Mechanisms, Institutions and Multi-Agent Systems

Jeffrey K. MacKie-Mason

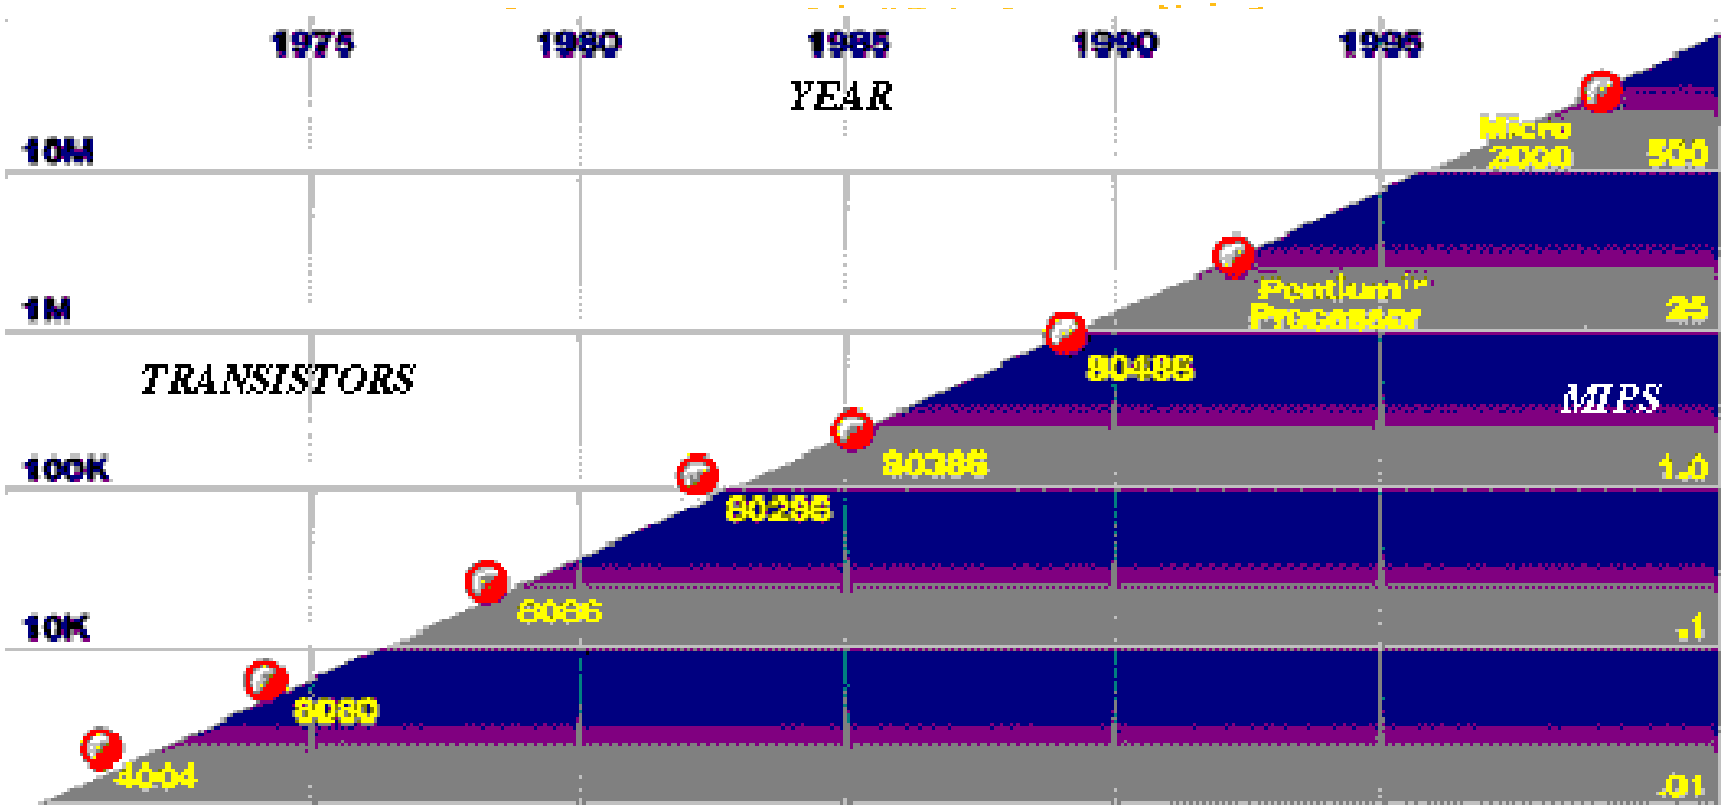
Dept. of Economics and School of Information

University of Michigan



“The problem of bounded rationality is diminished by the use of cheap silicon and sand.”

Silicon



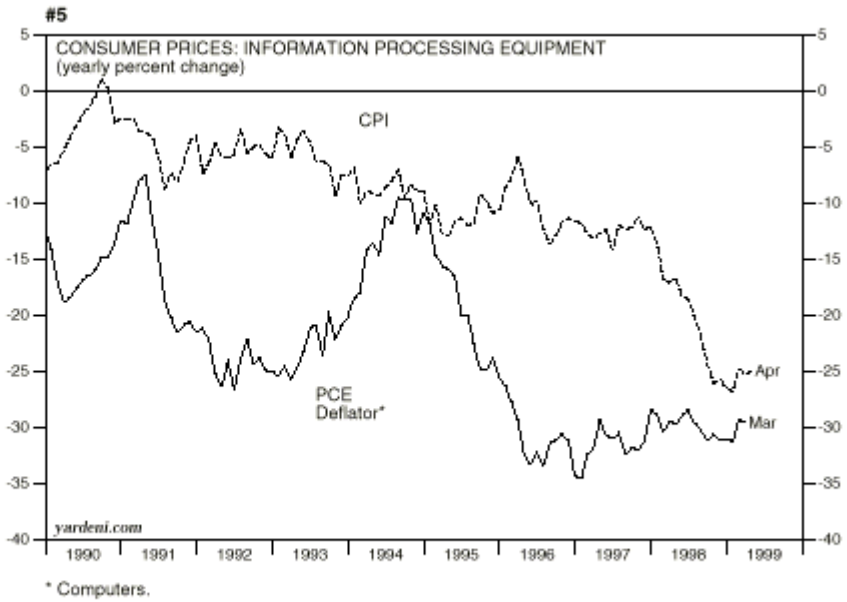
Transistor density doubling every 2 years since 1965

By Y2K, 100 million transistors, 1 billion instructions/sec

Source: Intel Corp.

Silicon Prices

Computer prices paid by consumers are deflating at a 30% annual rate since 1996.

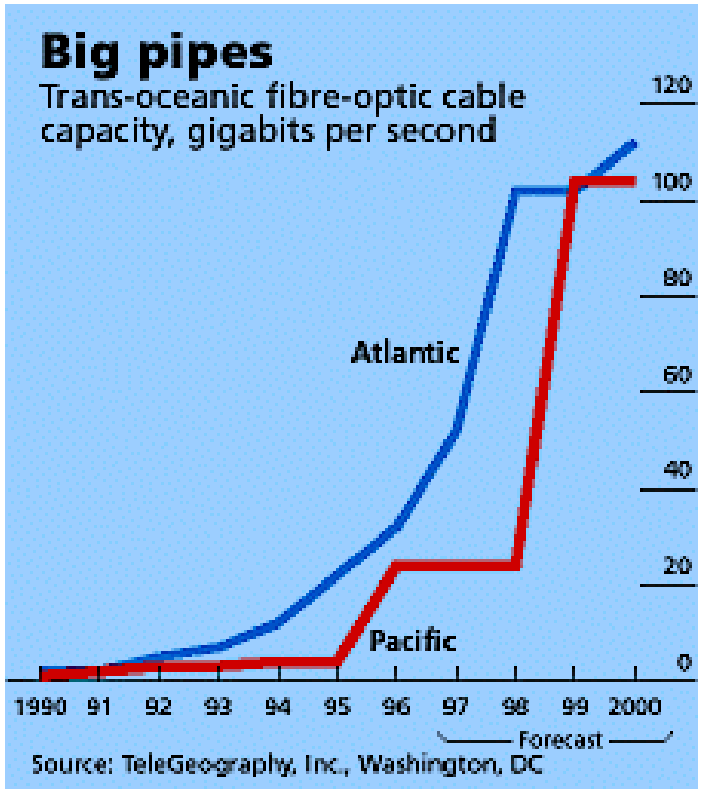


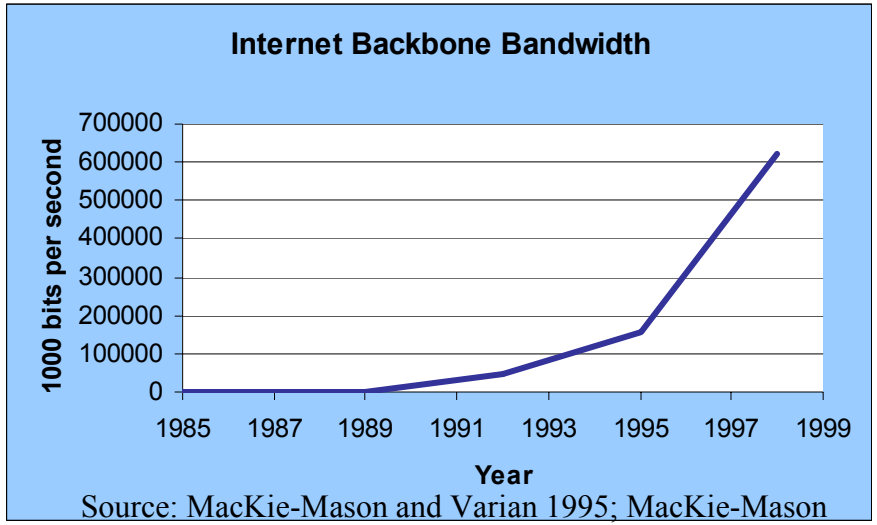


Computers: 30% price drop per year since 1996
GDP inflation: 3.25% without, 1.75% with computers

Source: Ed Yardeni, Deutsche Morgan Grenfell

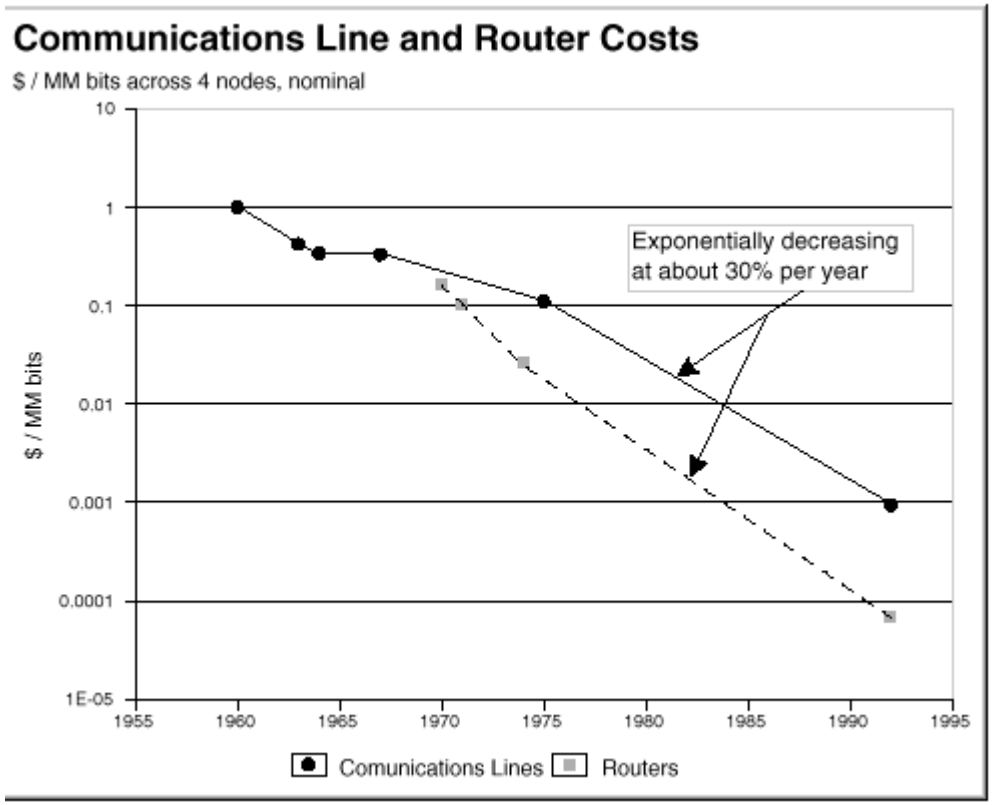
Sand





Qwest has a nationwide fiber backbone with more capacity than AT&T, MCI/Worldcom, Sprint combined. Also has 1.3 million miles of deployed fiber elsewhere.

Sand Prices



Source: MacKie-Mason and Varian 1995

Would you use free communications?

Internet traffic growth:

	Jan-88	Nov-94
packets	85 Million	86 Billion

6% / month since early 1990s,
continuing at least through 1997
This is an annual doubling!

Source: NSF; MCI

What do these numbers mean?

- If faculty salaries fell 30%/year, a $125,000 **business prof would make $6000** in 10 years
- At 6% growth/month, Gates > GE in 23 months

Bill Gates	GM	GE
$100B	$46B	$392B

As of 18 July 1999

What Does This Have to Do With Mechanism Design?

- Multiple autonomous agents, *i*, have some private **information**, θ_i , and **concerns**, $\underline{x}_i$. Agent preferences can be represented by $u(\underline{x}_i, \theta_i)$
- Social choice function maps info state into **preferred settlement** of concerns

$$f:\Theta_1 \times \Theta_2 \times \cdots \times \Theta_I \rightarrow X$$

Mechanism space

- A mechanism defines rules and concerns:

  Mechanism Space = Rules Space x Concerns Space

- Concerns: What are the goods, tasks, &c? that can be transacted?

- Rules:
  - Allowable actions (**strategies** $S_i \in S$)
  - Procedures for **settling concerns**

    $\Gamma = (S_1,\ldots,S_I,g(\bullet))$

    where g maps actions into settlements

    $g: S_1 \times \cdots \times S_I \rightarrow X$

So?

- Much of the mechanism design literature has focused on a subset of rules, taking as given
  - **concerns**
  - **other rules**
- Exponentially falling costs: **everything is in flux**
- Pay attention to the definition of concerns and to the formulation of the "other" rules: **institutions**

1. Designing Rules in MAS

- Designing interplay of **controlled and uncontrollable agents** to accomplish resource allocation goals
  - efficient network support for multiple QoS
  - automation and customization in e-commerce
  - electricity distribution
  - spectrum allocation
  - space lab / railroad / &c. scheduling

Harnessing multi-agent systems

- **First Welfare Theorem**: A competitive equilibrium is **Pareto efficient**
- **Second Welfare Theorem**: Any Pareto efficient allocation can be supported by a **competitive equilibrium**

Harnessing market mechanisms

- Make "**economically-smart**" agents who converge to competitive equilibrium
- If desired, move endowments to obtain preferred allocation from Pareto set

Not too fast...

- Decentralized multi-agent systems provide **poor solutions** for distressingly fundamental problems
- Many fundamental problems have distressingly high computational **complexity**

Problem I: Poor solutions

- Competitive equilibria often **don't exist**, or **aren't efficient**
  - externalities
  - transaction costs
  - increasing returns
  - integer constraints
  - asymmetric information
- **Negative results** for more general mechanisms
- E.g., (Myerson & Satterthwaite) For 2-person exchange, **no mechanism** exists that is always
  - efficient
  - non-coercive
  - no-deficit

Mechanism design theory

- What is the **goal**?
  - efficiency?
  - revenue maximization?
- What are the **constraints**?
  - information
  - voluntary participation
  - assumed agent strategies are plausible
  - low computation and complexity costs
- **Solve** for set of mechanisms that achieve goal subject to constraints

Pragmatic mechanism design

- What **relaxed set of constraints** admits optimal solutions?
- Alternatively, how close do **admissible mechanisms** come to optimality?

Example: Ascending k-Bundle Auction

- Peter Wurman (1999, UM dissertation)
- Suppose N goods, preferences defined on any bundle of 1,…,N goods (monotone on the lattice). An efficient allocation f* exists.
- Theorem: A **payment equilibrium** exists that supports f*: Nash equilibrium in nonlinear (payment,allocation) pairs. (Price equilibria don't always exist.)

k-Bundle payment equilibria

- Consider the bundles that are assigned in f*. Efficiently allocating this is the **assignment problem** (Shapley & Shubik), with known equilibrium payment vector bounds π_{lower} and π_{upper}
- Wurman: Finds a π_0 for unassigned items. Shows that k π_{lower} + (1-k) π_{upper} is a **payment equilibrium for all k $\in$ [0,1]**

k-Bundle Auction

- Iterative bid with multiple exclusive (bundle,payment) pairs
- Information cost: Need not bid on all possible bundles
- Simulations so far with myopic best-response agents
  - 1000 random bundle problems
  - In 918 trials optimal allocation reached
  - Average efficiency 99.8%

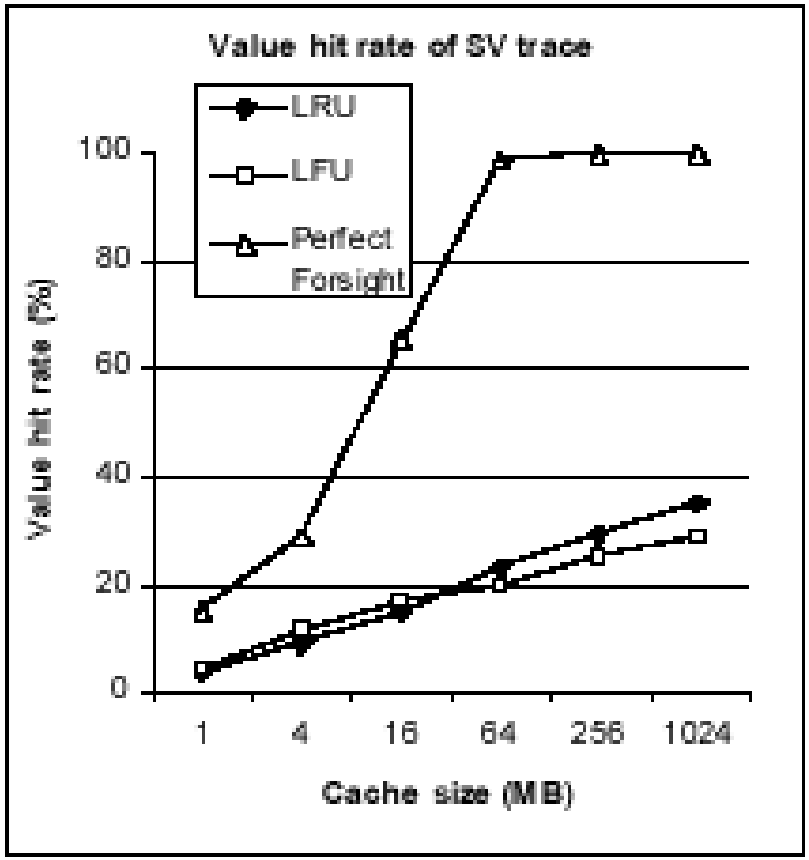
Example: Web Caching

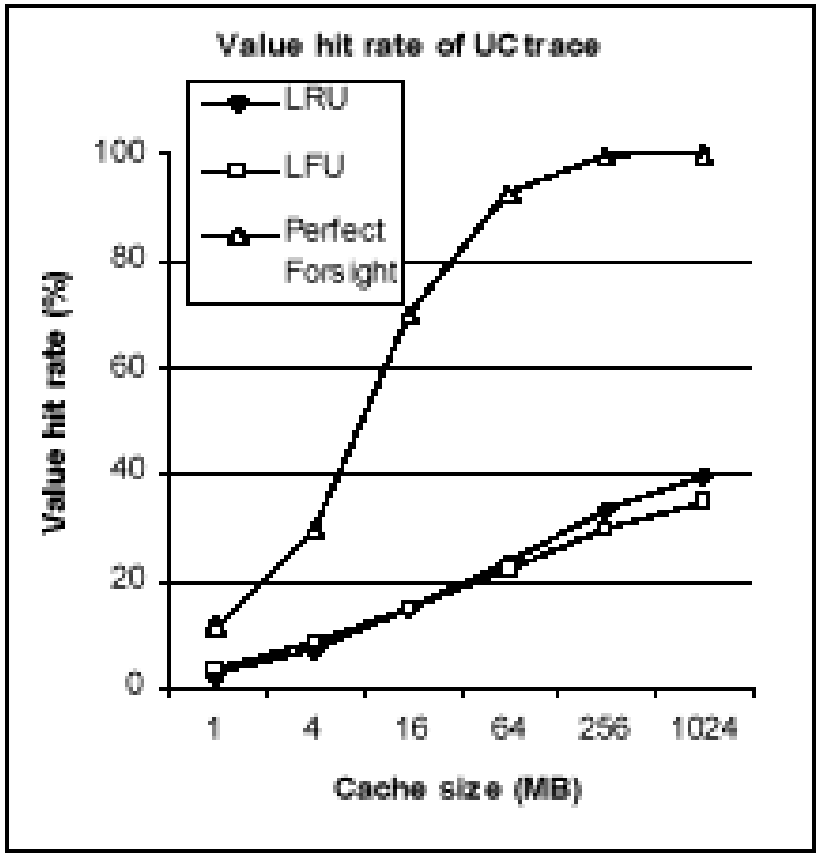
- Network QoS is **unreliable** and **insufficiently differentiated**
- QoS for data transport is one possibility
- **Dynamic distributed file storage** to reduce transport loads is another
- How well might a market-based file caching system work?
- What are the **incentives problems**?

Basic idea: QoS in distributed file services

- File storage characterized along several dimensions:
  - latency
  - reliability (incl. recency)
- Fundamental economic trade-off
  - floppy, HD, tape, CDROM, RAM….LAN, WAN...
- The more flexibly **user needs** for data access can be accommodated, the **higher the value**

Potential for market-based Web caching





Market for retention priority

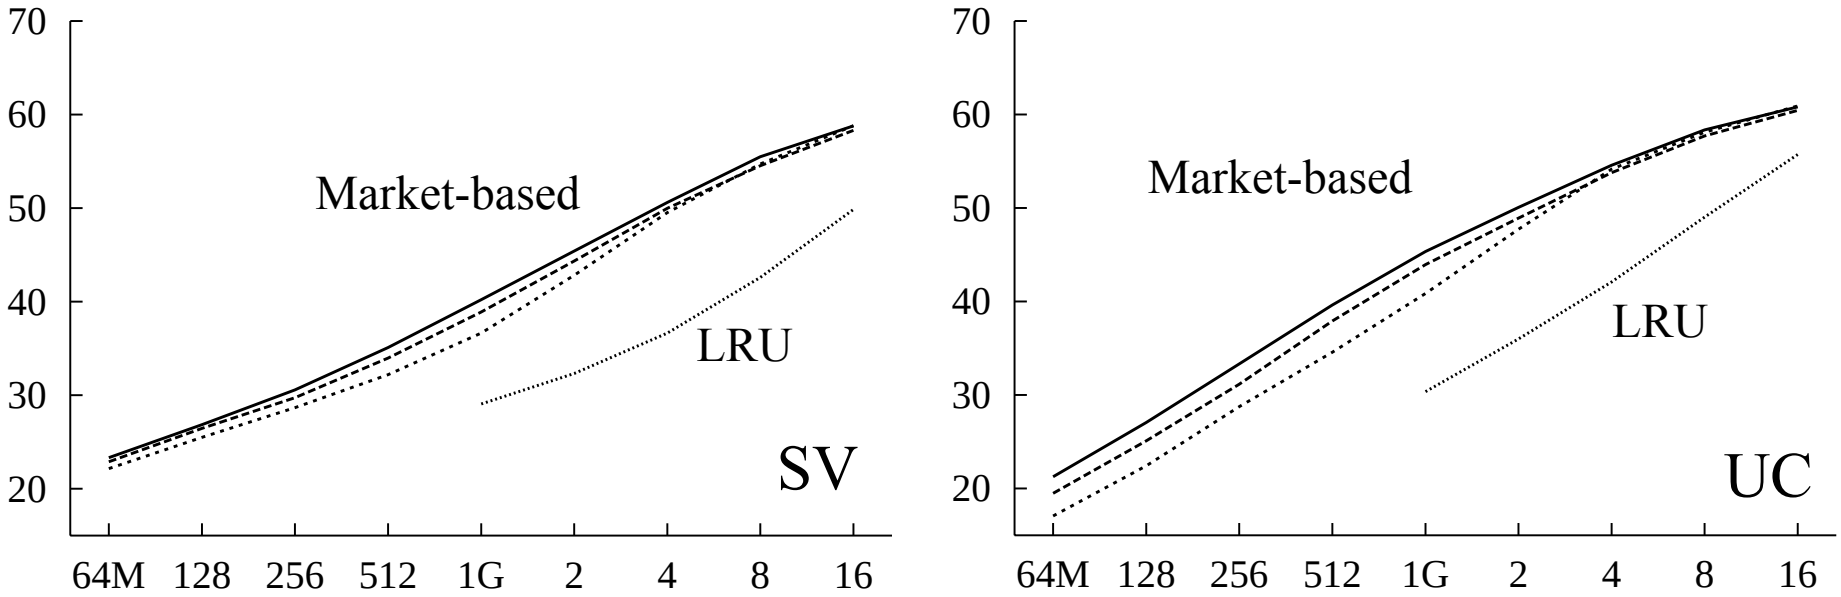
- Simple, **easy-to-implement** extension to standard LFU caching
  - weight frequencies by server-declared value per byte delivered
- Provides **variable QoS**: higher valued file access obtains higher hit rate
- Provides **higher aggregate user value**

Experimental data

- 6 Web cache request streams from NLANR, March 1999
- Example: Champaign-Urbana

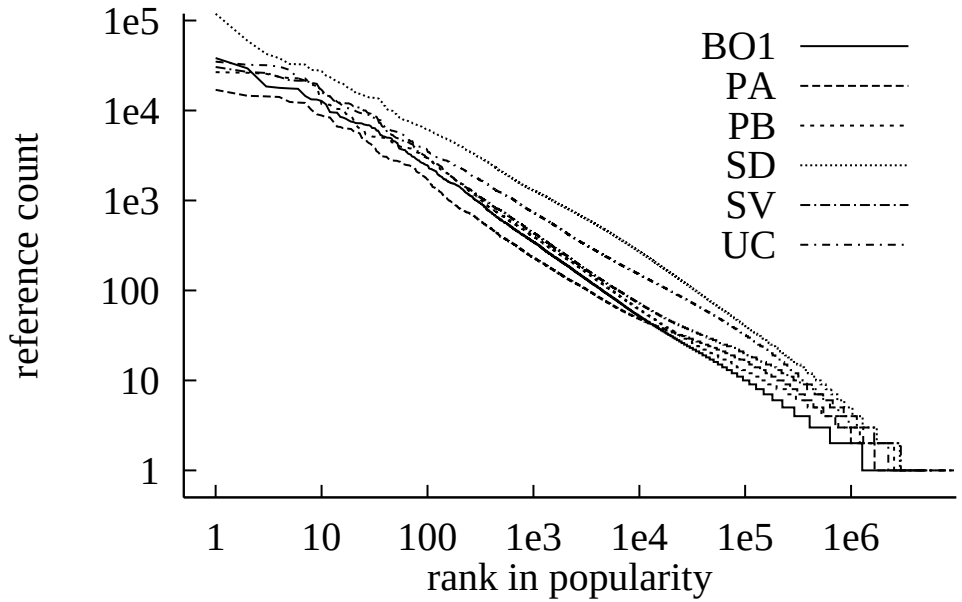
	Univ Ill. Urbana-Champaign	
# requests	**26,024,662**	
# servers	**250,484**	
# documents	**7,615,462**	
bytes requested	**397,548,684,913**	
mean # requests	**3.4**	
Infinite Cache		
unique bytes	**150,119,984,279**	**(150 GB)**
max. hit rate	**70.7%**	
max. byte HR	**64.4%**	

Results



Value hit rate as function of cache size for three market-based variants and LRU for comparison.

Problem: Exponential decline in frequency of requests



- **Zipf-like** distribution: Number of requests to ith most popular is proportional to $1/i^{\alpha}$, $\alpha \approx 0.7$
- **Need wide variance** in valuations to outweigh frequency dominance

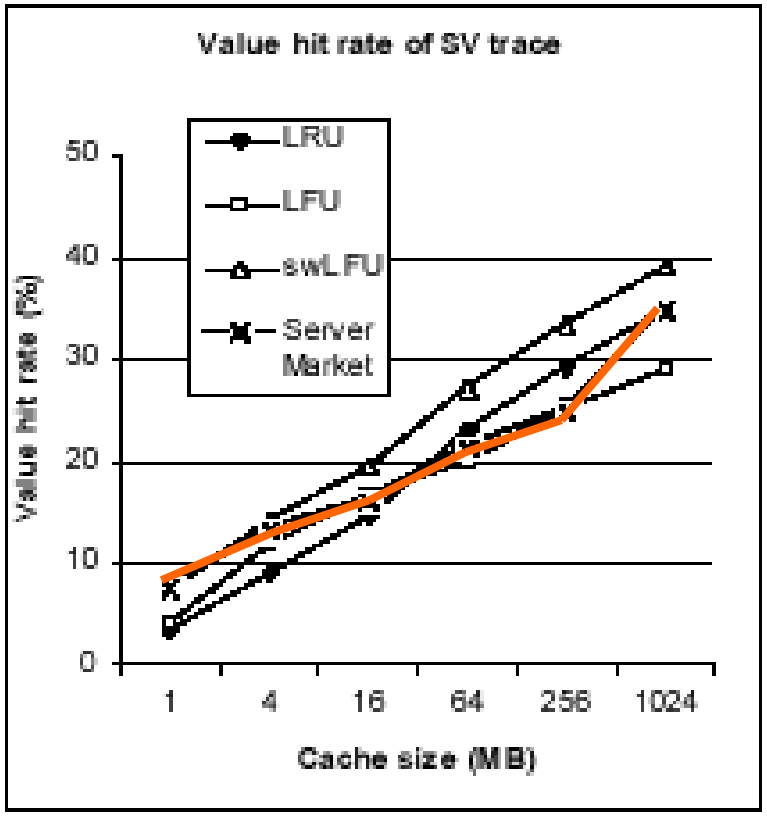
Problem: Incentive compatibility

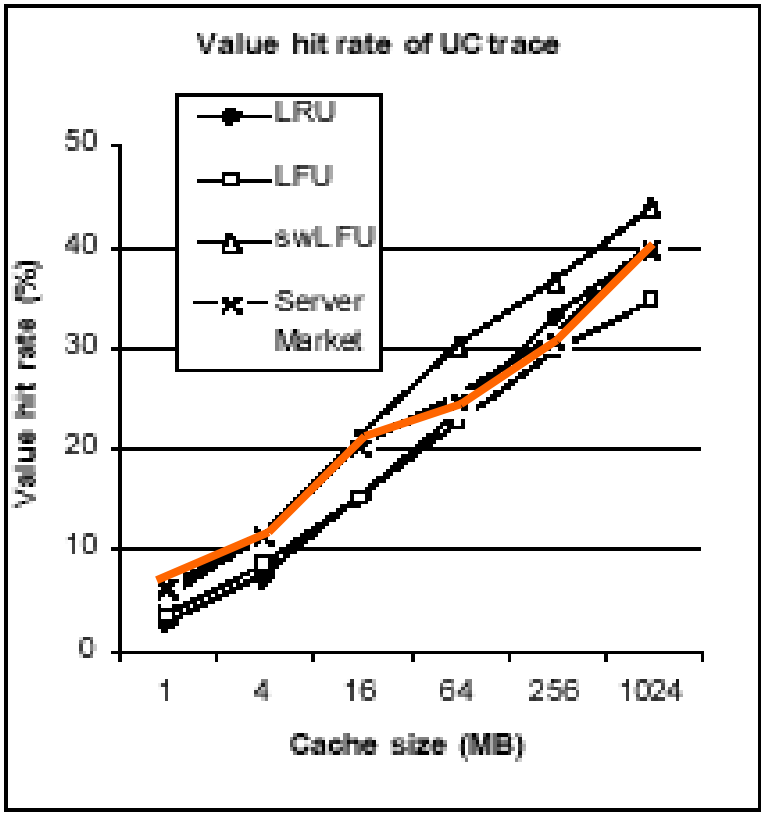
- Our results **assume servers truthfully report** their values.
  - Persistence of weighted reference counts undoes simple incentive-compatible schemes
  - E.g., "Vickrey" price: Suppose we use bid of 1st evicted document
    - current bid determines how long something stays in the cache
    - if expected current price < value, then overbid to extend time in cache

Bounding cost of incentives: market for disk space

- Chan, Womer, MacKie-Mason, Jamin (1999)
  - auction off write access (universal read access)
- One-shot 20 min. clearing, agents bid to control weights on document placement in cache
  - **incentive compatibility possible** with, e.g., 2nd price auction

Performance of disk space market





For small cache sizes, can match A-swLFU algorithm of Kelly et al., with incentive compatibility

Other Examples from MARX

- **Facility scheduling**: Wellman, Walsh, Wurman, MacKie-Mason (1998)
- **Packet handling**: Jamin, Calloway, MacKie-Mason (1999)
- **Task networks**: Walsh and Wellman (1999)

What to do absent analytics?

- Human-subject **experiments**
  - costly
  - incentive alignment hard
  - limited to simpler mechanisms
- Agent **simulations**
  - search for "reasonable" strategic behavior
  - study replicator dynamics to find performance with successful strategies

More “what to do”

- **Learning agents**:
  - implicit modeling (e.g., neural nets) of rational strategies
  - Brooks, Fay, Das, MacKie-Mason, Kephart & Durfee (1999)
    - In electronic marketplace for info goods, what price and bundle strategies does monopolist learn when consumer preferences are unknown?
    - Next: Duopolists need to learn each other’s strategies

2. Concerns in MAS

- Over which goods, tasks, plans are we negotiating an allocation?
- A mechanism may max $\Sigma u_i(\underline{x})$ with respect to $\underline{x}$ as defined, but higher aggregate utility possible for some $\underline{x}'$
- Example: **PEAK Project**
  - We can deliver 1200 Elsevier journals as about 250,000 separate articles, or about 10 million separate pages

Example: Learning in Dynamic Environment

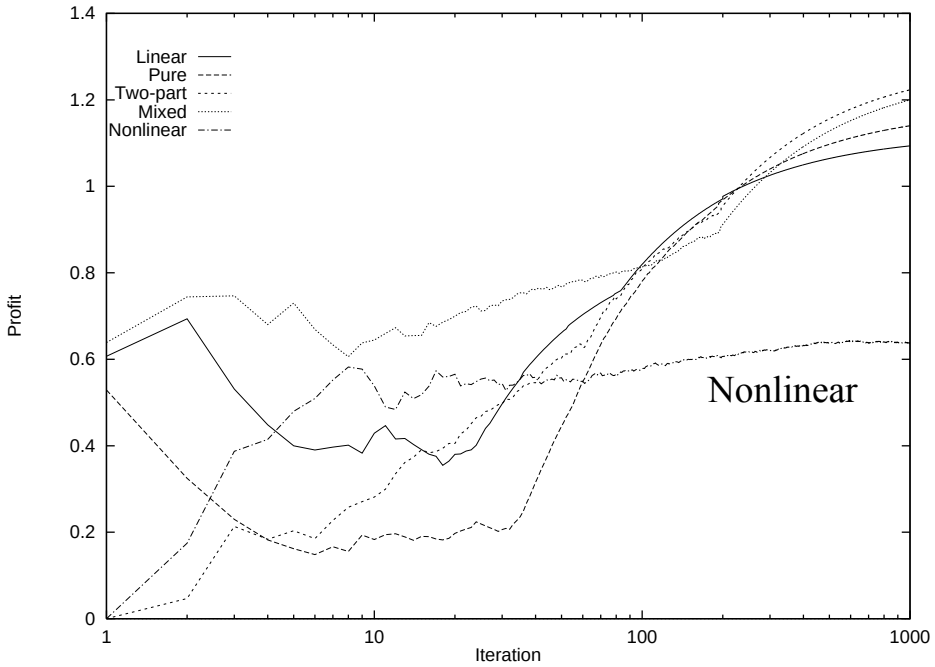
- Brooks, Fay, Das, MacKie-Mason, Kephart, Durfee (EC-99):
  - monopolist with 10 information objects that are periodically updated (news stories)
  - large consumer population with **heterogeneous preferences** for objects
  - anonymous pricing (cannot track who reads what from period to period)

Complete information possibilities

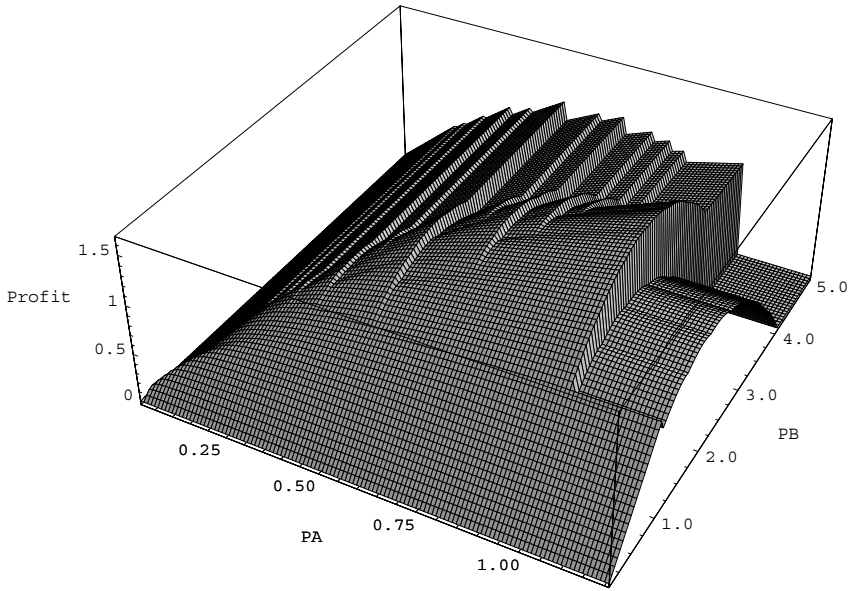
Pricing Scheme	Parameters	Profit
Pure bundling	1	0.875
Linear pricing	1	0.875
Two-part tariffs	2	1.037
Mixed bundling	2	1.037
Block pricing	3	1.094
Non-linear pricing	N	1.152

See Brooks et al. (1999) for assumptions about preferences. N=10 items.

Profits with Learning



Average profits as producer learns preferences



Profit landscape for mixed bundling

Even after 1000 transactions, nonlinear pricing has reached **less than 50%** of its profit potential.

3. Institutions and MAS

- In complex, dynamic agent environments, a mechanism configuration will be at best **approximately satisfactory**, and **temporary**
- Need to consider mechanism adaptations and meta-dynamics

Mechanisms are Approximately Satisfactory

- Gibbard and Satterthwaite assumed merely
  - (1) agents pursue dominant strategies if they have them; and
  - (2) preferences are rational (ordinal, transitive and complete)
- Then, the only implementable social choice functions are **dictatorial**
- What can we do to implement more desirable social choice functions?

Approximation due to Complexity, too

- Silicon and sand are getting cheap fast, but bounded rationality is still important, especially for real time decision-making
- Example: "Smart market" for gateway priority (MacKie-Mason and Varian 1994):
  - when congested, service the packet with the highest value for immediacy
  - agents bid their value for immediate service
  - M+1st price auction

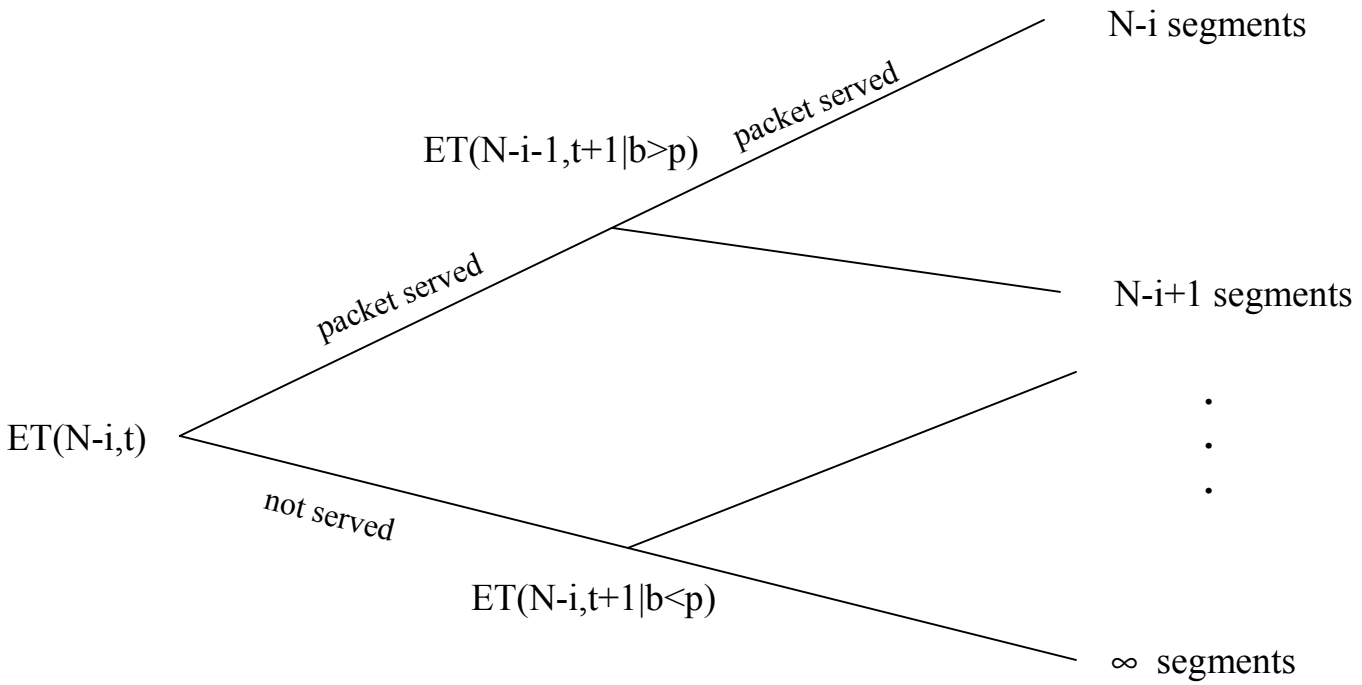
Example: Bidding for priority

- Example: "Smart market" for gateway priority (MacKie-Mason and Varian 1994):
  - when congested, service the packet with the highest value for immediacy
  - agents bid their value for immediate service
  - M+1st price auction

Mapping service value to priority

- How do **agents figure out** their valuations?
  - Suppose intrinsic value is for **completion time** (*T*) for N-packet flow
  - Need map from E(T(N,t)) to reservation value for immediate service on first packet

Optimal bidding



Unbounded # of paths, with some branches of unbounded length. Each depends on updated beliefs about congestion, and on evaluation of sequence of continuation bids.

Problem solving as search

- Uninformed search algorithms are $O(b^d)$ in time (b=branching factor, d=depth)

Depth	Nodes	Time
0	1	1 millisecond
2	111	0.1 seconds
4	11111	11 seconds
6	1.0E+07	18 minutes
8	1.0E+09	31 hours
10	1.0E+11	128 days
12	1.0E+13	35 years
14	1.0E+15	3500 years

Assumes 1000 nodes/second, b=10

Does "informed" economic structure reduce complexity?

- Many problems are **intractable**: $comp(\varepsilon,d) = O((1/\varepsilon)^d)$
  - utility maximization
  - expected utility maximization
  - finding competitive equilibria
  - finding rational expectations equilibria
  - option valuation
  - finding Nash equilibria
  - see Rust 1997, and Traub et al. 1988

Obtaining approximate optimality

- We focus on mechanisms that implement social choice functions in more **restrictive** environments:
  - agent **preferences are restricted** to a subset (e.g., quasilinear)
  - agent **strategies are restricted** more narrowly (e.g., Bayesian-Nash, myopic, &c.)
- Or we **relax the goal** or its constraints (e.g., design mechanisms that are not fully efficient)

Mechanisms are Temporary

- **Deviation** between outcomes of a feasible mechanism and the desired system goal depend on, e.g.:
  - number of agents
  - distribution of beliefs about other agents' private information
  - technological possibilities
  - match between definitions of concerns and actual agent concerns
- All of these can and do **change over time**

Mechanism Design in Context

- **Approximate and temporary** satisfaction of design goals is **intrinsic** to MAS system design

- Therefore, a **learning** and **re-configuration feedback loop** might improve performance

Context: Institutions

- Real economies can manage complexity…
  - What are they doing, and how?
- **Institutions** are essential
  - not all problems can be simultaneously endogenized
  - a dynamic form of **decomposition**
  - institutions are resource allocation mechanisms that **adapt at lower frequency**

Example of overlapping adaptive mechanisms

- English auction for Tori Amos bootlegs
  - Specific instance of family of auctions supported by auction server (good definition, reserve price, bid increment, closing date…)
    - governed by trust conventions, reputation, language
      - governed by UCC and case law
        - governed by Constitution
        - &c.

Can't build an MAS without institutions

- Institutions **often implicit**
- To make progress on dealing with complexity (better problem solutions, higher profits, &c.) **need to make institutions explicit**
  - understand their role
  - design and plan for institutional adaptation

When is slow adaptation appropriate?

- some adaptations have **higher direct costs** (siting a nuclear plant)
  - transaction costs
  - coordination costs
- some environments are **more stable**:
  - transactions are routinized (cash withdrawals)
- some designs are more **tightly coupled**
  - other processes depend on them (e.g., law of contracts)
  - change takes time to propagate, and adjustment is costly: limit rate of change
- some design **errors** are more consequential

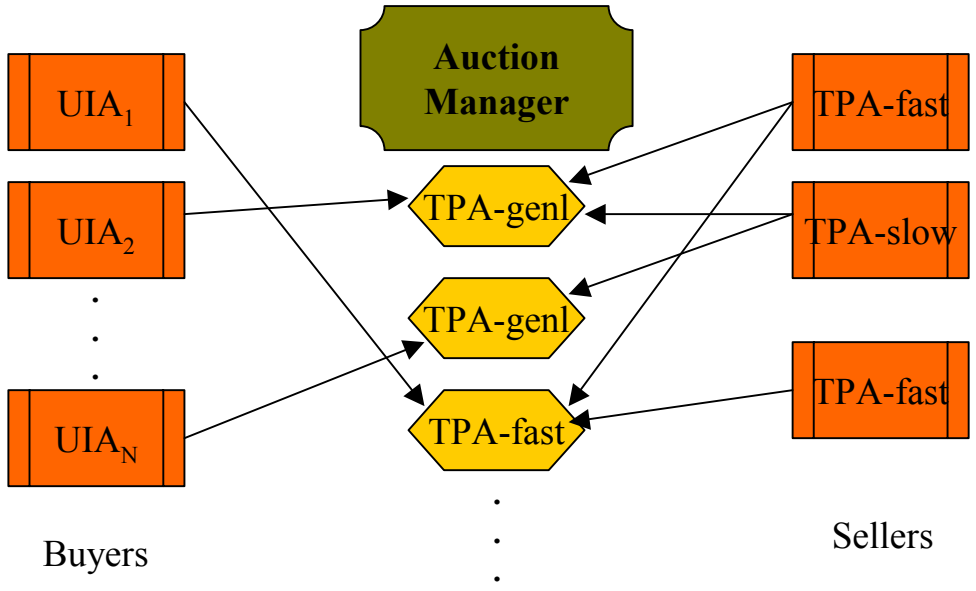
Roles for institutions in electronic commerce

- **bound the concerns space**
  - technical standards
  - conventional product configurations
  - default rules (UCC)
  - enforcement mechanisms
- facilitate computational **feasibility**
  - reduce the complexity of the problem space
  - create block structure, shared memory

Examples

- Why are most prices in US non-negotiable?
- Why do consumers trust credit cards over the phone (or at all!)?
- So what if it's a Federal Reserve Note?

Example: UMDL society of agents



- **Who makes decisions about:**
  - starting/ending auctions?
  - participating in which auctions?
  - defining scope of auctions

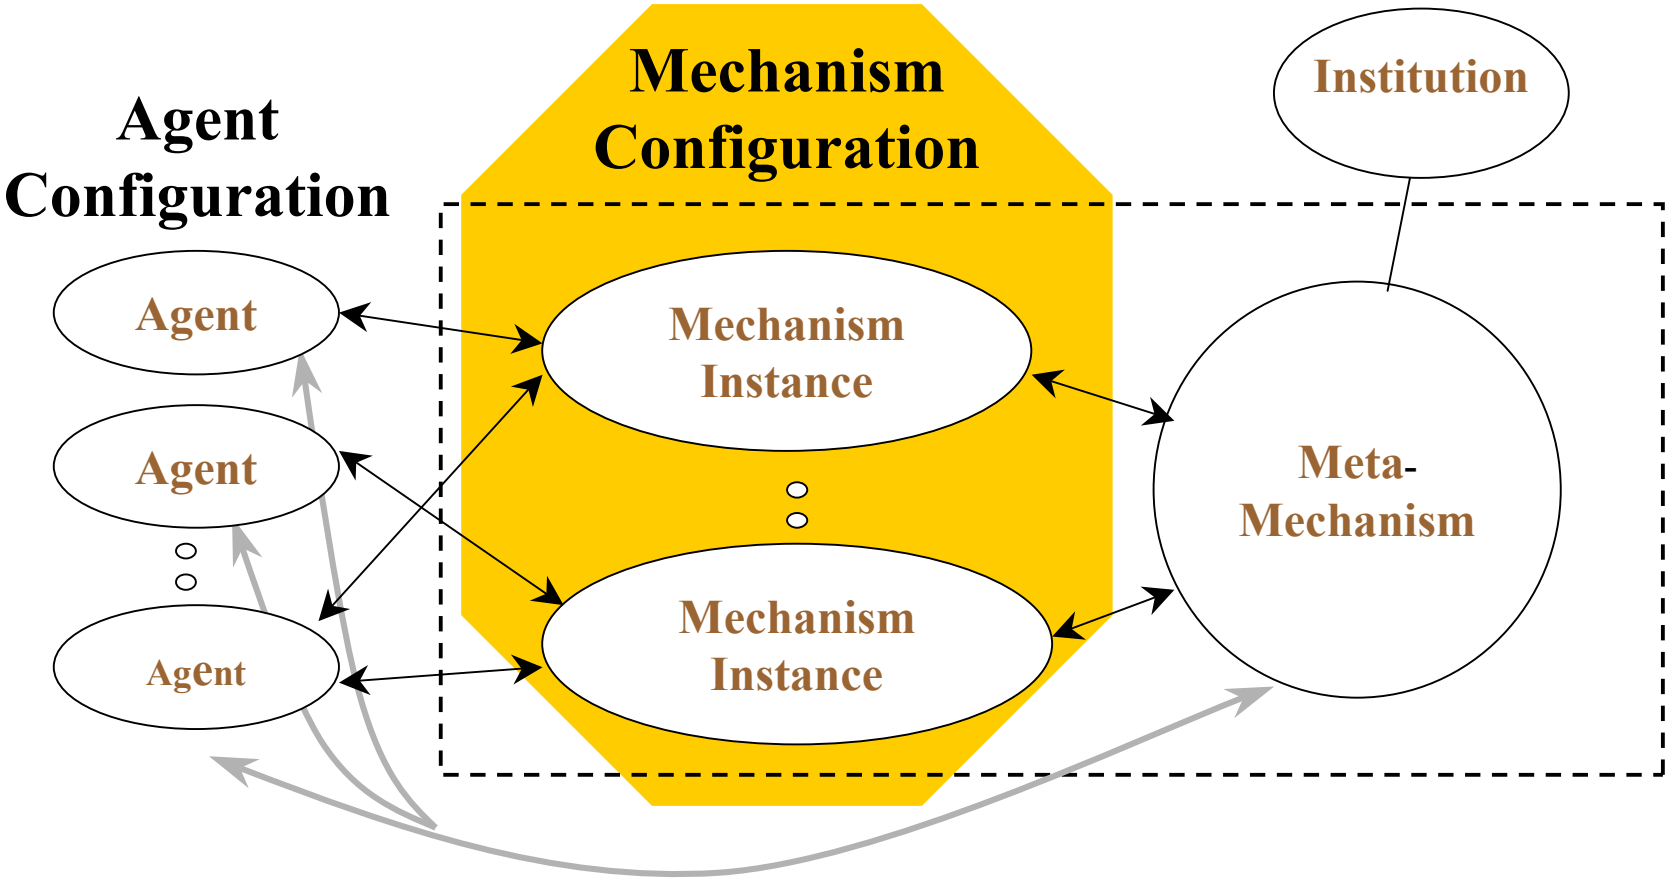
Which decisions are endogenous, which are not?

Many TPAs, compete to buy auction resources from AMA and sell to UIAs?
AMA buys disk, bandwidth from...? (Competing AMAs or central planner?)
Disk and bandwidth brokers buy resources *they* need from...?
Does AMA enforce allocations?
Does AMA purchase enforcement inputs from police and court agents?

Endogenizing the institutions

- Good problem solving requires adaptive institutions
- How do we implement the mechanism design process **internally and dynamically** within an MAS?
- want to design **metamechanisms**: mechanisms that reconfigure mechanisms

MAS structure



Metamechanisms

- Agents **interact** through **instances** of mechanisms
- A metamechanism can **create and modify** mechanism instances
- These relationships can be defined **recursively** (metamechanisms take place within governance structures, which can reconfigure metamechanisms, &c.)

A grand adventure...

- With cheap silicon and sand we can solve previously intractable allocation problems
- Vast opportunities for **low-friction** commerce, social problem solving

...With grand challenges

- Existing institutions evolved for expensive computation and communications
- How to design mechanisms **and their evolution** with everything in flux:
  - **concerns**
  - **rules**
  - **institutions?**

Resources

- Program for Research on the Information Economy (PRIE): **www.si.umich.edu/~prie/**
- MARX project: **ai.eecs.umich.edu/MARX/**
- My papers: **www-personal.umich.edu/~jmm/**