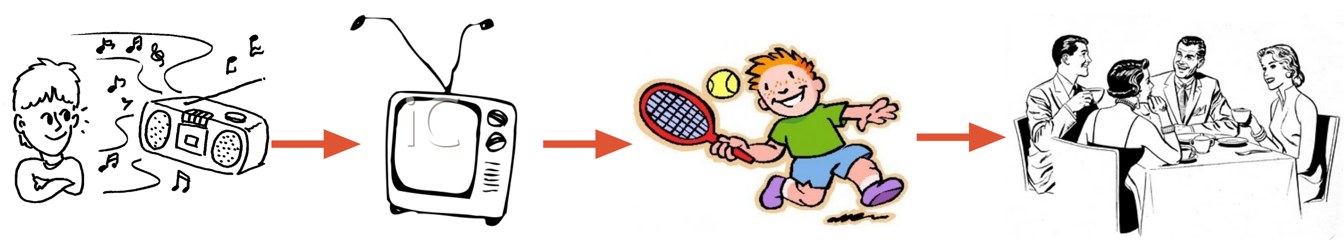

COMBINING PREFERENCES

SC: Individual Preferences $\rightarrow$ Global Preferences

Some reasonable properties:

1) Consistent: If everyone prefers some activity, that's the top choice.

2) No Dictator: no one guy decides based on his preferences alone.

3) Independence of Irrelevant Alternatives: The global order of TV and Tennis depends only on their order in the individual choices, not on the presence of other activities.

Arrow: Impossible if at least two people and three tasks!

COMBINING PREFERENCES

Arrow: Impossible if at least two people and three tasks!

"Proof:"

Assume:

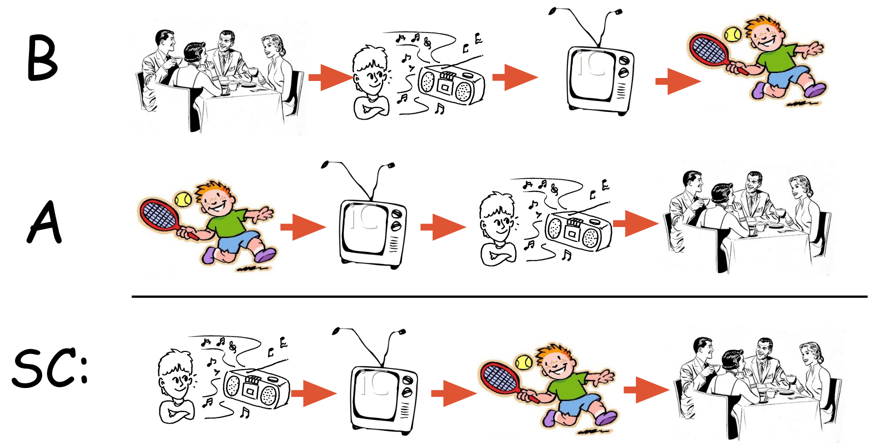

Modify A by shifting Tennis up until SC ranks Tennis highest.

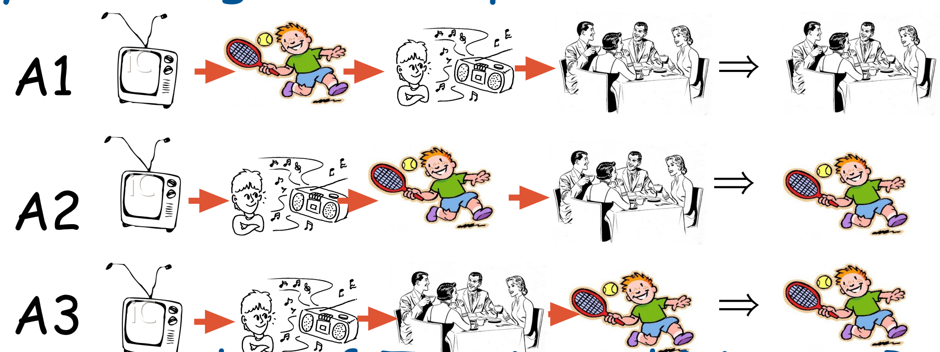

See [PW12] for details.

In A1, A2: same order of Tennis and Dinner. Remove everything else: the SC order of Tennis and Dinner cannot switch from (A1, B) to (A2, B). So switch must be at A3. A is a dictator.

COMBINING UTILITIES

Can one combine utilities?

10 15 20 30

Bob:

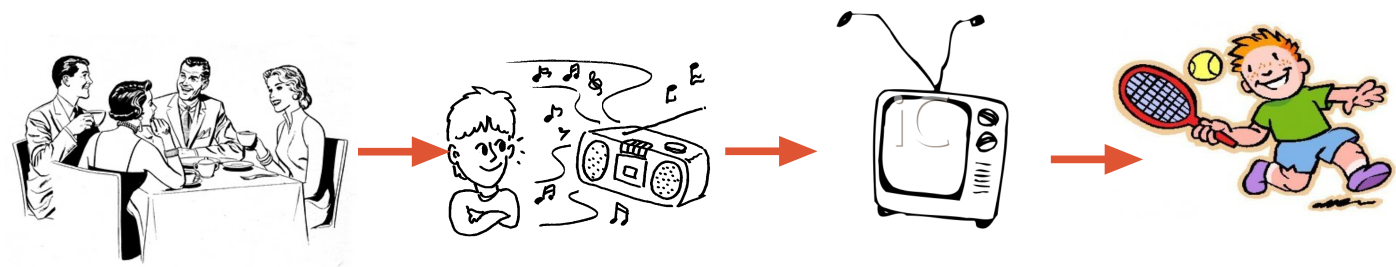

Alice:

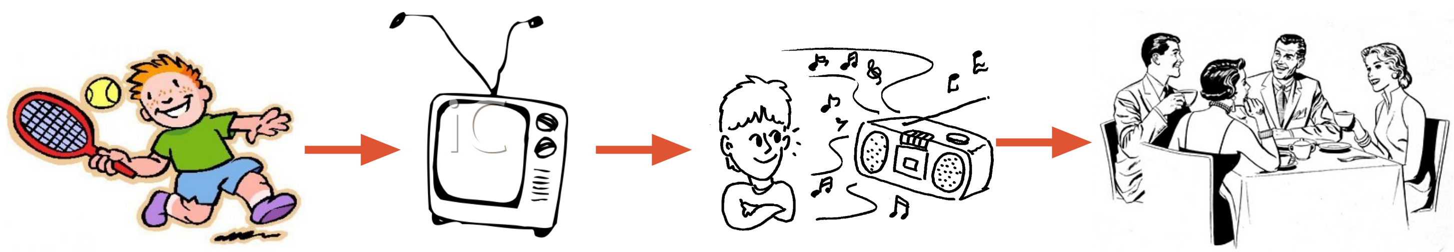

5 10 30 40

Say Bob and Alice must choose one common activity.

Does it make sense to maximize the <u>sum</u> of their utilities?

⇒ Dinner with friends

COMBINING UTILITIES

Can one combine utilities?

10 15 20 30

Bob:

Alice:

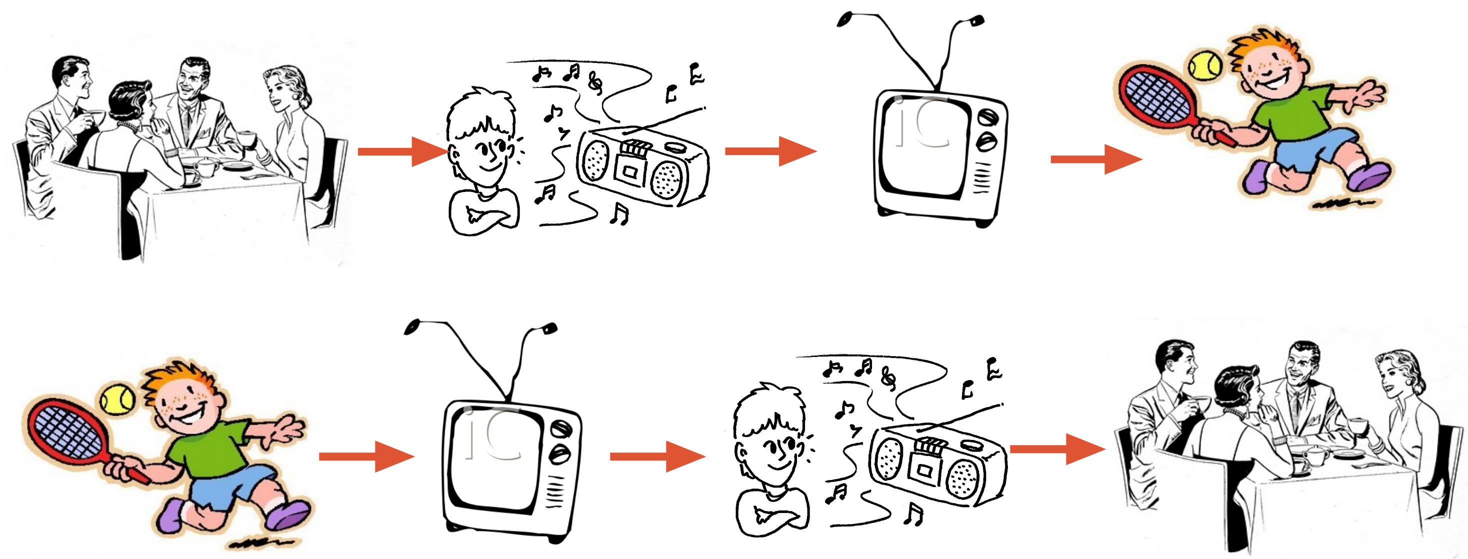

5 10 30 40

Say Bob and Alice must choose one common activity.

Nash says we should maximize the <u>product</u> of the utilities

⇒ Listen to music

COMBINING UTILITIES

Nash says we should maximize the <u>product</u> of the utilities

Insensitive to how Alice and Bob scale their utility.

However, still requires some cardinal measure for each!

Alice:  → 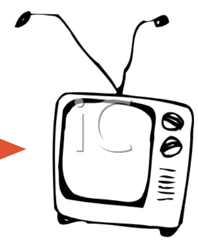 → 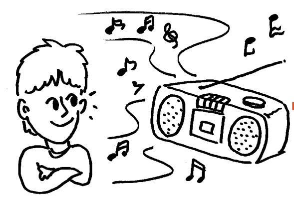 → 

	5	10	30	40
equivalent to	50	100	300	400
but not to	5	6	7	8

COMBINING UTILITIES

In the case of monetary utilities, maximizing the sum is sensible only if one can compensate the users (e.g., by exchanging money).

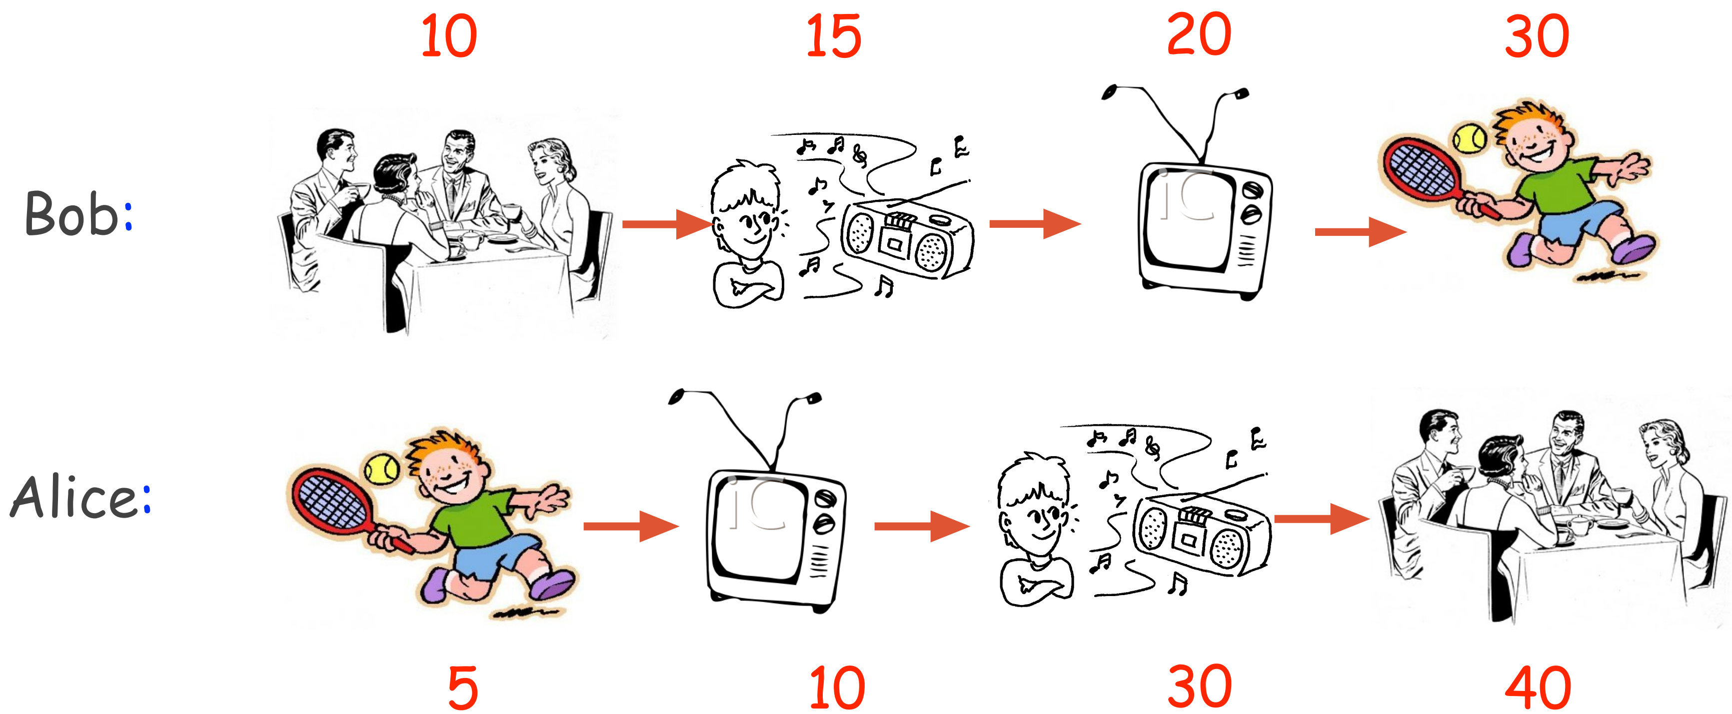

--> Dinner with friends, and Alice gives $15 to Bob.

2. UTILITIES & CHOICE - 2.2 SOCIAL CHOICE

WHAT IF MONEY EXCHANGE IS NOT POSSIBLE?

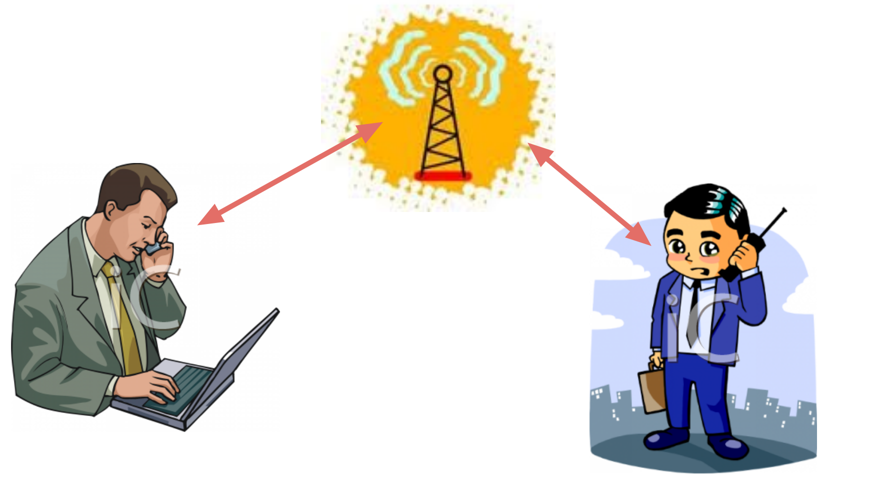

or

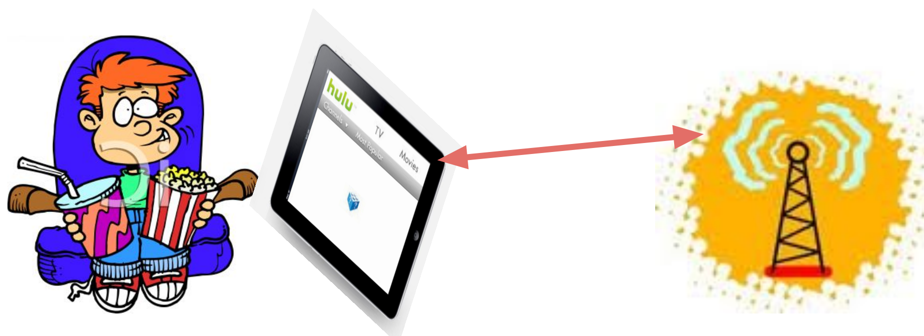

Possible Approaches:

Whoever pays the most: Maximize revenue

Some social sense of fairness: Maximize user welfare

Total benefit to society: Maximize welfare of all parties - cost

CONCAVITY?

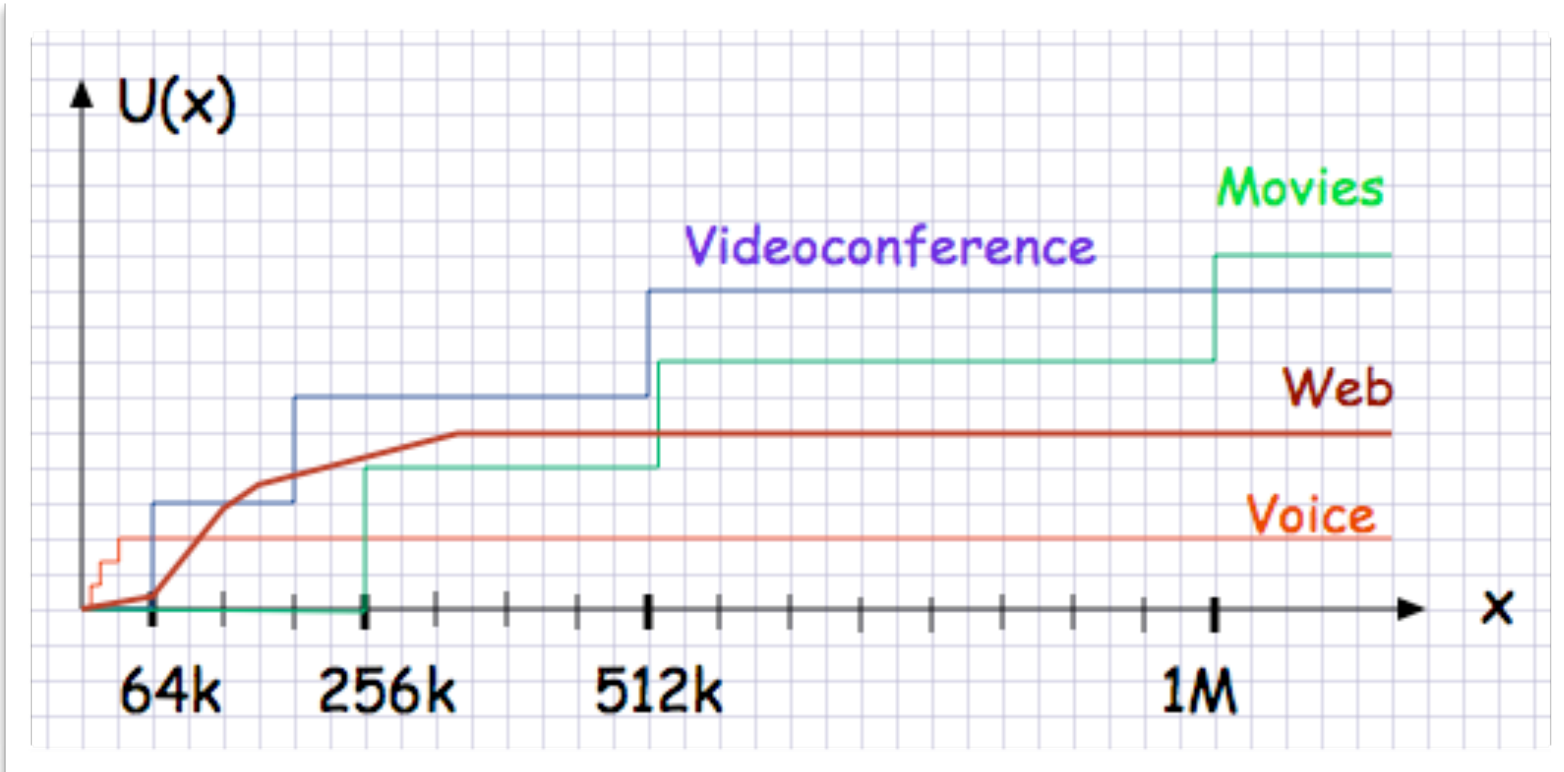

Hypothetical utilities, say per unit of time.
Utilities depend on the user, obviously.
The jumps come from the discrete set of codec rates.

A Good Goal?

Assume user i has some utility U_i.

There are many vectors $(U_1, \ldots, U_N)$ that are possible.

Which one is preferable?

- Maximize the sum $\Sigma_i U_i$: MAX SUM
- Maximize the minimum $\min_i U_i$: MAX MIN
- Maximize the product $\Pi_i U_i$: Proportionally Fair
- Maximize $\Pi_i (u_i)^{1-\alpha}$: α-Fair

EFFICIENCY & FAIRNESS:

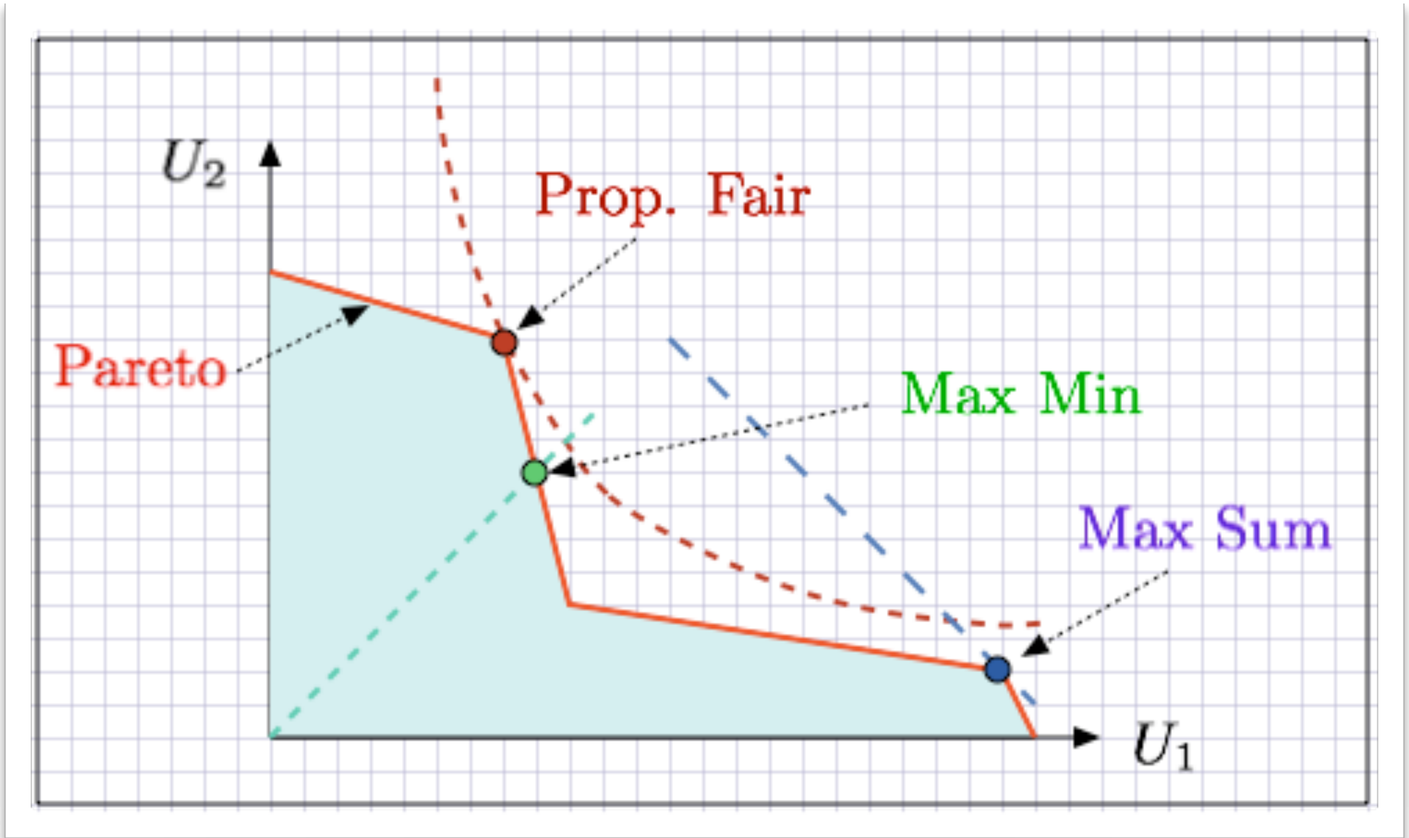

NOTES:

- Not all Pareto points are maximizers of $\Sigma_i a_i U_i$ for some a_i's.
- For some applications, one may "convexify" by time sharing.

EFFICIENCY & FAIRNESS:

$$U_i(x_i) = x_i$$
$$x_1 + 2x_2 \leq 1$$

$$\text{Maximize } \sum_i (x_i)^{1-\alpha}$$

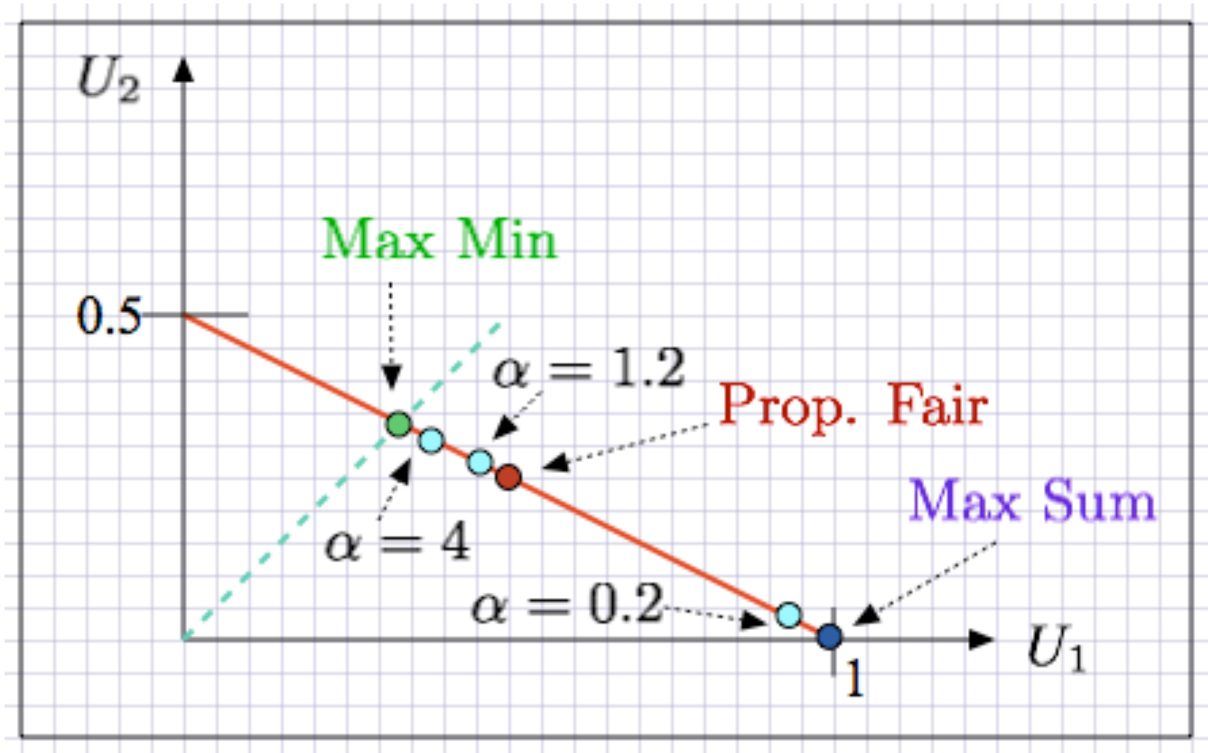

NOTE:

- One finds $U_1 = [1 + 2^{1-\alpha^{-1}}]^{-1}$

[MO00]

Goal: Sharing resources

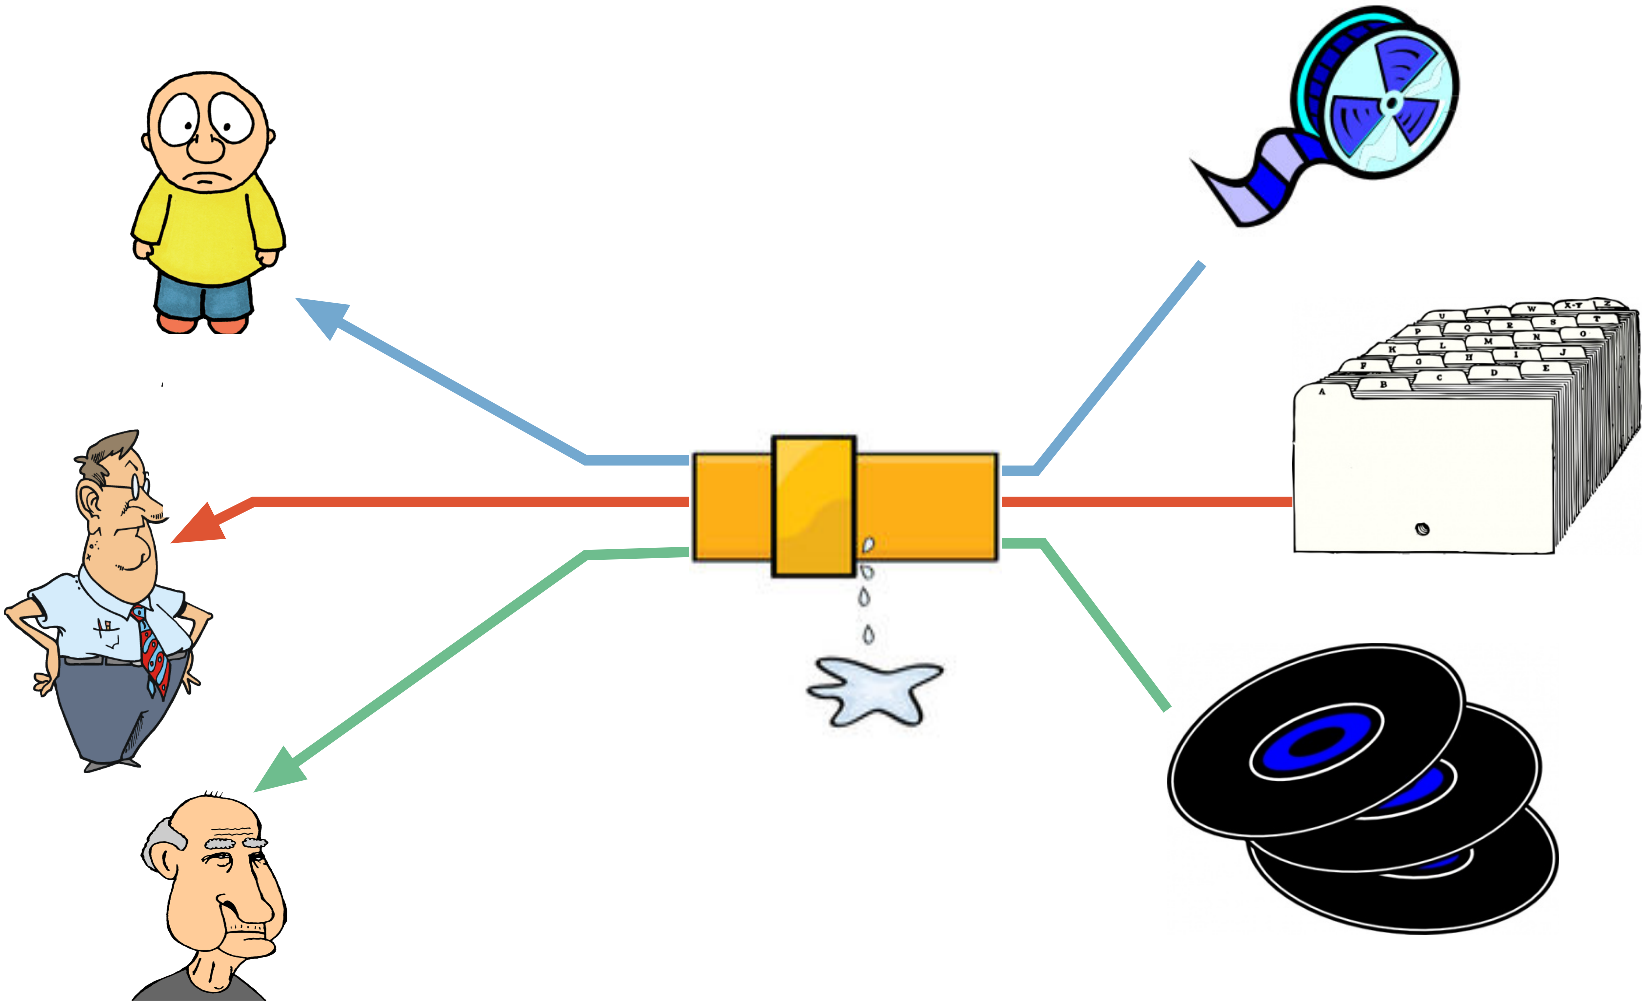

Centralized:

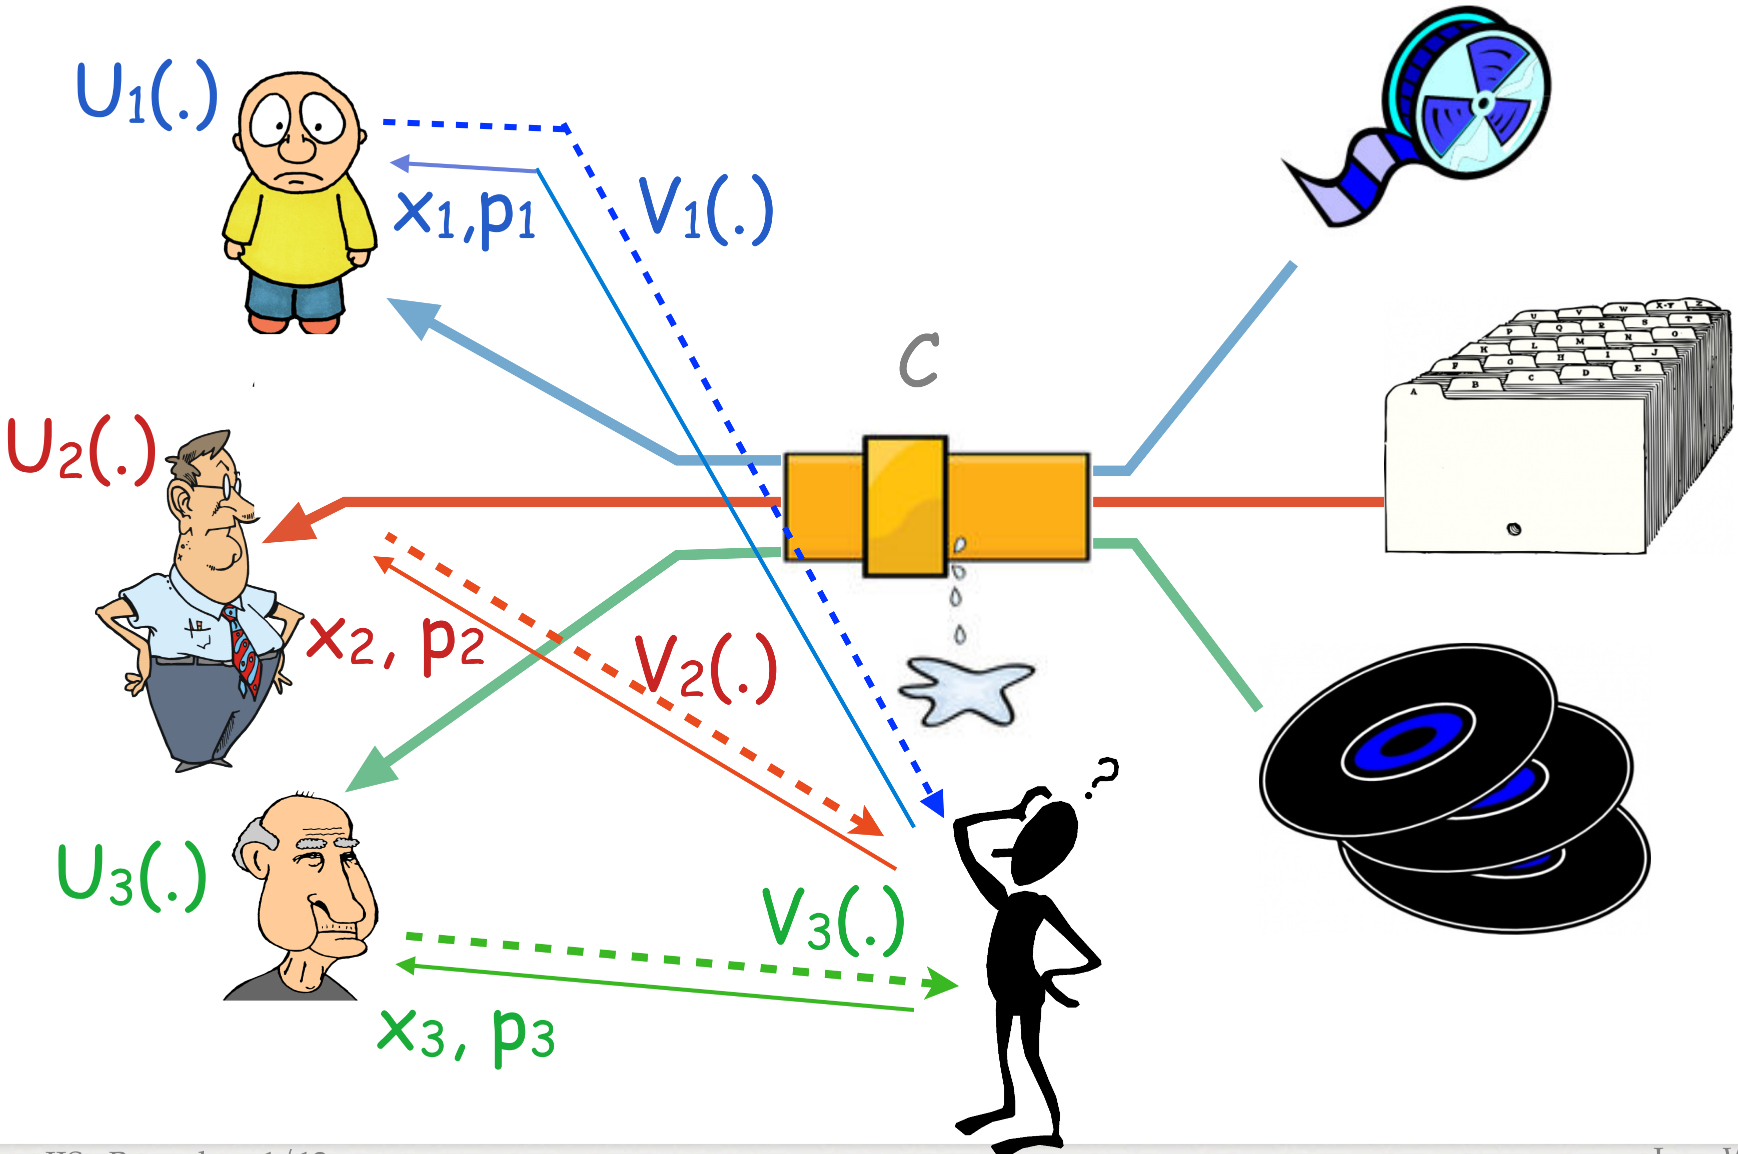

Centralized:

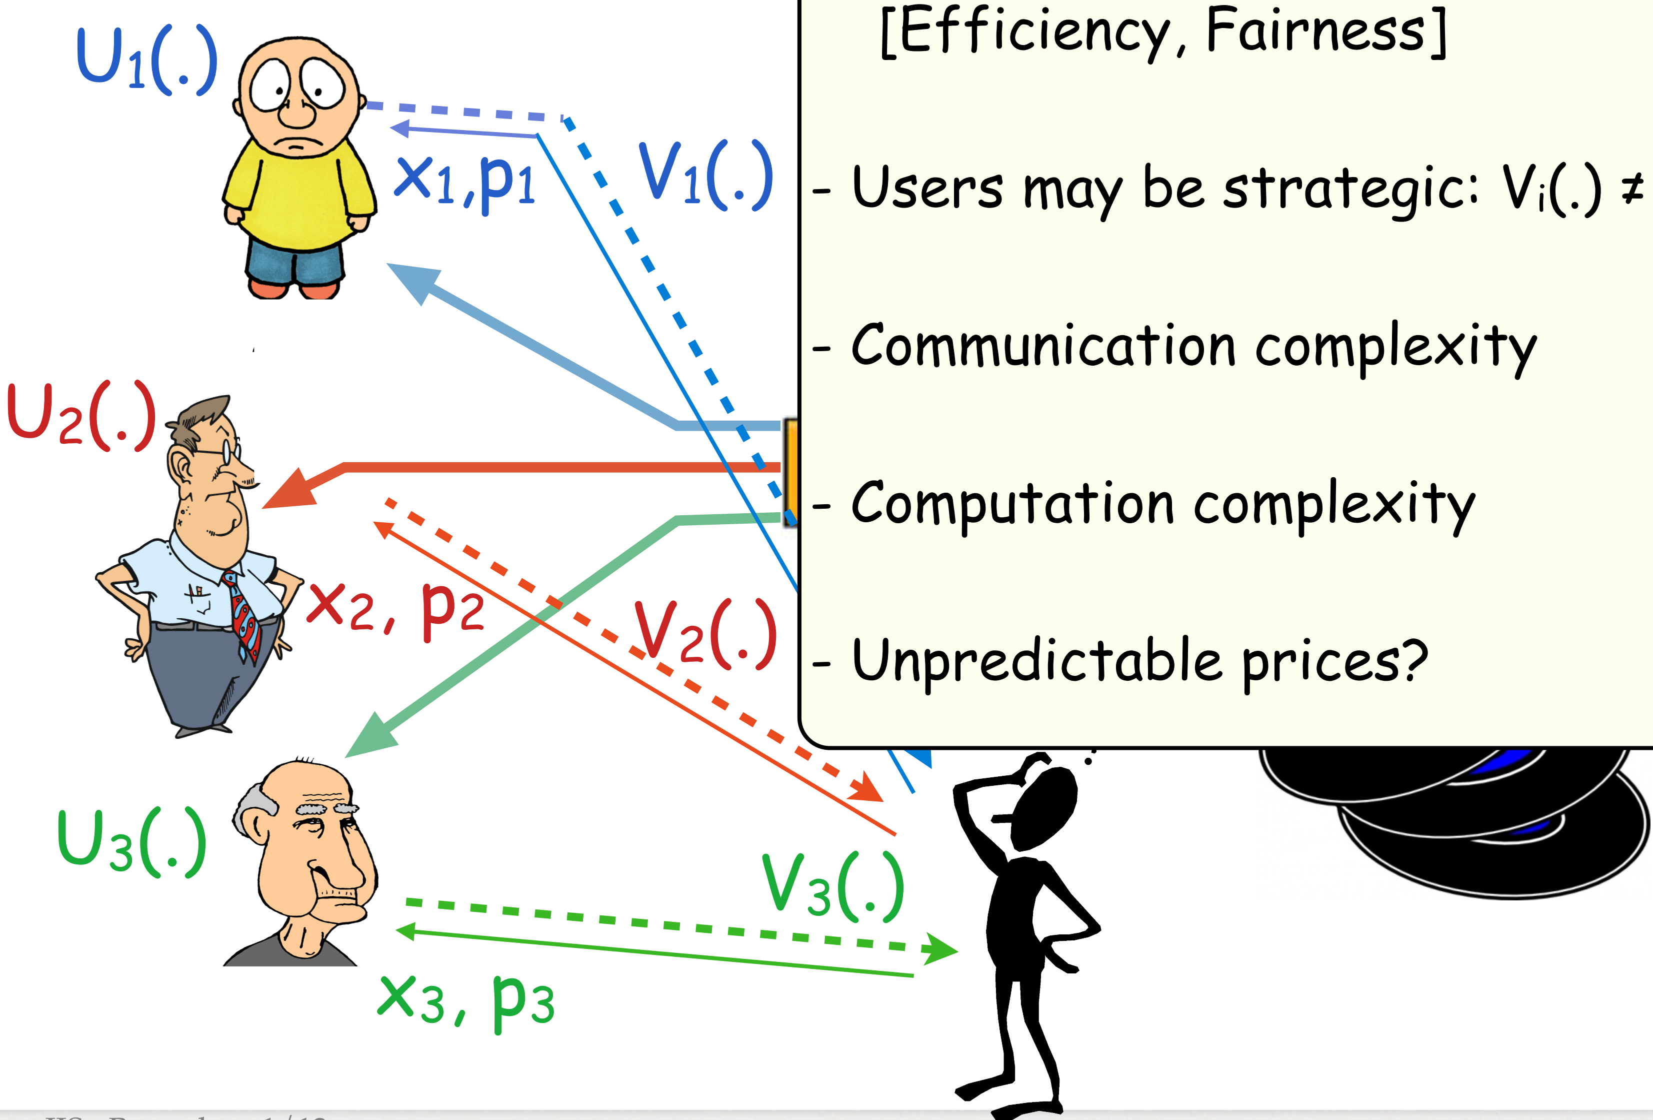

- What is a good allocation?
  [Efficiency, Fairness]
- Users may be strategic: $V_i(.) \neq U_i(.)$
- Communication complexity
- Computation complexity
- Unpredictable prices?

Decentralized:

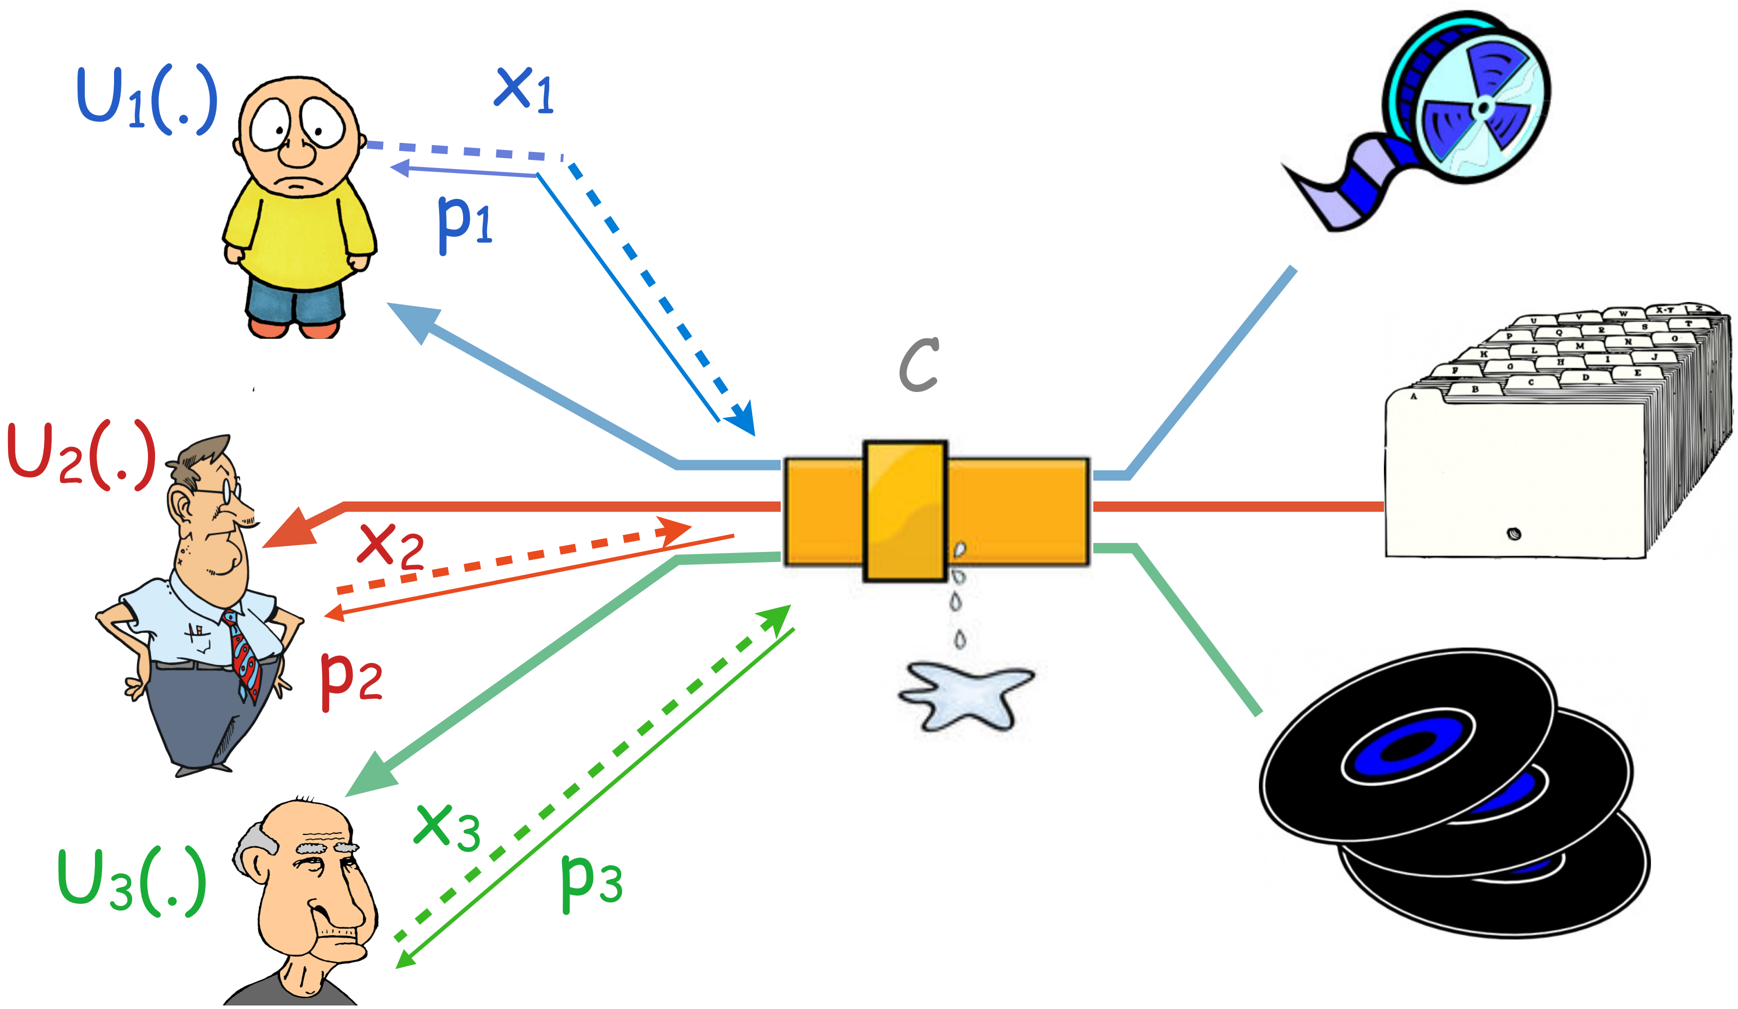

Decentralized:

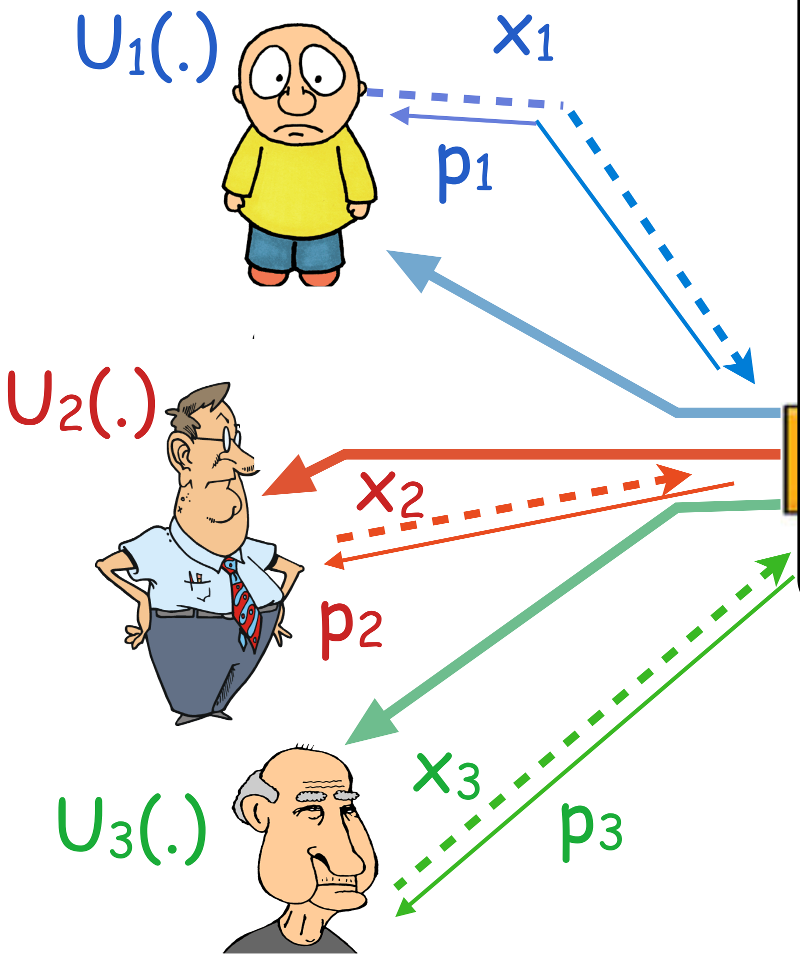

- OK if real prices; if only signals, then not strategy-proof.

- Extends to combinatorial versions

- Stability with delays?

- Convergence speed?

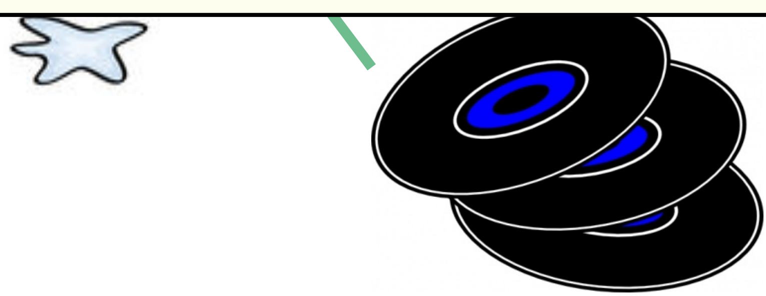

2. UTILITIES & CHOICE - 2.4 SUMMARY

- Users Share Resources
  - Practical issues: demand outstrips supply
  - New types of problems and new capabilities (4G, P2P, Cloud)
- Goal?
  - Satisfy the user preferences: "Impossibility Theorem"
  - Maximize a function of their utilities (e.g., α - fair), or
  - Maximize profit, or
  - Maximize social welfare ....
- Challenges
  - Scalabilitiy (complexity of communication & computation)
  - Strategic users
  - Robustness (stability, convergence speed)

TABLE OF CONTENTS

3. NETWORK ECONOMICS

GOALS

- Understand the interaction of users, providers and network
  - Prices and quality affect demand
  - Demand and investments affects quality (congestion)
  - (Demand)x(Price) affects revenue
  - Revenue affects investments
- Understand the strategic behavior of users and providers
  - Users know that their behavior affects prices
  - Providers compete and try to learn users' willingness to pay

3. NETWORK ECONOMICS - 3.1. OVERVIEW

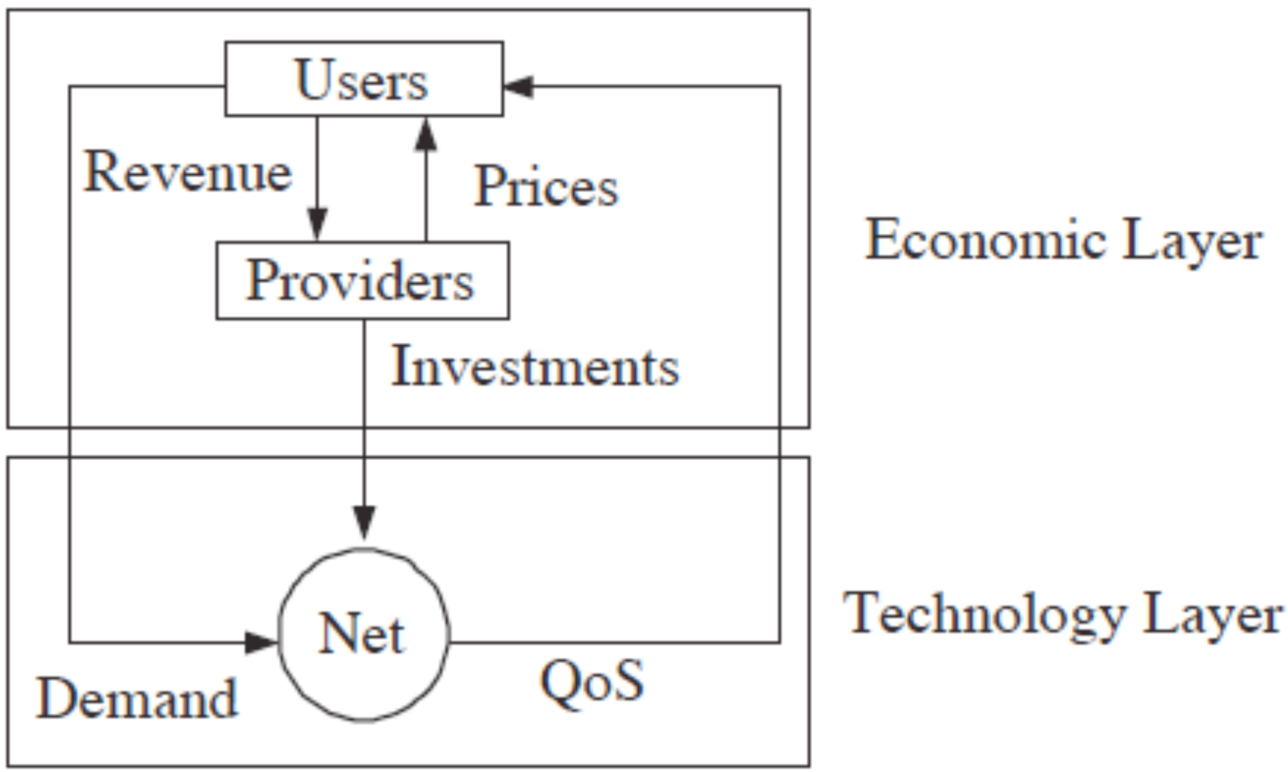

Users and providers respond to economic incentives and affect the network.

[WAL08b]

BASIC EFFECTS

- Network users and providers affect one another's utility
  - The usage of one user degrades the quality of service of other users → Congestion
  - The ability to connect to other users increases the network's value → Network Value
  - Security investments of one user (generally) improves the security of other users → Security Externality
  - Investments by content providers improve the value of the network for users → Content Value
  - Investments by transport providers improve the value of the network for content providers → Transport Value

Definition of Externality:

Externality: When the increase of one agent's utility corresponds to a change in the utility of other agents not reflected in the price

- Positive if increase; negative if decrease

[FEU99]

Consequences of Externality:

Positive Externality often results in free-riding

Negative Externality often results in over-consumption

3. NETWORK ECONOMICS - 3.2. EXTERNALITY

Externality in Networks:

	Users	Content	Transport
Users	+ Contacts - Congestion +/- Security	+ Interest	+ Access
Content	+ Revenue	+ Appeal - Competition	+ Access
Transport	+ Revenue	+ Traffic	+ Access - Competition

[WAL08b]

Tragedy of the Commons:

Utility of one user depends on the choices of other users, through congestion.

Example: n users sharing a network

$$U(x_1) - [x_1 + x_2 + \cdots + x_n]$$

Disutility due to congestion

Utility of user 1 for activity level x_1

Tragedy of the Commons:

$$U(x_1) - [x_1 + x_2 + \cdots + x_n]$$

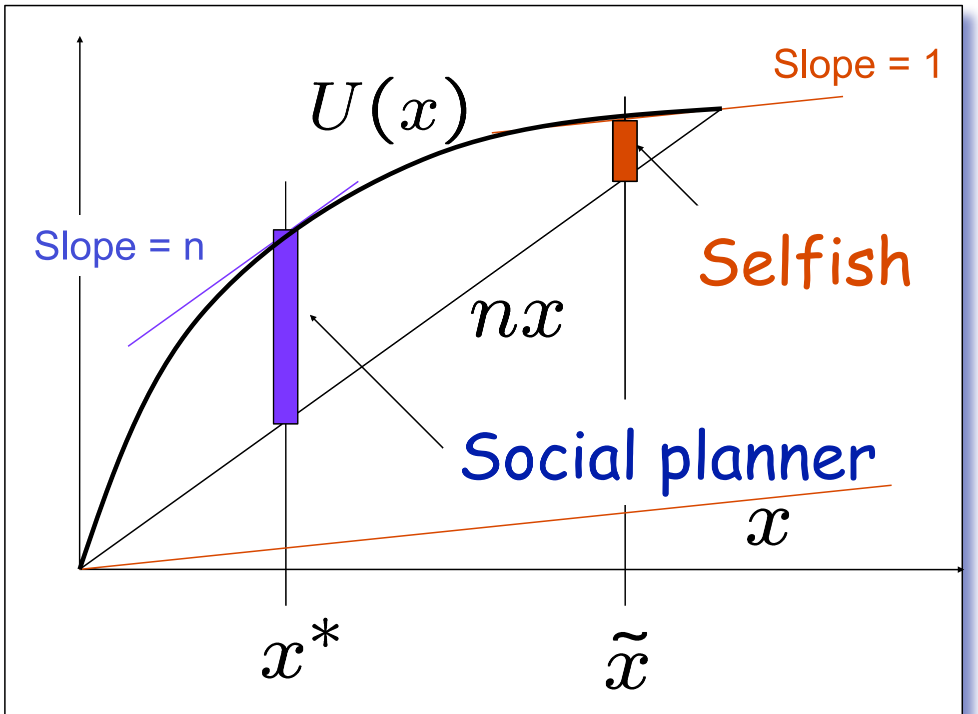

Selfish users over-consume. They neglect their impact on others, and they all hurt each other!

Routing Games:

MAIN POINTS:

- If users "selfishly" choose the fastest available path, all users may end up facing a longer travel time.
- The factor by which the average delay is increased is called the Price of Anarchy.
- It may happen that increasing the capacity of the network increases the travel times (Braess' Paradox)

Routing Games:

Example 1:

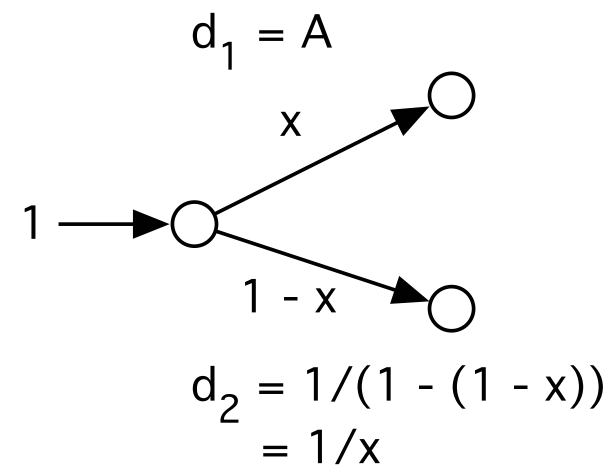

Many small users (non-atomic)

- Selfish choices: d1 = d2 --> x = 1/A, d = A
- Centralized choice: d = min {xA + (1 - x)/x} --> $A = 1/x^2$ --> $x = A^{-0.5}$ --> $d = 2A^{0.5} - 1$
- Hence, PoA = $A/(2A^{0.5} - 1) \to \infty$ as $A \to \infty$
- Thus, PoA is unbounded even in this simple example.

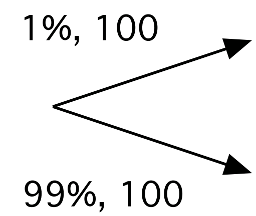

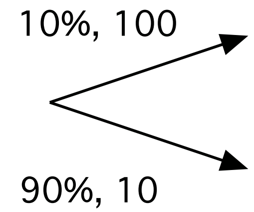

$A = 100$

Routing Games:

Example 2:

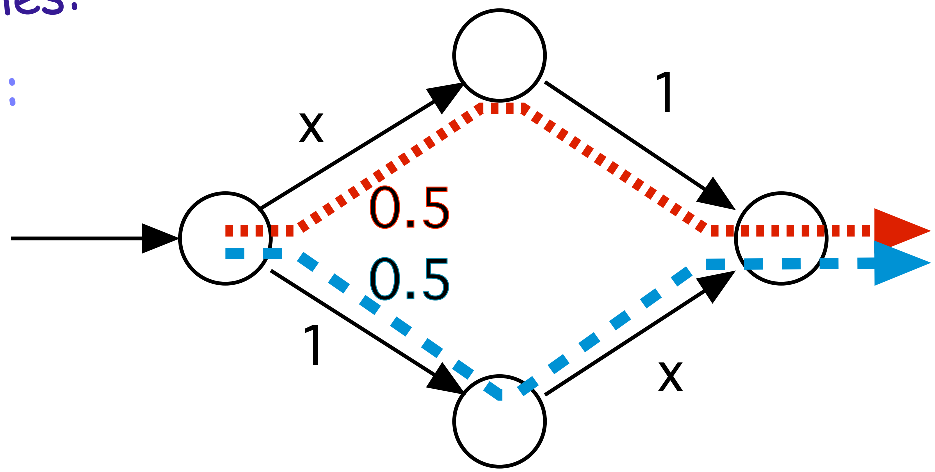

Selfish Choice: 50% choose top path, 50% choose bottom path.

All the users face a delay equal to 1.5.

This is also the social optimum.

PoA = 1

[ROU02]

Routing Games:

Example 2':

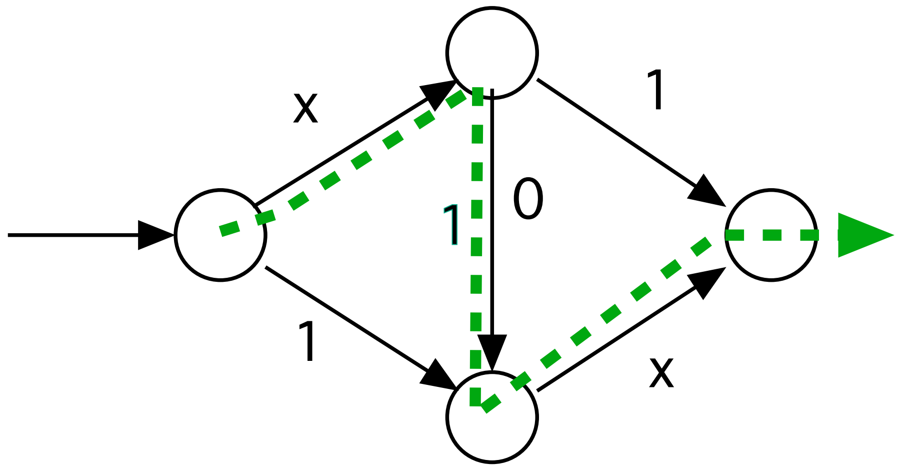

Now one adds a link with zero latency.

Selfish Choice: Users go along indicated path --> delay = 2

Social Optimum: Ignore new link --> delay = 1.5

Price of Anarchy: 2/1.5

Fact: This is a worst case for networks with affine latencies.

[ROU02]

Braess' Paradox:

Adding resources may make the system worse.
In system below, weight goes up after cutting the 3/8 string.

Extension:
1/2 + 3/8 + 1/2
= 1 + 3/8

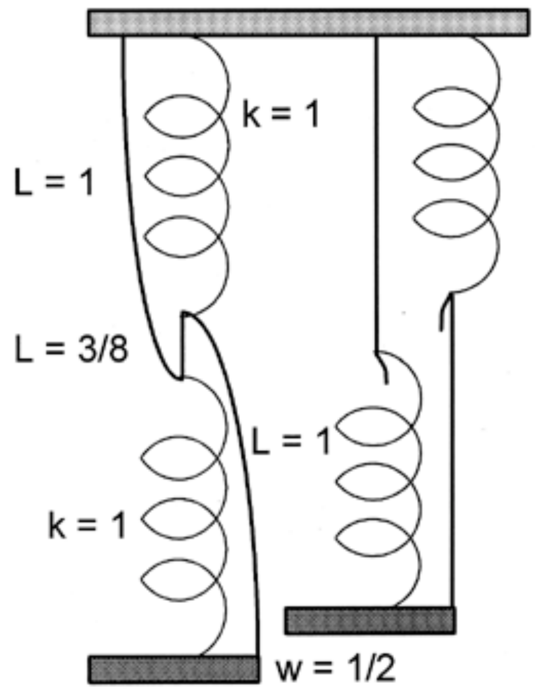

Extension:
1 + 1/4

Example due to Joel Cohen, U. Rochester; illustration from Ivar's Peterson Math Trek

Many Aspects:

- Open access to all
  - content
  - applications
  - devices
- No unreasonable traffic discrimination by
  - source
  - user

Notes:

- Necessary traffic management allowed if rules are disclosed
- Most rules do not apply to wireless access

These rules became law on 11/20/11, but are regularly challenged.

Question: Should ISPs be allowed to charge content providers?

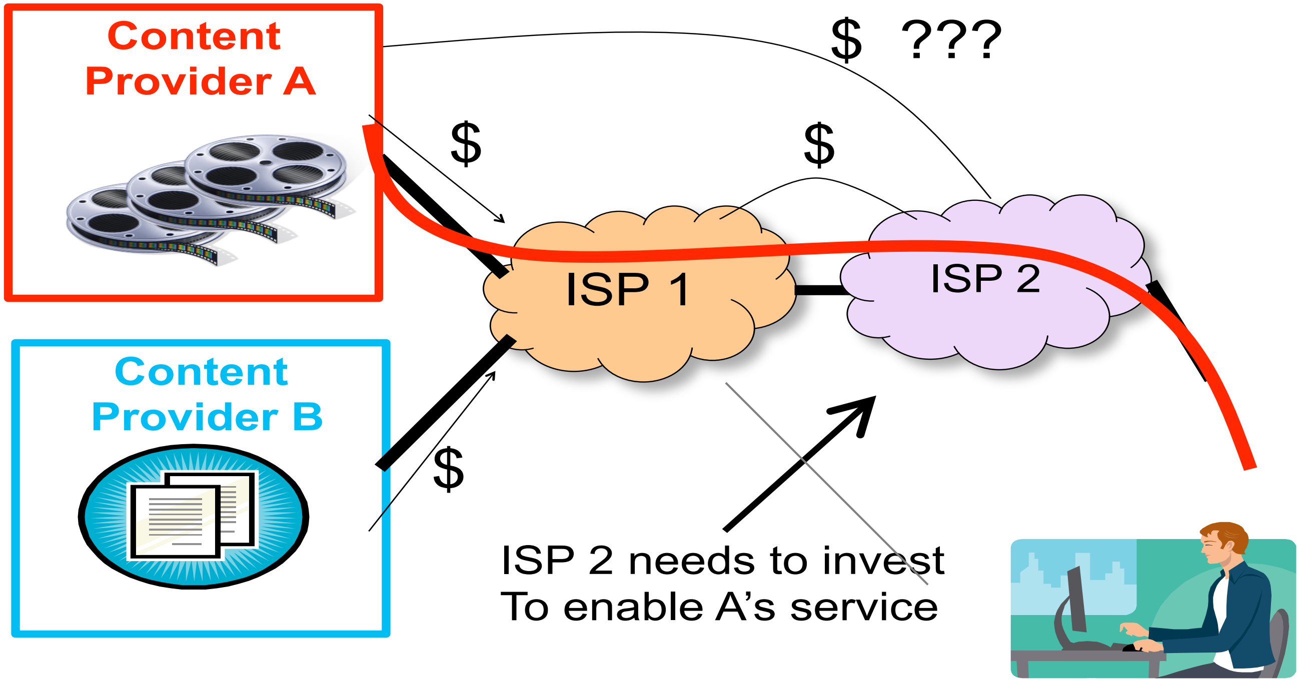

- Would allowing 2 to charge A
  - encourage 2 to invest?
  - discourage A to invest?
- What revenue sharing mechanisms should new Internet have?

This is user discrimination: non-neutral

[MUS08]

Model Structure: Leader-Follower Game

(Content Investment, ISP Investment) --> Usage

Usage --> (Ad Revenue to Content, User Revenue to ISPs)

ISP

NEUTRAL → *Each Chooses:*
- Investment level
- User Price

NOT NEUTRAL →
- Investment level
- User Price
- Price for each content provider

Content Prov.

→ *Each Chooses*
- Investment level

Simple Model:

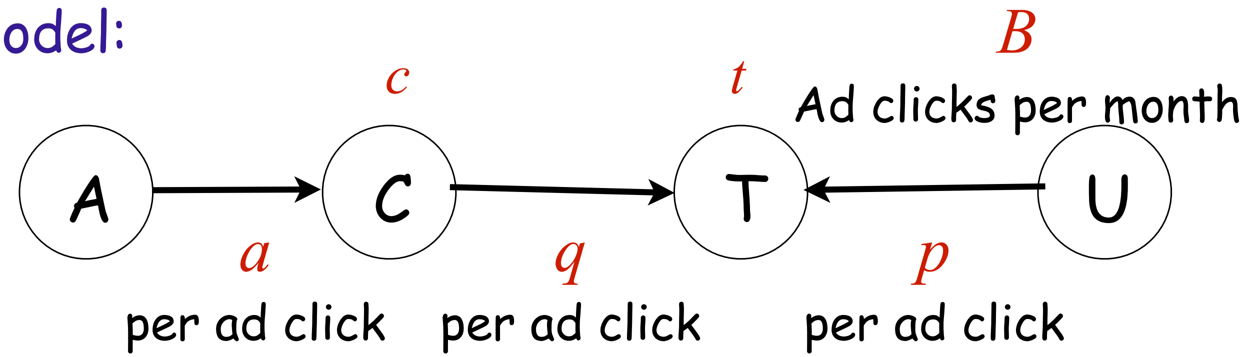

$$B = c^v t^w e^{-p/\theta} \quad 0 < v, w; v + w < 1$$

$$R_C = (a - q)B - \alpha c$$

$$R_T = (q + p)B - \beta t$$

A = advertisers
C = content provider
T = transport provider
U = regular users

Question: $q \uparrow \Rightarrow R_T \uparrow$

or $q \uparrow \Rightarrow R_C \downarrow \Rightarrow c \downarrow \Rightarrow B \downarrow \Rightarrow R_T \downarrow$

Assume T chooses (t, p, q). Then C chooses c to max

$$R_C = (a - q)c^c t^w e^{-p/\theta} - \alpha c$$

Given this c(t, p, q), T then chooses (t, p, q) to max

$$R_B = (p + q)c^c t^w e^{-p/\theta} - \beta t$$

Analysis:

Result 1:

The revenues per click and ROIs are the same under both regimes for content and transport providers.

$$\left(\frac{B(neutral)}{B(non\!-\!neutral)}\right)^{1-v-w}$$

$$\left(\frac{t(neutral)}{t(non\!-\!neutral)}\right)^{1-v-w}$$

$$\left(\frac{R_T(neutral)}{R_T(non\!-\!neutral)}\right)^{1-v-w}$$

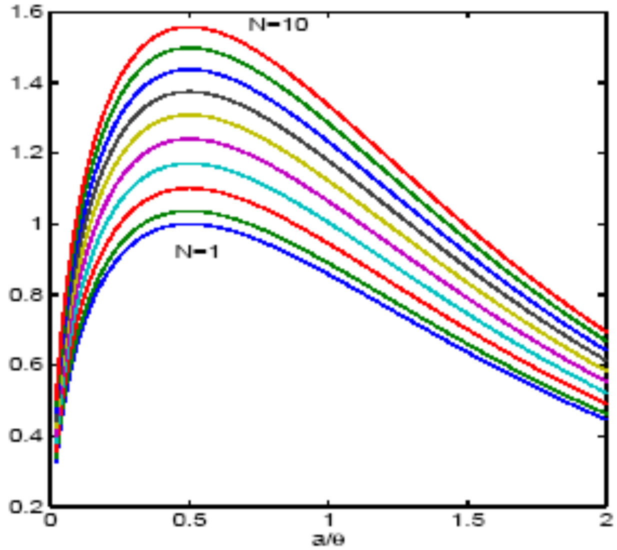

Result 2:

The size of the market is larger in the neutral case only if a/θ is neither very large nor very small.

$$\left(\frac{c(neutral)}{c(non\!-\!neutral)}\right)^{1-v-w}$$

$$\left(\frac{R_C(neutral)}{R_C(non\!-\!neutral)}\right)^{1-v-w}$$

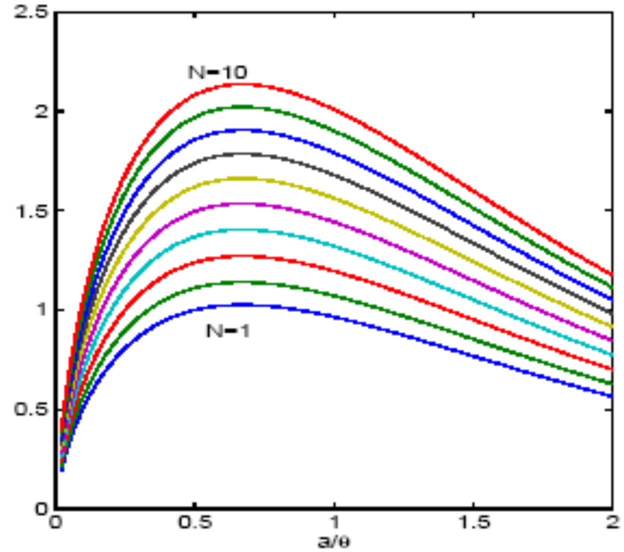

FOUR SECURITY MODELS

Price of Anarchy for Security Investments:

Goals of study:

- How bad is selfish investment?
  - Is regulation necessary?
  - How to improve incentives?
- Important factors that determine POA?
  - Network topology, players' inter-dependency
  - Players' heterogeneous cost functions
  - Strategic-form (one-shot) game or Repeated game

Price of Anarchy for Security Investments:

Model:

- There are n users in the network.
- User i invests x_i in security
- The "security cost" of user i is

$$g_i(\mathbf{x}) = f_i(\mathbf{x}) + c_i x_i, i = 1, 2, \ldots, n$$

- The function f_i -- the risk -- is decreasing and convex
- Thus, positive externality
- We expect free-riding

Assume

$$f_i(\mathbf{x}) = V_i(\sum_{j=1}^{n} \beta_{ji} x_j)$$

$$\beta_{ii} = 1, \beta_{ji} \geq 0$$

$$V_i(.) \geq 0, \downarrow, \text{ convex}$$

[JIA08b]

$$g_i(\mathbf{x}) = V_i(\sum_{j=1}^{n} \beta_{ji} x_i) + c_i x_i$$

Price of Anarchy: $$\rho := \frac{\bar{G}}{G^*} = \frac{\sum_i g_i(\bar{\mathbf{x}})}{\sum_i g_i(\mathbf{x}^*)}$$

$$\bar{\mathbf{x}} = \text{ Nash Equilibrium}; \mathbf{x}^* = \text{ Social Optimum.}$$

Proposition:

$$\rho \leq \max_k \{1 + \sum_{i \neq k} \beta_{ki}\} \qquad \text{(tight)}$$

--> PoA depends on most influential node

Example 1

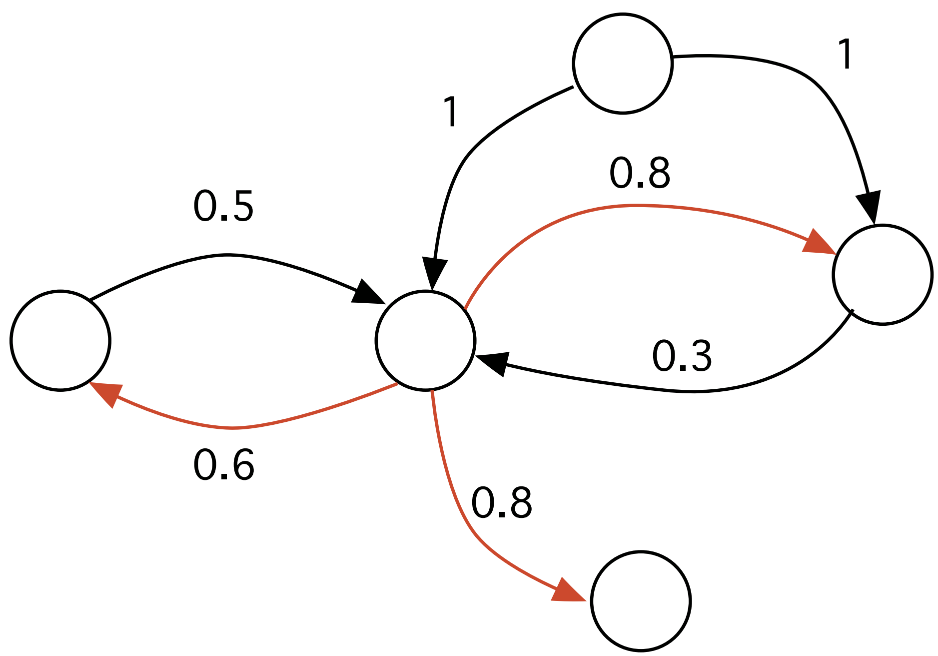

$$\rho \leq 1 + 2.2 = 3.2$$

Example 2

$$\text{If } \beta_{ij} = 1, \forall i, j, \text{ then } \rho \leq n$$

Example 3

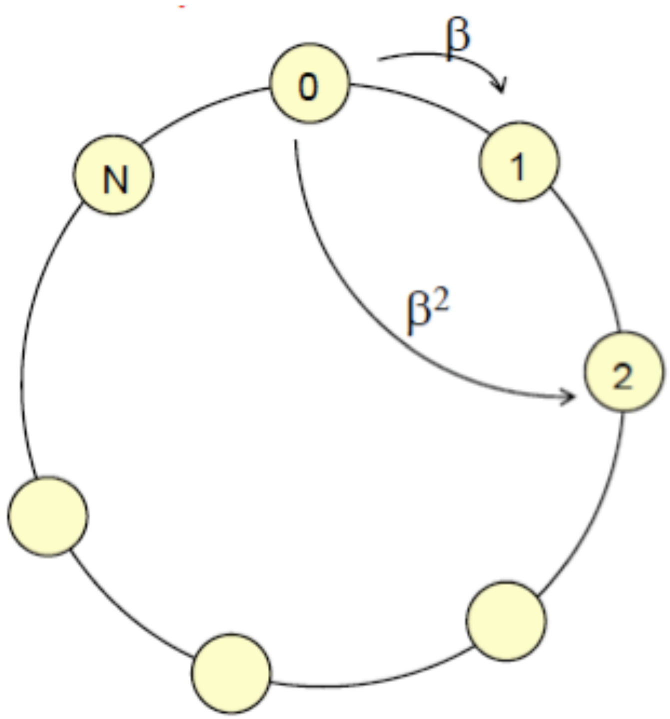

$$V_i(v) = \delta e^{-\lambda v}$$

$$\beta_{ij} = \beta^{d(i,j)}$$

Minimum number of hops

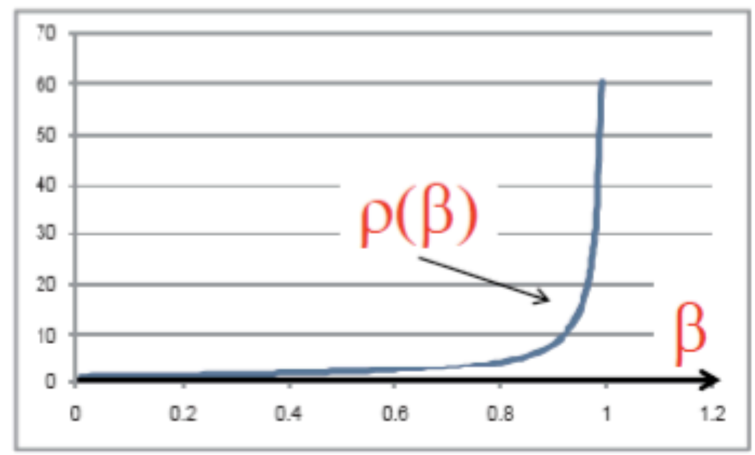

Security as a Game: 1. Intruder Game

Scenario:

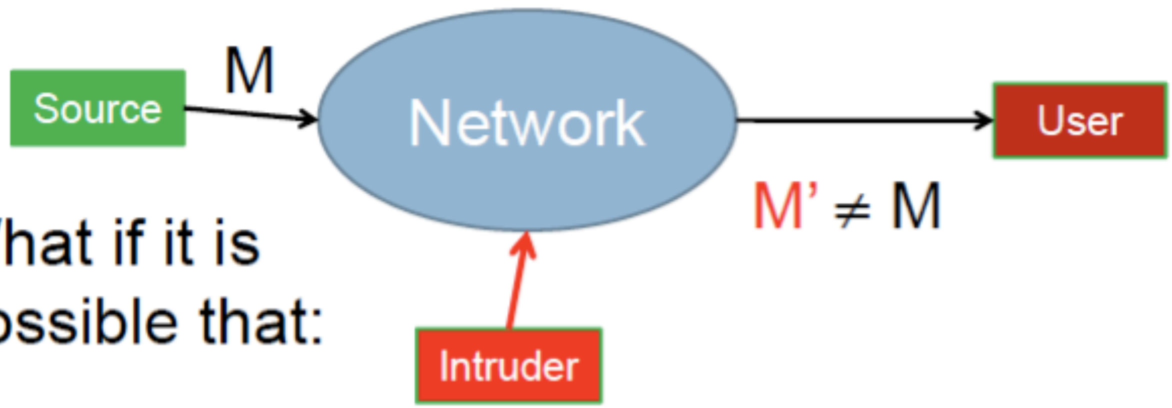

What if it is possible that:

Encryption is not always practical ….

Formulation: Game between Intruder and User

[GUE08]

Security as a Game: 1. Intruder Game

Model

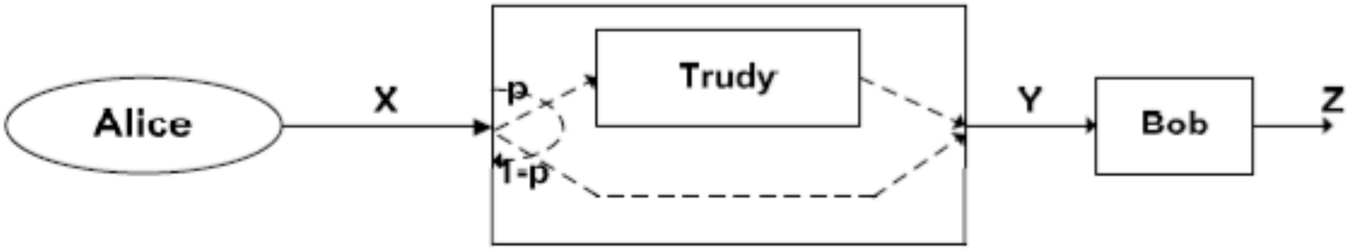

$$X \sim f_X(x) \qquad Y = g(X, W) \qquad Z = h(Y)$$

Goal: $\max_g E[\Phi(X,Z)] \quad \max_h E[\Psi(X,Z)]$

- One-shot Game
  - p and $f_X(x)$ common knowledge
- Two models: with and without challenge
- Nash Equilibria

Security as a Game: 1. Intruder Game

Simple Example:

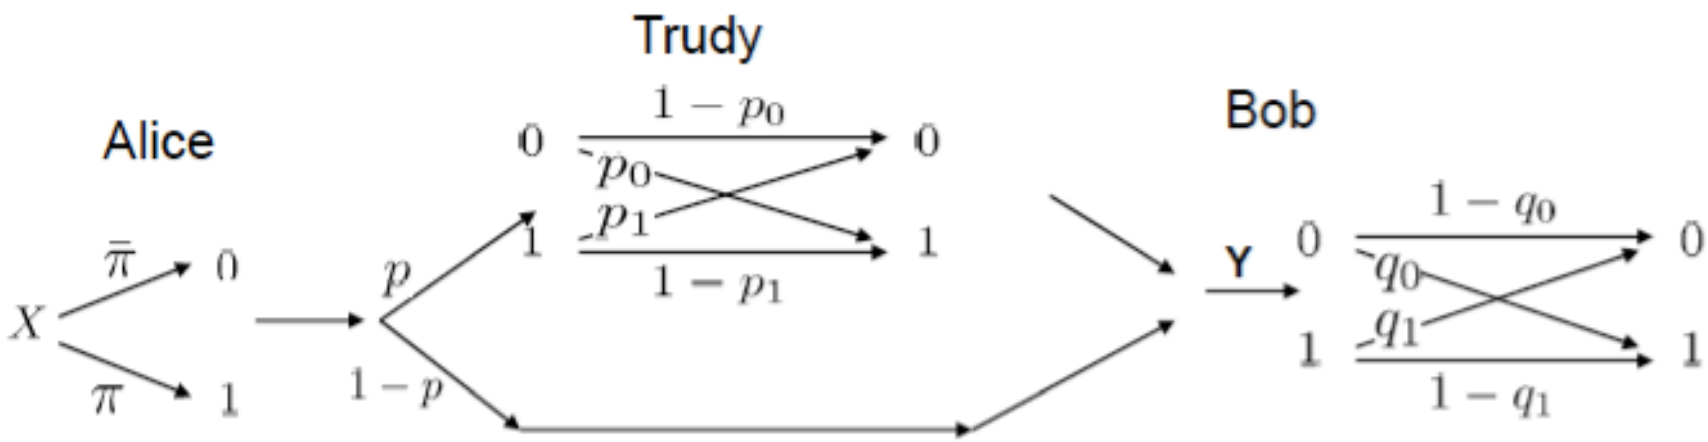

- Payoffs:

$$Trudy: \quad \Phi(Z,X) = A * 1_{(Z=1,X=0)} + B * 1_{(Z=0,X=1)}$$

$$Bob: \quad \Psi(Z,X) = -\Phi(Z,X)$$

- Bayesian, Zero-sum game ⇔ NE=MaxMin

Security as a Game: 1. Intruder Game

Result:

$\bar{\pi}A \leq \pi B$

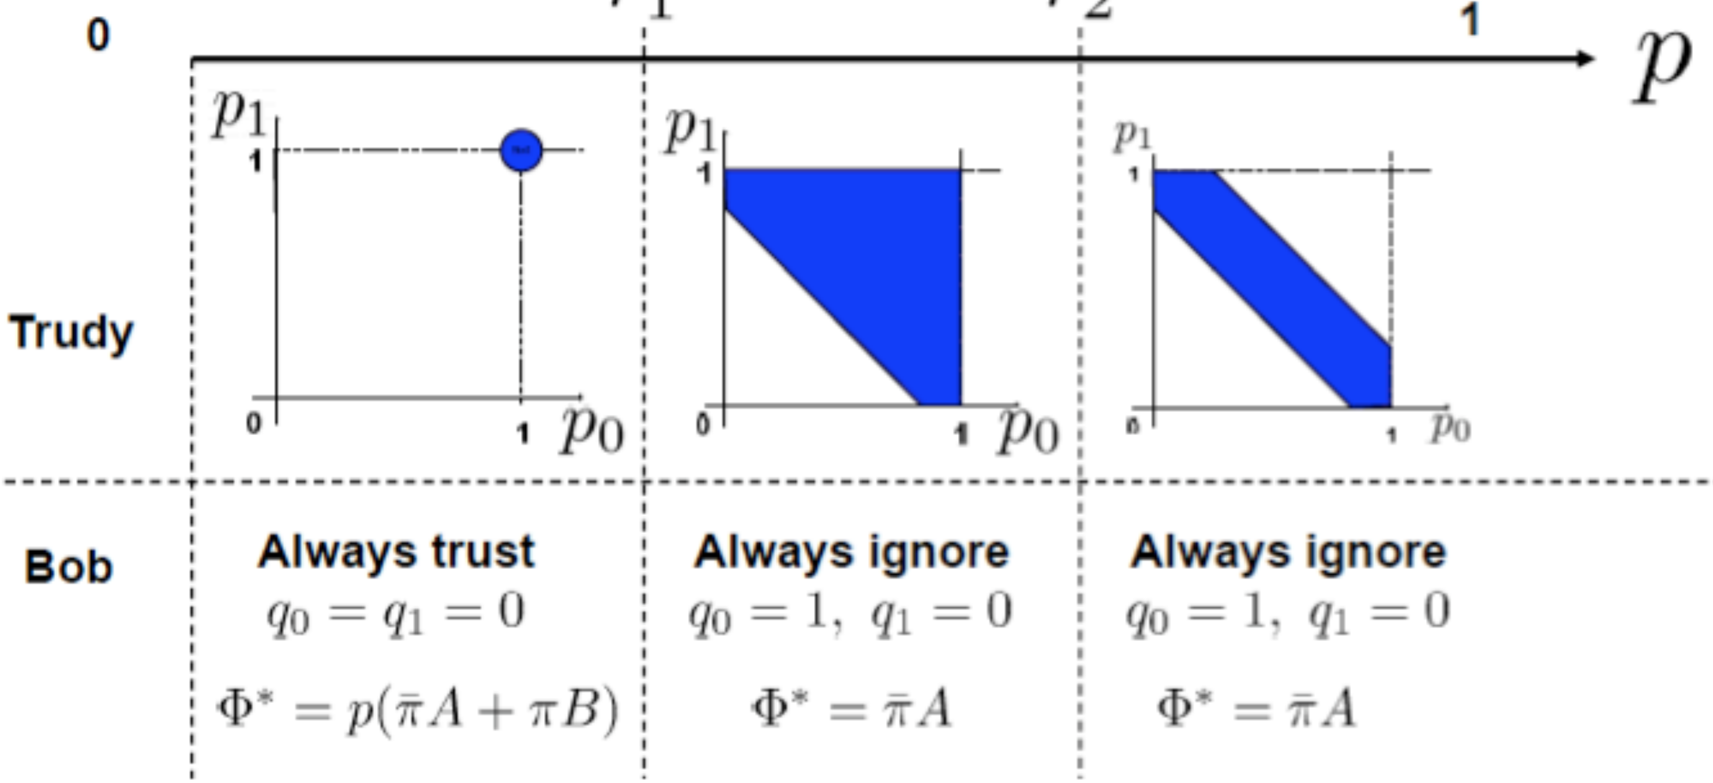

$$r_1 = \frac{\bar{\pi}A}{\bar{\pi}A + \pi B} \qquad r_2 = \frac{\pi B}{\bar{\pi}A + \pi B}$$

Security as a Game: 1. Intruder Game

Challenge:

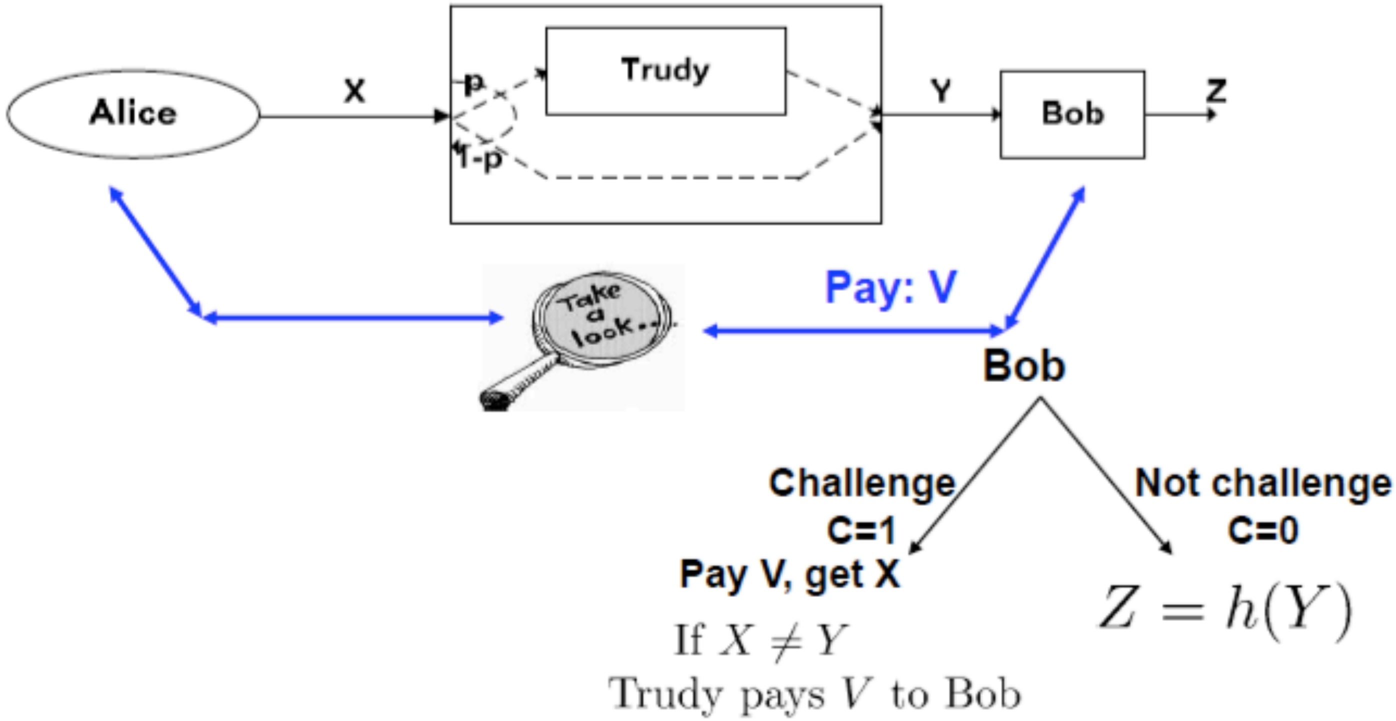

Goal: $\max_g E\left[\Phi(X, Z, C)\right] \quad \max_h E\left[\Psi(X, Z, C)\right]$

- Intrusion Game: NE
  - If P(Intruder present) large, ignore data
  - If challenge is possible, reduces chance of corruption

Security as a Game: 2. Virus Detection

Scenario 1: Intrusion Detection System

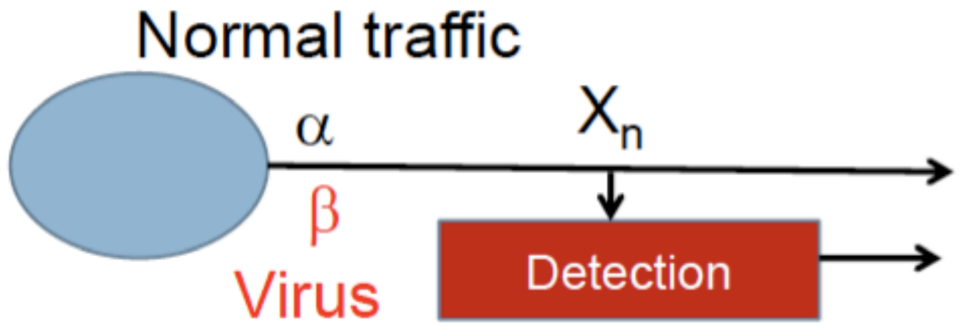

If $X_n > \lambda$ => Alarm, ….

Assume (α, β) known
Choose λ to minimize average cost of infection (spread) + clean up

But, smart designer picks very large β, so that the cost is always high ….

[GUE08]

Security as a Game: 2. Virus Detection

Scenario 2: Intrusion Protection System

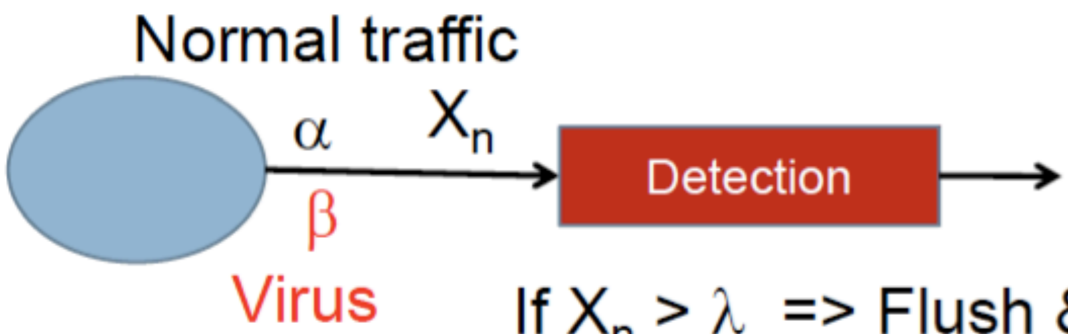

If $X_n > \lambda$ => Flush & Alarm

Now a game between Virus (β) and Detector (λ)

[GUE08]

Security as a Game: 2. Virus Detection

Scenario 2: Model Details

- Observe traffic at discrete time n
  - Normal traffic is Unif(0,α); With Virus Unif(β,α+β)
  - Virus arrives at any time with rate μ, Infect with rate γ

$$X_n > x, \text{ then flag, else continue.}$$
$$p_0 = P(Flag|NoVirus), \quad p_1 = P(Flag|Virus)$$

Markov Chain model:

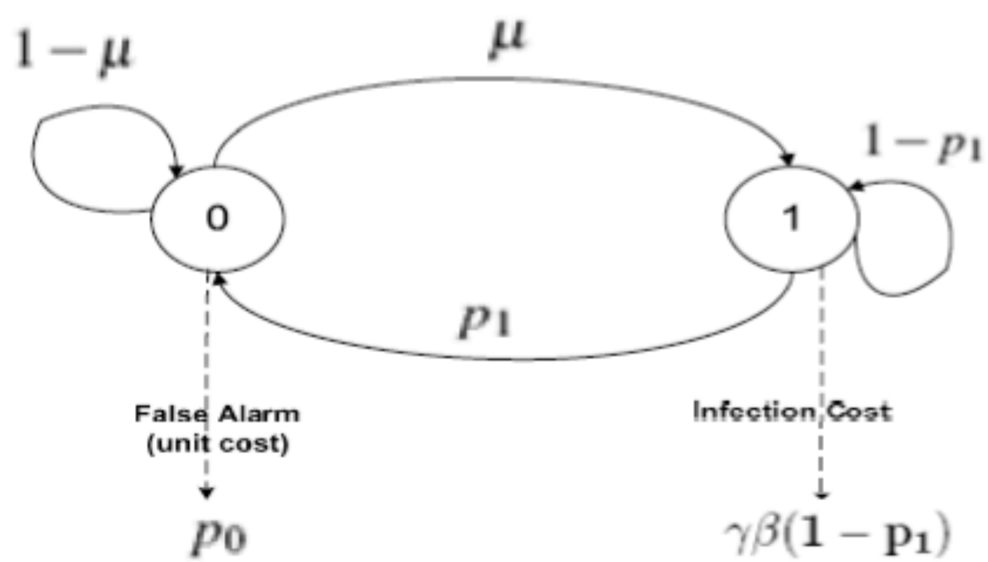

[GUE08]

Security as a Game: 2. Virus Detection

Scenario 2: Results

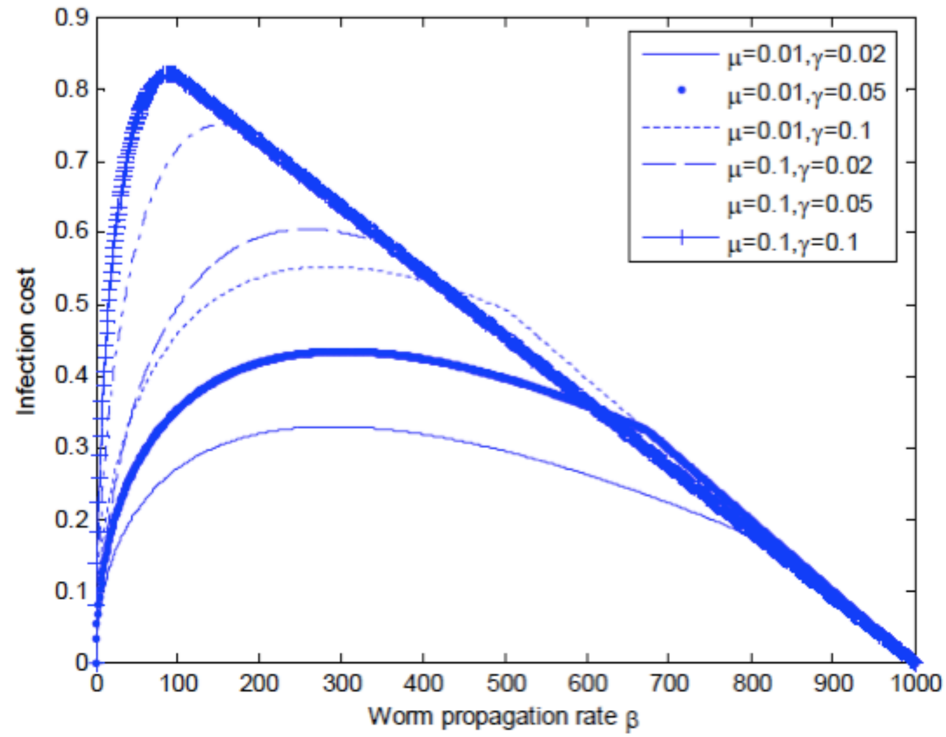

- Virus Design: With IPS (buffer before letting traffic through)
  - Optimal propagation rate!
  - By tuning to that optimal rate, one can guarantee a certain security level/upper bound on cost?

[GUE08]

FOUR SECURITY MODELS

1. POA of Security Investments
2. Intrader
3. Virus
4. Graph Attack

Security as a Game: 3. Graph Attack

Motivation:

- Attack on resources
  - Communication networks
  - Infrastructure
- Modeled as a static game
- Question: Structure of NE
  - Typically PPAD complex
  - When are NE simpler to find?
  - Design: Which set of resources is less vulnerable?
- Result: Class of games with polynomial algorithm

[GWA11]

Security as a Game: 3. Graph Attack

Model:

- S is a set of resources
- Alice uses a subset A of resources, $A \in \underline{A} \subset 2^S$
- Bob attacks a subset B of resources, $B \in \underline{B} \subset 2^S$
- Alice loses $\Lambda(A, B)$
- Bob wins $\Lambda(A, B) - \mu(B)$
- Almost zero-sum
- Mixed strategy: α on A, β on B
- Alice minimizes $E[\Lambda(A, B)]$, Bob maximizes $E[\Lambda(A, B) - \mu(B)]$

Security as a Game: 3. Graph Attack

Example: Cutting a Spanning Tree

- S is the set of links of a graph
- Alice uses any spanning tree *A*
- Bob attacks any single link b
- Alice loses $\Lambda(A, b) = 1\{b \in A\}$
- Bob wins $\Lambda(A, b) - \mu(b)$

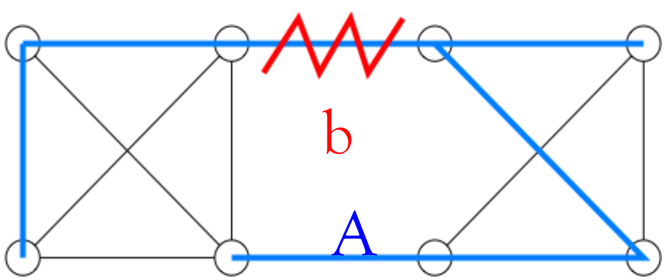

Alice loses 1

Bob wins $1 - \mu(b)$

Note: > 700 spanning trees

Security as a Game: 3. Graph Attack

Example: Cutting a Spanning Tree

DEFINITION The set E of links is critical if it maximizes the vulnerability

$$V(E) := \min \{ |E \cap A| \text{ s.t. } A = \text{spanning tree}\}/|E|$$

THEOREM

1) The following is a NE: Let E be critical;

Bob attacks the links b of E with equal probabilities

Alice chooses the spanning tree A at random among those that have a minimal intersection with E, so that

$$P(e \in A) = V(E), \text{ for all } e \in E$$
$$\leq V(E), \text{ for all } e \notin E$$

2) There is a polynomial algorithm to find E

Security as a Game: 3. Graph Attack

Example: Cutting a Spanning Tree

For some graphs, E is not a minimum cut-set:

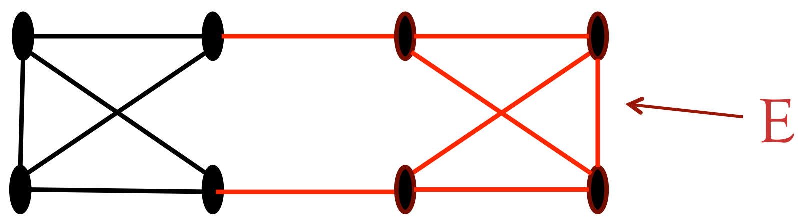

$$V(E) = \frac{\min\{\ |E \cap A| \text{ s.t. } A = \text{spanning tree}\}}{|E|} = 4/7 > 1/2$$

E depends on μ:

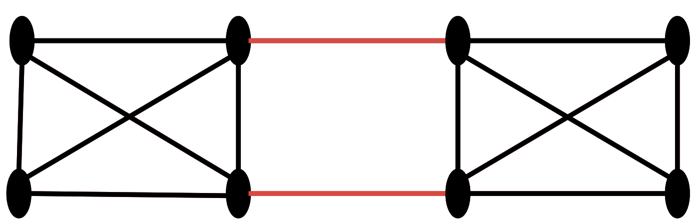

μ (b) = 1, all b

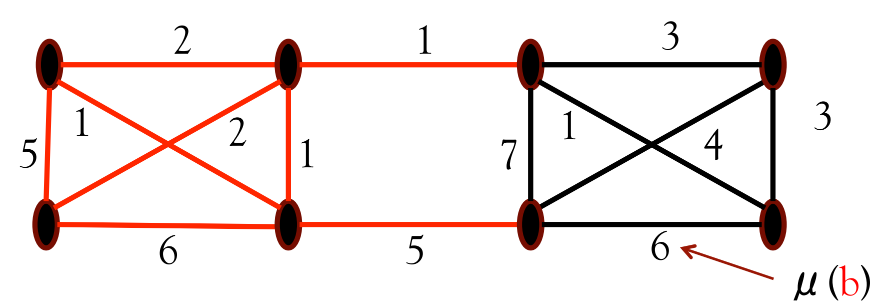

Security as a Game: 3. Graph Attack

Proof technique:

Structure of NE:

Characterization of Spanning Tree Polyhedron

Polynomial Algorithm:

Maximization of supermodular function

Security as a Game: 3. Graph Attack

Structure of NE:

Consider spanning tree problem with $\mu = 0$

If there is a choice $\underline{\alpha}$ such that

$P(e \in A) = V(E)$, for all $e \in E$

$\leq V(E)$, for all $e \notin E$

then, this is a NE (with a uniform attack on E)

$\leftrightarrow$ $\underline{\alpha}\, \Lambda \leq \underline{w}$, $\underline{w} := V(E)\underline{1}$

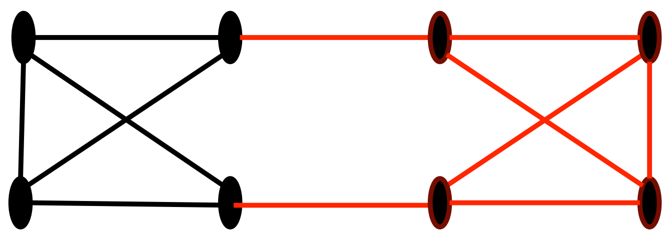

4/7

Indeed: Alice cannot do better: $P(e \in A) = |E \cap A|/|E| \geq V(E)$

Bob cannot either: $P(e \in A)$ is maximized for $e \in E$

Security as a Game: 3. Graph Attack

Structure of NE:

Key: Show $\max\{\underline{a}.\underline{1} \mid \underline{a} \geq \underline{0},\ \underline{a}\Lambda \leq \underline{w}\} \geq 1$

Then, let $\underline{\alpha} = \underline{a}/(\underline{a}.\underline{1})$ where $\underline{a}$ achieves max. [Indeed $\underline{\alpha}\Lambda \leq \underline{w}$]

Definitions: Let Λ be a ≥ 0 matrix.

$P(\Lambda)$ = co{ rows of Λ} + R_+^n = polyhedron of Λ

$P(\Lambda)^* = \{ \underline{x} \geq 0 \mid \underline{y}.\underline{x} \geq 1, \forall \underline{y} \in P(\Lambda)\}$ = blocker of $P(\Lambda)$

$P(\Lambda)^* = P(\Lambda^*)$ for some matrix $\Lambda^* \geq 0$

Theorem 1: $(P(\Lambda)^*)^* = P(\Lambda)$

Theorem 2 (Fulkerson '71):
If Λ = incidence matrix of spanning trees, then

$P(\Lambda) = \{\underline{x} \geq 0 \mid x(F) \geq |F| - 1, \forall F \leftrightarrow$ proper partition of graph$\}$

where $x(F) = \sum_{e \in F} x(e)$

Security as a Game: 3. Graph Attack

Structure of NE:

Show $\max \{ \underline{a}.\underline{1} \mid \underline{a} \geq \underline{0},\ \underline{a}\Lambda \leq \underline{w}\} \geq 1$

Argument: Dual is $\min \{\underline{w}.\underline{1} \mid \underline{y} \geq \underline{0},\ \underline{y}\Lambda \geq \underline{1}\}$

Now, $\{ \underline{y} \geq \underline{0},\ \underline{y}\Lambda \geq \underline{1}\} = P(\Lambda)^*$

Also, by T1,

$\underline{w}.\underline{1} \geq 1$ for all $\underline{y}$ in $P(\Lambda)^*$ iff $\underline{w} \in (P(\Lambda)^*)^* = P(\Lambda)$

i.e., by T2, iff $\underline{w}$ is such that

$$w(F) \geq |F| - 1,\ \forall F \leftrightarrow \text{proper partition of graph}$$

But, since E is critical,

$$w(F) = V(E)\underline{1}(F) = V(E)\,|F| \geq |F| - 1 = \min \{|A \cap F|,\ A = \text{S.T.}\}$$

Security as a Game: 3. Graph Attack

Submodular Function:

Key observation (Cunningham '85):

$$f(E) = \min\{|A \cap E| \text{ s.t } A = \text{spanning tree}\}$$

is <u>supermodular</u>.

That is, $f(E + e + f) - f(E + e) \geq f(E + f) - f(E)$

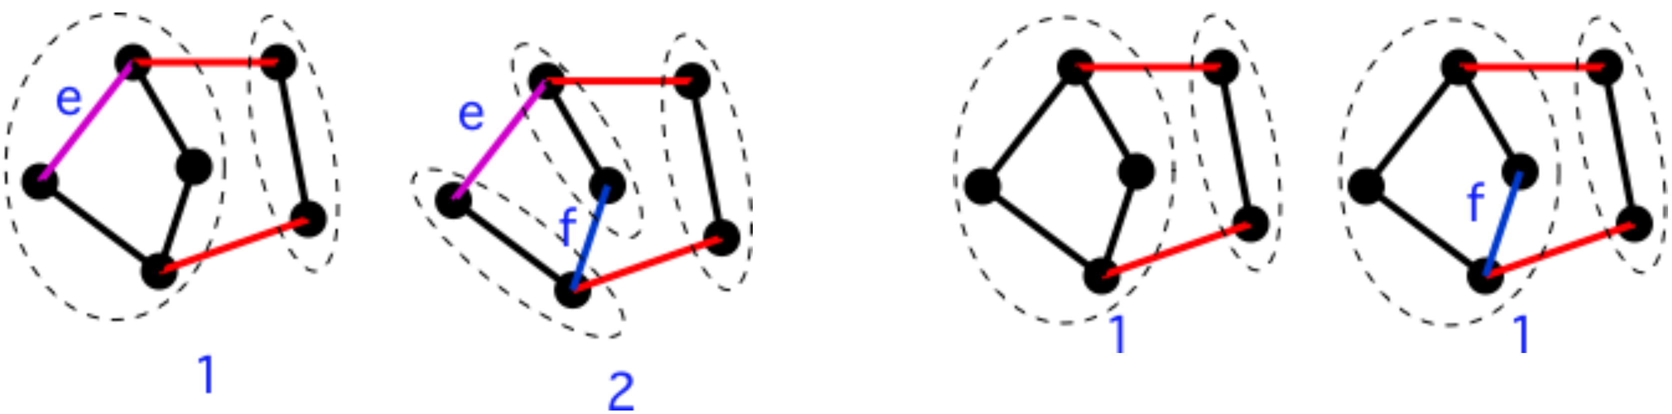

Security as a Game: 3. Graph Attack

Summary:

Resource attack games:

$\min_{\alpha} \alpha \Lambda \beta$; $\max_{\beta} \alpha \Lambda \beta - \mu . \beta$

Nash Equilibrium characterized by extreme points of P(Λ)*

For Spanning Tree Game,

Attack is uniform on a critical set

Critical set found by minimizing a supermodular function:

polynomial

Vulnerability of a graph --> guidelines for design

[GWA11]

When should an operator upgrade from 3G to 4G?

Tradeoff:

- Early upgrade
  - is more expensive (cost decreases)
  - starts with few users
- Late upgrade
  - results in smaller market share
  - delays 4G revenues

Model:

- Timing game between provider
- Model must capture
  - decreasing cost of technology
  - dynamic of users adopting 4G and switching providers

[DHW12]

3. NETWORK ECONOMICS - 3.6. UPGRADES

When should an operator upgrade from 3G to 4G?

Model: Customer migrations

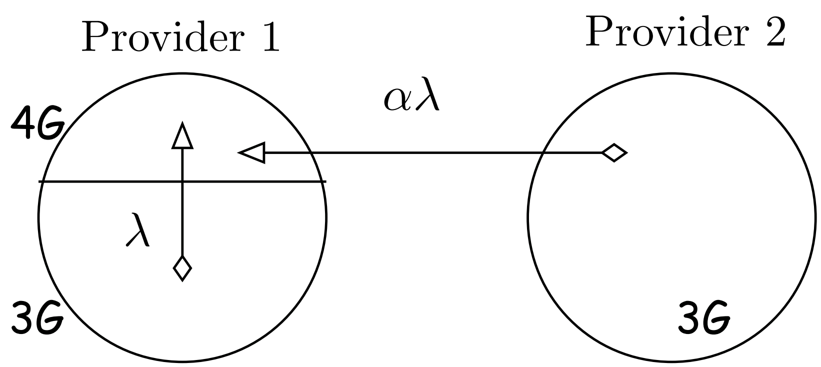

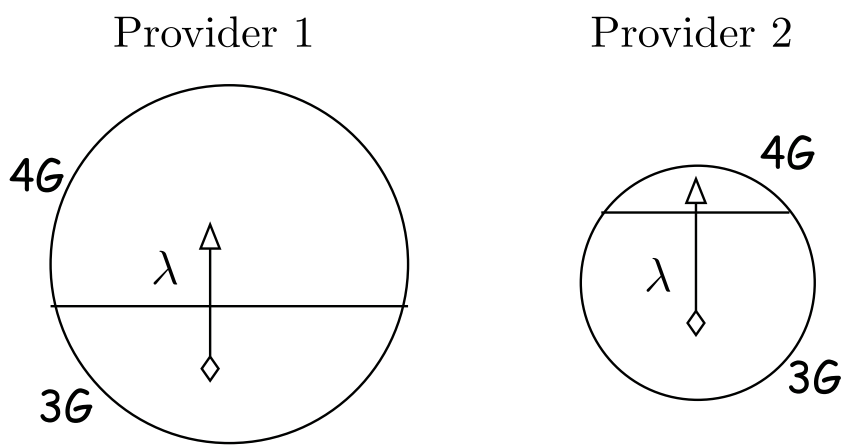

Customers of one provider upgrade to 4G at rate λ.
Customers switch providers to get 4G, at rate $\alpha\lambda$, $\alpha < 1$.

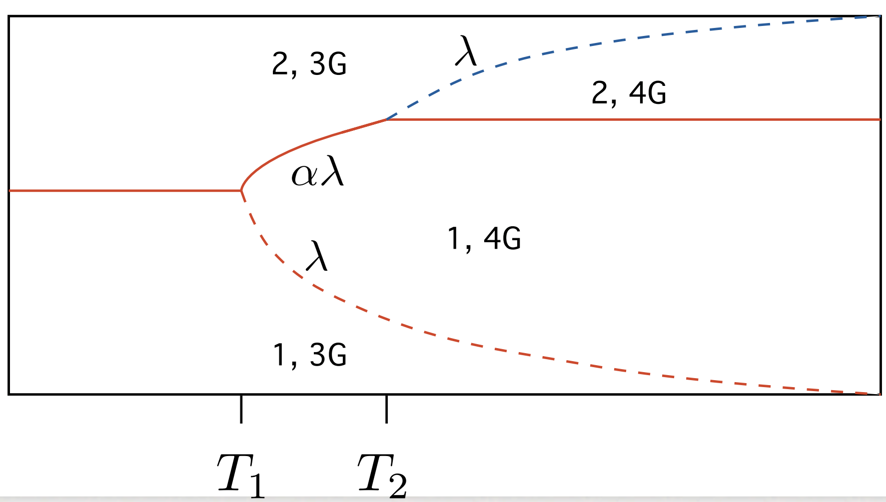

When should an operator upgrade from 3G to 4G?

Model: Calls

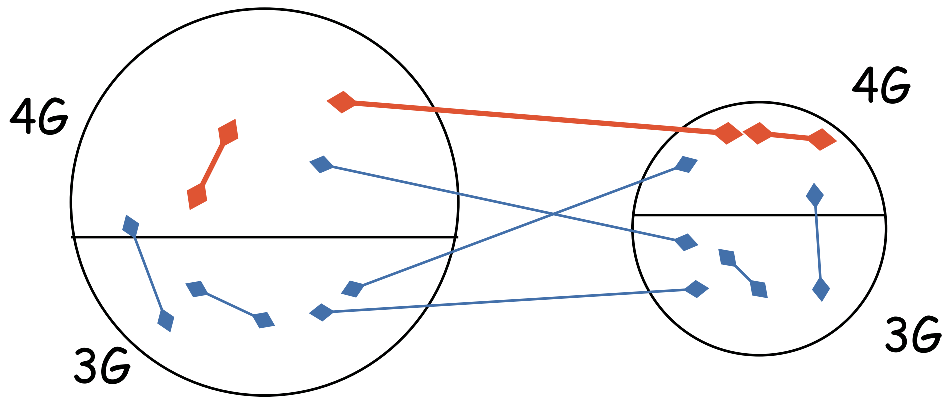

Each user generates calls at rate ρ.

A fraction A/N are to users in a set with A users.

The rate of calls from a set with A users to a set with B users is

$$\rho \frac{AB}{N} = AB \text{ (after normalizing)}$$

When should an operator upgrade from 3G to 4G?

Model: Revenue

4G calls cost 1

Other calls cost $\gamma < 1$

$$R_i(t) = \gamma N_t^i N + (1-\gamma) N_t^{i*} N_t^*$$

$N_t^i :=$ number of users of provider i
$N_t^{i*} :=$ number of 4G users of provider i
$N_t^* := N_t^{1*} + N_t^{2*}$

Profit

$$\pi_i = \int_0^\infty e^{-\beta t} R_i(t) dt - K e^{-U T_i}$$

$$U > \beta + \alpha\lambda$$

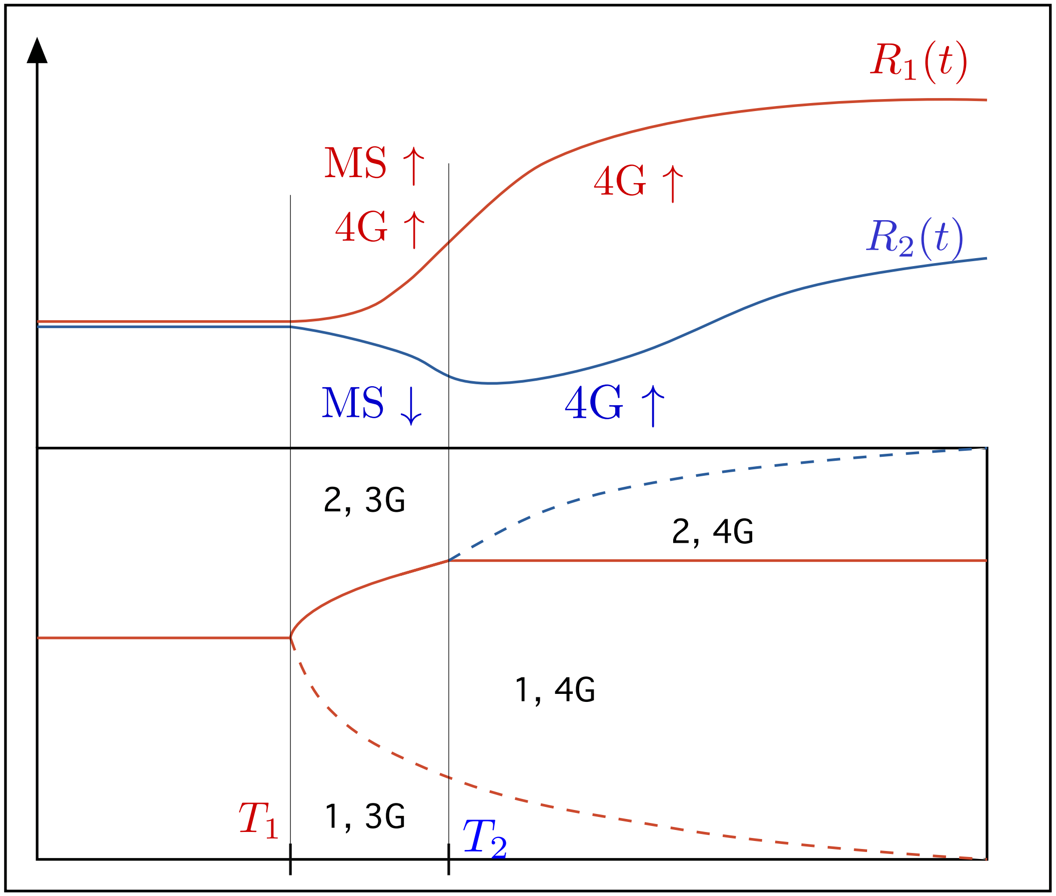

U = decrease rate of technology cost
β = discounting rate

When should an operator upgrade from 3G to 4G?

Timing Game:

User i chooses T_i to maximize $\pi_i(T_1, T_2)$.

Results:

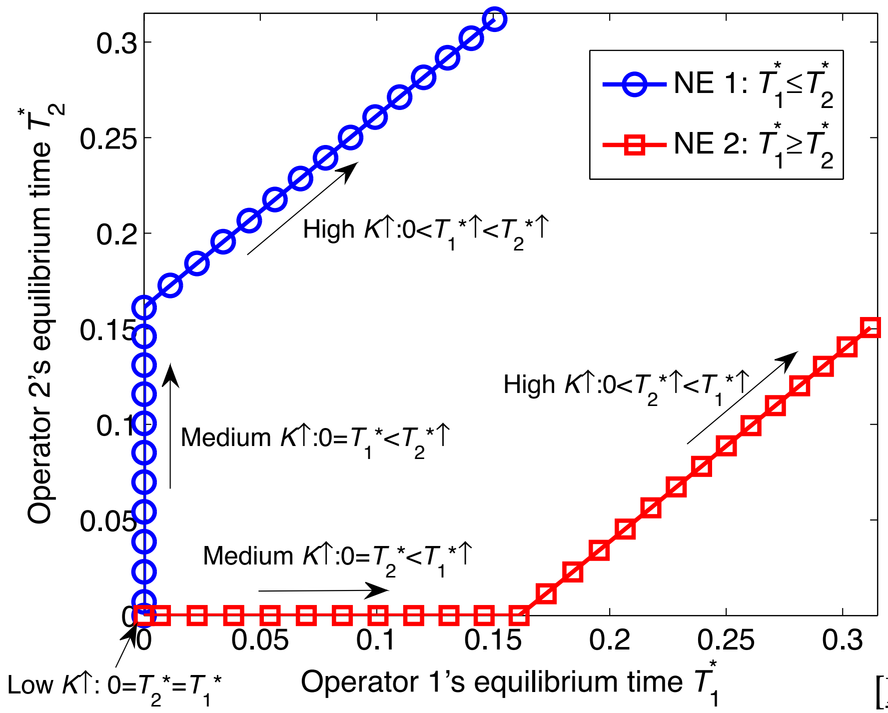

[DHW12]

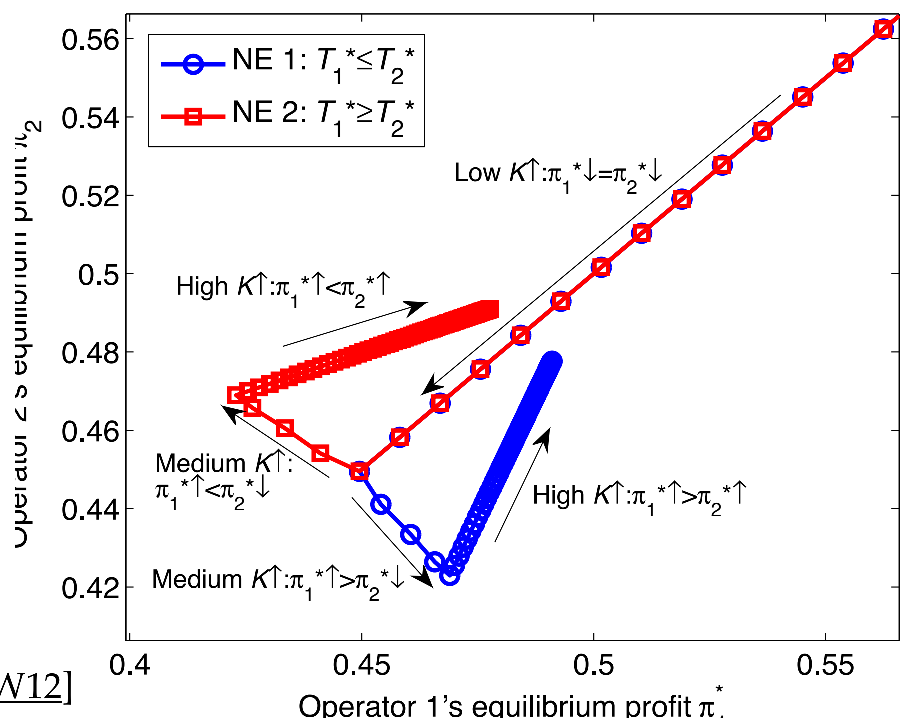

- Externality reflects the impact of a user on other users. Positive externality results in free-riding.

- The price of anarchy is the reduction in social welfare caused by selfish agents.
  - Examples: Routing (Braess), Security Investments

- Suitable pricing results in socially optimal behavior
  - Examples: VCG, Congestion Pricing, Sharing Revenue

- Security is a game between attackers and defenders.
  - The strategic attacker anticipates the defender's strategy.

- Upgrade decisions are modeled as a timing game.

TABLE OF CONTENTS

4. DISTRIBUTED ALGORITHMS

4. DISTRIBUTED ALGORITHMS 4.1 OVERVIEW

- Users share resources: links, cloud computers, data centers, spectrum
- Everything is distributed: decision makers, ownership, users, resources

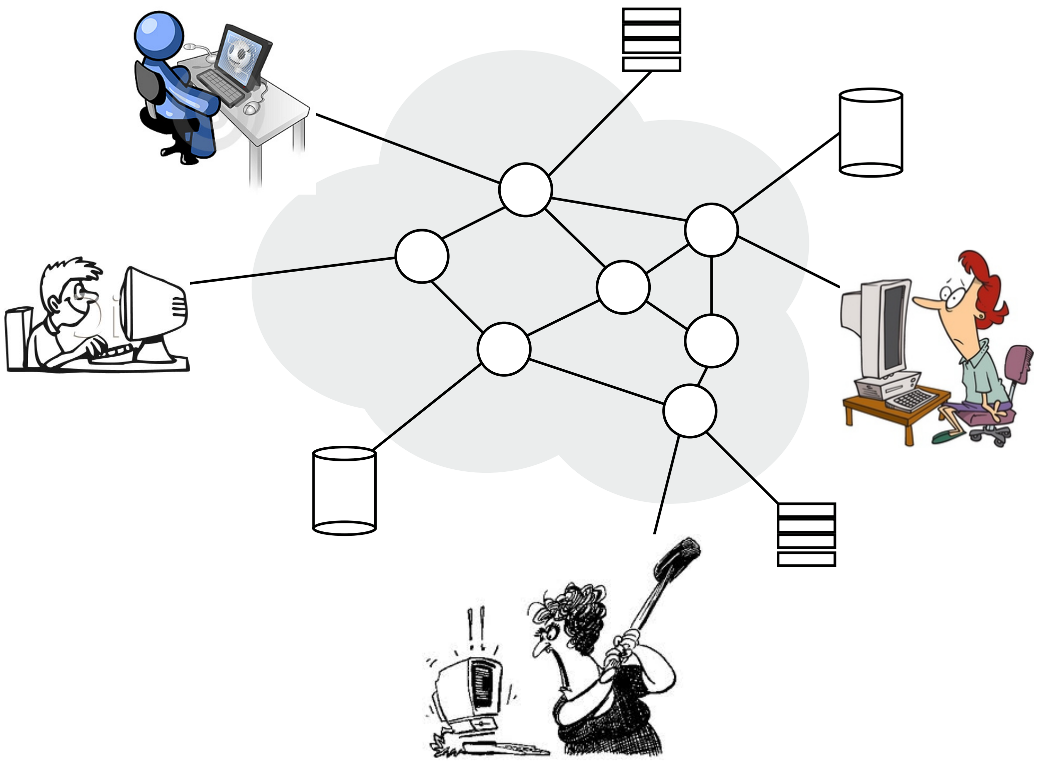

- Goals:
  - Distributed algorithms that "work": Stable network
  - Hopefully, that work well: Efficient and fair

Basic Ideas:

- Formulate goal as an optimization problem with constraints:

  Maximize utility of users
  Subject to utilization ≤ capacity

- Decouple problem by considering the dual:

  Maximize utility of users - shadow price x (utilization - capacity)
  ⇒ individual problem for each user: Maximize utility - price

- Compute shadow price with gradient algorithm:

  Δ(shadow price) = utilization - capacity
  ⇒ shadow price ≅ backlog at resource: decomposed

- Extension to indivisible resources

History: Arrow-Debreu (54), Arrow-Hurwitz (58), Kelly (97), Low (99), ...

TCP: Transmission Control Protocol

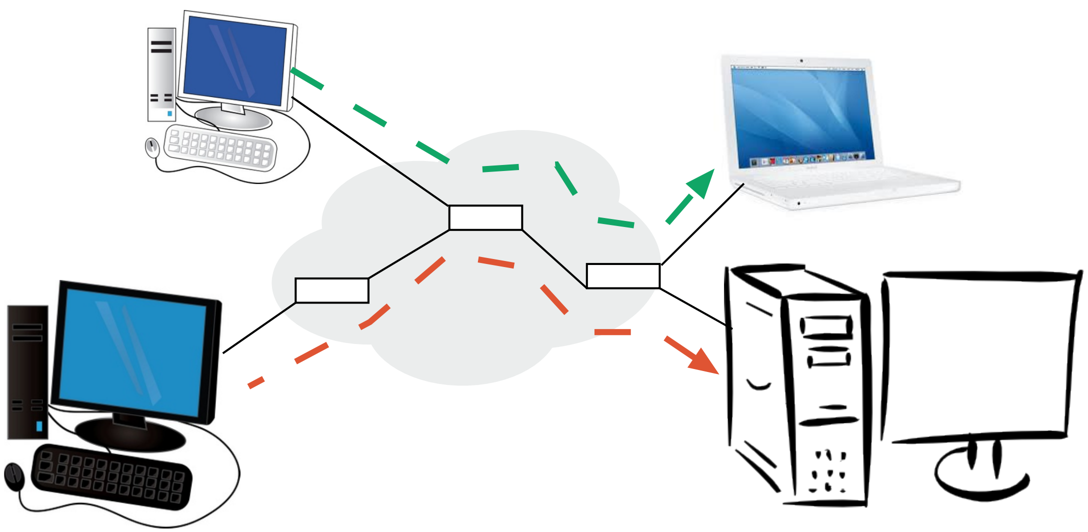

Packets

Packets

Problem: How to control the transmissions to share the network links fairly and efficiently?

TCP: Transmission Control Protocol

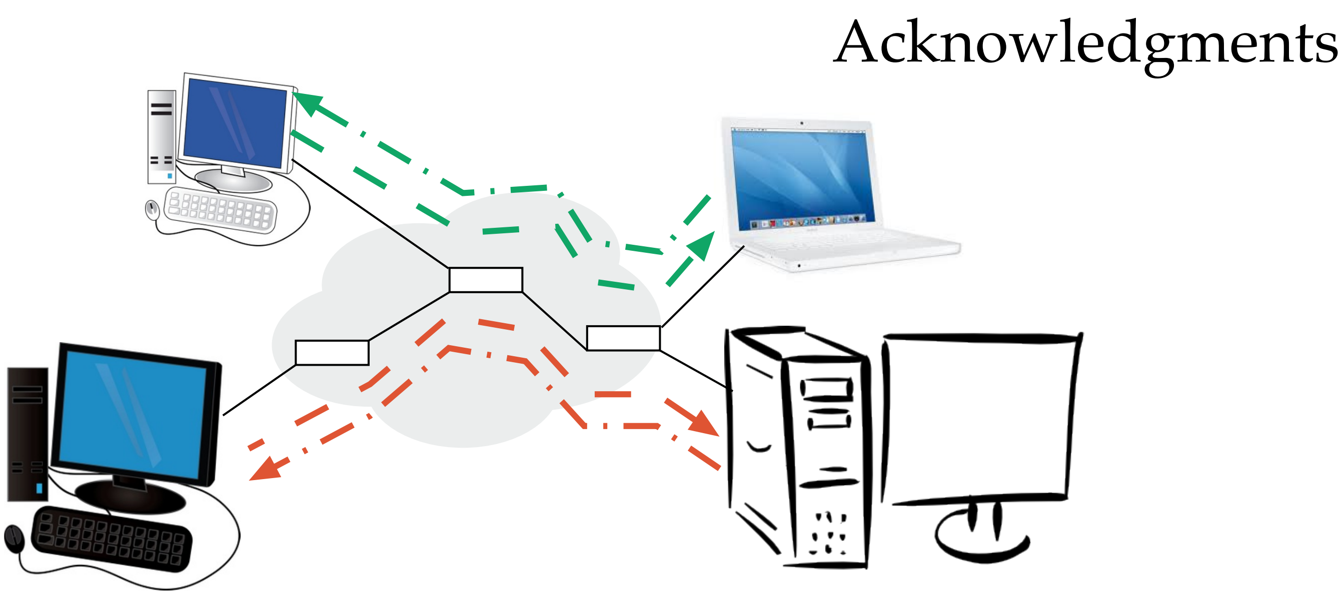

TCP approach:

- Increase rate as long as there is no congestion
- Slow down when network gets congested.

Congestion? Missing ACKs Routers drop packets when congested

TCP: Transmission Control Protocol

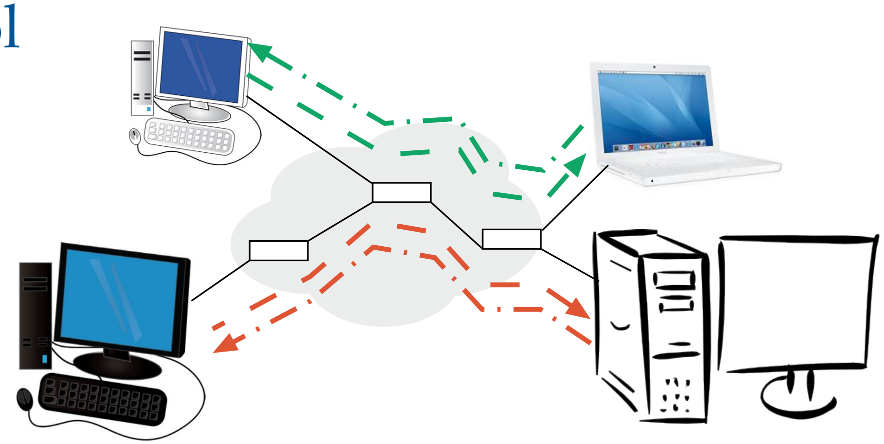

TCP approach:

- Increase rate as long as there is no congestion
- Slow down when network gets congested.

Additive Increase - Multiplicative Decrease

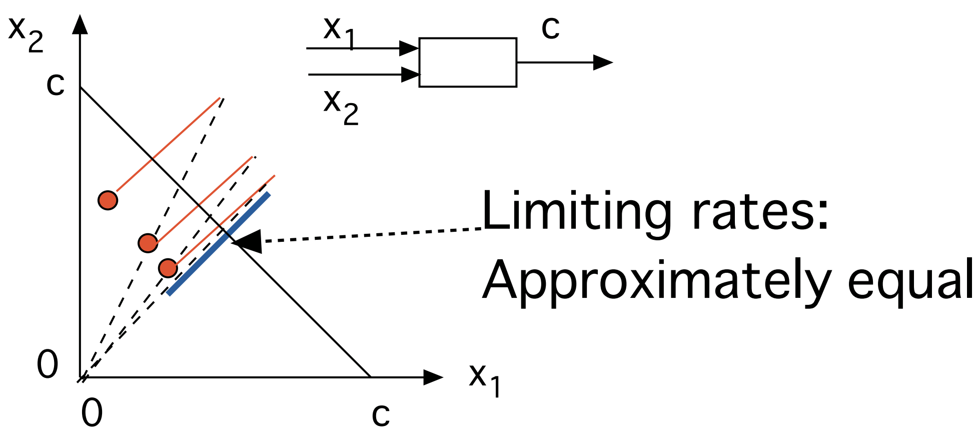

Works when sharing one link.
What about
- multiple links?
- objective other than equal rates?

Chiu-Jain '88, V. Jacobson '88

TCP: Transmission Control Protocol

Mathematical Model:

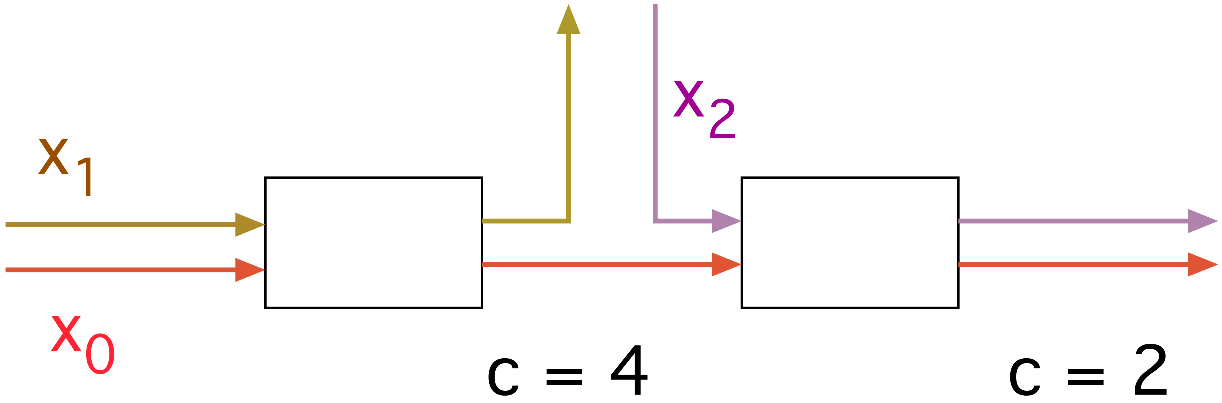

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$

Utility of users (= function of rate)

$$\text{Subject to } x_0 + x_1 \leq 4 \text{ and } x_0 + x_2 \leq 2$$

Capacity constraints

Kelly 96, see KEL98

TCP: Transmission Control Protocol

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$

Solution:

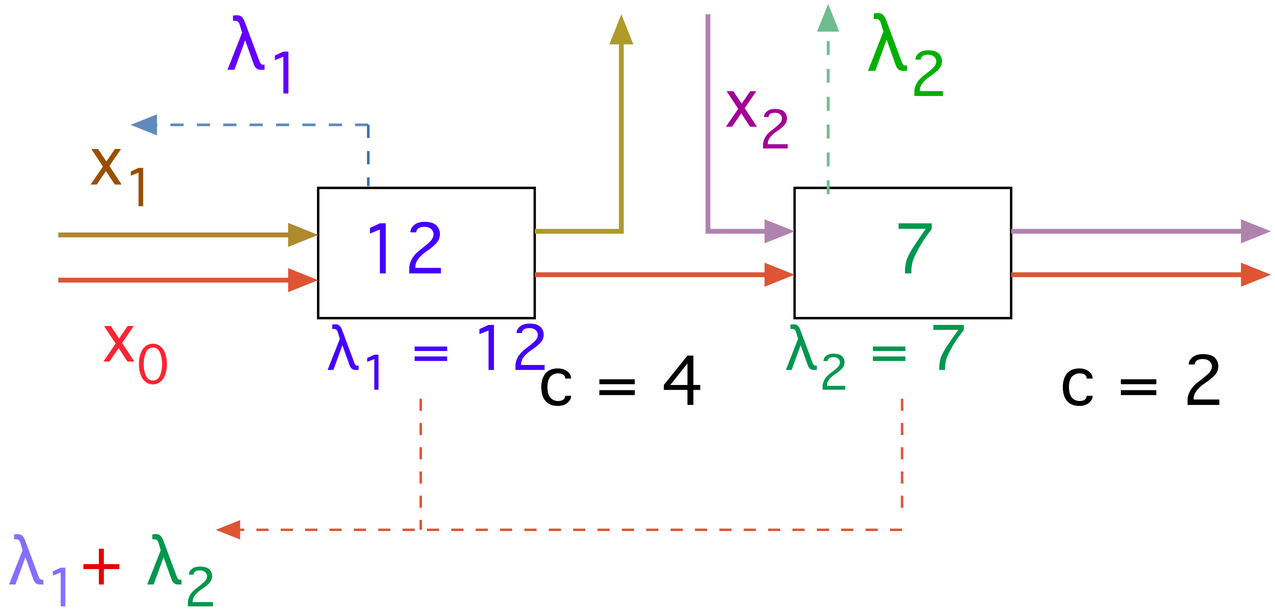

$$x_0 : \text{ Maximize } 2\log(x_0) - 19x_0$$

$$x_1 : \text{ Maximize } 3\log(x_1) - 12x_1$$

$$x_2 : \text{ Maximize } \log(x_2) - 7x_2$$

TCP: Transmission Control Protocol

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$

Simulation:

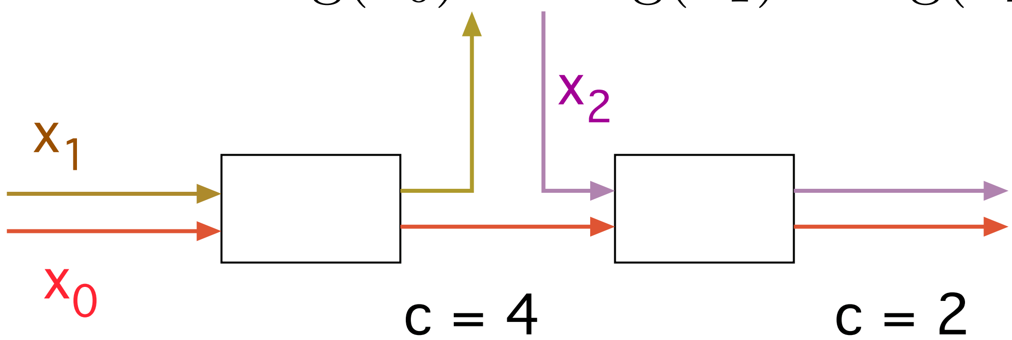

x_0 : Maximize $2\log(x_0) - (q_1 + q_2)x_0$

x_1 : Maximize $3\log(x_1) - q_1 x_1$

x_2 : Maximize $\log(x_2) - q_2 x_2$

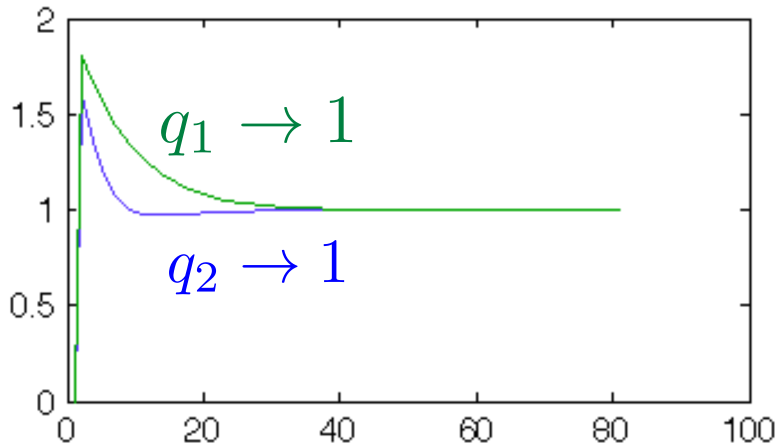

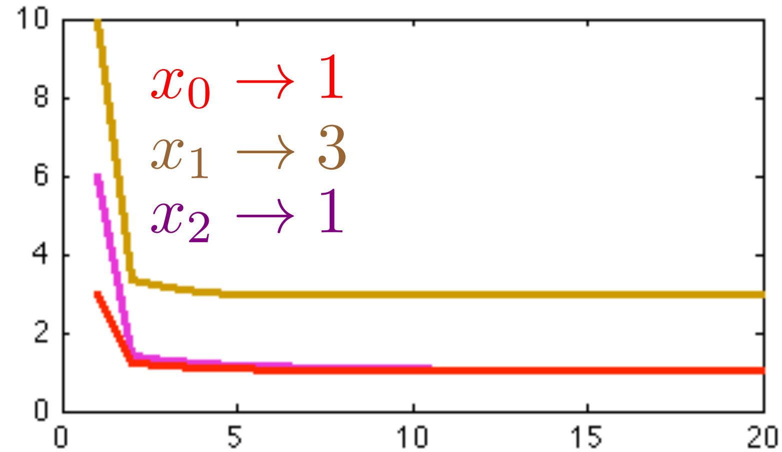

TCP: Transmission Control Protocol

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$

Simulation:

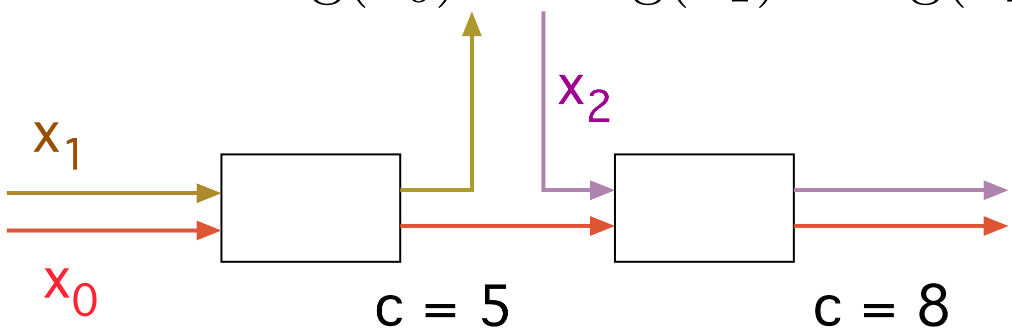

x_0 : Maximize $2\log(x_0) - (q_1 + q_2)x_0$

x_1 : Maximize $3\log(x_1) - q_1 x_1$

x_2 : Maximize $\log(x_2) - q_2 x_2$

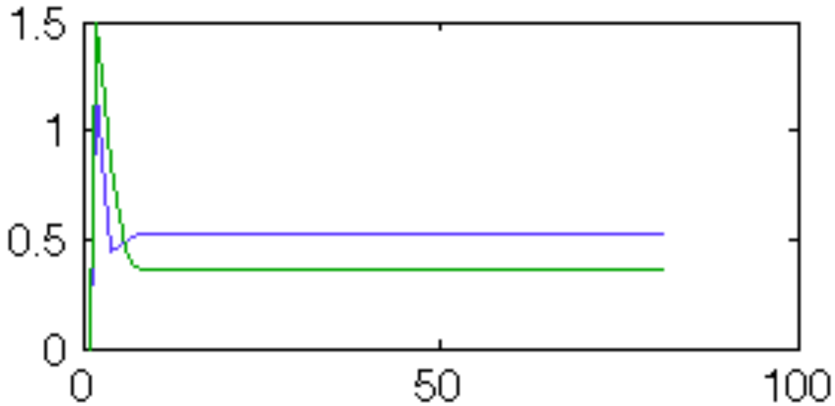

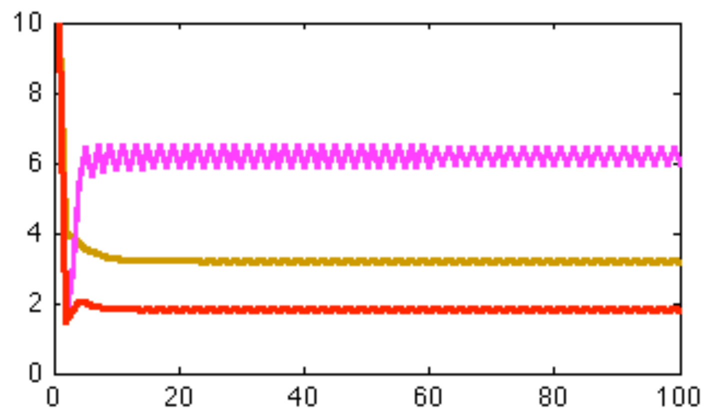

TCP: Transmission Control Protocol

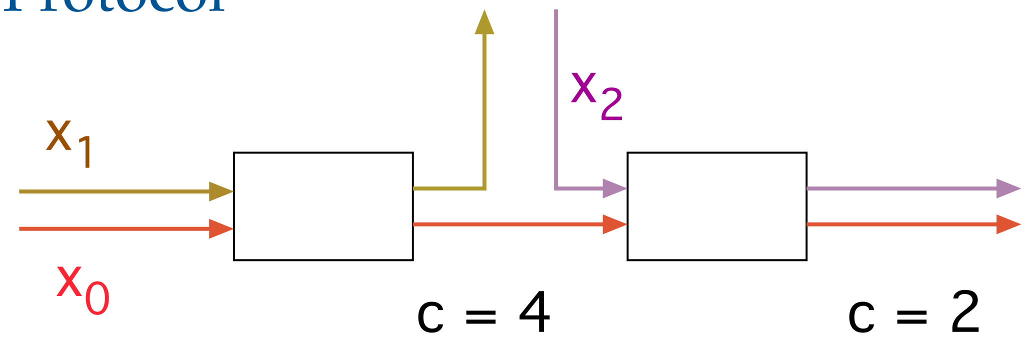

Primal:

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

Subject to $x_0 + x_1 \leq 4$ and $x_0 + x_2 \leq 2$

Lagrangian:

$$L(\lambda, \mathbf{x}) = 2\log(x_0) + 3\log(x_1) + \log(x_2) - \lambda_1(x_0 + x_1 - 4) - \lambda_2(x_0 + x_2 - 2)$$

Replace constraints by "penalty"

Theorem:

If $(\lambda^*, \mathbf{x}^*)$ is a *saddle point*, then $\mathbf{x}^*$ solves Primal.

Saddle Point:

$$\min_{\lambda} L(\mathbf{x}^*, \lambda) = L(\mathbf{x}^*, \lambda^*) = \max_{\mathbf{x}} L(\mathbf{x}, \lambda^*)$$

Digression: Constrained Optimization

Primal: Maximize $f(\mathbf{x}) = \log(x_1) + \log(x_2)$
Subject to $g(\mathbf{x}) = 2x_1 + x_2 - 10 \leq 0$

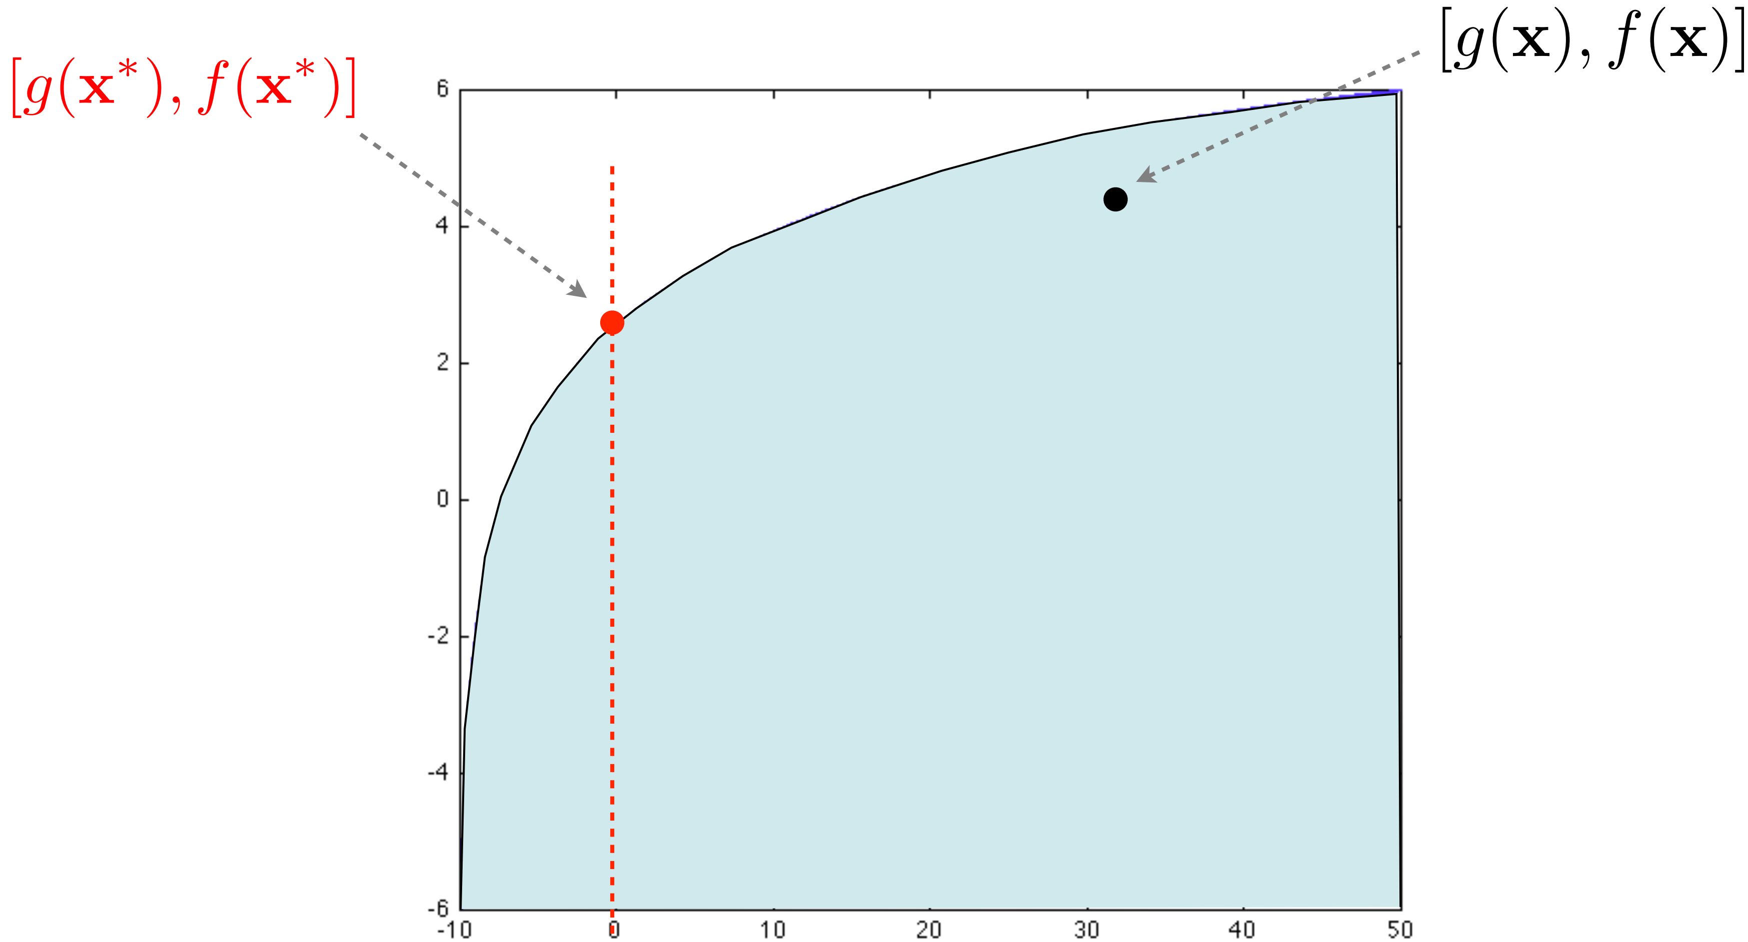

Digression: Constrained Optimization

Primal: Maximize $f(\mathbf{x}) = \log(x_1) + \log(x_2)$

Subject to $g(\mathbf{x}) = 2x_1 + x_2 - 10 \leq 0$

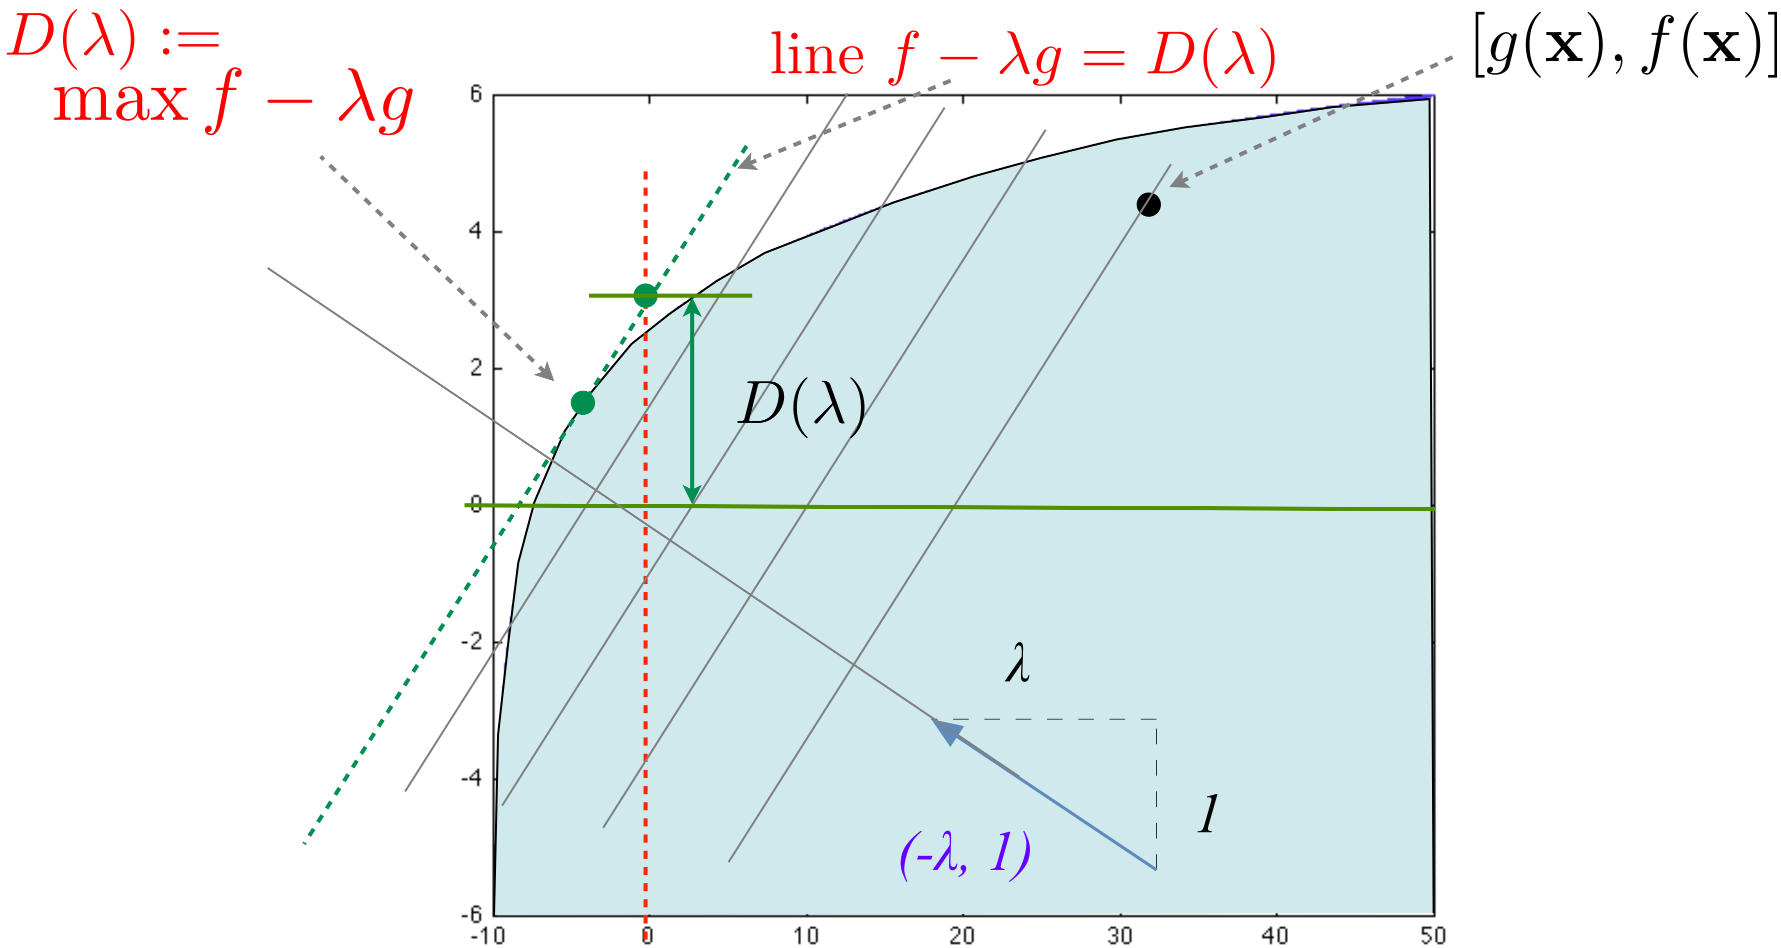

Digression: Constrained Optimization

Primal: Maximize $f(\mathbf{x}) = \log(x_1) + \log(x_2)$

Subject to $g(\mathbf{x}) = 2x_1 + x_2 - 10 \leq 0$

$D(\lambda) = \max_{\mathbf{x}} f - \lambda g$

$[g(\mathbf{x}), f(\mathbf{x})]$

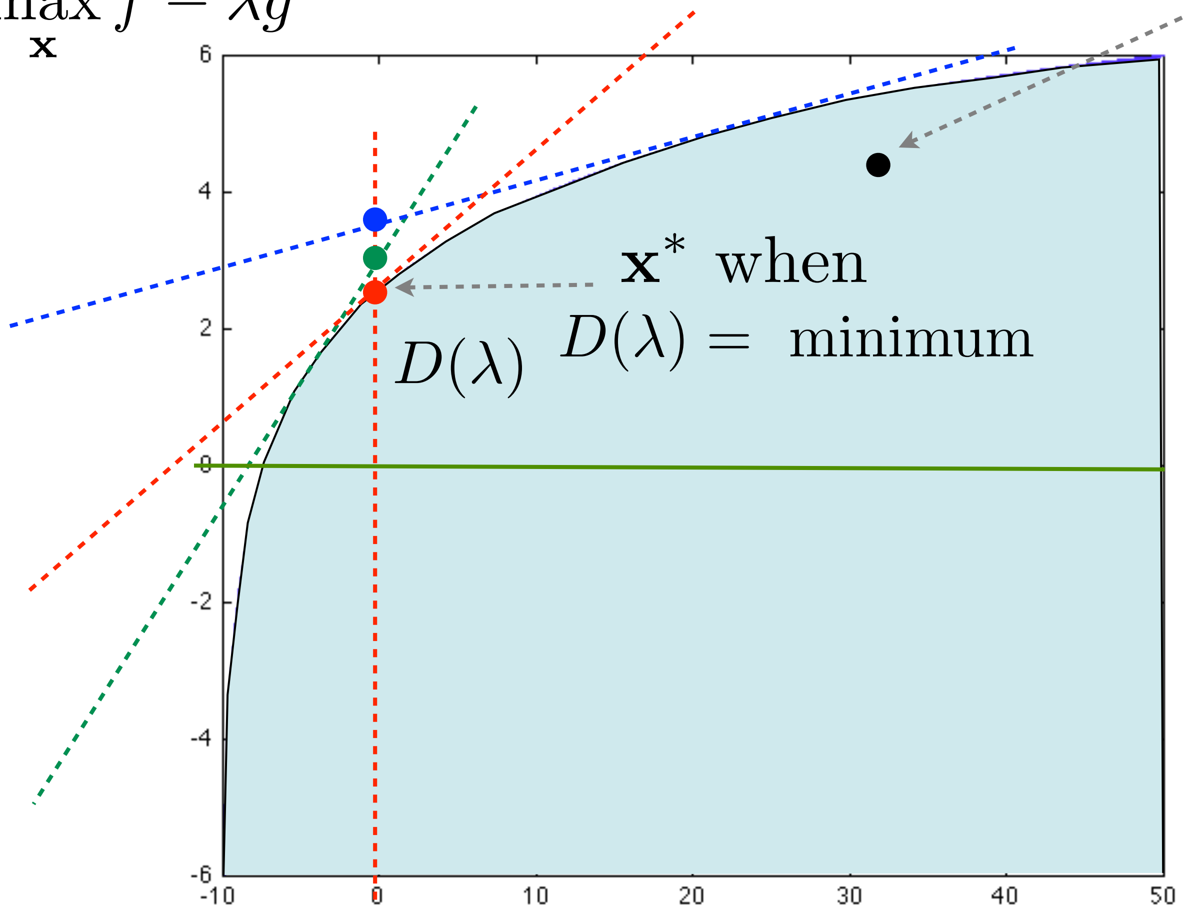

Digression: Constrained Optimization

Summing up:

Primal: Maximize $f(\mathbf{x}) = \log(x_1) + \log(x_2)$

Subject to $g(\mathbf{x}) = 2x_1 + x_2 - 10 \leq 0$

Lagrangian: $L(\mathbf{x}, \lambda) = f(\mathbf{x}) - \lambda g(\mathbf{x})$

Dual:

$$D(\lambda) = \max_{\mathbf{x}} L(\mathbf{x}, \lambda) = L(\mathbf{x}(\lambda), \lambda)$$

$$D(\lambda^*) = \min_{\lambda} D(\lambda) = \min_{\lambda} \max_{\mathbf{x}} L(\mathbf{x}, \lambda)$$

$$\Rightarrow \mathbf{x}^* = \mathbf{x}(\lambda^*)$$

Saddle Point: $\min_{\lambda} L(\mathbf{x}^*, \lambda) = L(\mathbf{x}^*, \lambda^*) = \max_{\mathbf{x}} L(\mathbf{x}, \lambda^*)$

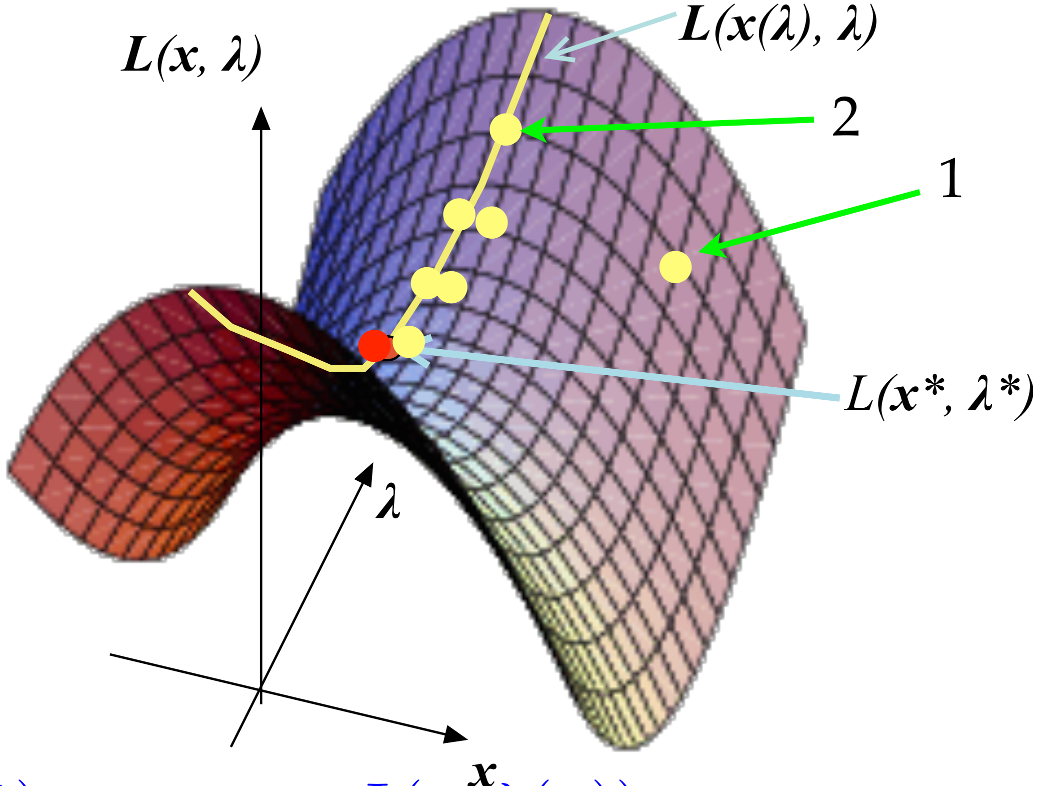

Dual:

$$\mathbf{x}(n+1) = \arg\max_{\mathbf{x}} L(\mathbf{x}, \lambda(n))$$

$$\lambda(n+1) = [\lambda(n) - \alpha(n)\nabla_{\lambda} L(\mathbf{x}(\mathbf{n}), \lambda(n))]^{+}$$

$\alpha(n) =$ step size (e.g. $1/n$)

TCP: Transmission Control Protocol

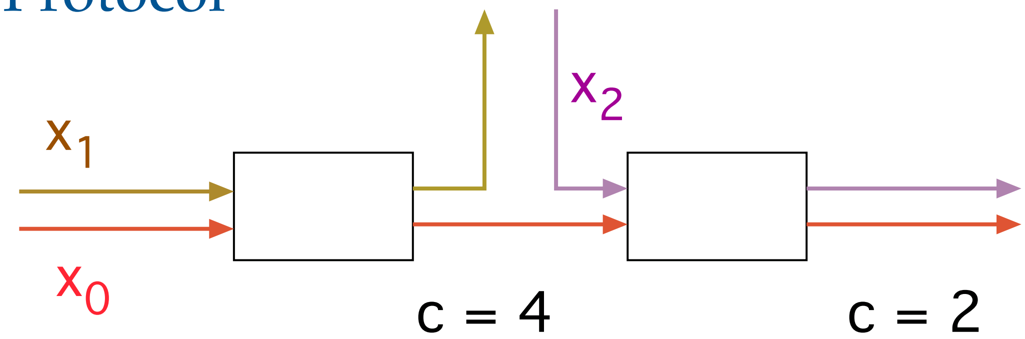

Primal:

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

Subject to $x_0 + x_1 \leq 4$ and $x_0 + x_2 \leq 2$

Lagrangian:

$$L(\lambda, \mathbf{x}) = 2\log(x_0) + 3\log(x_1) + \log(x_2)$$
$$-\lambda_1(x_0 + x_1 - 4) - \lambda_2(x_0 + x_2 - 2)$$

Replace constraints by "penalty"

Users:

Given $\lambda = (\lambda_1, \lambda_2)$,

$x_0 \rightarrow$ Maximize $2\log(x_0) - (\lambda_1 + \lambda_2)x_0$

$x_1 \rightarrow$ Maximize $3\log(x_1) - \lambda_1 x_1$

$x_2 \rightarrow$ Maximize $\log(x_2) - \lambda_2 x_2$

Kelly 96, see KEL98

TCP: Transmission Control Protocol

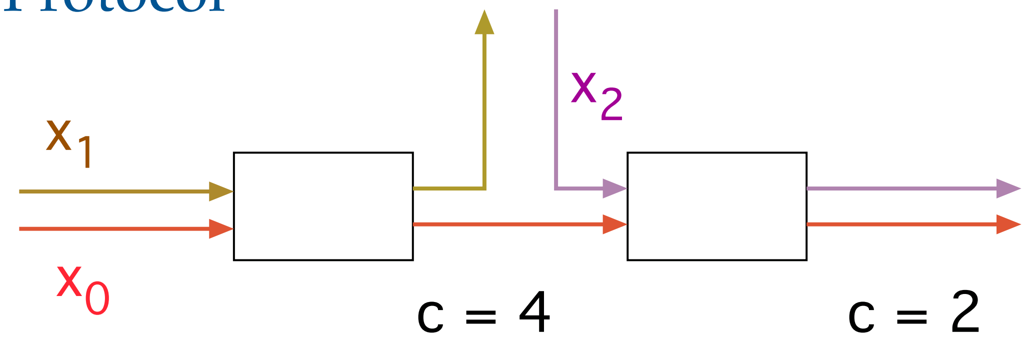

Primal:

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

Subject to $x_0 + x_1 \leq 4$ and $x_0 + x_2 \leq 2$

Replace constraints by "penalty"

Lagrangian:

$$L(\lambda, \mathbf{x}) = 2\log(x_0) + 3\log(x_1) + \log(x_2) - \lambda_1(x_0 + x_1 - 4) - \lambda_2(x_0 + x_2 - 2)$$

Routers:

$$\lambda_1(n+1) = [\lambda_1(n) + \alpha(n)(x_0 + x_1 - 4)]^+$$

$$\lambda_2(n+1) = [\lambda_2(n) + \alpha(n)(x_0 + x_2 - 2)]^+$$

Note: $\lambda_i \approx$ backlog in router i

Low-Lapsley, see LOW99

Backpressure Congestion Control

Example:

Per flow queues

Service rates = control variables.

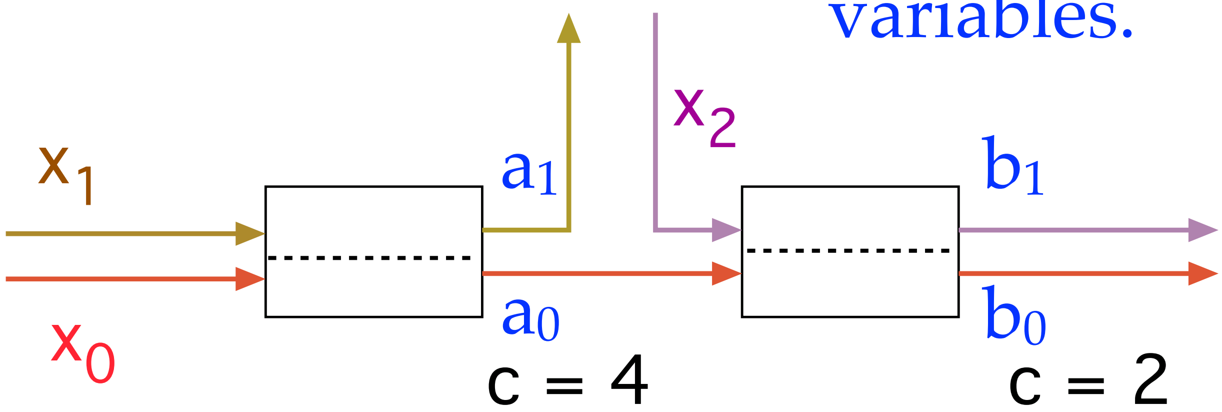

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

Subject to $x_0 \le a_0, x_1 \le a_1, a_0 \le b_0, x_2 \le b_1$

$$a_0 + a_1 \le 4, b_0 + b_1 \le 2$$

TAS92 NEE05

Backpressure Congestion Control

Example:

Per flow queues

Service rates = control variables.

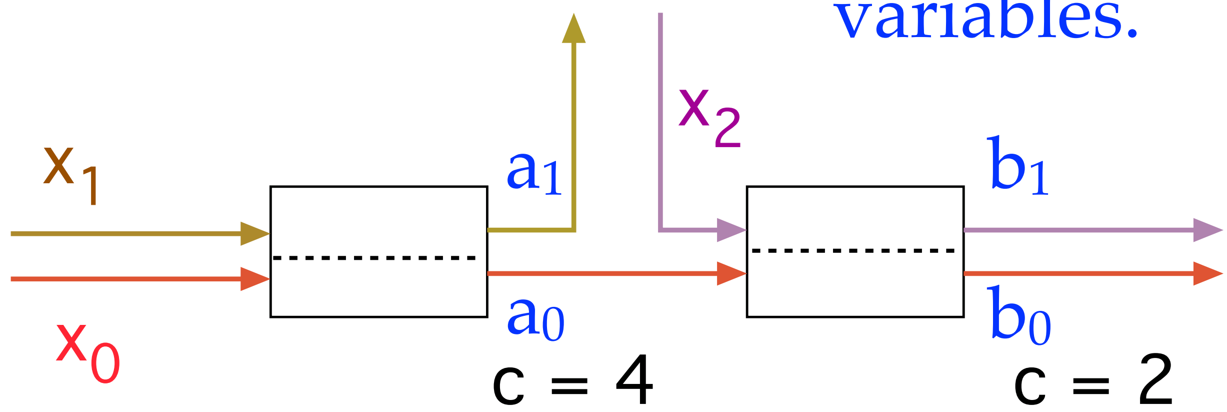

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$
$$-\lambda_{10}(x_0 - a_0) - \lambda_{11}(x_1 - a_1) - \lambda_{20}(a_0 - b_0) - \lambda_{22}(x_2 - b_1)$$
$$\text{s.t. } a_0 + a_1 \leq 4, b_0 + b_1 \leq 2$$

Backpressure Congestion Control

Example:

Per flow queues

Service rates = control variables.

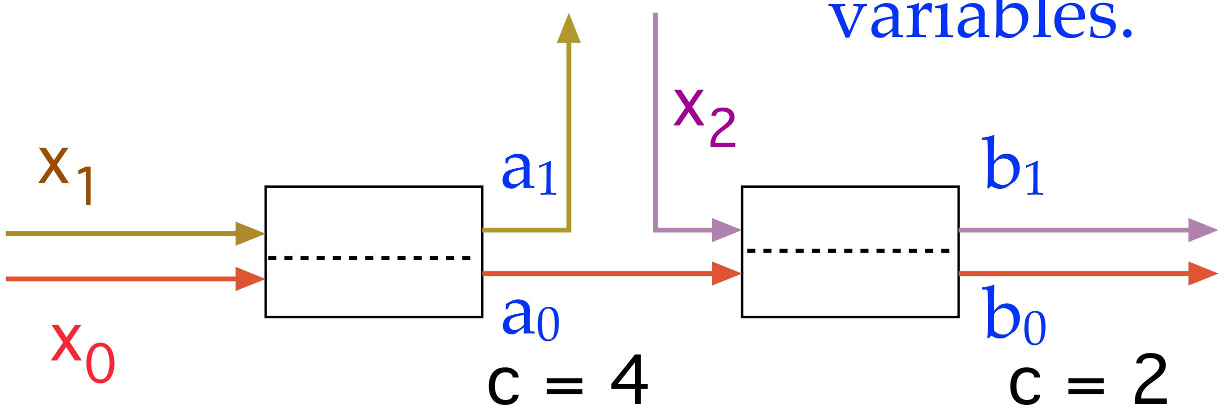

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

$$-\lambda_{10}(x_0 - a_0) - \lambda_{11}(x_1 - a_1) - \lambda_{20}(a_0 - b_0) - \lambda_{22}(x_2 - b_1)$$

s.t. $a_0 + a_1 \leq 4, b_0 + b_1 \leq 2$

x_0 maximizes $2\log(x_0) - \lambda_{10}x_0 \leftarrow$ price = ingress backlog

$$\lambda_{10}(n+1) = [\lambda_{10}(n) + \alpha(n)(x_0 - a_0)]^+$$

Backpressure Congestion Control

Example:

Per flow queues
Service rates = control variables.

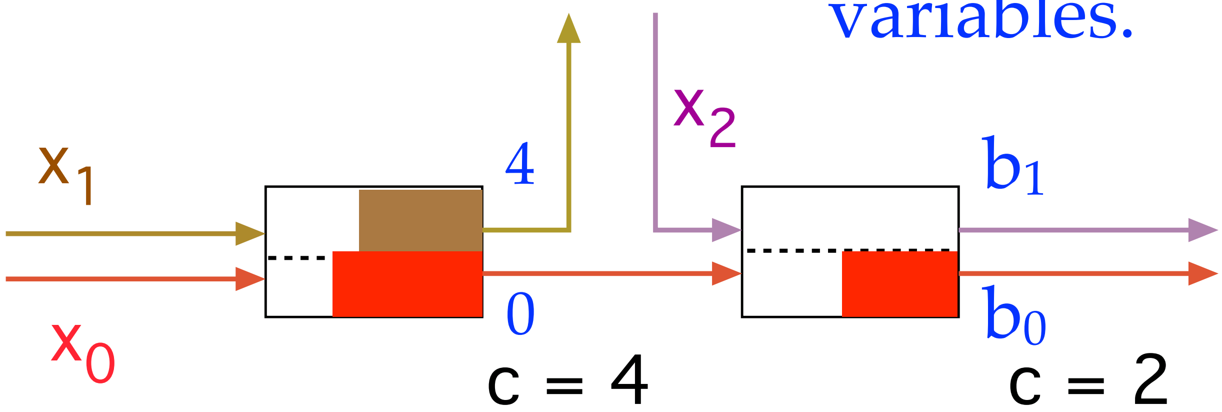

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$

$$-\lambda_{10}(x_0 - a_0) - \lambda_{11}(x_1 - a_1) - \lambda_{20}(a_0 - b_0) - \lambda_{22}(x_2 - b_1)$$

s.t. $a_0 + a_1 \leq 4, b_0 + b_1 \leq 2$

If $\lambda_{10} - \lambda_{20} > \lambda_{11}$, then $a_0 = 4, a_1 = 0$
else, $a_0 = 0, a_1 = 4$: maximum backpressure

Backpressure Congestion Control

Example:

Per flow queues

Service rates = control variables.

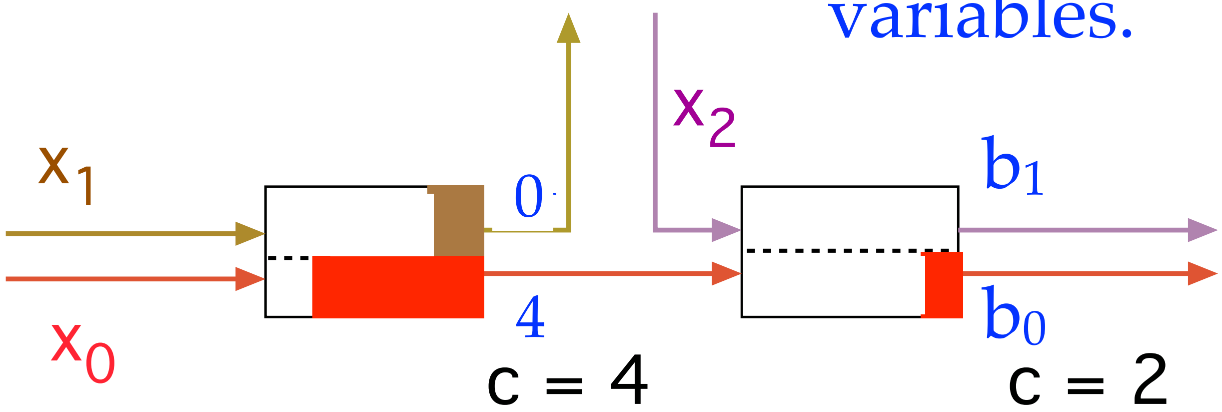

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$
$$-\lambda_{10}(x_0 - a_0) - \lambda_{11}(x_1 - a_1) - \lambda_{20}(a_0 - b_0) - \lambda_{22}(x_2 - b_1)$$
$$\text{s.t. } a_0 + a_1 \leq 4, b_0 + b_1 \leq 2$$

If $\lambda_{10} - \lambda_{20} > \lambda_{11}$, then $a_0 = 4, a_1 = 0$
else, $a_0 = 0, a_1 = 4$: maximum backpressure

Backpressure Routing

Example:

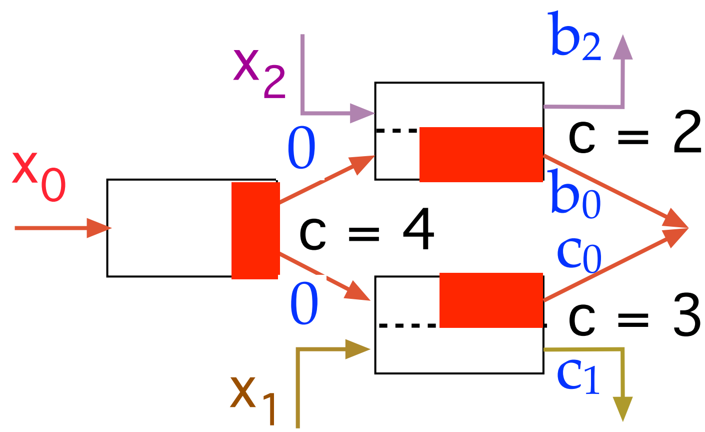

Per flow queuing

Service rates are control variables

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$
$$-\lambda_0(x_0 - 4) - \lambda_{12}(x_2 - b_2) - \lambda_{10}(a_1 - b_0) - \lambda_{21}(x_1 - c_1) - \lambda_{20}(a_2 - c_0)$$
$$\text{s.t. } a_1 + a_2 \leq 4, b_0 + b_2 \leq 2, c_0 + c_1 \leq 3$$

$\rightarrow \lambda_{ij} = \text{ backlog}$

If $\lambda_0 < \max\{\lambda_{10}, \lambda_{20}\}$, then $a_1 = 0, a_2 = 0,$
else, if $\lambda_{10} < \lambda_{20}$, then $a_1 = 4, a_2 = 0,$
else, $a_1 = 0, a_4 = 4.$

TAS92 NEE05

Backpressure Routing

Example:

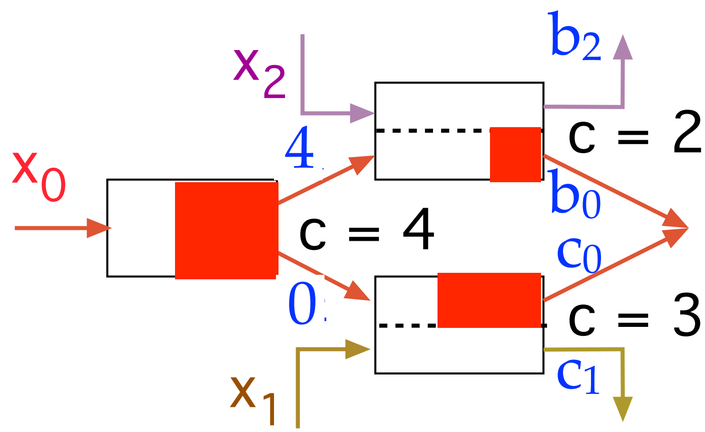

Per flow queuing
Service rates are control variables

Maximize $2\log(x_0) + 3\log(x_1) + \log(x_2)$
$-\lambda_0(x_0 - 4) - \lambda_{12}(x_2 - b_2) - \lambda_{10}(a_1 - b_0) - \lambda_{21}(x_1 - c_1) - \lambda_{20}(a_2 - c_0)$
s.t. $a_1 + a_2 \leq 4, b_0 + b_2 \leq 2, c_0 + c_1 \leq 3$

$\rightarrow \lambda_{ij} =$ backlog

If $\lambda_0 < \max\{\lambda_{10}, \lambda_{20}\}$, then $a_1 = 0, a_2 = 0,$
else, if $\lambda_{10} < \lambda_{20}$, then $a_1 = 4, a_2 = 0,$
else, $a_1 = 0, a_4 = 4.$

Backpressure Routing

Example:

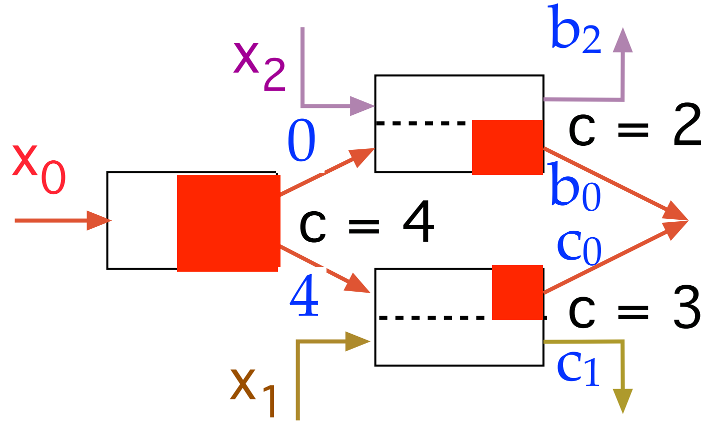

Per flow queuing

Service rates are control variables

$$\text{Maximize } 2\log(x_0) + 3\log(x_1) + \log(x_2)$$
$$-\lambda_0(x_0 - 4) - \lambda_{12}(x_2 - b_2) - \lambda_{10}(a_1 - b_0) - \lambda_{21}(x_1 - c_1) - \lambda_{20}(a_2 - c_0)$$
$$\text{s.t. } a_1 + a_2 \leq 4, b_0 + b_2 \leq 2, c_0 + c_1 \leq 3$$

$\rightarrow \lambda_{ij} =$ backlog

If $\lambda_0 < \max\{\lambda_{10}, \lambda_{20}\}$, then $a_1 = 0, a_2 = 0$,
else, if $\lambda_{10} < \lambda_{20}$, then $a_1 = 4, a_2 = 0$,
else, $a_1 = 0, a_4 = 4$.

Which packets to
send at any given time?

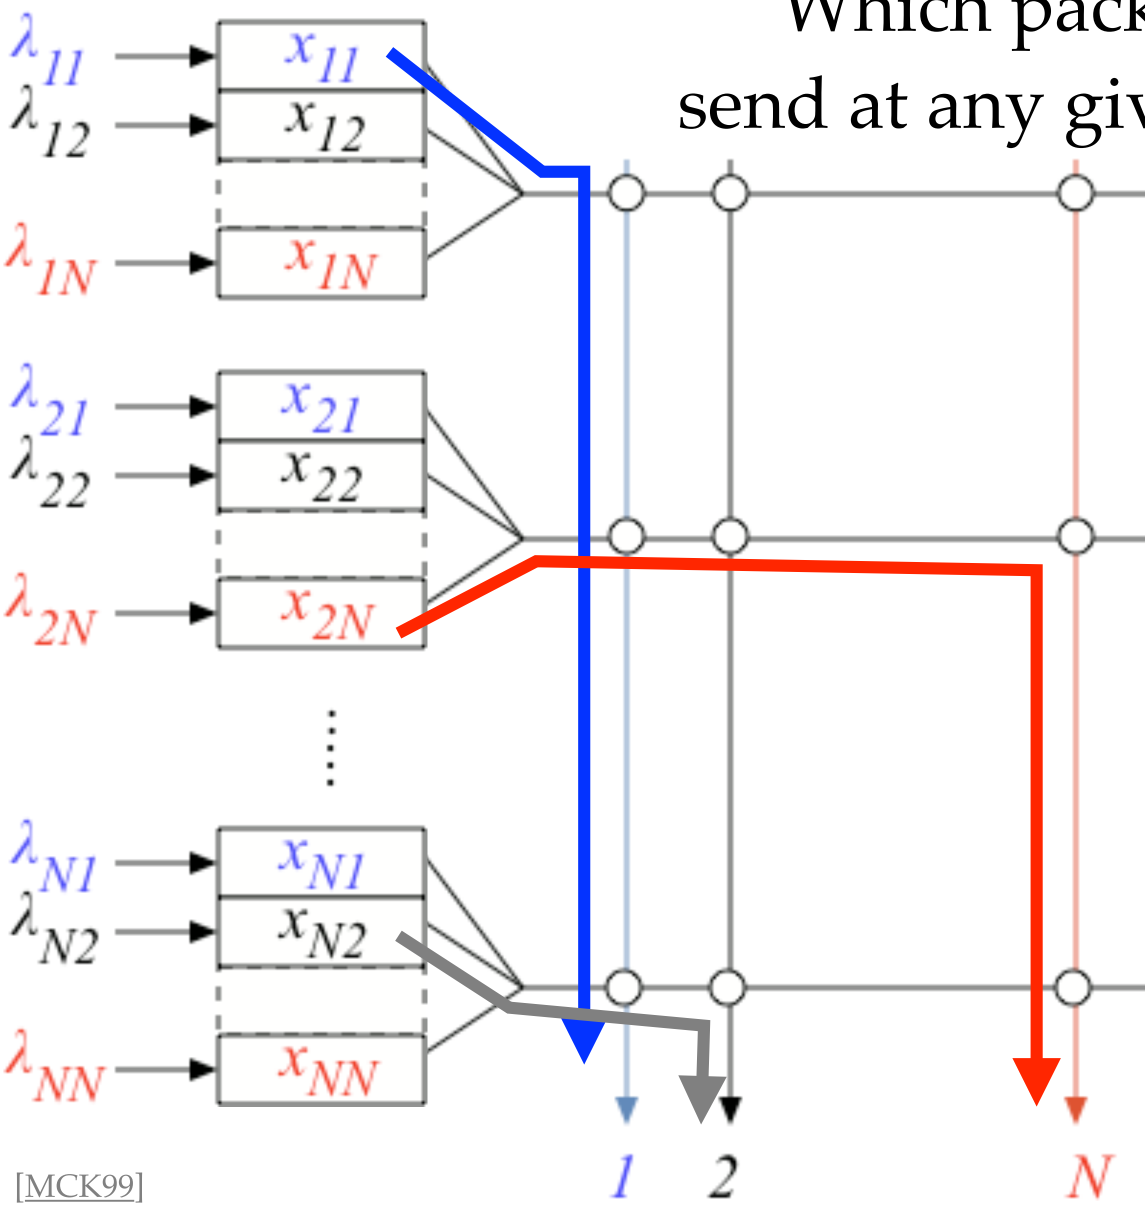

[MCK99]

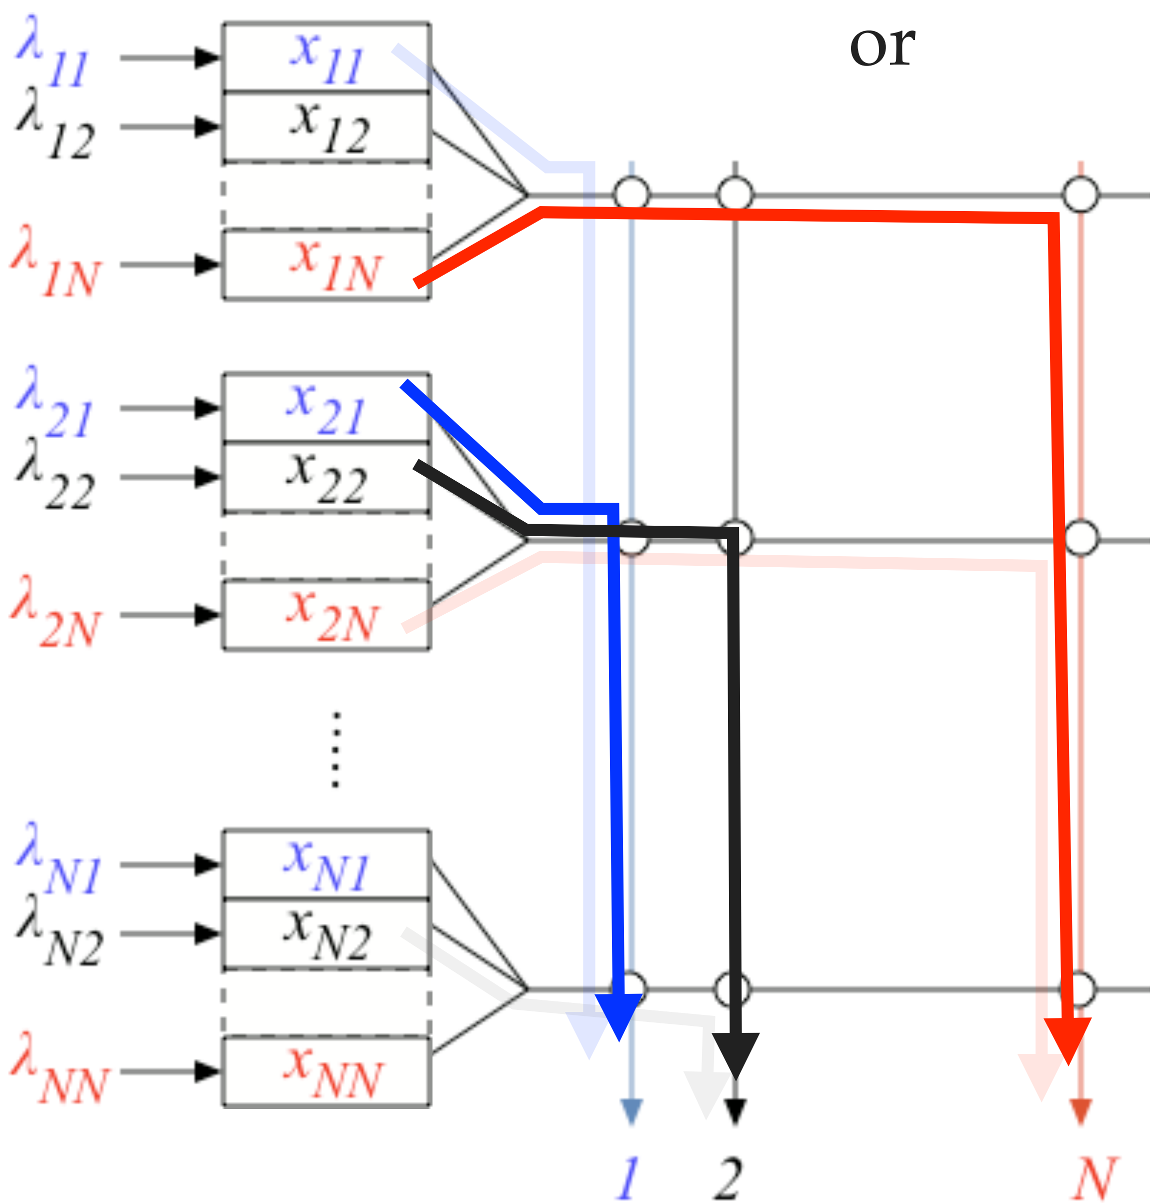

λ_{11}
λ_{12}
λ_{1N}
λ_{21}
λ_{22}
λ_{2N}
λ_{N1}
λ_{N2}
λ_{NN}
x_{11}
x_{12}
x_{1N}
x_{21}
x_{22}
x_{2N}
x_{N1}
x_{N2}
x_{NN}
or
1
2
N

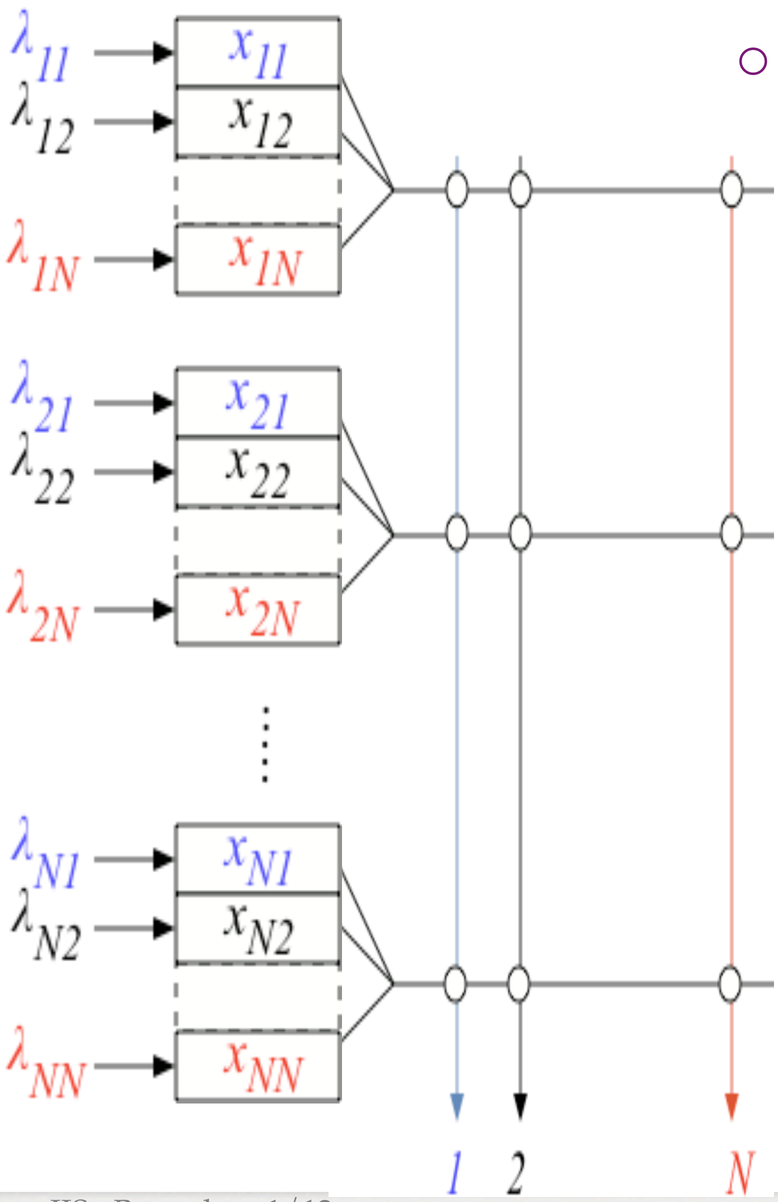

- Goals?
  - High Throughput
  - Fairness
  - Low Delays
- Classical Answer:
  - Maximum Weighted Matching

  Much Too Complex!
- Simpler Answer:
  - Q-CSMA

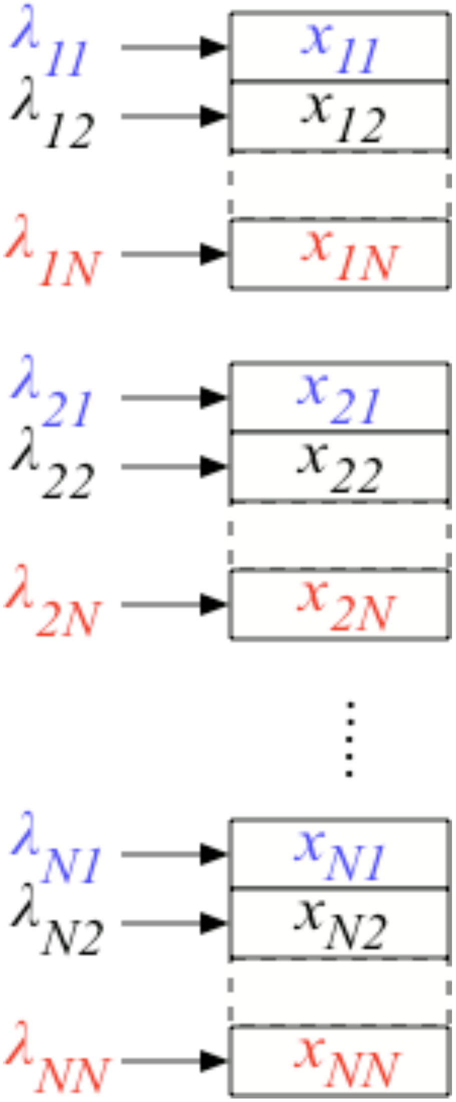

Q-CSMA:

- Input 1 : Select random delay with mean $\exp\{-\alpha X_{1j}\}$ for every j
- If minimum delay is for j, input 1 checks if output j is busy
  - If not, it sends a packet to j
  - If yes, it repeats
- Same for the other inputs
- Basic Idea: Favor larger backlogs

[JW10]

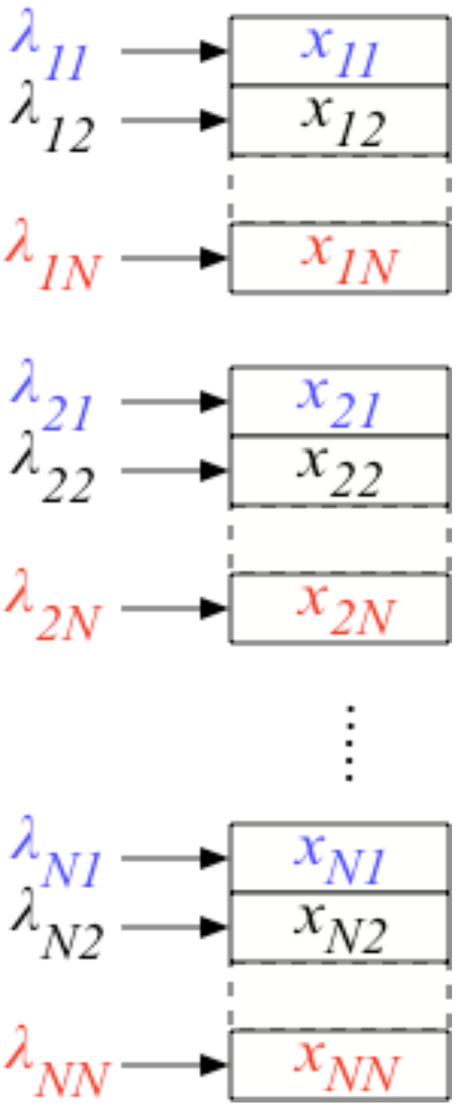

○ Results:

- Essentially 100% throughput
- Delays can be controlled if we accept a small throughput reduction
- Works with variable packet lengths

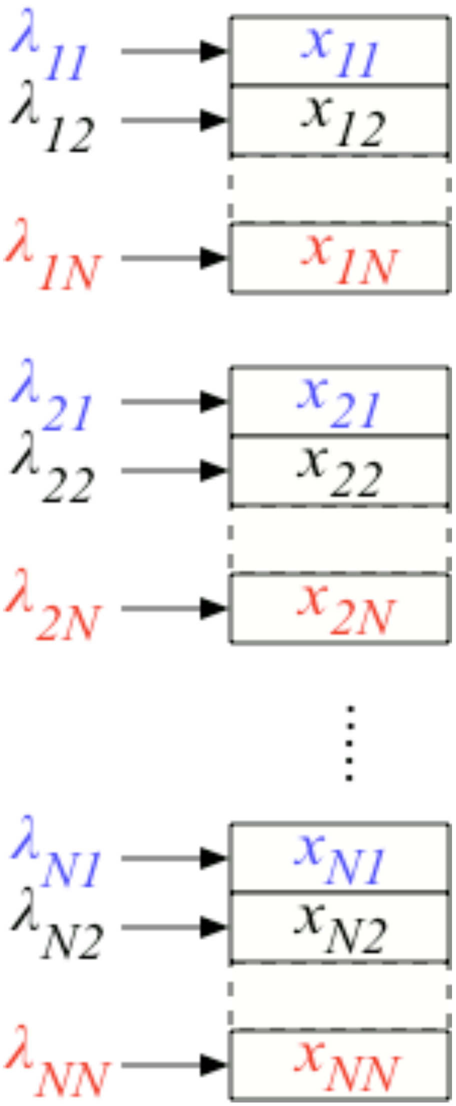

- Fairness:
  - Requires congestion control
  - Input ij reduces λ_{ij} if x_{ij} increases
  - Choose λ_{ij} to maximize

$$u_{ij}(\lambda_{ij}) - \beta x_{ij}\lambda_{ij}$$

- Result:
  - Essentially maximizes $\Sigma u_{ij}(\lambda_{ij})$

Formulation:

$$\text{Maximize} \sum_{i,j} u_{i,j}(\lambda_{i,j})$$

$$s.t.\ \lambda_{i,j} \le \sigma_{i,j}(\pi) := \sum_{\mathbf{z}\in\mathcal{Z}|(i,j)\in\mathbf{z}\}} \pi(\mathbf{z}).$$

$$\mathcal{Z} = \text{ possible simultaneous transfers}$$
$$\pi = \text{ some p.m. on } \mathcal{Z}$$

Relaxation:

$$\text{Maximize} \sum_{i,j} u_{i,j}(\lambda_{i,j}) + \beta H(\pi)$$

$$\text{subject to } \lambda_{i,j} \le \sigma_{i,j}(\pi)$$

$$\text{and } \sum_{\mathbf{z}\in\mathcal{Z}} \pi(\mathbf{z}) = 1, \pi(\mathbf{z}) \ge 0, \forall \mathbf{z} \in \mathcal{Z}.$$

$$H(\pi) := -\sum_{\mathbf{z}\in\mathcal{Z}} \pi(\mathbf{z}) \log(\pi(\mathbf{z}))$$ ← Entropy

$$\text{Maximize } \sum_{i,j} u_{i,j}(\lambda_{i,j}) + \beta H(\pi)$$
$$\text{subject to } \lambda_{i,j} \leq \sigma_{i,j}(\pi)$$
$$\text{and } \sum_{\mathbf{z}\in\mathcal{Z}} \pi(\mathbf{z}) = 1, \pi(\mathbf{z}) \geq 0, \forall \mathbf{z} \in \mathcal{Z}.$$

$$H(\pi) := -\sum_{\mathbf{z}\in\mathcal{Z}} \pi(\mathbf{z}) \log(\pi(\mathbf{z}))$$

Lagrangian:

$$L(\lambda, \pi, \mu) = \sum_{i,j} u_{i,j}(\lambda_{i,j}) + \beta H(\pi) - \sum_{i,j} \mu_{i,j}[\lambda_{i,j} - \sigma_{i,j}(\pi)] - \gamma[\sum_{\mathbf{z}\in\mathcal{Z}} \pi(\mathbf{z}) - 1]$$

Maximize over λ:

$$\lambda_{i,j} \text{ maximizes } u_{i,j}(\lambda_{i,j}) - \mu_{i,j}\lambda_{i,j}$$

Maximize over π:

$$\pi(\mathbf{z}) = A\exp\{\beta^{-1}\sum_{i,j}\mu_{i,j}1\{(i,j)\in\mathbf{z}\}\}, \mathbf{z}\in\mathcal{Z}$$

Minimize over μ:

[JW10]

$$\mu_{i,j}(t+1) = [\mu_{i,j}(t) + \alpha(t)\{\lambda_{i,j} - \sigma_{i,j}(\pi)\}]^{+} \approx q_{i,j}$$

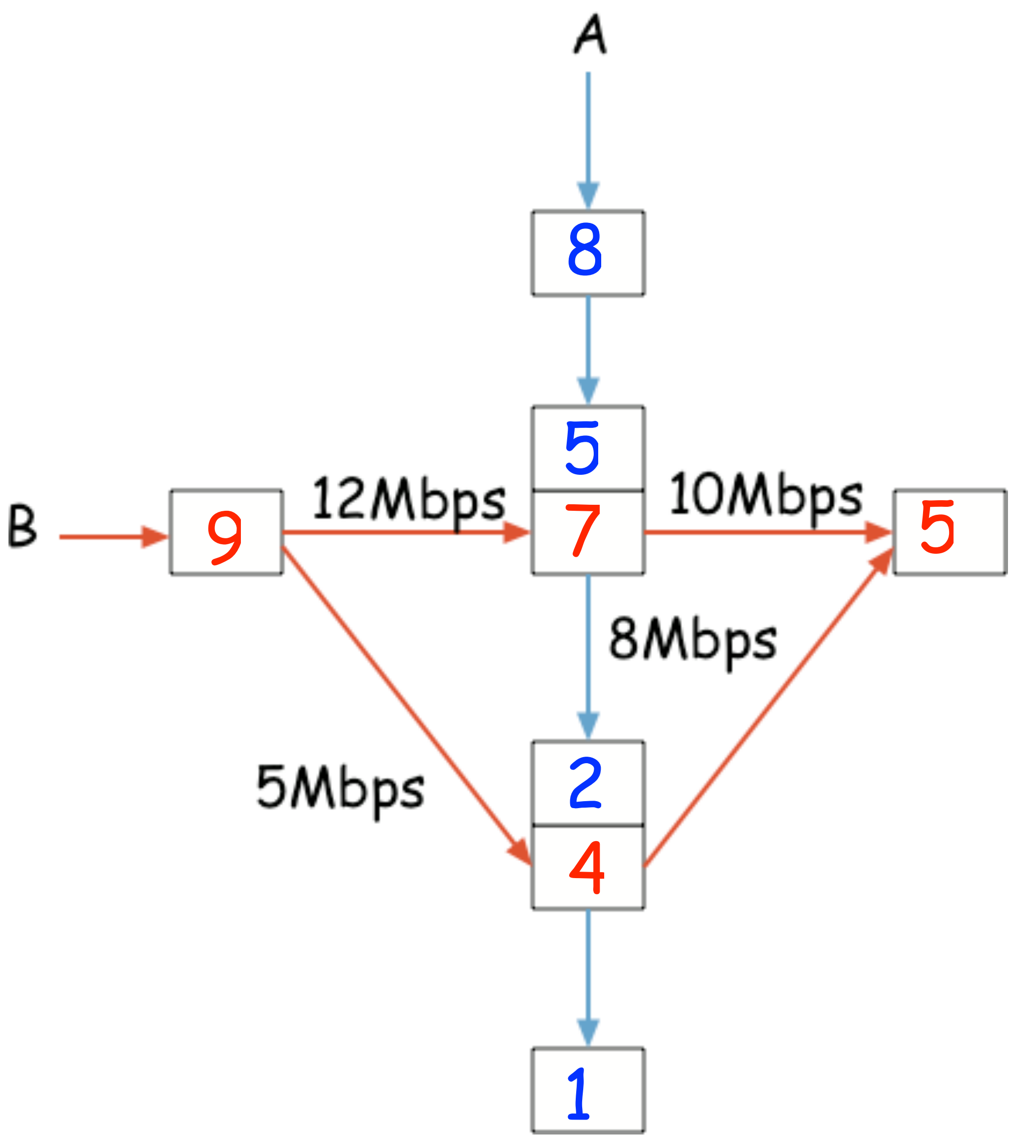

[JW10]

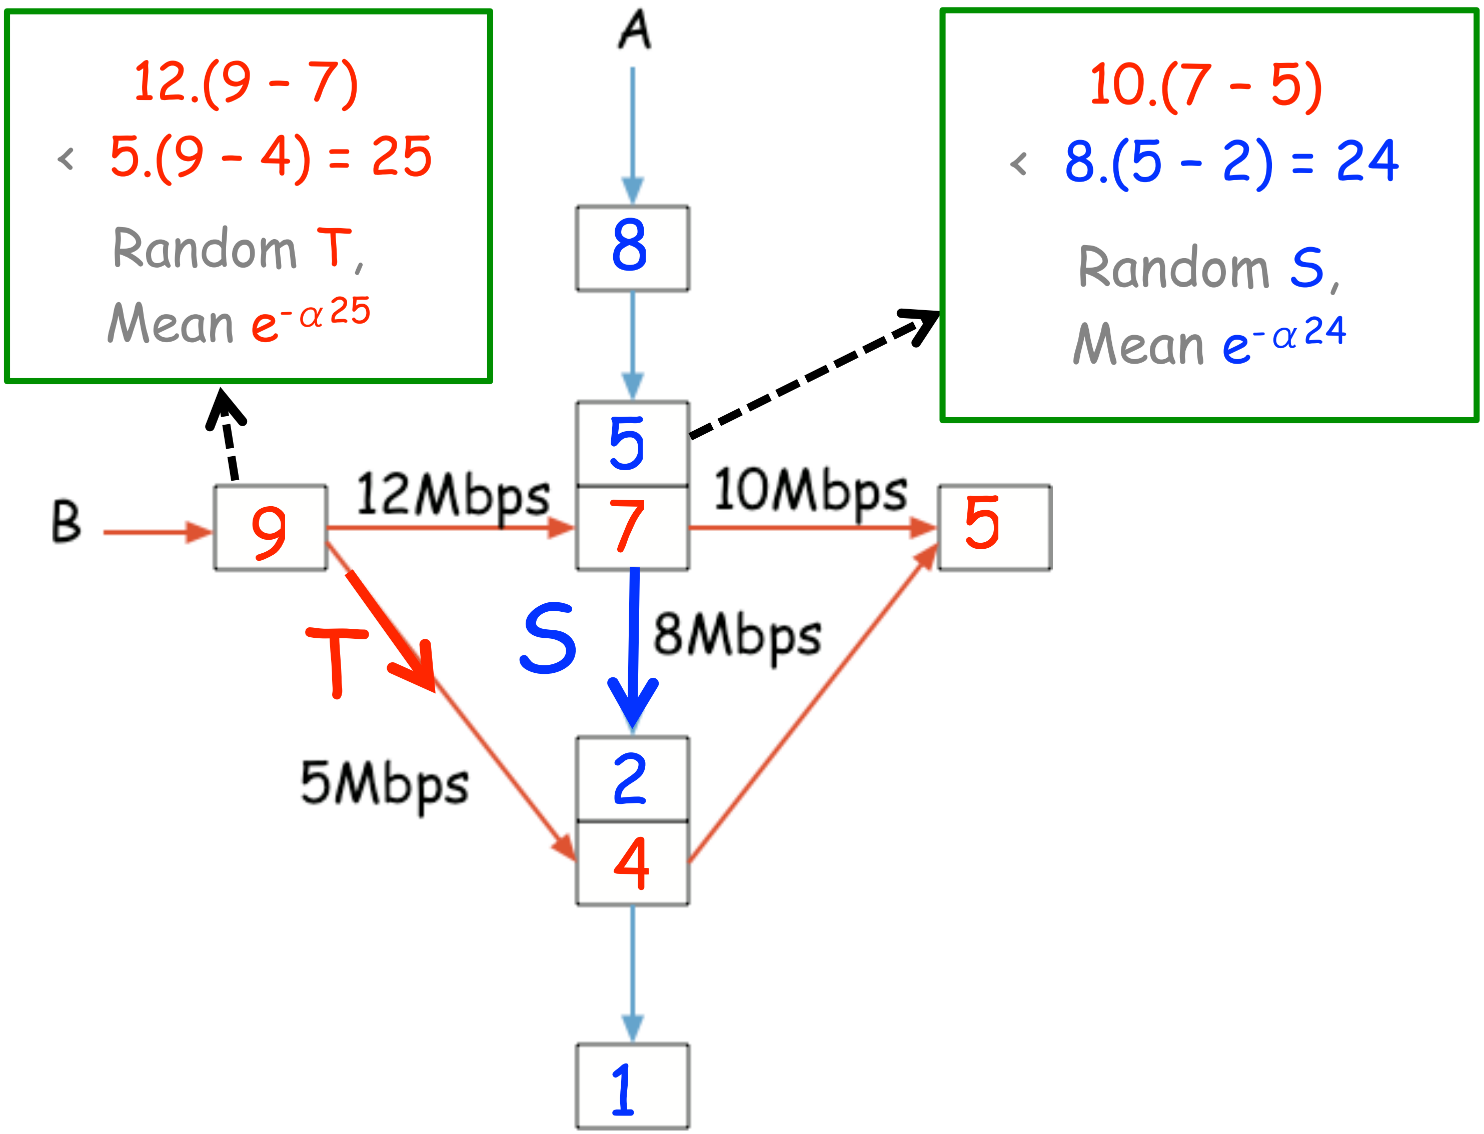
12.(9 - 7)
< 5.(9 - 4) = 25
Random T,
Mean e^{-α25}
A
8
10.(7 - 5)
< 8.(5 - 2) = 24
Random S,
Mean e^{-α24}
5
B
9
12Mbps
7
10Mbps
5
T
S
8Mbps
5Mbps
2
4
1

[JW10]

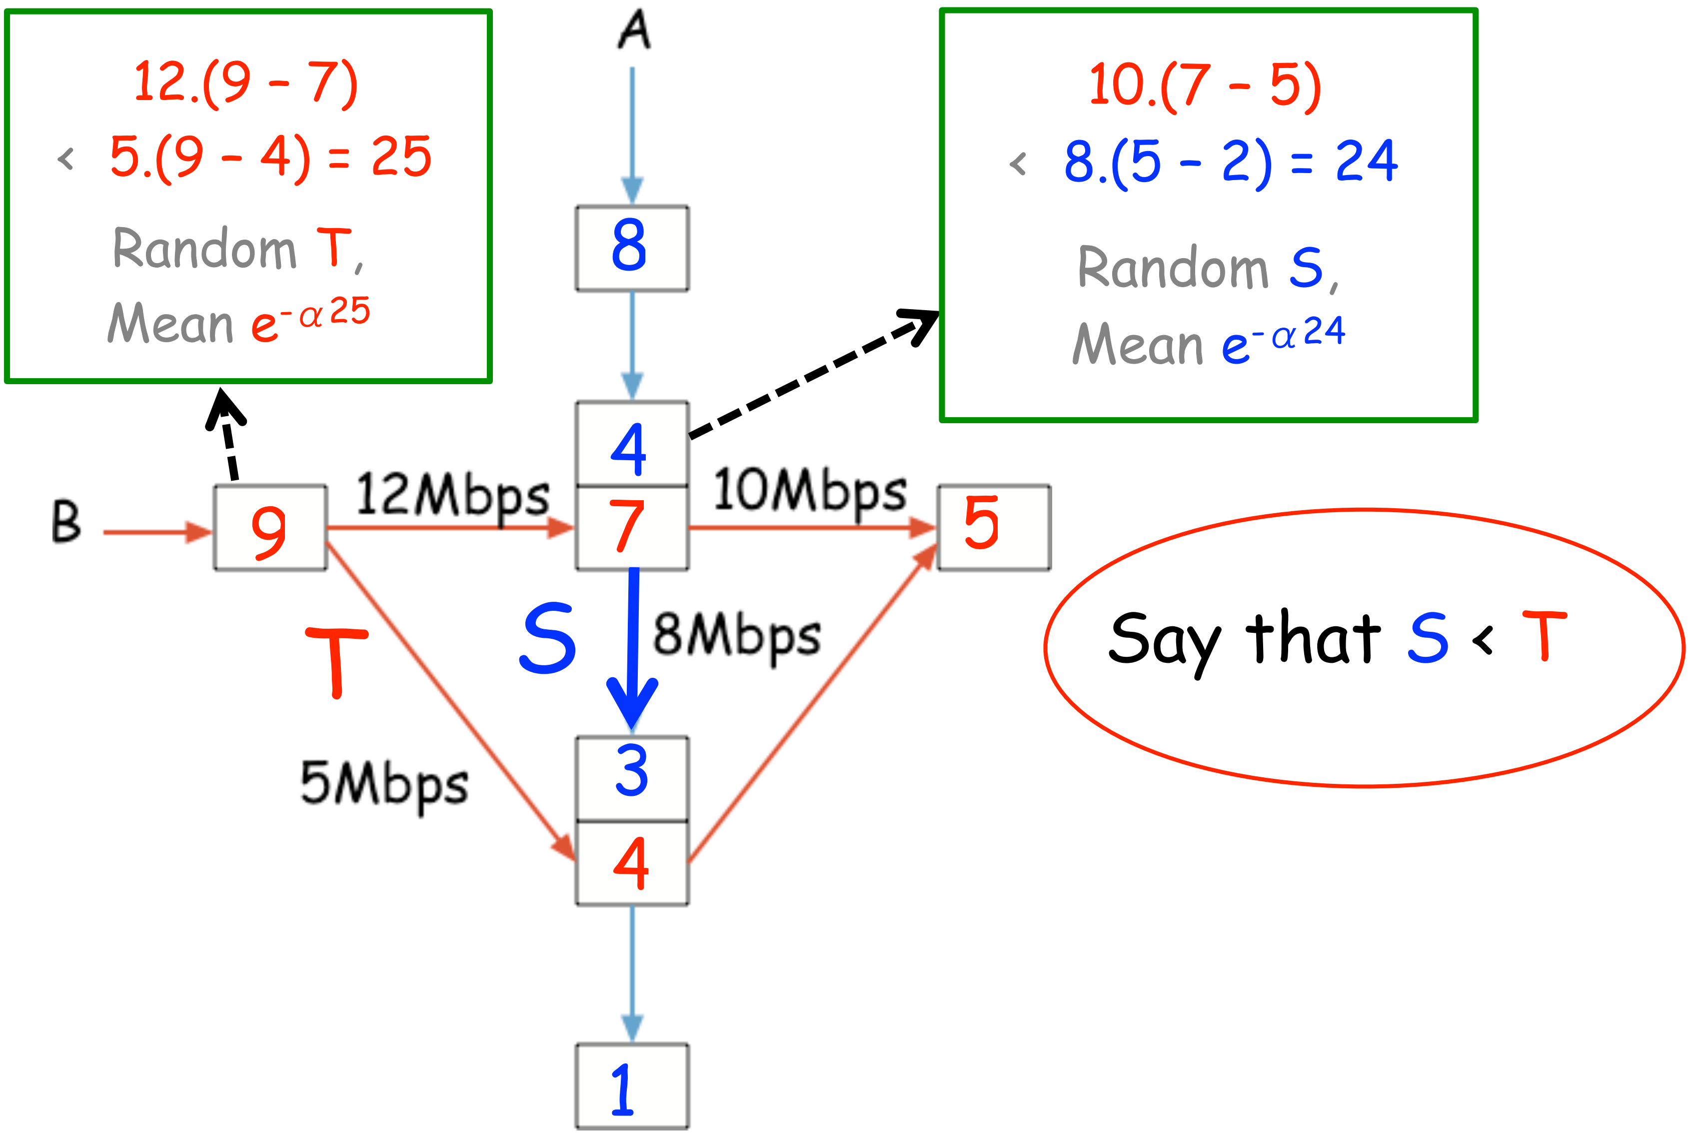
12.(9 - 7)
< 5.(9 - 4) = 25
Random T,
Mean e^-α25
A
8
10.(7 - 5)
< 8.(5 - 2) = 24
Random S,
Mean e^-α24
4
B
9
12Mbps
7
10Mbps
5
T
S
8Mbps
Say that S < T
5Mbps
3
4
1

[JW10]

Admission Control:

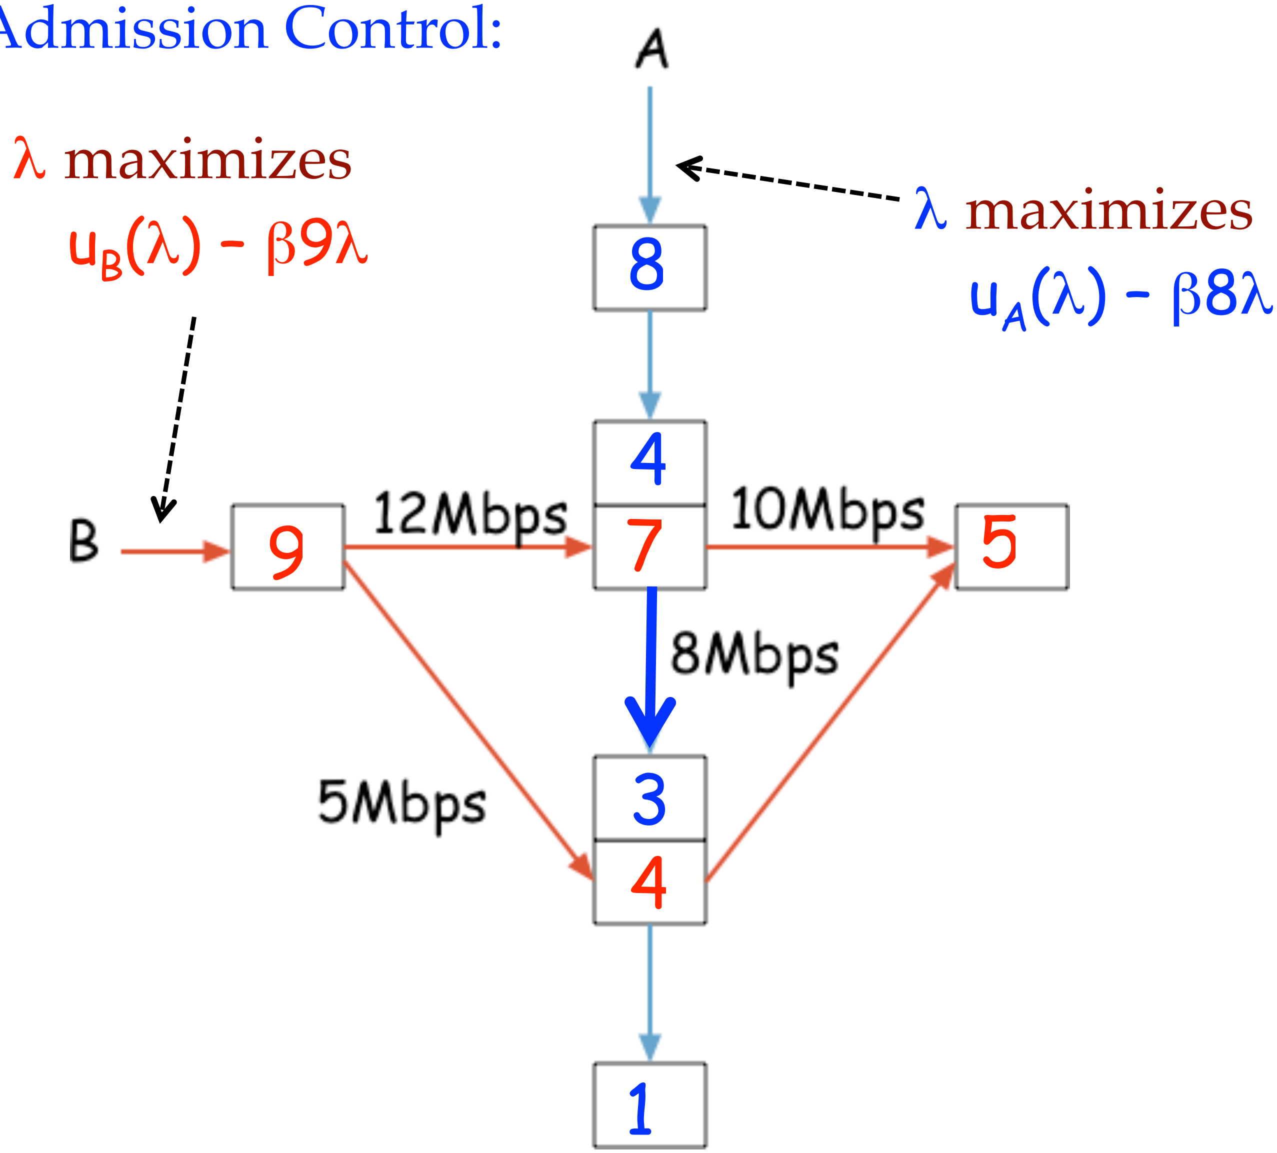

[JW10]

Result:

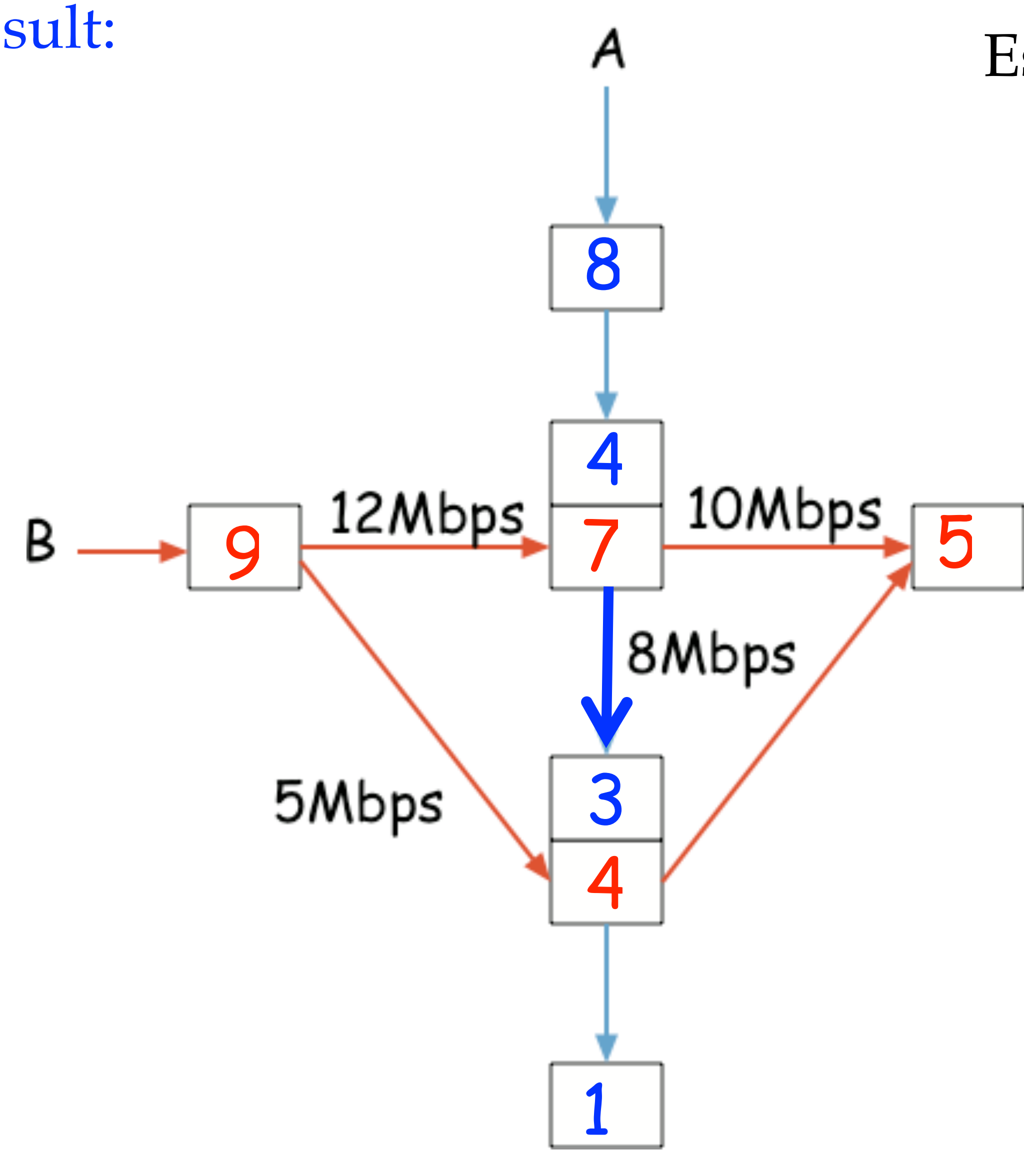

Essentially maximizes the sum of flow utilities

Note: Integrates
- congestion control
- routing
- MAC scheduling

[JW10]

Example 1:

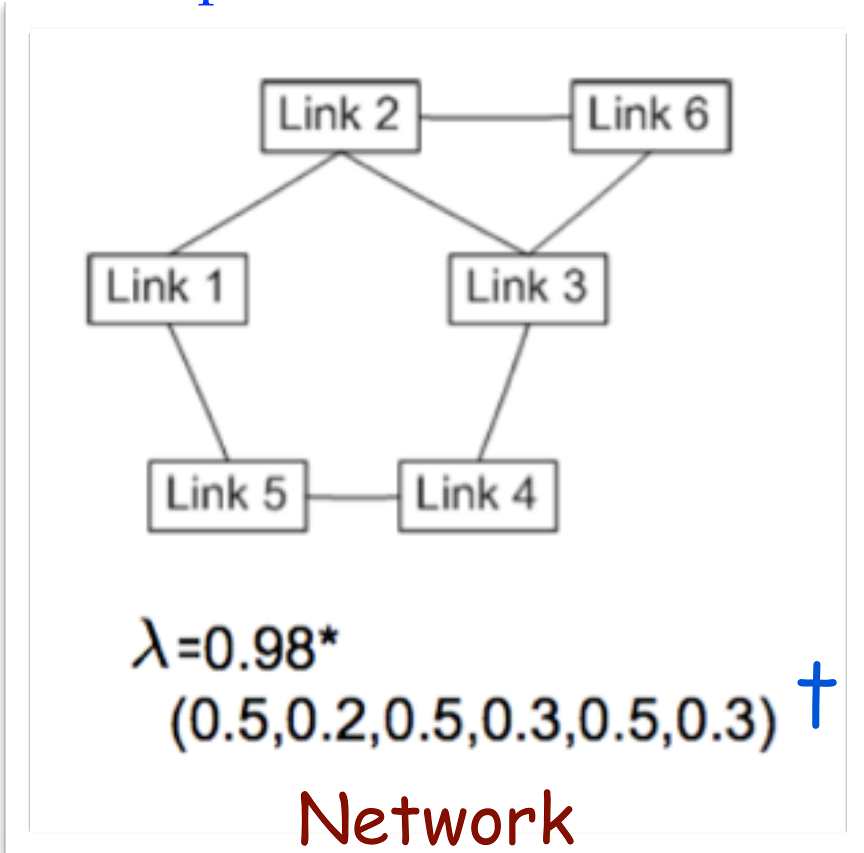

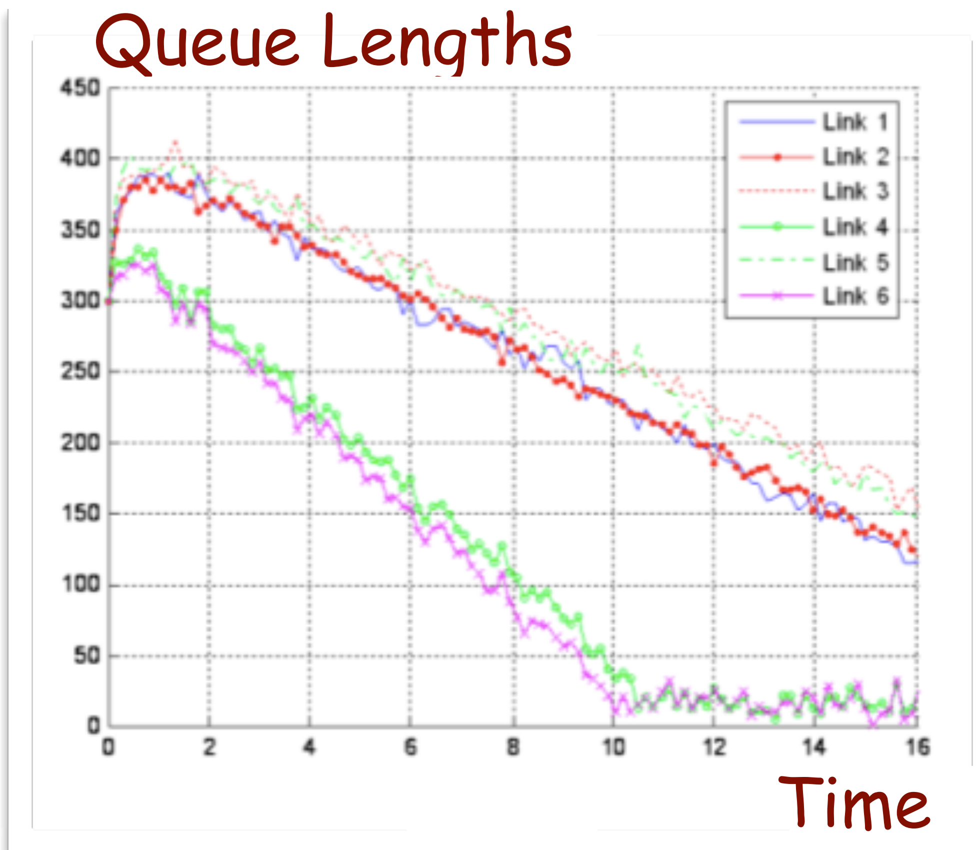

† λ = 0.98*(convex combination of maximal independent sets)

0.2*{1, 3} + 0.3*{1, 4, 6} + 0.3*{3, 5} + 0*{2, 4} + 0.2*{2, 5}

[JW10]

Example 2:

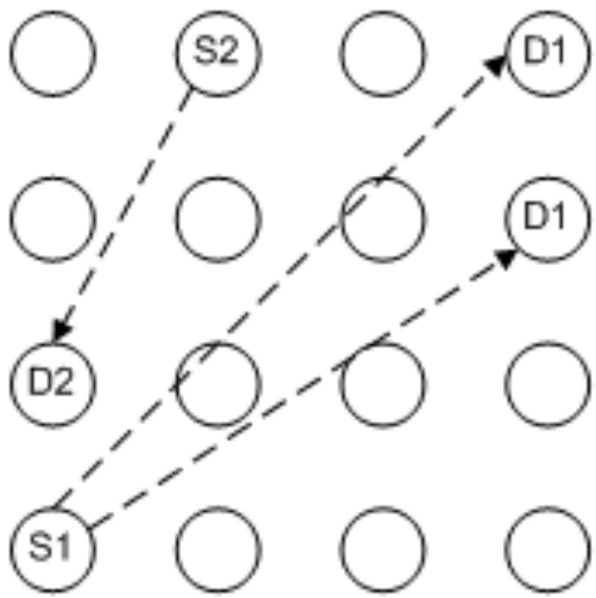

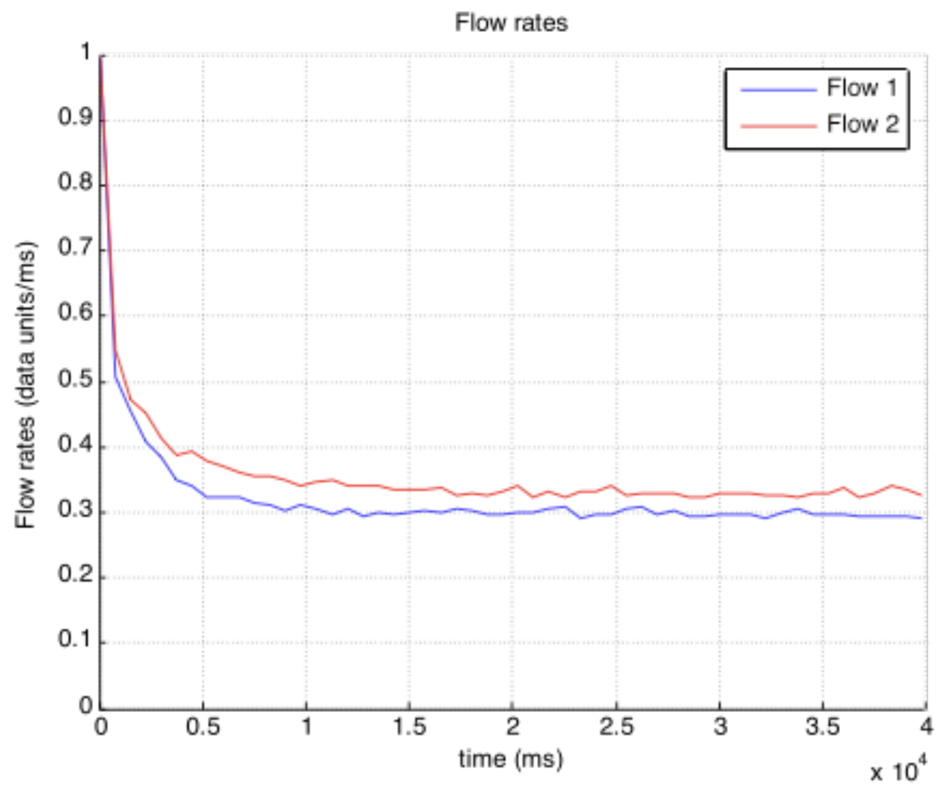

Multipath routing allowed

Unicast S2 -> D2

Anycast S1 to any D1

[JW10]

- Many users compete for resources
  - CPU, Memory in Cloud
  - Energy
  - Wireless Channels
- For scalability, the protocols must be distributed
- The protocols should be efficient and strategy-proof
- Optimal allocation is NP-hard and requires full knowledge

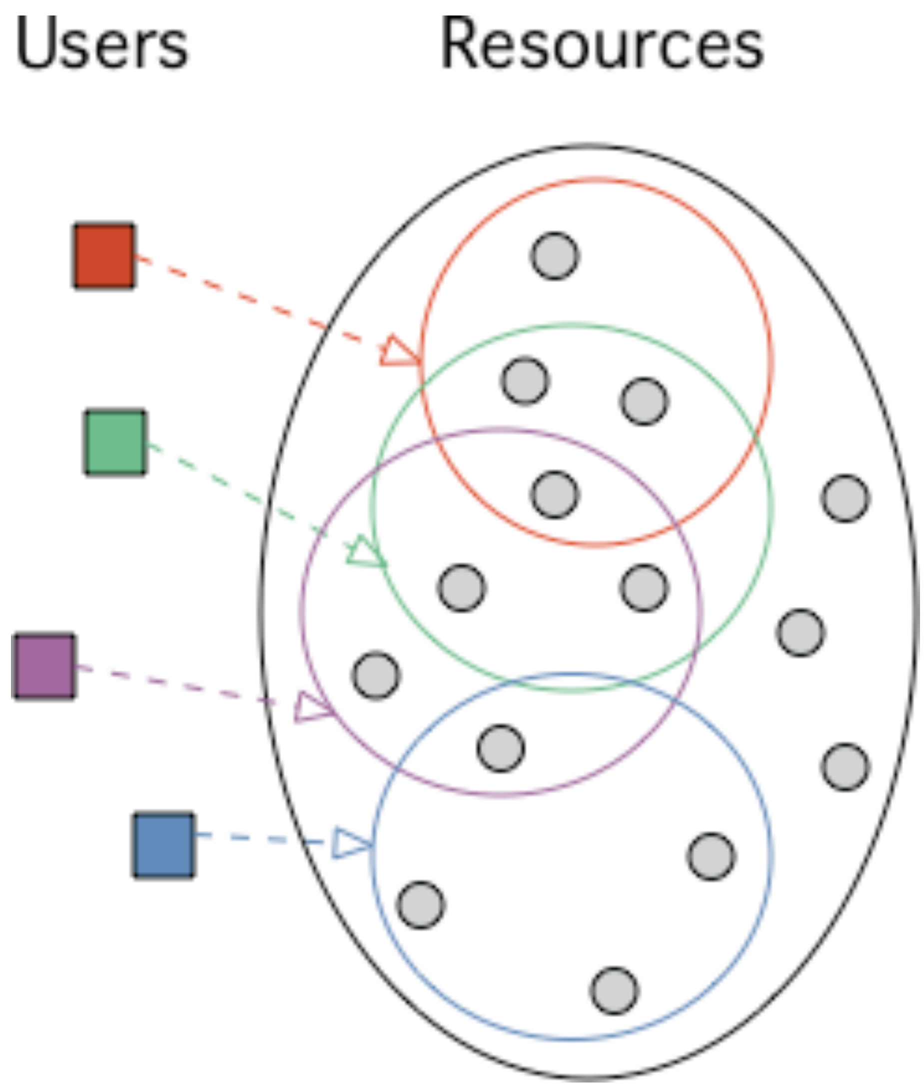

- Replace

$$\text{MAX } \Sigma_i u_i(x_i)$$

by

$$\text{MAX } \Sigma_i u_i(x_i) + \beta H(p)$$

H = entropy of allocation

- Magic:

From NP-hard, the problem becomes

- Distributed
- Easy

The solution is

$O(T/\beta)$-optimal

T = mixing time ….

Bounds on T based on topology of resource conflicts.

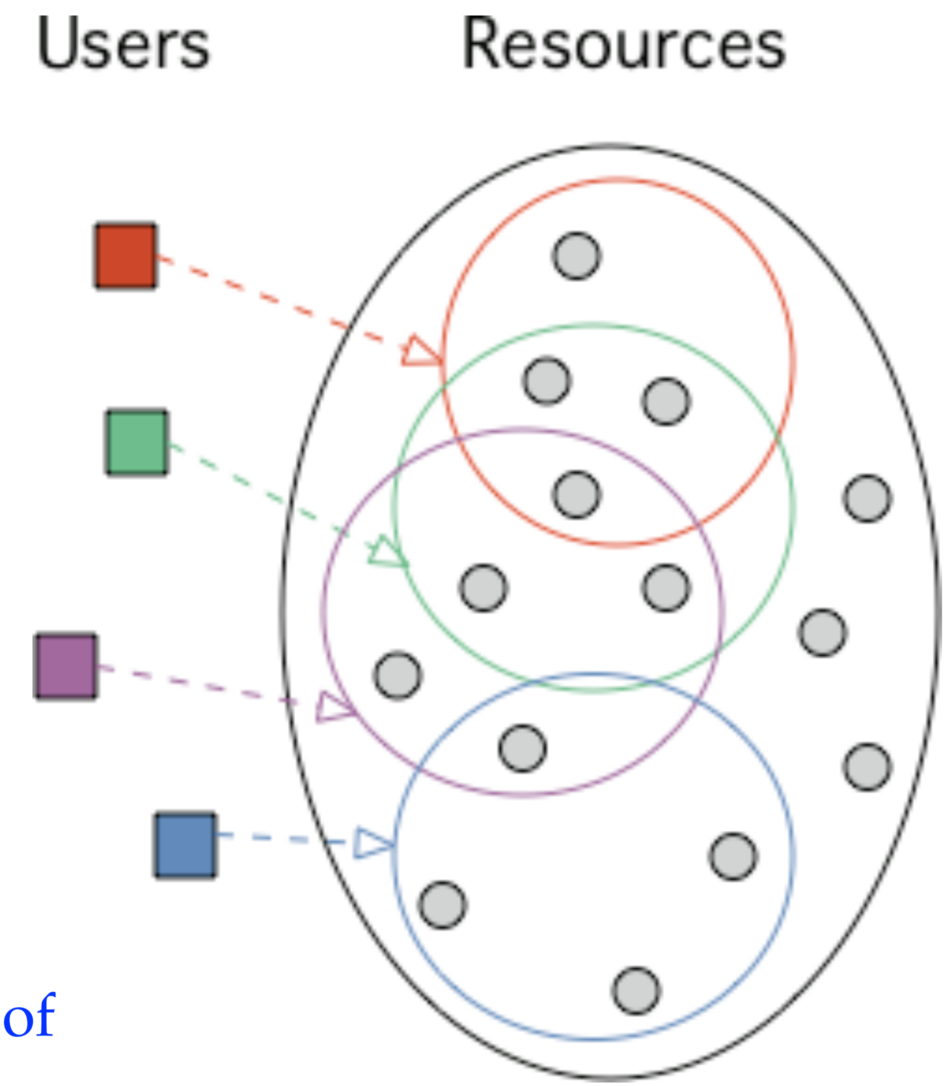

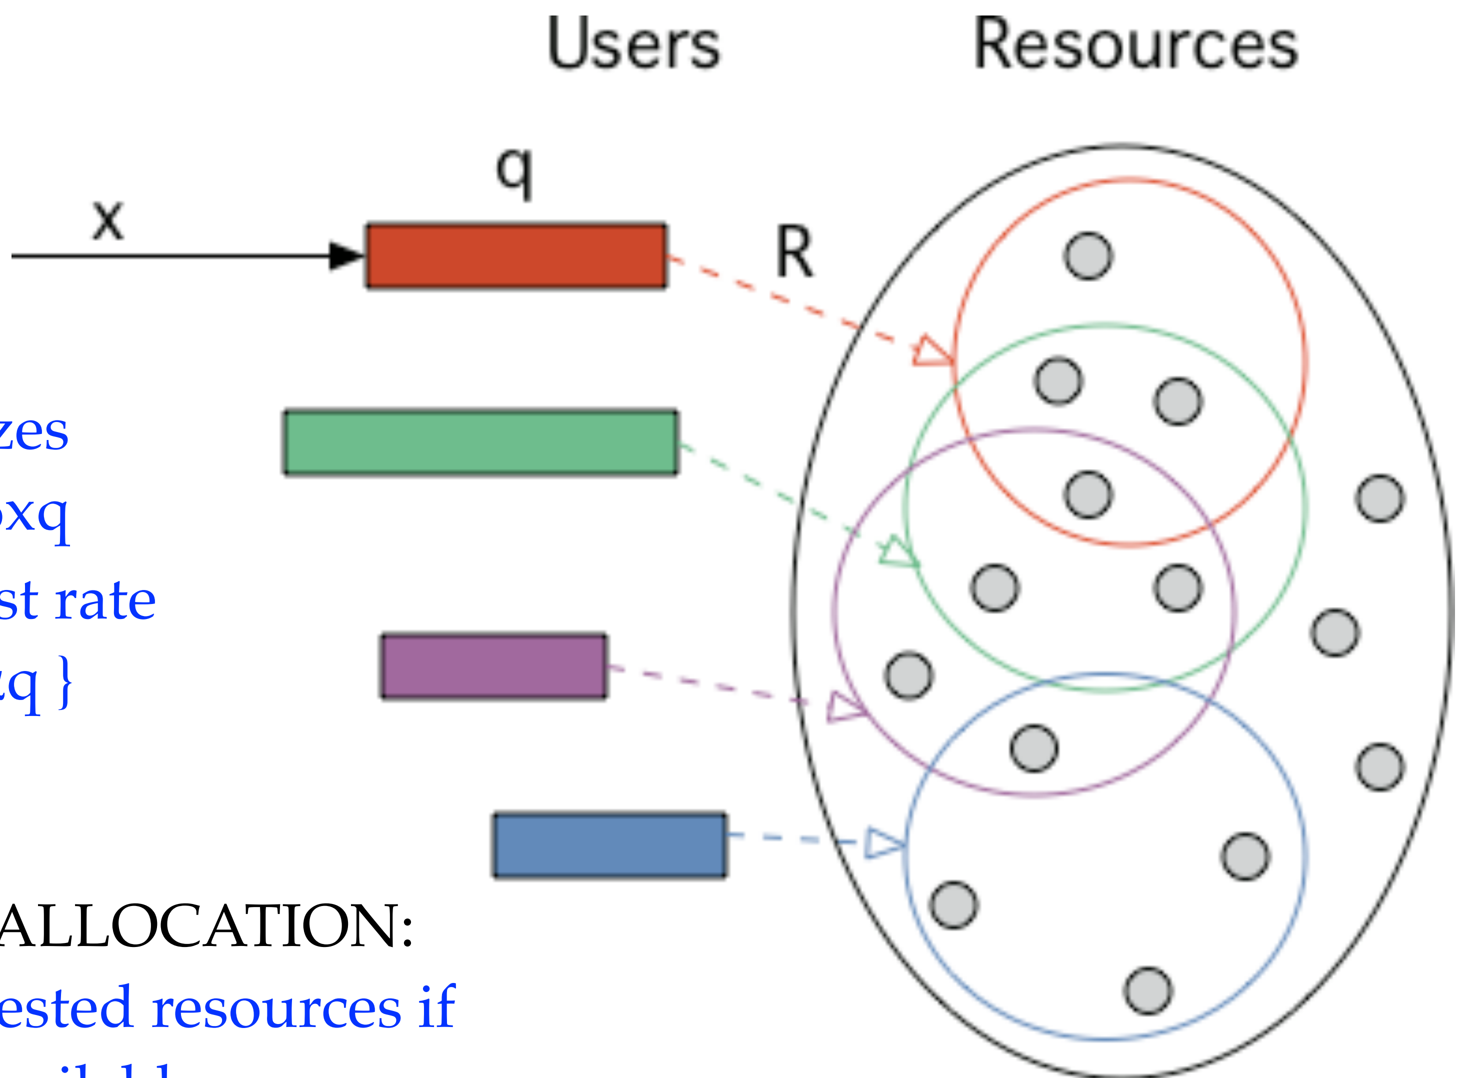

USER i:

- x maximizes $u_i(x) - \beta x q$
- R = request rate $\approx \exp\{\alpha q\}$

RESOURCE ALLOCATION:

- Grant requested resources if they are all available

What about strategic users?

USER i:

- Charge βxq

Intuition:

- Greed is expensive

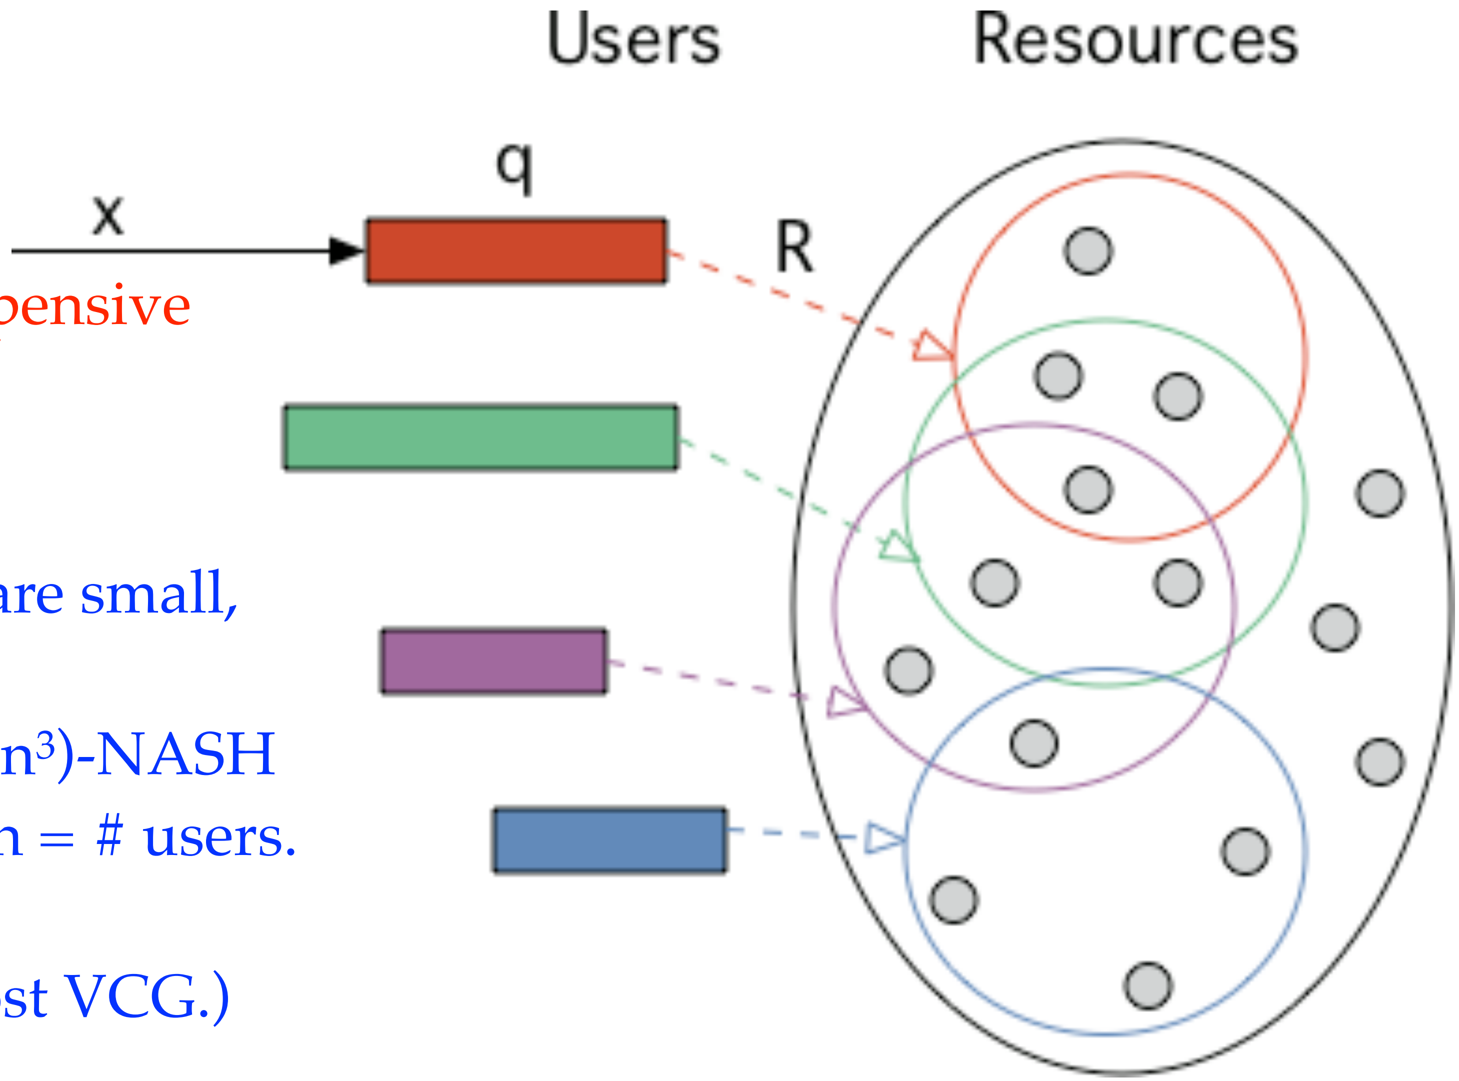

RESULT:

- If all users are small,

scheme is $(1/n^3)$-NASH equilibrium; n = # users.

(Price is almost VCG.)

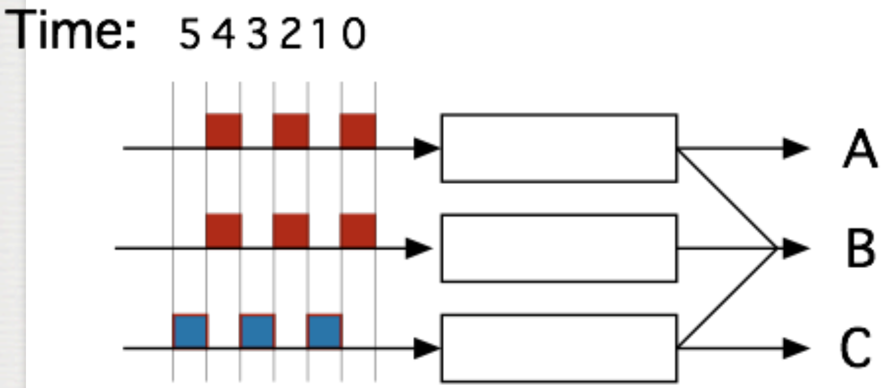

Task: 1 from queue 1;
Task B: 1 from all queues;
Task C: 1 from queue 3

MWM Maximum Weighted Matching is not stable.

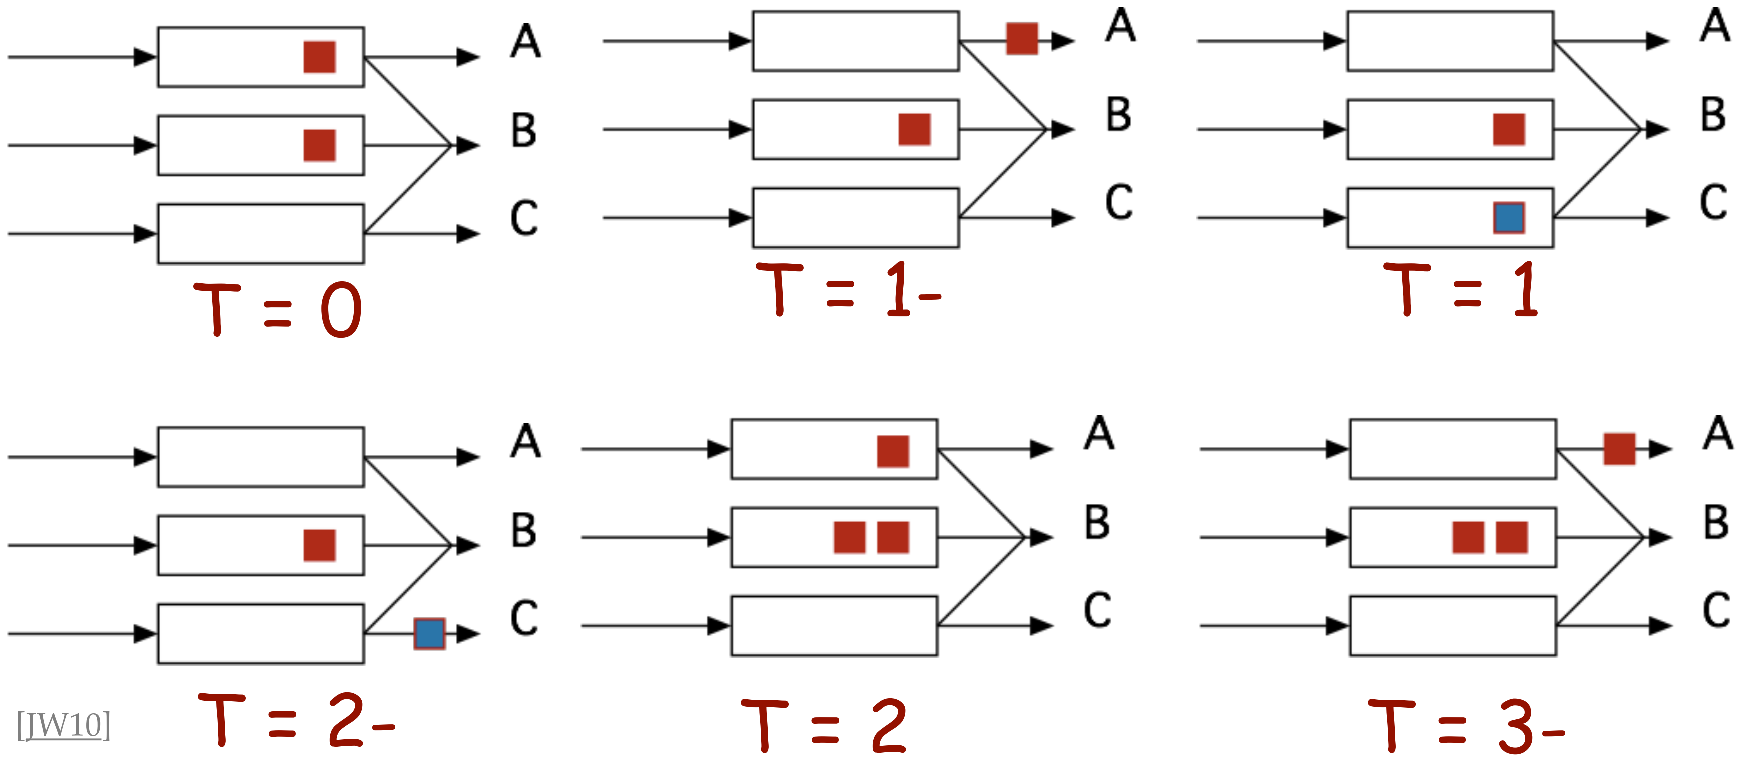

[JW10]

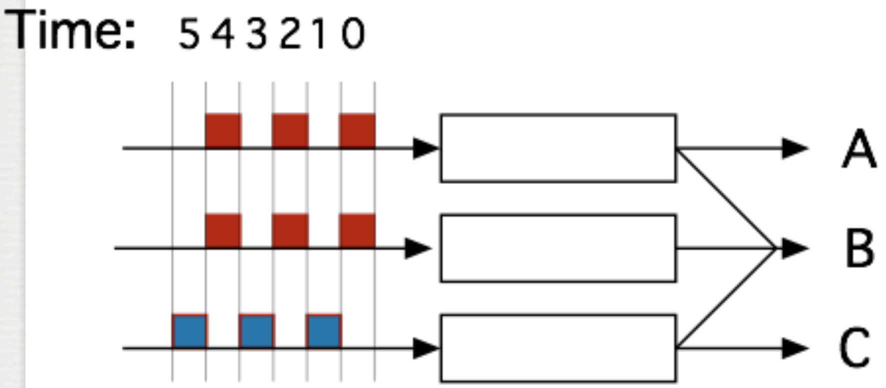

Task: 1 from queue 1;
Task B: 1 from all queues;
Task C: 1 from queue 3

Deficit Maximum Weighted Matching is stable.

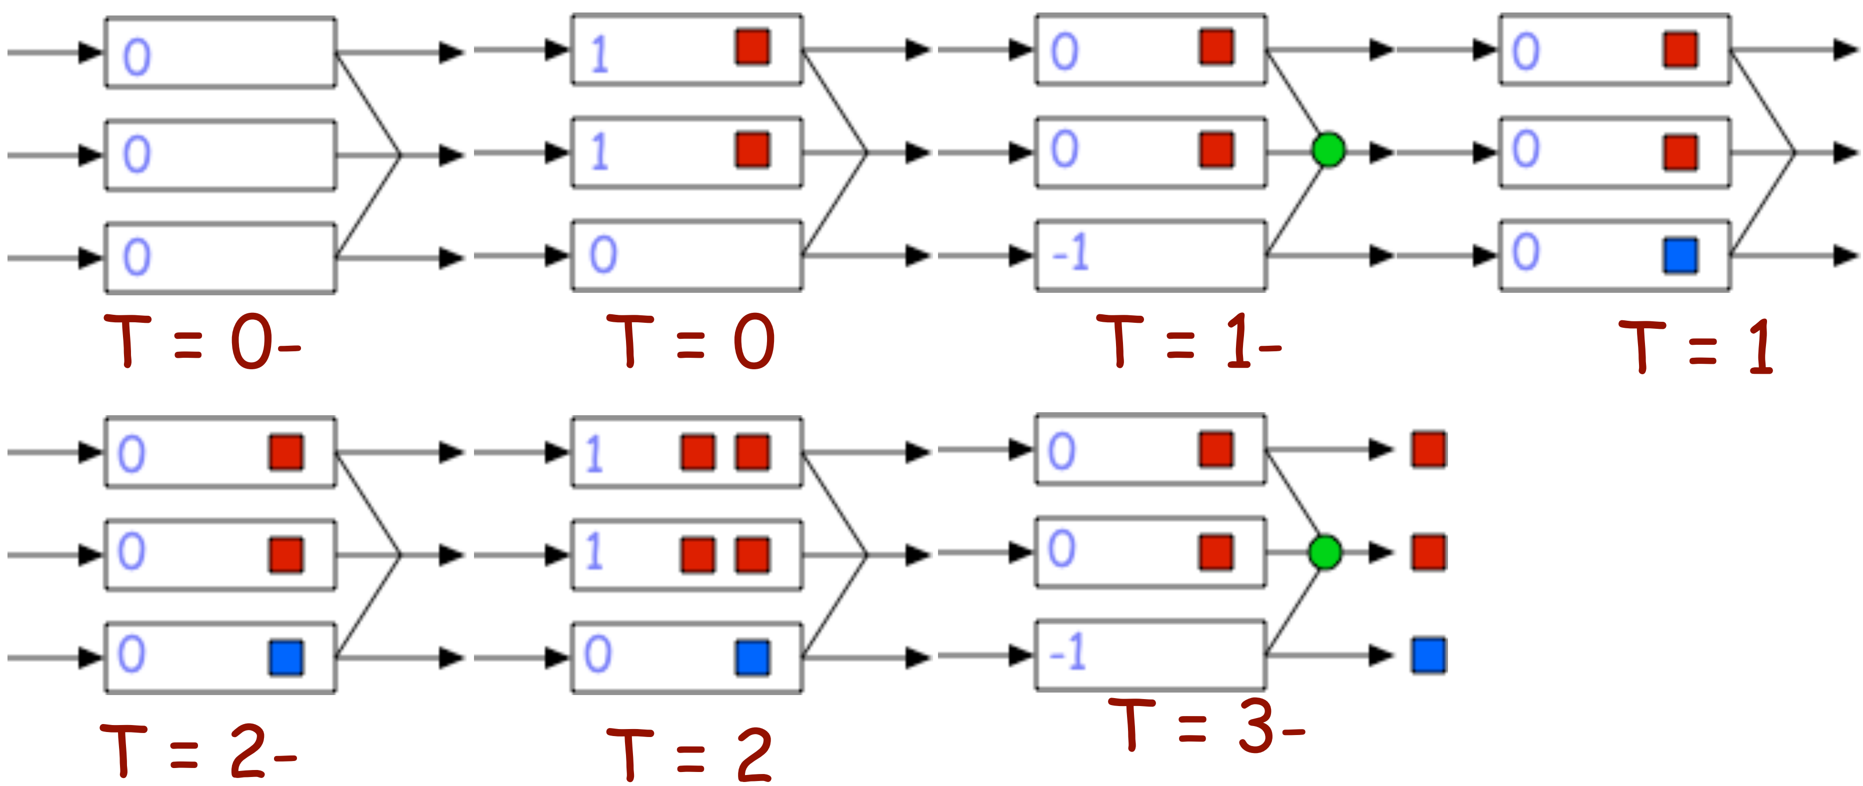

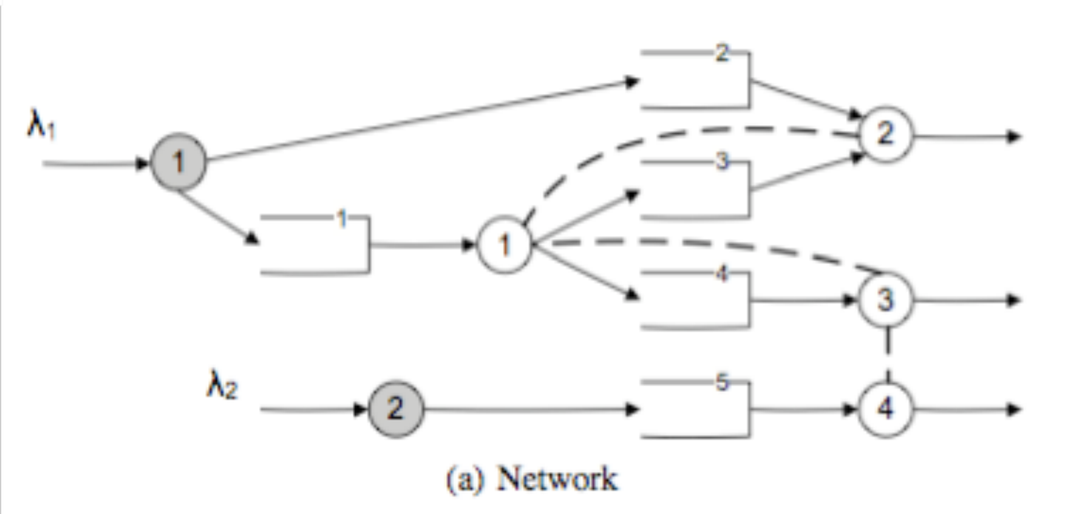

[JW10]

Parts arrive
at 1 & 2 with rate λ_1
and at 5 with rate λ_2

Task 2 consumes one part from 2 and one from 3; ...

Tasks 1-2, 1-3, 3-4 conflict

Algorithm stabilizes the queues and achieves the max. utility

* Formulate allocation as a utility maximization problem with constraints
* The dual algorithm decomposes into individual user problems (max. net utility) and resource problems (price = backlog)
* TCP: user price = sum of link prices
* Backpressure: service rate = control variable. User price = ingress node price. Router serves packet with max. backressure

* Random allocations with adaptive requests rates are ε-optimal in utility
  - The request rates increase with the backlog
  - This price make the scheme almost strategy-proof in a large system
- Processing networks are scheduled based on virtual queues
  - These queues can become negative

TABLE OF CONTENTS

5. PRICING - 5.1. OVERVIEW

Wide range of applications
- different utilities
- different resource requirements

Pricing is designed to maximize
- provider revenue
- user welfare
- social welfare

Pricing should reflect externality and user utility

Congestion Pricing

$$U(x_1) - [x_1 + \cdots + x_n] - px_1$$

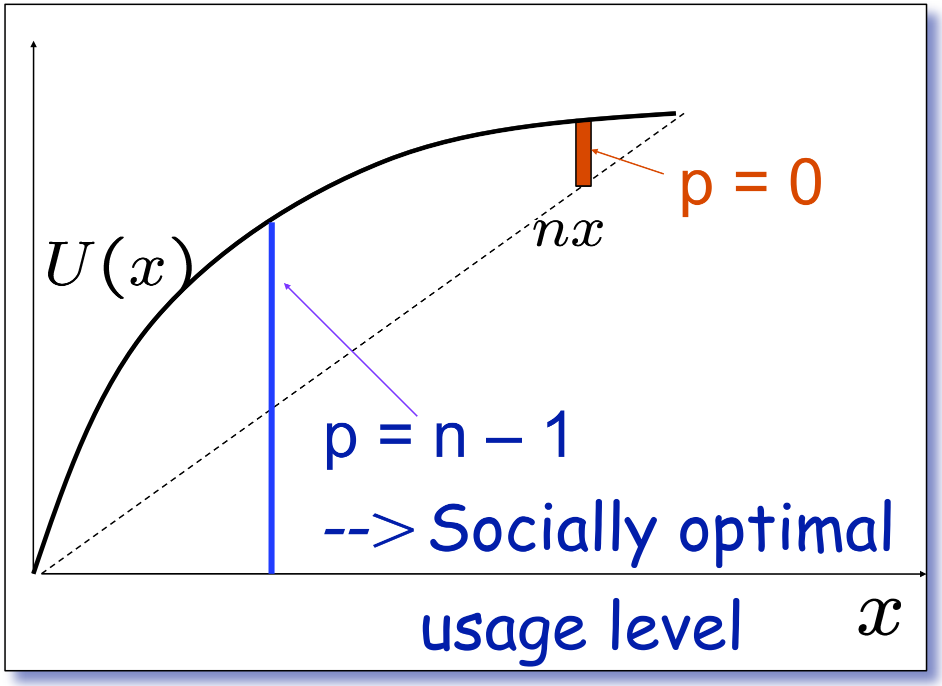

Charge the cost that the user imposes on others:

"Internalize the externality"

[WAL08b]

When to use the network?

Long delay

6 pm

Short delay,
but inconvenient time

8 pm

How can one increase toll to improve social welfare?

[JIA08c]

When to use the network?

Assume $x_i(t)$ users of class i use the network at time $t \in \{1, \ldots, T\}$.

The user "cost" is $\mathbf{x}_i[\mathbf{g}_i + \mathbf{d} + \mathbf{p}]$
where

$g_i(t) =$ inconvenience of time t for user i
$d(t) = d(N(t)) =$ delay at time t
$N(t) = \sum_j x_j(t) =$ load at time t
→ $p(t) = N(t)d'(N(t)) =$ externality cost at t.

> If $d(\cdot)$ is convex increasing, then
> the selfish equilibrium is socially optimal.

Note: The price does not depend on the user inconvenience.

[JIA08c]

Token Pricing

Goals:

Improve social welfare by shifting usage
Maintain a fixed monthly bill

[LMW11]

Token Pricing

Goals:

Improve social welfare by shifting usage
Maintain a fixed monthly bill

Approach:

Users get tokens at a constant rate
To get a better QoS, a user consumes tokens

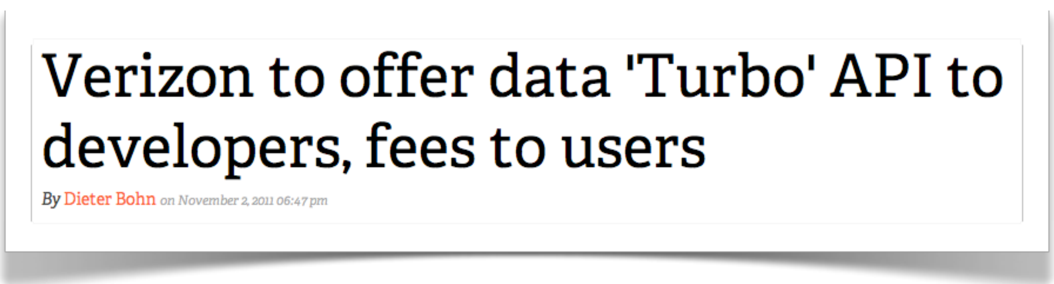
Verizon to offer data 'Turbo' API to developers, fees to users

By Dieter Bohn on November 2, 2011 06:47 pm

Token Pricing

1. Two service classes: H, L
2. Get 1 token every epoch
3. To use H, need K tokens
   To use L, no token needed

Results:

1. Token Scheme improves user welfare (over flat pricing)

2. Improvement larger when capacity is scarce (4G?)

5. PRICING - 5.3. TOKEN

Model:

M tokens per day. “Services” arrive as Poisson(λ).
K tokens to use H; 0 token to use L.

H-service has random added value g(x, p) over L-service
where p = P(user uses H-service) = measure of congestion.

Google x = 1

Skype x = 5

YouTube x = 4

iTunes x = 3

PANDORA x = 2

Model:

M tokens per day. "Services" arrive as Poisson(λ).
K tokens to use H; 0 token to use L.

H-service has random added value g(x, p) over L-service where p = P(user uses H-service) = measure of congestion.

To maximize the average utility of the network,
use H if x > a(s, p) where s = number of saved tokens.

$$p = P(x > a(s,p)); \lambda pK = M$$

Model:

Throughput Service:

$$g(x,p) = Ax - \frac{\lambda p}{c}$$

Latency-Sensitive Service:

$$g(x,p) = Ax - \frac{1}{c - \lambda p}$$

5. PRICING - 5.3. TOKEN

Results:

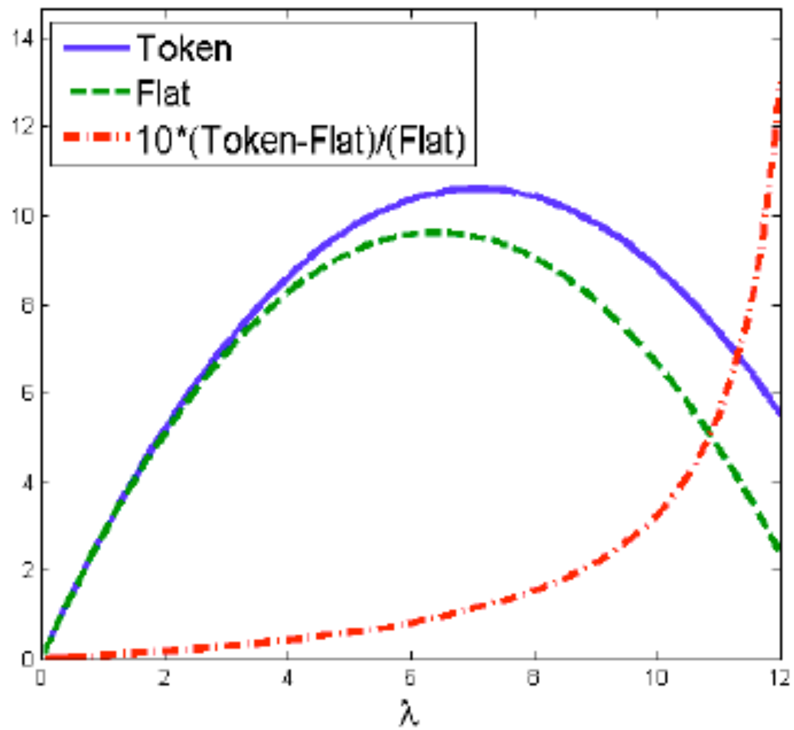

“throughput” services

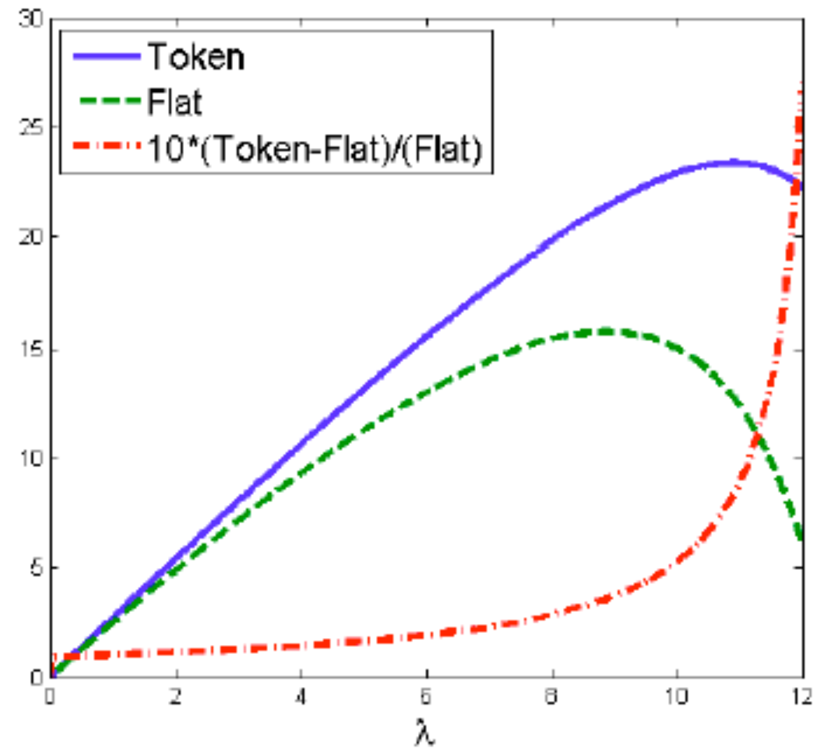

“latency” services

Service Differentiation: Paris Metro Pricing

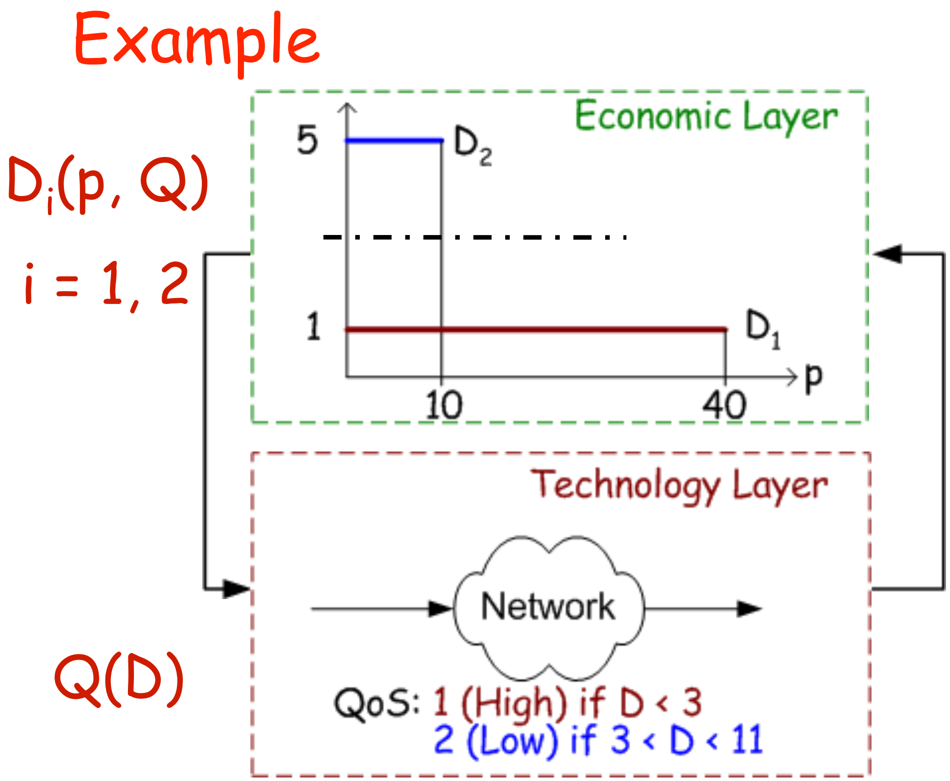

p --> D_1, D_2

--> Revenue

Question:

How to price?

[ODL98]

Service Differentiation: Paris Metro Pricing

Solution 1:

$p = 10$ --> D_2

--> Revenue = 5x10

(Note: $p = 40$ --> D_1 --> 1x40)

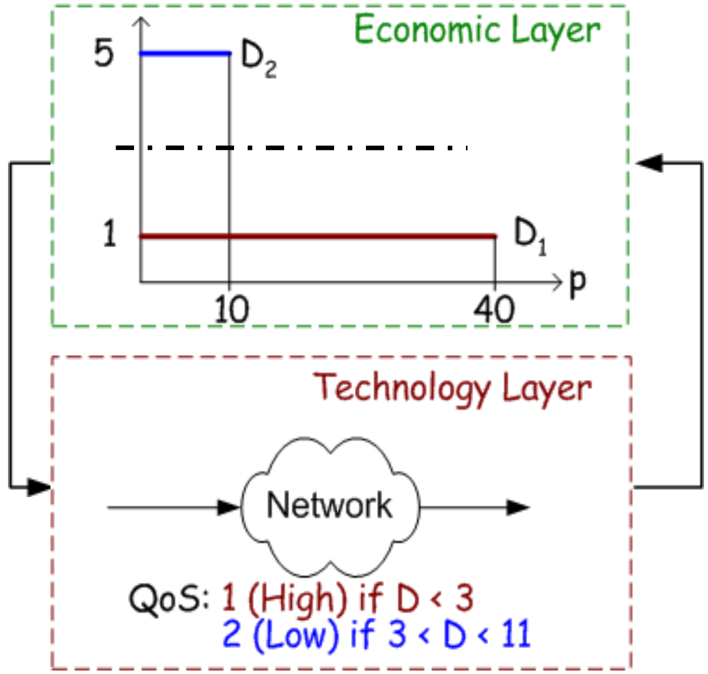

Service Differentiation: Paris Metro Pricing

Solution 2:

$p_1 = 40$ --> D_1

--> Revenue = 1x40

$p_2 = 10$ --> D_2

--> Revenue = 5x10

Total Revenue: 90

Note: QoS achieved
by pricing
(not by QoS mechanisms)

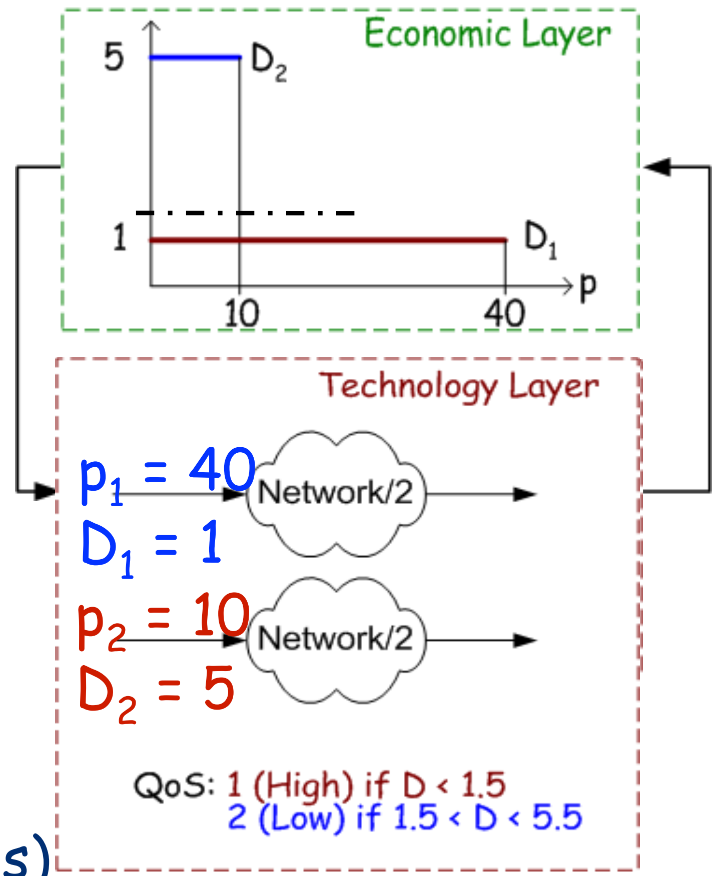

Service Differentiation: Paris Metro Pricing

"Paris Metro Pricing" (A. Odlyzko)

Cheap, but crowded

First class, comfort because of price

Pricing with Competition:

Two competing providers:
A and B

Best response:

$11 \leq p_1 \leq 40$ --> $p_2 = 10$

$p_1 = 10$ --> $p_2 = 9$

$p_1 = 9$ --> $p_2 = 8$ or 40

$p_1 \leq 8$ --> $p_2 = 40$

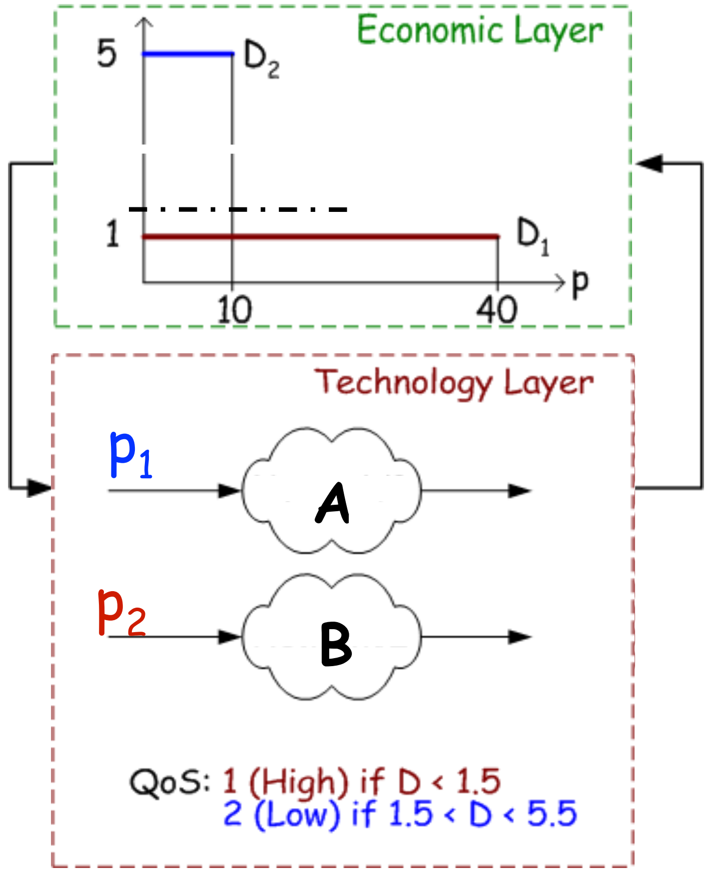

[WAL08b]

Pricing with Competition:

Two competing providers:
A and B

Best response:

$11 \leq p_1 \leq 40$ --> $p_2 = 10$

$p_1 = 10$ --> $p_2 = 9$

$p_1 = 9$ --> $p_2 = 8$ or 40

$p_1 \leq 8$ --> $p_2 = 40$

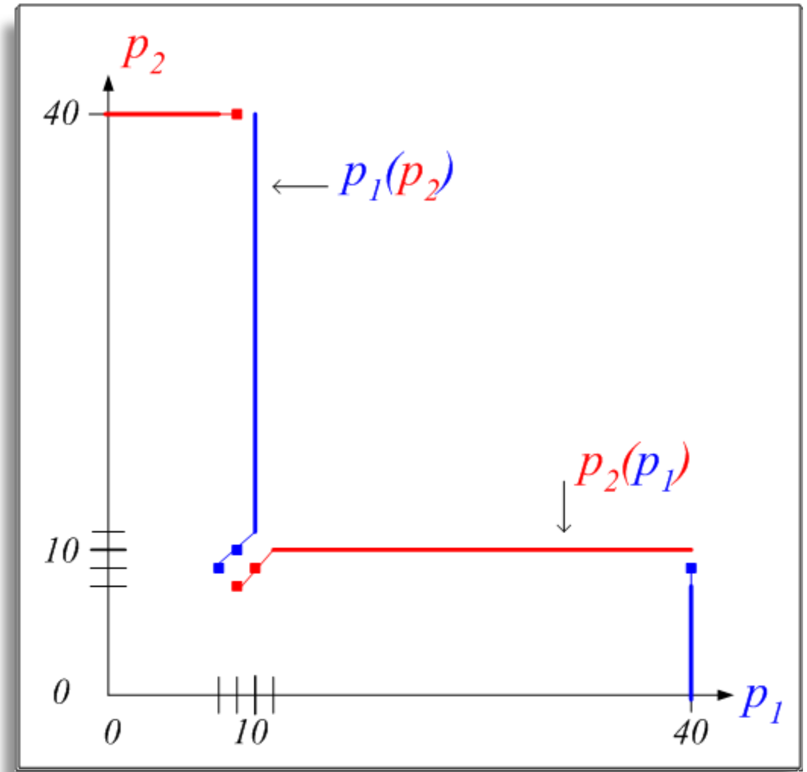

No pure
Nash Equilibrium

Why sell spectrum?

Revenue

Fairness

Efficiency

How to sell spectrum?

Auction: reveals user's utility

Simple Case: Licensed Spectrum, Single Block

Bidders:	at&t	Sprint	T-Mobile	verizon
Valuations:	Y_1	Y_2	Y_i	Y_n
Bids:	X_1	X_2	X_i	X_n

Second Price Auction

Outcome: Spectrum goes to highest bidder

Payment: Second highest bit

Simple Case: Licensed Spectrum, Single Block

Bidders:	at&t	Sprint	T-Mobile	verizon
Valuations:	10	6	8	5
Bids:	9	5	6	7

Second Price Auction - Example

Outcome: Spectrum goes to highest bidder Bidder 1

Payment: Second highest bit 7

Simple Case: Licensed Spectrum, Single Block

Bidders: at&t, Sprint, T-Mobile, verizon

Valuations: Y_1 Y_2 Y_i Y_n

Bids: X_1 X_2 X_i X_n

Fact - Dominant Strategy: $X_i = Y_i$

Proof: $(Y_i - Z_i)1\{Y_i > Z_i\} \geq (Y_i - Z_i)1\{X_i > Z_i\} \quad \forall X_i, Z_i$

$$Z_i := \max_{j \neq i} X_j$$

[VIC61]

(Clear if $Y_i > Z_i$ and if $Y_i \leq Z_i$.)

Mixed Case: Spectrum is either licensed or unlicensed
Single Block

Bidders:

Valuations:	Y	Y_1	Y_2	Y_i	Y_n
Bids:	X	X_1	X_2	X_i	X_n

AT&T bids for exclusive licensed spectrum

All others bid for unlicensed spectrum

Possible Auction:

$X > Z := X_1 + \cdots + X_n \Rightarrow$ AT&T gets spectrum, pays Z

$X < Z \Rightarrow$ Unlicensed, i pays $X\frac{X_i}{Z}, i = 1, \ldots, n$

Mixed Case: Spectrum is either licensed or unlicensed
Single Block

Results (so far ...):

Assume we know Y = 1 and $\{Y_1, \ldots, Y_n\}$ i.i.d. f(.)

One can calculate a symmetric Nash Equilibrium where

$$X_i = g(Y_i)$$

For example:

$$Y_i =_D a + bW$$

$$W =_D \text{ exp. dist.}$$

Note: Underbidding

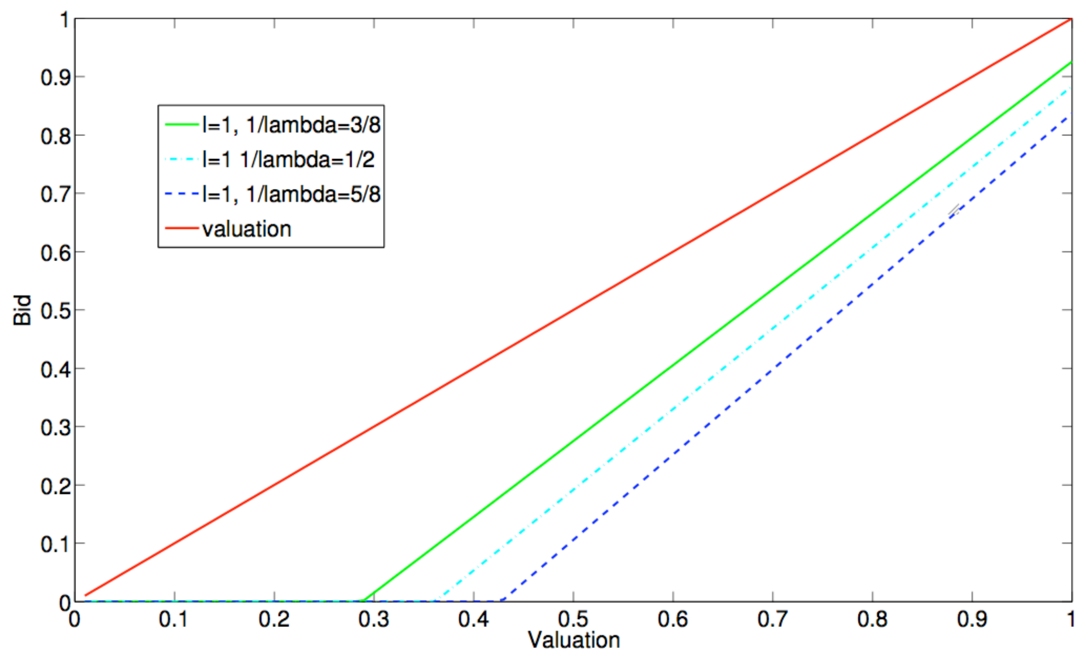

[Silva, Beltran, W., in preparation]

Model

Reward offer
R

Client

Agents

n accept offer
and
collaborate

Utility $U(n) - nR$

or

Accept if $R > K_i$

or if

[D12]

$U(n) - R$	$\frac{R}{n} > K_i$

Analysis

Assume $U(n) = A.1\{n \geq n_0\}$.

Client reward $= A.1\{n \geq n_0\} - R$

$$\text{Agent } i \text{ reward} = \begin{cases} \frac{R}{n} - K_i, & \text{if } n \geq n_0 \\ 0, & \text{otherwise.} \end{cases}$$

Nash Equilibrium:

Collaborate if $K_i \leq \gamma$ where γ maximizes

$E[$ reward of agent i | each collaborates w.p. $P[K_j > \gamma]]$

Client chooses R to maximize his expected reward.

5. PRICING - 5.7. SUMMARY

Pricing can "internalize the externality."

Pricing should reflect externality and user utility, as in congestion pricing.

Token pricing provides an incentive to congest the network only when the user utility is high.

Paris Metro pricing explains why price differentiation leads to quality differentiation and may be self-sustaining.

How should one share revenue among collaborators?

What are suitable incentives to interest potential collaborators?

INDEX

INDEX

REFERENCES

[ACE08] D. Acemoglu, M. A. Dahleh, Ilan Lobel, and A. Ozdaglar, “Bayesian Learning in Social Networks,” MIT 2008

[AHL00] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung. Network Information Flow. In IEEE Transactions on Information Theory, 2000.

[BOY04] S. Boyd and L. Vandenberghe, “Convex Optimization”, Cambridge University Press, 2004.

[CHU05] Shang-Tse Chuang, Sundar Iyer, Nick McKeown: Practical Algorithms for Performance Guarantees in Buffered Crossbars, Proceedings of IEEE INFOCOM 2005.

[COU03] Courcoubetis, C and R.R. Weber(2003). Pricing Network Services, Springer Verlag.

[DIM05] Antonis Dimakis and Jean Walrand, “Sufficient Conditions for Stability of Longest Queue First Scheduling: Second Order Properties Using Fluid Limits,” Journal of Applied Probability, May 2005.

[DHW12] Lingjie Duan, Jianwei Huang, Jean Walrand, “Economic analysis of 4G network upgrade,” JSAC 2012.

[D12] Lingjie Duan et al., “Motivating Smartphone Collaboration in Data Acquisition and Distributed Computing,” Infocom 2012.

[FUD91] D. Fudenberg and J. Tirole: *Game Theory*, MIT Press, 1991.

[GWA11] Assane Gueye , Jean C. Walrand , and Venkat Anantharam, “Design of Network Topology in an Adversarial Environment,” Gamenets 2011.

[GUE08] A. Gueye, and J. Walrand, “Security in networks: A game-theoretic approach, IEEE CDC 2008.

[LMW11] Dongmyung Lee, Jeonghoon Mo, and Jean Walrand, “A Token Pricing Scheme for Internet Services,” ICQT 2011.

[HE06] Linhai He and J. Walrand, “Pricing and Revenue Sharing Strategies for Internet Service Providers,” IEEE JSAC, May 2006.

[JAI04] Jain, R., and P. Varaiya. "Combinatorial Exchange Mechanisms for Efficient Bandwidth Allocation." *Communications in Information and Systems* 3, no. 4 (2004): 305-324.

[JW10] Libin Jiang and J. Walrand. Scheduling and Congestion Control for Wireless and Processing Networks. Morgan-Claypool 2010.

REFERENCES

[JIA08b] Libin Jiang, V. Anantharam and J. Walrand, “Efficiency of Selfish Investments in Network Security,” NetEcon’08.

[JIA08c] Libin Jiang, Shyam Parekh and Jean Walrand, "Time-dependent Network Pricing and Bandwidth Trading", IEEE International Workshop on Bandwidth on Demand (BoD 2008), in conjunction with IEEE/IFIP NOMS 2008.

[LYN96] Nancy Lynch. Distributed Algorithms. Morgan Kaufmann, 1996.

[JOH04] R.Johari AND J.Tsitsiklis, “Efficiency loss in a network resource allocation game”, Mathematics of Operations Research, 2004.

[KAT06] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, *J. Crowcroft, “XORs in The Air: Practical Wireless Network Coding,” SIGCOMM’06,* September 11–15, 2006, Pisa, Italy.

[KEL79] F. P. Kelly, Reversibility and Stochastic Networks, Wiley, 1979.

[KEL98] Kelly, F.P., Maulloo, A. and Tan, D. (1998) 'Rate control for communication networks: shadow prices, proportional fairness and stability', *Journal of the Operational Research Society*, Vol.49, No.3, pp.237-252.

[LEE09] J. Lee and J. Walrand, " Reliable Relaying with Uncertain Knowledge,” GameNets, Istambul, May 2009.

[LOW99] Steven H. Low and David E. Lapsley , “Optimization Flow Control, I: Basic Algorithm and Convergence,” IEEE/ACM Transactions on Networking, 1999.

[MCK99] N. McKeown, A. Mekkittikul, V. Anantharam, J. Walrand, “Achieving 100% throughput in an input-queued switch,” IEEE Transactions on Communications, 1999.

[MO00] Mo and Walrand "Fair End-to-End Window-based Congestion Control," *IEEE/ACM Trans. Networking* 8, 5 (Oct. 2000), Pages 556 - 567.

[MUS09] J. Musacchio, G. Schwartz, and J. Walrand, “Network Neutrality and Provider Investment Incentives,” Review of Network Economics, 2009.

REFERENCES

[NEE05] M. J. Neely, E. Modiano, C. P. Li, “Fairness and Optimal Stochastic Control for Heterogeneous Networks,” In Proceedings of IEEE Infocom, 2005.

[PW12] Abhay Parekh and Jean Walrand. Resource Allocation in Networks. Manuscript in preparation.

[ODL98] A. Odlyzko, “Paris Metro Pricing for the Internet,”ACM Conference on Electronic Commerce, 1998

[ROU02] T. Roughgarden and E. Tardos, “How bad is selfish routing?” Journal of the ACM (JACM), vol. 49, no. 2, pp. 236–259, 2002.

[SCH08] G. Schwartz, N. Shetty, and J. Walrand, “Impact of QoS on Internet User Welfare,” WINE 2008.

[SHA07] Srinivas Shakkottai and R. Srikant, "Network Optimization and Control," Foundations and Trends in Networking, NoW Publishers, 2007.

[SHA08] Srinivas Shakkottai, R. Srikant, Asuman Ozdaglar, and Daron Acemoglu, " The Price of Simplicity."JSAC, Vol. 26, No. 7, September 2008

[SHO06] A. Shokrollahi, "Raptor Codes," IEEE Transactions on Information Theory, vol. 52, pp. 2551-2567, 2006

[TAS92] L. Tassiulas and A. Ephremides, ``Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks,'' IEEE Transactions on Automatic Control, Vol. 37, No. 12, pp. 1936-1949, December 1992.

[VIC61] William **Vickrey,** Counterspeculation, Auctions, and Competitive Sealed Tenders, Journal of Finance, Volume 16, Issue 1 (Mar., **1961**), 8-37.

[WAL08] J. Walrand, “Entropy in Communication and Chemical Systems,” Isabel’08.

[WAL08b] Jean Walrand: Economic Models of Communication Networks, Sigmetrics Tutorial, 2008. Chapter 3 in Performance Modeling and Engineering. Liu, Z., Xia, C., Liu, Z., and Xia, C. (Eds), Springer Publishing Company, 2008. PDF

[WAN05] X. Wang and K. Kar, “Throughput Modelling and Fairness Issues in CSMA/CA Based Ad-Hoc Networks,” Proceedings of IEEE Infocom 2005, Miami, March 2005.

[YAN06] S. Yang and B. Hajek. VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. In 40th Annual Conference on Information Sciences and Systems (CISS06), Princeton, NJ, Mar 2006.