Design and Analysis of a Content-Oriented Internet

Jim Kurose
School of Computer Science
University of Massachusetts
Amherst MA USA

UNIVERSITY OF Massachusetts Amherst

Distinguished Faculty Lecture, April 2014

Overview

- introduction: "getting connected"
- Internet 101
- content-oriented networks:
  - "over the top" approaches
  - network-based approaches
- wrap-up

Goals:

- learn something about networking
- appreciate research challenges in content networking
  - learn about research here at UMass

Getting connected: a history

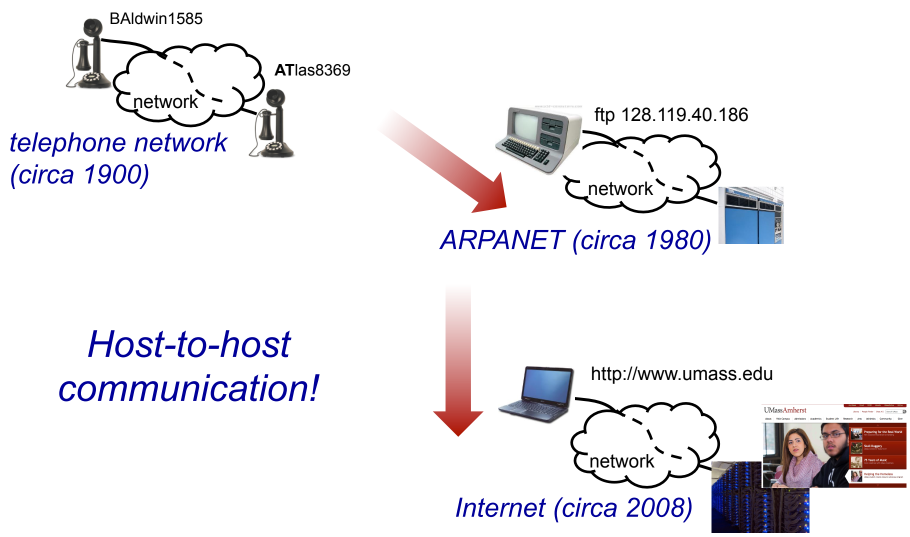

The Internet in 2014

- content is king!
  - video: 64% of all global consumer Internet traffic in 2012
  - TV, VOD, P2P: 80-90 % global consumer traffic by 2017
- *what* not *where*

> *Challenge:* can a network architecture designed for wired, host-host communication meet demands of today's (and tomorrow's) *content-oriented* mobile Internet?

Overview

- introduction: "getting connected"
- Internet 101
- content access:
  - "over the top" approaches
  - network-based approaches
- wrap-up

Internet 101

- hosts: clients, servers
- packets, routers
- Internet: a network of networks
- protocols

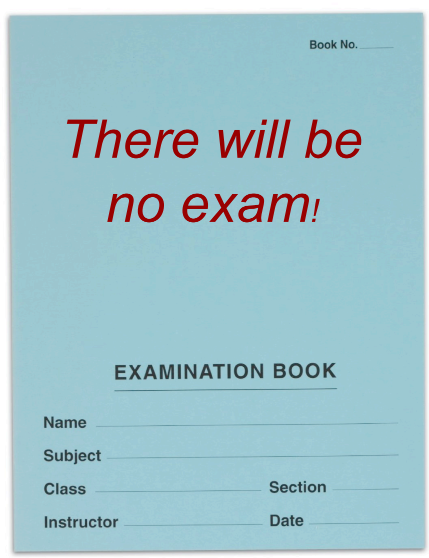

Internet 101

billions of connected computing devices: *hosts*

- running *network apps*
- clients, servers

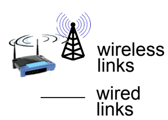

communication links

- wired, wireless
- transmission rate: bandwidth

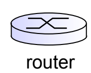

routers: switch/forward *packets* (chunks of data) from source to destination

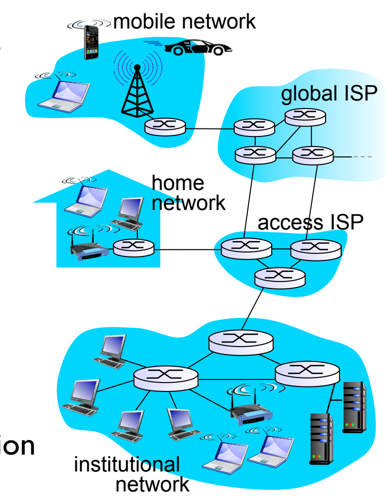

Internet 101: a network of networks

Question: given *millions* of access ISPs, how to connect them together?

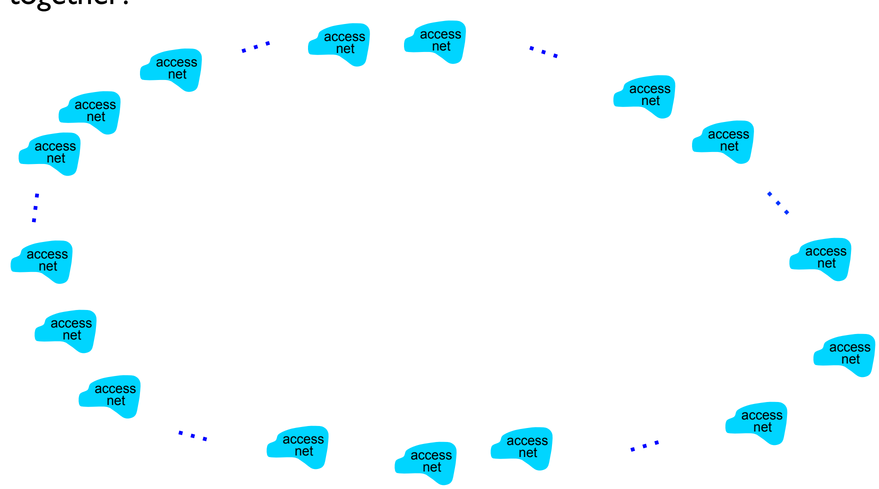

Internet 101: a network of networks

Option: *connect each access ISP to every other access ISP?*

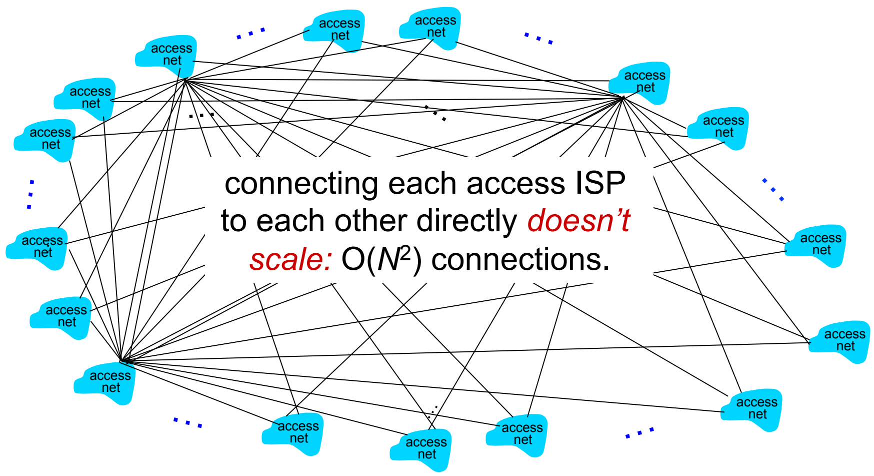

Internet 101: a network of networks

Option: *connect each access ISP to a global transit ISP?*

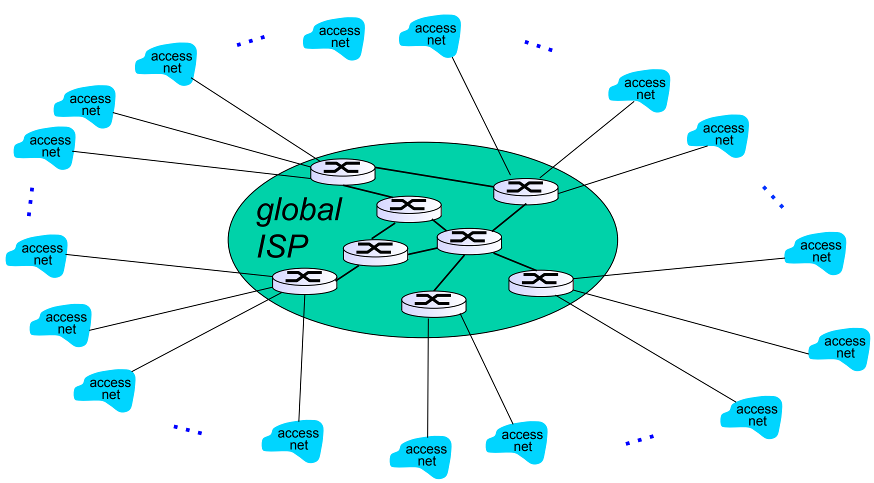

Internet 101: a network of networks

If one global ISP is viable business, there will be competitors ….

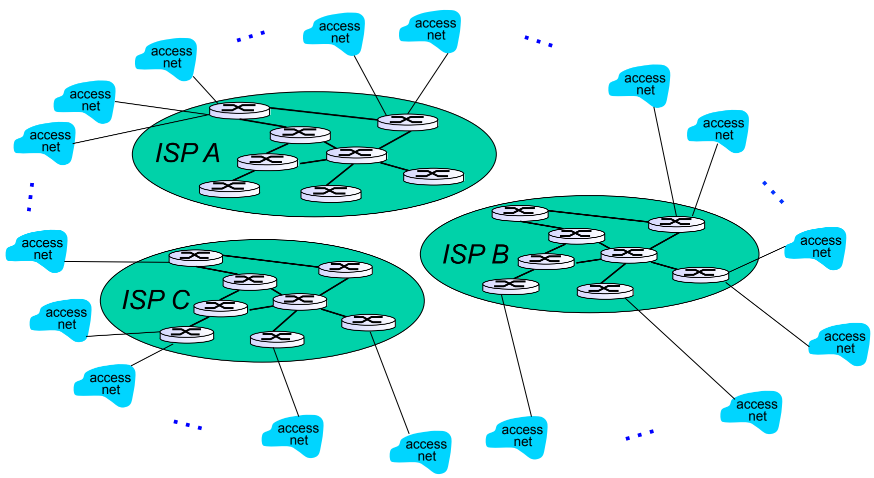

Internet 101: a network of networks

If one global ISP is viable business, there will be competitors ….

… which must be connected

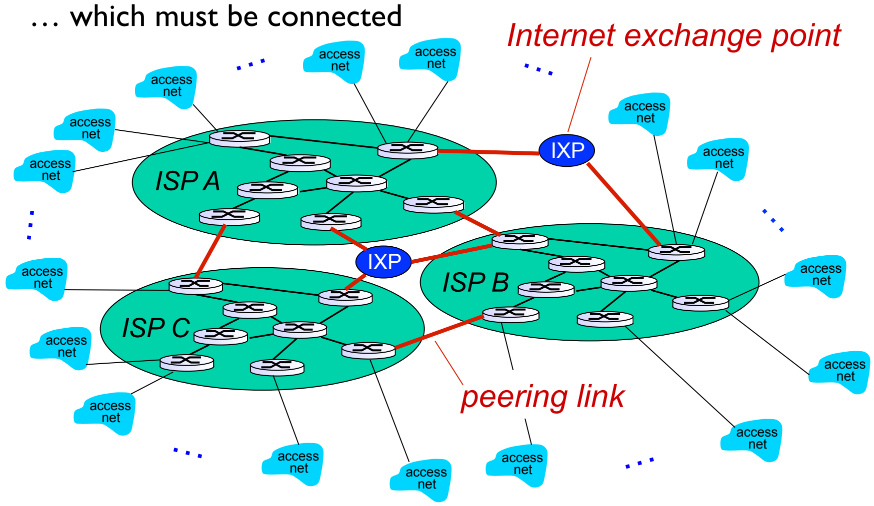

Internet 101: a network of networks

… and regional networks may arise to connect access nets to ISPs

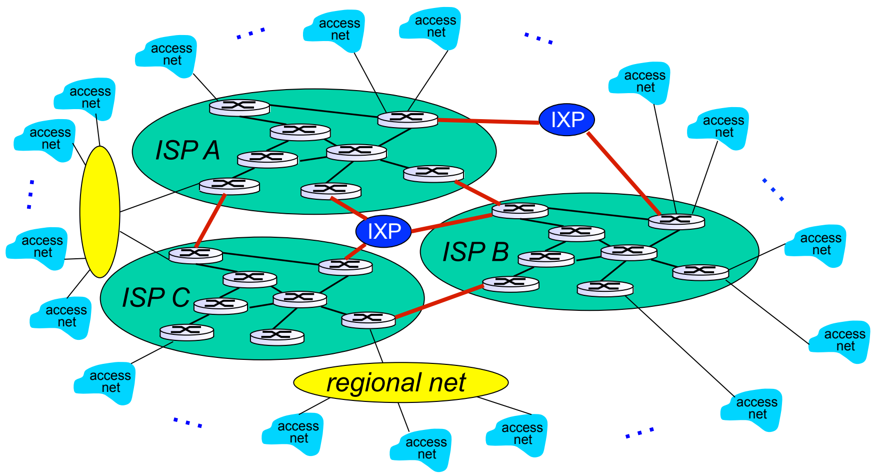

Internet 101: a network of networks

… and content providers may run their own network, to bring services, content close to end users

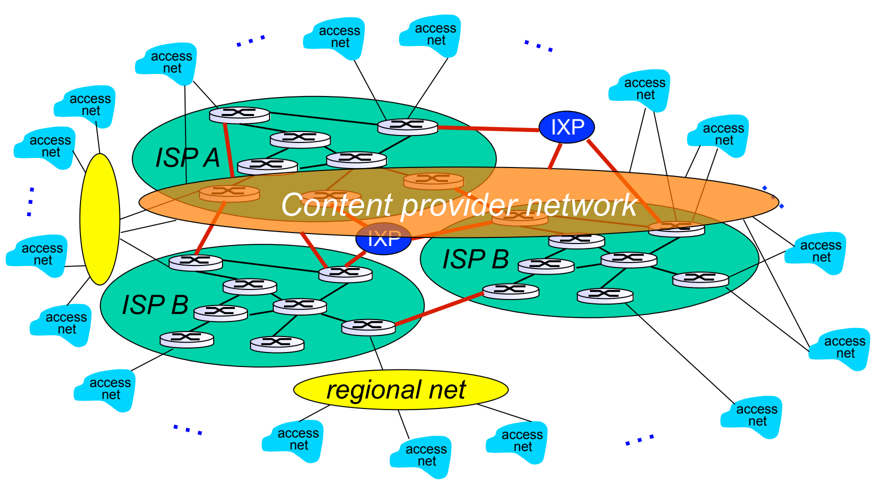

Internet 101: protocols

human protocols:

- "what's the time?"
- "I have a question"
- introductions

... specific messages sent
... specific actions taken when messages received

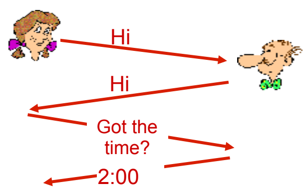

network protocols:

- hardware, software rather than humans
- define *format, meaning of messages* sent/received, *actions taken* on transmission/receipt
- all communication activity in Internet governed by protocols

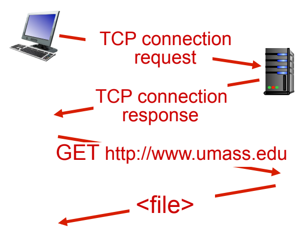

Overview

- introduction: "getting connected"
- Internet 101
- content-oriented networks:
  - "over the top" approaches
  - network-based approaches
- wrap-up

Challenge: can a network architecture designed for wired, host-host communication meet demands of today's (and tomorrow's) *content-oriented* mobile Internet?

Content Distribution Networks (CDNs)

- CDN: stores copies of content at CDN nodes
  - e.g. Netflix stores copies of MadMen
- subscriber requests content from CDN
  - directed to nearby copy, retrieves content
  - may choose different copy if network path congested

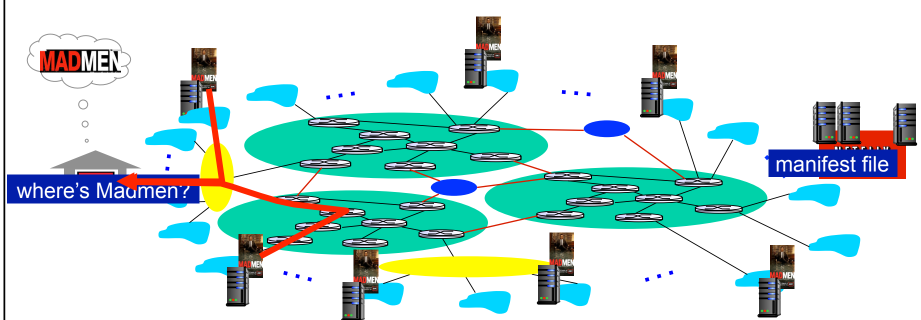

Content Distribution Networks (CDNs)

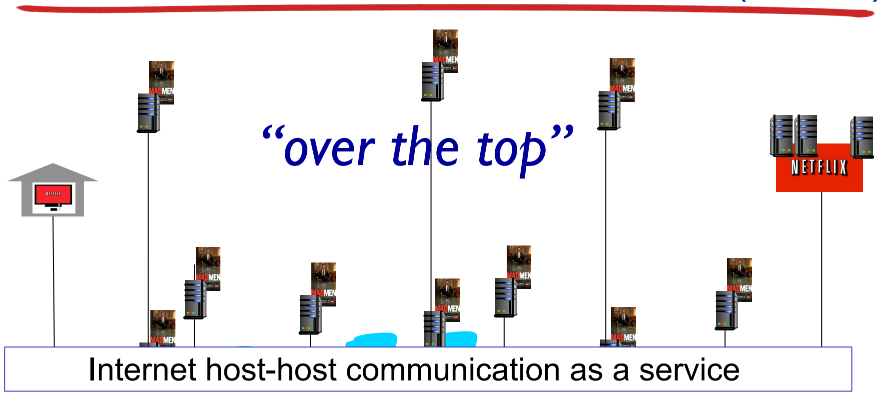

OTT challenges: coping with a congested Internet

- from which CDN node to retrieve content?
- viewer behavior in presence of congestion?
- what content to place in which CDN node?

CDNs: locating a "good" server

- CDNs (Akamai, Netflix): proprietary algorithms matching user with "nearby" server with high-bandwidth congestion-free path to client
  - "maps" of Internet topology
  - historical data (measurements) of Internet performance

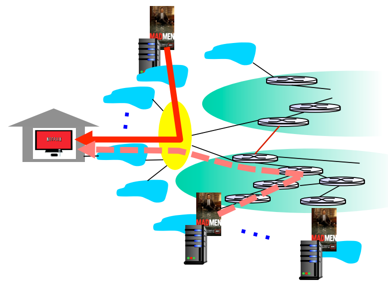

Nygren, Sitaraman,Sun, "The Akamai network: a platform for high-performance internet applications," SIGOPS Oper. Syst. Rev. 44, 3 (August 2010)

- *client control:* request lower/higher quality video encoding, using less/ more bandwidth
- *client control:* seek new server if performance drops below minimum

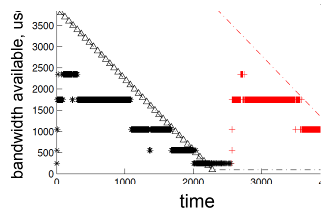

Adhikari, Guo, Hao, Varvello, Hilt, Steiner, Zhang, "Unreeling netflix: Understanding and improving multi-CDN movie delivery," IEEE INFOCOM, 2012.

CDN and user behavior

- *Question:* how does improved CDN performance affect user behavior?
  - abandonment, engagement, repeat viewership
  - quantitative, causal (avoid confounding factors)
- (big) data analysis: 6.7M users, 23M views, 100K unique videos

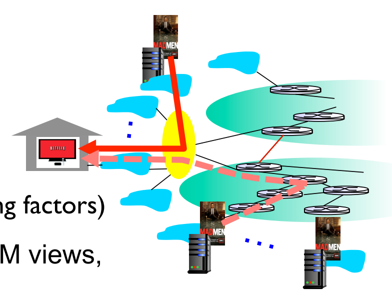

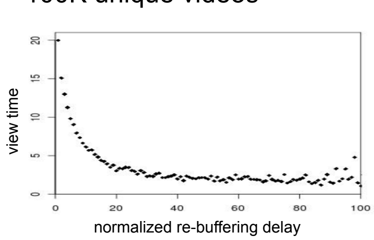

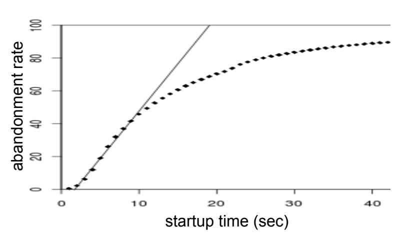

Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-Experimental Designs, Shunmuga, Sitaraman, ACM Internet Measurement Conf. (IMC), Boston, MA, Nov 2012

CDN: content placement in servers

- *Question:* which content to place in which servers?
  - Netflix: all servers identical
  - predict future content access using ML techniques: content viewing, popularity ranking by neighbors?
    - populate server, make recommendation accordingly?
    - degree of difference among regional servers?

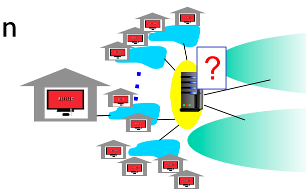

…. on going work (Dernbach, Kurose, Mahadevan) with Technicolor Research

DETOUR network neutrality

Question: to what extent can ISP regulate (or differentially charge) traffic flows within its network?

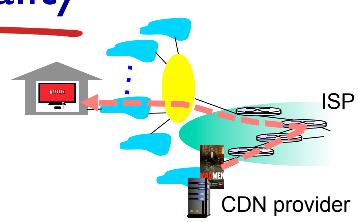

FCC "Open Internet Order" (2010):

- *[Transparency rule:]* … shall publicly disclose accurate information … sufficient for consumers to make informed choices … and for content, application, service, and device providers to develop, market, and maintain Internet offerings
- *[No blocking rule:]* … shall not block lawful content, applications, services, or non-harmful devices, subject to reasonable network management.
- *[Non discrimination rule:]* ... shall not unreasonably discriminate in transmitting lawful network traffic over a consumer's broadband Internet access service.

DETOUR network neutrality

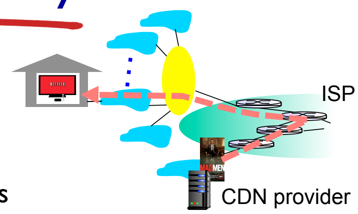

- *2000-2005: deregulation* – ISPs re-defined as "information service" providers rather than "telecommunication service" providers
- *2005:* ISP (Madison River) FCC consent degree: stop dropping VoIP packets
- *2008:* FCC order against Comcast for throttling BitTorrent file sharing. Overturned 2010: FCC lacks sufficient statutorily-mandated responsibility under Title I
- *Feb. 2014:* Verizon v FCC: vacates *Open Internet Order*
- *March 2014:* Netflix, noting degrading performance to Comcast subscribers, agrees to "pay-to-peer" with Comcast

The future?

Overview

- introduction: "getting connected"
- Internet 101
- content-oriented networks:
  - "over the top" approaches
  - network-based approaches
- wrap-up

Challenge: can a network architecture designed for wired, host-host communication meet demands of today's (and tomorrow's) *content-oriented* mobile Internet?

Architecture: form follows function

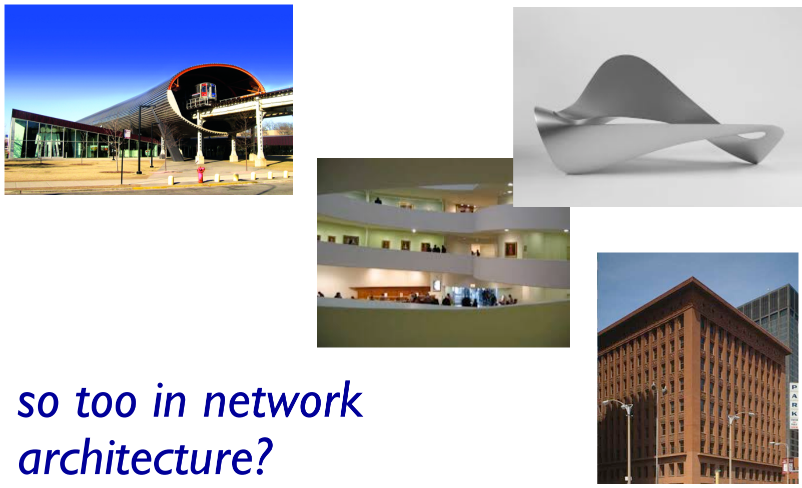

so too in network architecture?

In-network caching

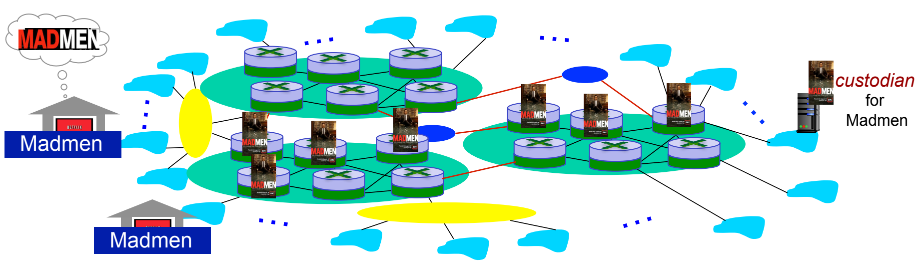

- routers equipped with massive storage
- client request routed towards custodian
- content returned, cached enroute
- client requests may be satisfied by in-network caches

In-network caching

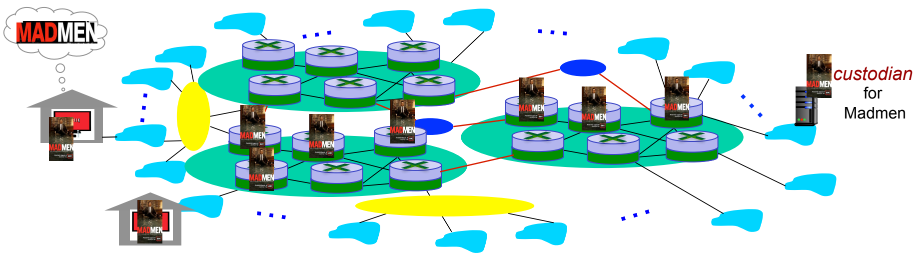

Many challenges!

- routing to custodian: based on name, or address?
  - how to search for content
- mobile users, content custodians
- cache management algorithms
- modeling, design, analysis of cache networks

Cache networks

- client *requests content*
- request *routed* (e.g., shortest path) to known *content custodian*
- en-route to custodian, *caches* inspect request
  - *hit:* return local copy
  - *miss*: forward request towards custodian
- during content download, *store* in caches along path
- content requests from different users interact: *cache replacement*

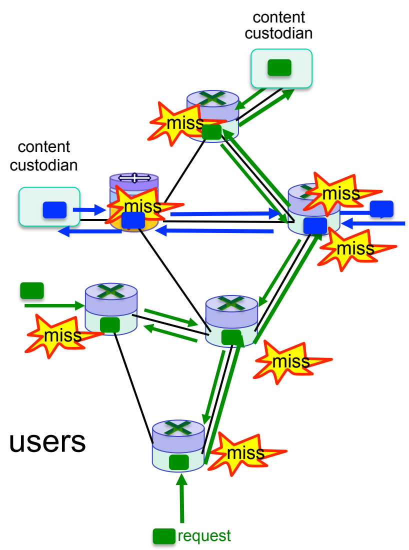

Challenge: networks of caches

- *network effect: interaction* among content request/ reply flows from different users:
  - content replacement: requested content by one user replaces content previously requested by others

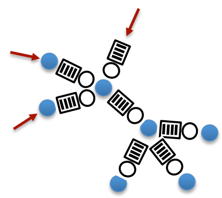

Packet-switching:
queueing networks
(Kleinrock, 1963)

Modeling a network of caches

- node i: *exogenous* (external) arrivals for content j: $\lambda_{i,j}$
- node i: *internal* arrivals *(miss stream)* for content j from downstream neighbors h: $m_{j,h}$
  - complex, correlated process
- $r_{i,j}$: aggregate rate of arrival requests at i for content j

$$r_{i,j} = \lambda_{i,j} + \sum_{\substack{\text{all downstream} \\ \text{neighbors, } h}} m_{j,h}$$

- ZDD: zero download delay assumption

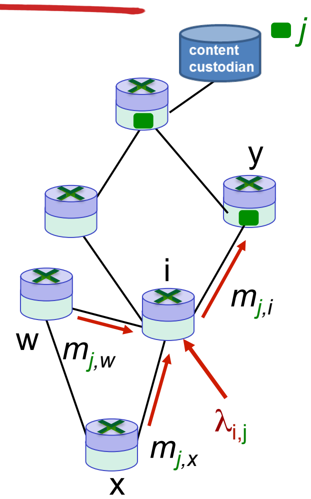

Cache networks: approximate analysis

- **SCA:** standalone cache *i* approximation algorithm: given *r(i,j),* compute miss rate for all content j
- Independence Reference Model (IRM) of incoming requests: $Pr(X_t = j \mid X_1,..,X_{t-1}) = Pr(X_t=j)$
- SCA approximation algorithm for LRU: [Dan 1985]

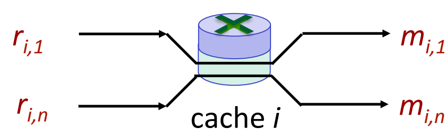

But we need $\{(r_{i,j}, m_{i,j})\}$ for a *network* of caches

Fixed-point iteration

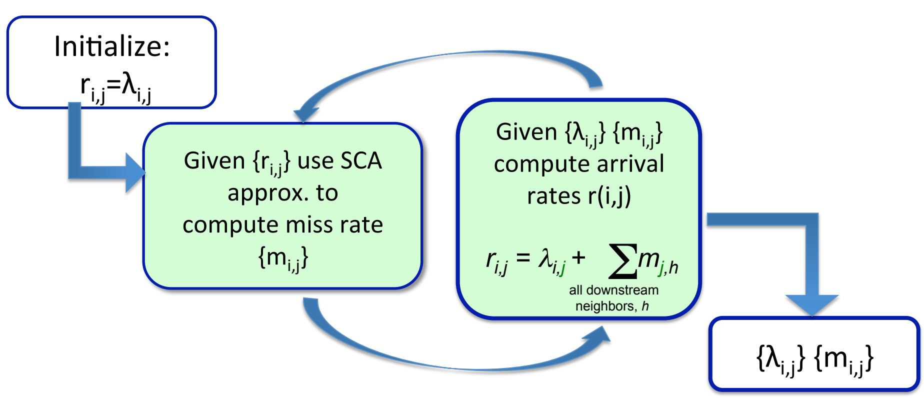

- fixed-point approximation relatively accurate: main source of error is independent reference model

"Approximate Models for General Cache Networks," Elisha J. Rosensweig, Jim Kurose and Don Towsley, IEEE NFOCOM 2010

32

Cache networks: bounding analysis

Question: can we provably *bound* characteristics of request flows between caches?

- $\rho_{v,y}^{j}$: *bound* on average arrival rate of requests for content item *j*, from cache *v* to cache *y*
- $\sigma_{v,y}^{j}$: *bound* on burstiness of requests for content item *j*, from cache *v* to cache *y*

$$\int_{t_1}^{t_2} r_{y,z}^{j}(t)dt < (t_2 - t_1)\rho_{v,y}^{j} + \sigma_{v,y}^{j}$$

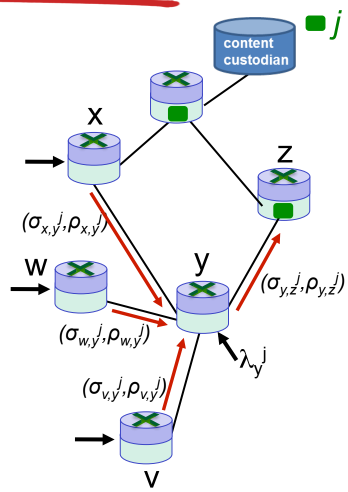

From $(\sigma_I^{in}, \rho_I^{in})$ to $(\sigma_I^{out}, \rho_I^{out})$

network calculus for cache flows:

- given input bounds, compute output bounds

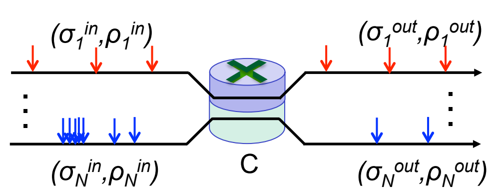

Theorem: For a cache of size *C*:

$$\rho_i^{out} = \min(\rho_i^{in}, M_i(\rho_1^{in}, \dots, \rho_n^{in}, C))$$

Can calculate ρ_i^{out} from $\{\rho_i^{in}\}$, σ_i^{out} also

E. Rosensweig, J. Kurose, **"A Network Calculus for Cache Networks,"** IEEE Infocom 2013.
E. Rosensweig, D. Menasche, J. Kurose, **"On the Steady-State of Cache Networks,"** IEEE INFOCOM 2013 (Best paper award)

In-network caching

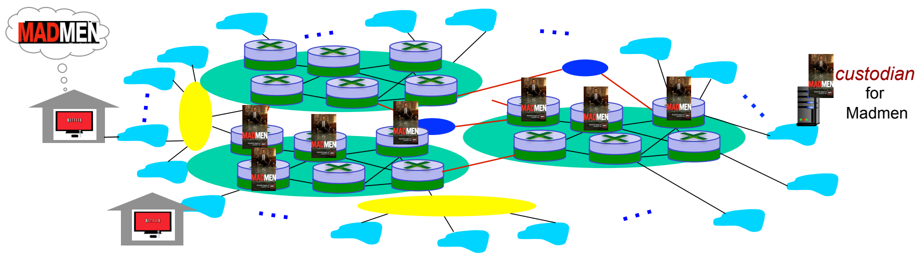

Many challenges!

- routing to custodian: based on name, or address?
  - searching for content
- mobile users, content custodians
- cache management algorithms
- modeling, design, analysis of cache networks

Overview

- introduction: "getting connected"
- Internet 101
- content access:
  - "over the top" approaches
  - network-based approaches
- wrap-up

Getting connected

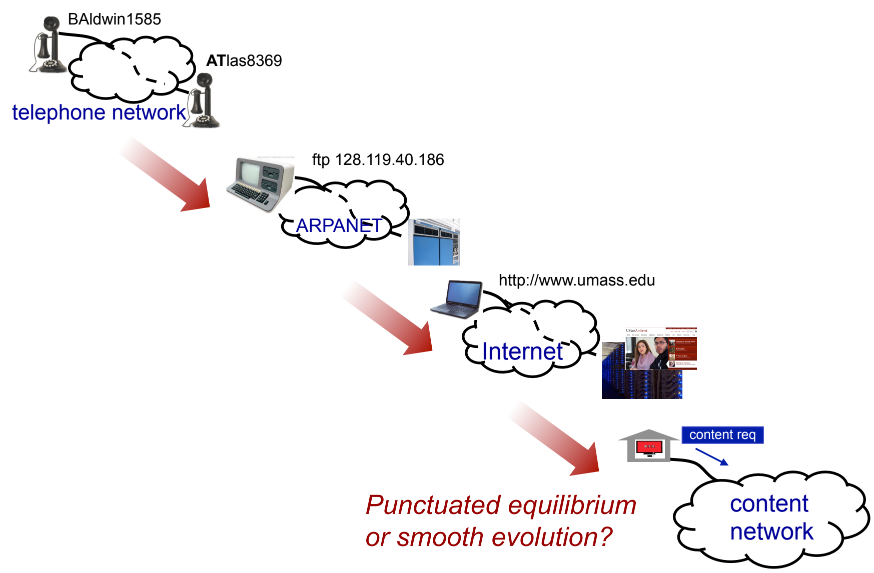

Punctuated equilibrium or smooth evolution?

Thanks !

UMass faculty collaborators

… and staff and friends

M. Adler
R. Adrion
W. Burleson
L. Clarke
L. Gao
D. Goeckel
W. Gong
N. Immerman
D. Irwin
W. Israel
C. Krishna
V. Lesser
B. Levine
S. Mahadevan
D. McLaughlin
G. Miklau
R. Moll
E. Riseman
P. Shenoy
J. Stankovic
D. Towsley
A. Venkataramani
B. Woolf
M. Zink

Thanks !

UMass students and postdocs

S. Abdallah	S. Heimlicher	A. Nyzio	M. Somasundaram
M. Badov	B. Horling	J. Padhye	J. Steinberg
S. Bhattacharyya	Y. Huang	J. Partan	J. Steshenko
M. Bradshaw	R. Hwang	D. Pepyne	M. Stern
C. Casetti	T. Ireland	S. Pingali	K. Suh
V. Chaganti	S. Jaiswal	S. Pingali	X. Tie
S. Chen	P. Ji	R. Ramjee	H. Uppal
Y. Chen	B. Jiang	B. Ribeiro	G. VanLeemput
W. Chen	S. Kasera	E. Rosensweig	S. Vasudevan
R. Chipalkatti	R. Khalili	D. Rubenstein	B. Wang
M. Dehghan	R. Koodli	S. Sahu	K. Watts
J. Dey	V. Lakamraju	J. Salehi	W. Wei
B. Donovan	H. Lee	H. Schulzrinne	D. Westbrook
D. Figueiredo	B. Liu	A. Seetharam	M. Yajnik
V. Firoiu	Y. Liu	S. Sen	S. Yang
T. Friedman	F. LoPresti	P. Serrano	D. Yates
S. Gangam	E. Lyons	J. Shapiro	X. Zhang
L. Gautier	V. Manfredi	A. Sharma	C. Zhang
Z. Ge	D. Menasche	C. Shen	Z. Zhang
M. Ghaderi	A. Misra	A. Shrivastava	C. Zhang
M. Girkar	S. Moon	R. Simha	H. Zhang
Y. Gu	R. Nagarajan	M. Sims	C. Zhang
Y. Guo	E. Nahum	S. Singh	T. Zhu
D. Gyllstrom	G. Neglia	S. Singh	Z. Zhu

To learn more:

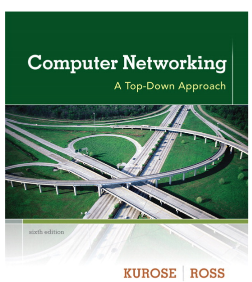

- CS 290nw: A networked world
- CS 453: Computer networks