Market-Biased Replacement Policies for QoS Web Caching



Jeffrey K. MacKie-Mason

School of Information
& Dept of Economics

University of Michigan

What is the problem?

- Rapid growth of convergent **multi-service** networks
- Highly variable, **unreliable**, stochastic service quality
- Increasing demand for multiple qualities-of-service (**QoS**)

How might economics help?

- Fundamentally a **resource allocation problem**. Who gets to use which network resources, with what priority or guarantees?
- Fundamentally a **distributed problem**. User values for different volumes and QoS are private information.

What is scarce?

- Low-delay, low-jitter network transport (among other things!)
- cf. Breese and Mullen (1998) on distributed file systems for mobile users
- Replicated "local" storage can improve mean, max and variance of delay

Basic idea: QoS in distributed file services

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

What have we accomplished?

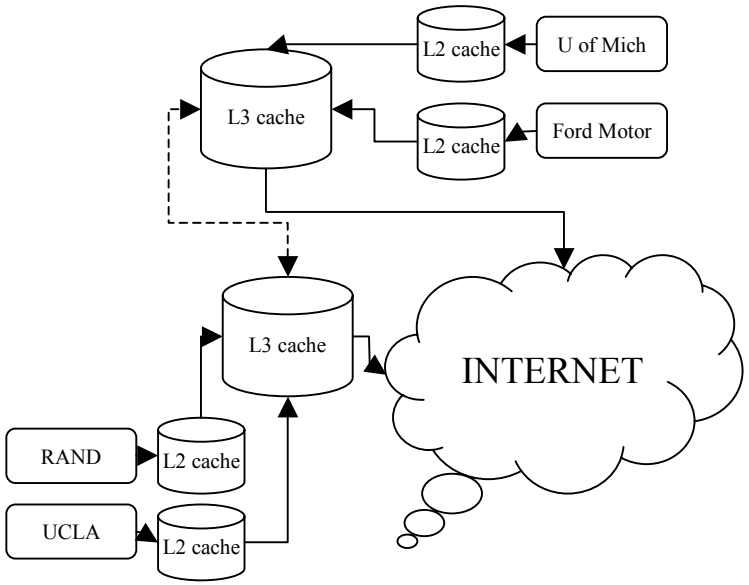
- Simple, **easy-to-implement** extension to standard LFU caching
- Provides **variable QoS**: higher valued file access obtains higher hit rate
- Provides **higher aggregate user value**
- Now implementing a deterministic QoS alternative

Resources

- This paper:
  - "Biased Replacement Policies for Web Caches: Differential Quality-of-Service and Aggregate User Value"
  - T. Kelly, Y. M. Chan, S. Jamin and J. K. MacKie-Mason
- My papers: www-personal.umich.edu/~jmm/
- Program for Research on the Information Economy (PRIE): www.si.umich.edu/~prie/
- MARX project: http://ai.eecs.umich.edu/MARX/

Caching and QoS

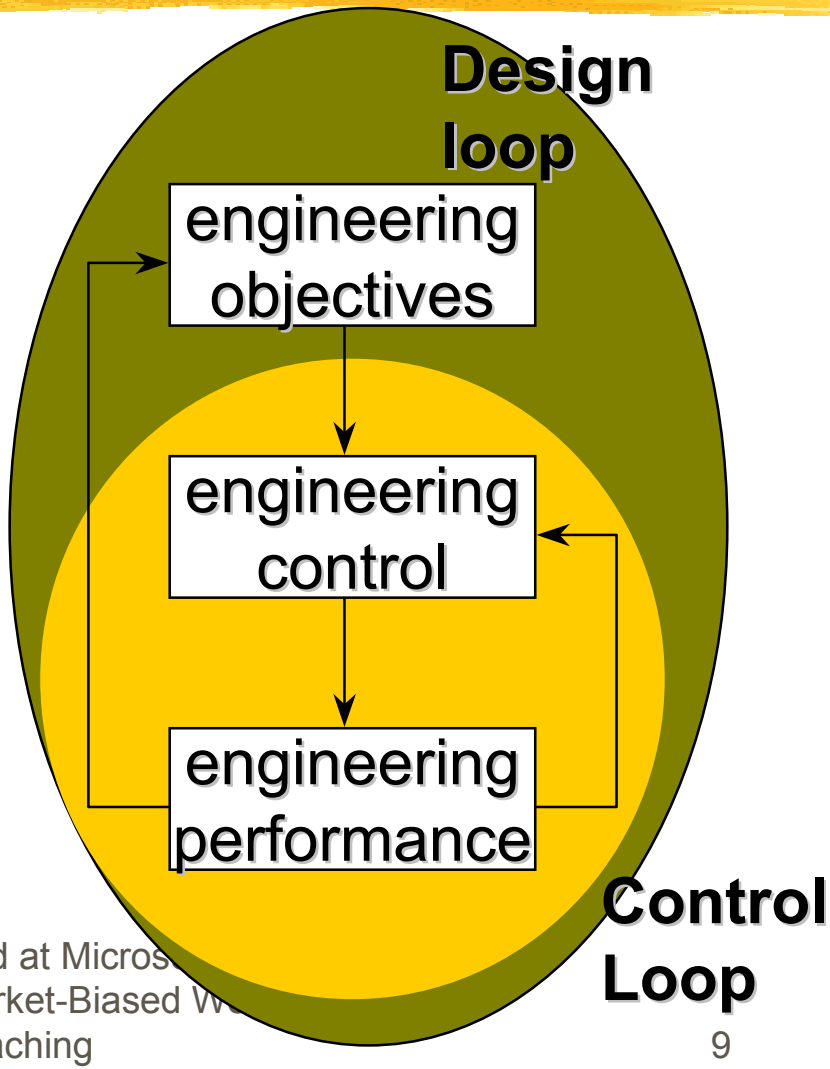
- View of a caching hierarchy:



- Common service measure:
  - hit rate or **byte hit rate**
  - measurable and "objective"
  - but, **doesn't necessarily measure** value of service to users
- Users place **heterogeneous** value on network services
  - private nets vs. xDSL vs. ISDN vs. coax vs. modem...

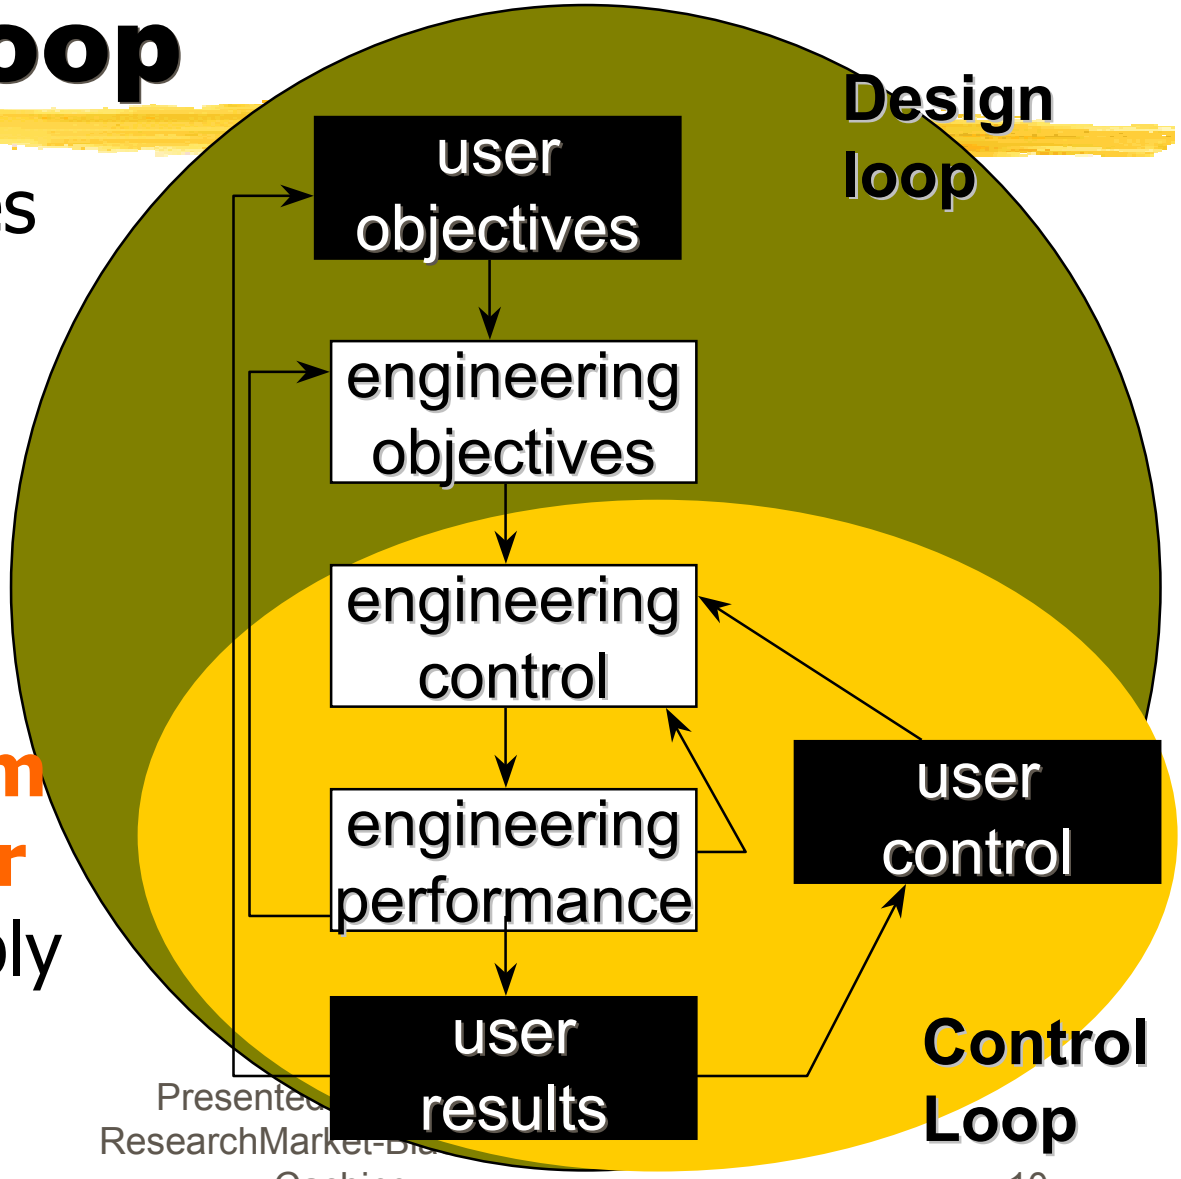
User centric view

- A networked resource is **only as good as its users think it is**
- Users not only the driver, but they are ***smart devices***, not ***dumb terminals***.



Putting users into the feedback loop

- Let user preferences determine:
  - which objects
  - located where
  - for whom
- Design goal: **maximize the sum of aggregate user valuations** (possibly subject to some constraints)



Cache Management Policies

- Users send messages to cache "manager"
  - object request
  - usage and forecast reports
  - bids and payments
- Cache implements several decisions as function of message:
  - **Replacement**
  - **Service**

Replacement Policy

- Which objects are retained, which replaced? Fundamental scarcity.
- Implicitly a ***sorting algorithm***: "remove URLs from tail of sort list"
  - LRU: sort on time of last access
  - LFU: sort on number of references
- Thus for every URL *i* there is p=Pr[*i* will be present during (t, t+Δt)]
- Users have ***heterogeneous*** values over p

Service policy

- **What services are provided** to users?
  How do messages *m* affect services?
  - Effect on replacement probability, *p*?
    (Equivalently, the cache loading profile)
  - Effect on object in cache being delivered?
    What quality of delivery?
- Example: If server *s* buys right to load object *u* into cache, can it also restrict the list of authenticated users of *u*?

Service policy problem

- Let quality of service increase in *m*
  - Absent disincentive all users report highest *m*
  - If quality is scarce, need to induce differentiated m_i to provide allocation with maximal aggregate user value
- ***Preference revelation*** problem. Prices, e.g., reveal willingness-to-pay values

User valuations: servers

- Servers, s = 1, 2, …, S
  - W_u = value to s per byte when cache requests for u are fulfilled
  - N_u = number of references since it entered cache

$$\frac{\text{contribution of } u \text{ to aggregate value}}{\text{unit size}} = \frac{W_u \times \text{size}(u) \times N_u}{\text{size}(u)} = W_u \times N_u$$

- Suggests weighted-LFU (swLFU):
  - primary sort key is W_u x N_u

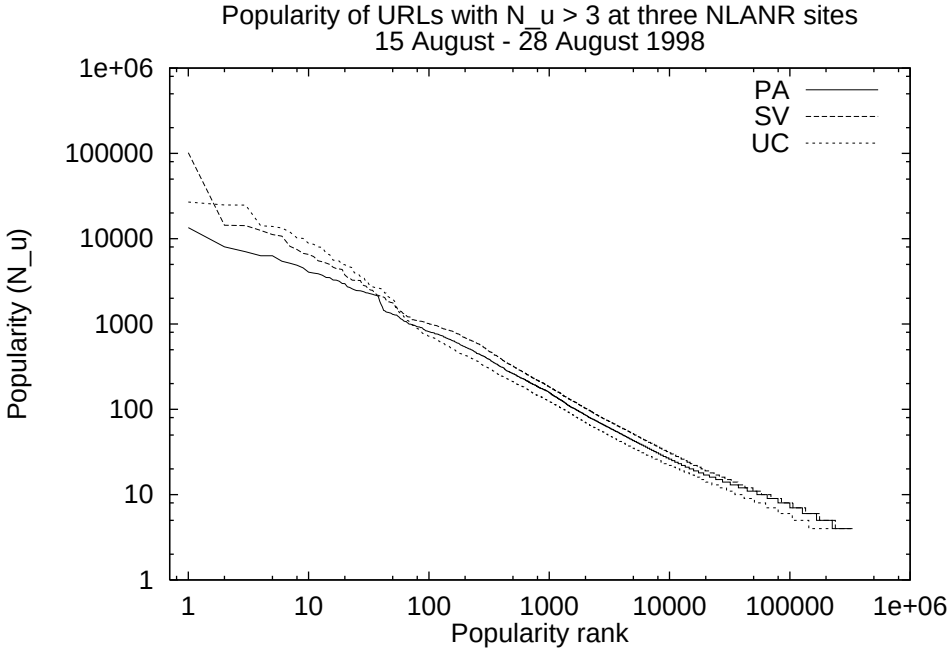
Experimental data for swLFU

- 3 actual L3 Web cache request streams from NLANR project
- Palo Alto, Silicon Valley and University of Illinois (UIUC), 15-28 Aug 98

	Palo Alto	Silicon Valley	Univ Ill. Urbana-Champ
# servers	114,381	124,698	105,710
# URLs	3,412,105	3,744,274	2,884,598
# requests	7,011,622	7,897,659	5,568,112
Bytes requested	131,665,275,664	161,620,444,331	127,346,723,989
Value requested	264,025,996,908,451	319,466,493,484,667	264,106,487,028,562
Infinite Cache			
size (bytes)	60,037,623,775	66,976,225,688	51,825,514,504
hit rate	51.3364	52.5901	48.1943
byte HR	54.4013	58.5596	59.3036
value HR	48.5422	57.467	56.5745

Assigning W_u

- Don't have direct information on hit value distribution
  - Simulate a distribution that could result in performance effect, measure potential effect on real data
  - Zipf-like request distribution requires logarithmic value distribution



- **Assign**: $W_s \in \{1,10,100,1000,10000\}$

Simulations

- We simulate caches of 1, 4, 16, 64, 256, 1024 MB
  - LRU, LFU, swLFU
  - Metrics: hit rate (HR), byte hit rate (BHR), value hit rate (VHR)
  - VHR = percent of all possible value that is hit

Results

- swLFU:
  - **higher value** than LRU, LFU
    - VHR 1/3 to 1/2 higher than LFU
    - due to serial correlation in data, performance improves with recency weighting (LRU)
  - **better QoS** to servers that value it more
    - UIUC anomaly: value-biased LFU not *guaranteed* to perform better
  - worse on HR, BHR: not all hits created equal, so swLFU **trades off** BHR for higher aggregate user value

Implementation 2: Market for write access

- swLFU provides ***stochastic*** QoS
- "Write access" caching provides ***deterministic*** QoS
  - users bid for write access over *y* MB of cache space
  - periodic auction allocates write access to winning servers, clients
  - read access can be free or allocated
- In general, offer ***multiple QoS*** model: some swLFU, some WA, some best effort

Write Access vs. swLFU

- Pros of Write Access
  - **deterministic** (?)
  - **flexible**: could implement custom replacement policies in each sub-cache
  - users might be able to **forecast** better than naïve policies
- Cons
  - more **complex** to explain
  - **computational** (forecasting) **burden** on users

User valuations: Clients

- Clients might want to influence replacement policy, too
  - heterogeneous values for particular URLs (swLFU)
  - want own usage to influence replacement more than other users (WA)
- But:
  - high burden unless aggregators (corps, ISPs)
  - law of large numbers undoes effect unless correlated values across consumers

Economics and Computer Science



Why an integrated research program?

Computing as problem solving and decision support

- Different modalities:
  - centralized vs. decentralized
  - rule/heuristic vs. negotiation/contract
- Economics:
  - rigorous and well-adapted approaches to **decentralized, negotiation-based** decision making
  - changing boundaries between modalities require system design research that is informed by decision theory (esp. economics)

Why are boundaries changing?

- Silicon and sand revolution changes the underlying cost structure for communications and computation
- **Undercuts existing boundaries**, institutions
  - more organizations managing with internal market-like systems (profit centers)
  - previously internalized functions are externalized: e.g., auto design and parts manufacturing

Good systems should be smart

- Respond to and represent **users**
- Economics provides well-understood, powerful approaches:
  - utility maximization paradigm for preference representation and choice-making
  - competitive market allocation for distributed decision making (1st and 2nd Welfare theorems)

Areas for research

- Models and methods for **computational resource allocation**
- Computational infrastructure for **digital goods** and services
- Changing **social rules** and policies

Program for Research on the Information Economy



Jeffrey K. MacKie-Mason, Director

What is PRIE? http://www.si.umich.edu/~prie/

- About 20 research faculty and grad students with overlapping interests in the information economy, computer science, management, policy
  - economics, computer science, AI, networking, management, organization theory, law
- Unusual array of interdisciplinary, funded projects and graduate courses

Research Projects

- Content economics
  - PEAK and Bundling
  - Digital Libraries (UMDL)
  - Media Financing
  - Evaluation Brokers
  - Copyflex
- Network economics
  - Smart Reservations
  - Internet Congestion
  - Service Architecture
  - Web Caching
- New media and society
  - Competition, antitrust and network externalities
  - Anonymity, Trust, and Collective Action in Public Spaces
  - Cyberbalkanization
  - Do the Info Rich Get Richer?
  - Community Support Systems
- Computational markets for distributed problem solving
  - AuctionBot
  - MARX