

Network Architecture Research in SPYCE CIP/URI Project


Jonathan M. Smith
University of Pennsylvania
3/31/03

Highlights

- **Network support of diffuse computing**
 - High-performance active networking for monitoring and measurement
 - Ultra-large-scale peer-to-peer architecture for simulations and multiplayer games
- **Diffuse computing** network measurement infrastructure - "cing"
- *Transcoding* and *proxies* to support **diffuse computing** on media streams and images in large-scale heterogeneous nets

The 70s Internet is no more!

- New Roles for Hosts (Hosts, WWW, P2P)
- Routers *plus* RFC3234 "middleboxes"
- Other Diffuse Computing Elements



Network Challenges

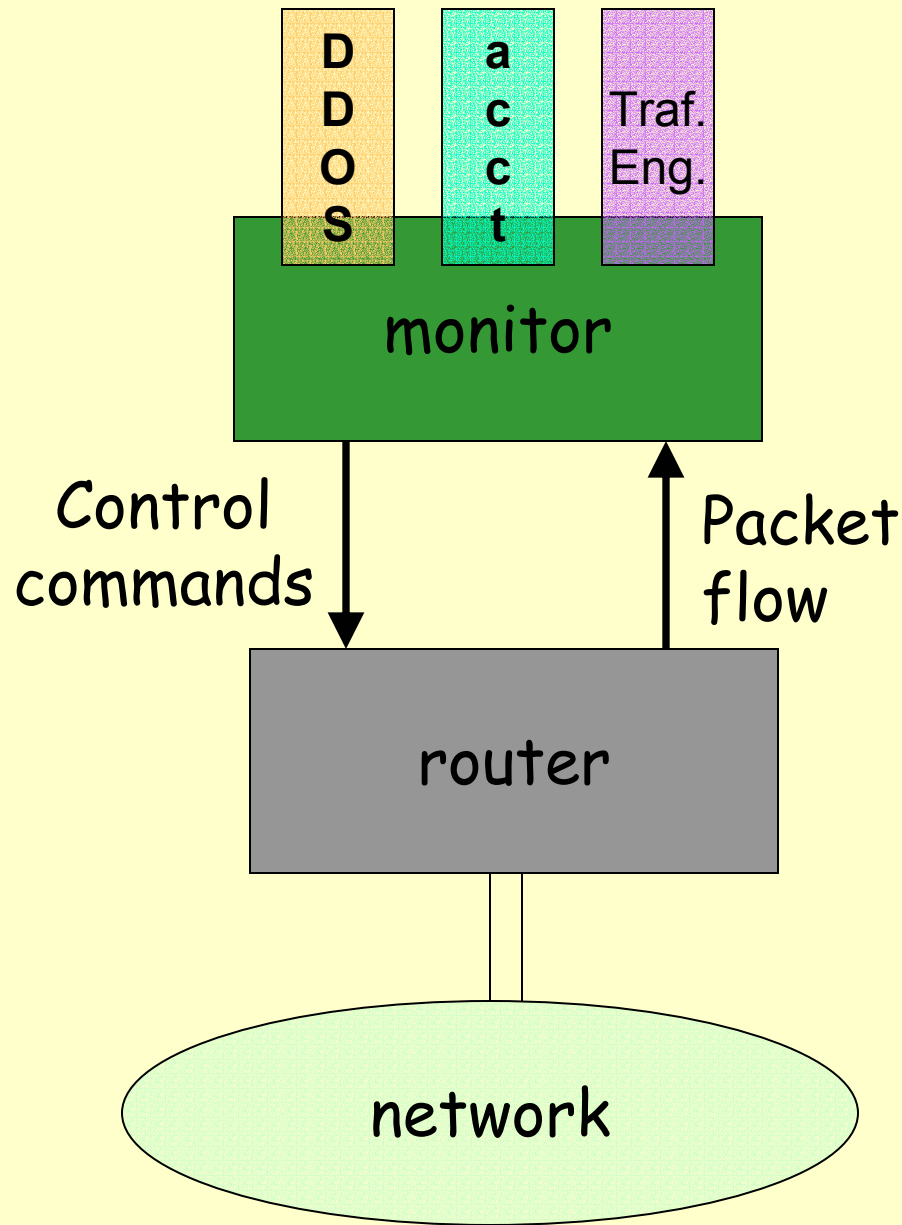
- Diagnosis in spite of complex network conditions, heterogeneity
- Reconfigurability for high availability
- Accommodating new applications (rapidly)
- Accommodating new tech rapidly (wireless)
- Scalability; Metcalfe's Law phenomenon
- Need for global security even *with* local failures and subversions

Achievements/Prototypes

- **LAME/FLAME*** active networking
 - fast/safe/extensible NW monitoring
 - NOMS 2002
- **CING*** network measurement tool
 - characterize Internet paths (Performance'02)
 - Indirect measurements (Infocom'03)
- **Transcoding proxies***
- **P2P massive multiplayer games***
 - Under submission
 - Applications to distributed simulation
- 3 slides on FLAME; rest on latter two systems

* Done with CIP/URI support

Diagnosis: AN-based monitors:



Uses:

IDS

DDoS/virus detection

performance debug

traffic engineering

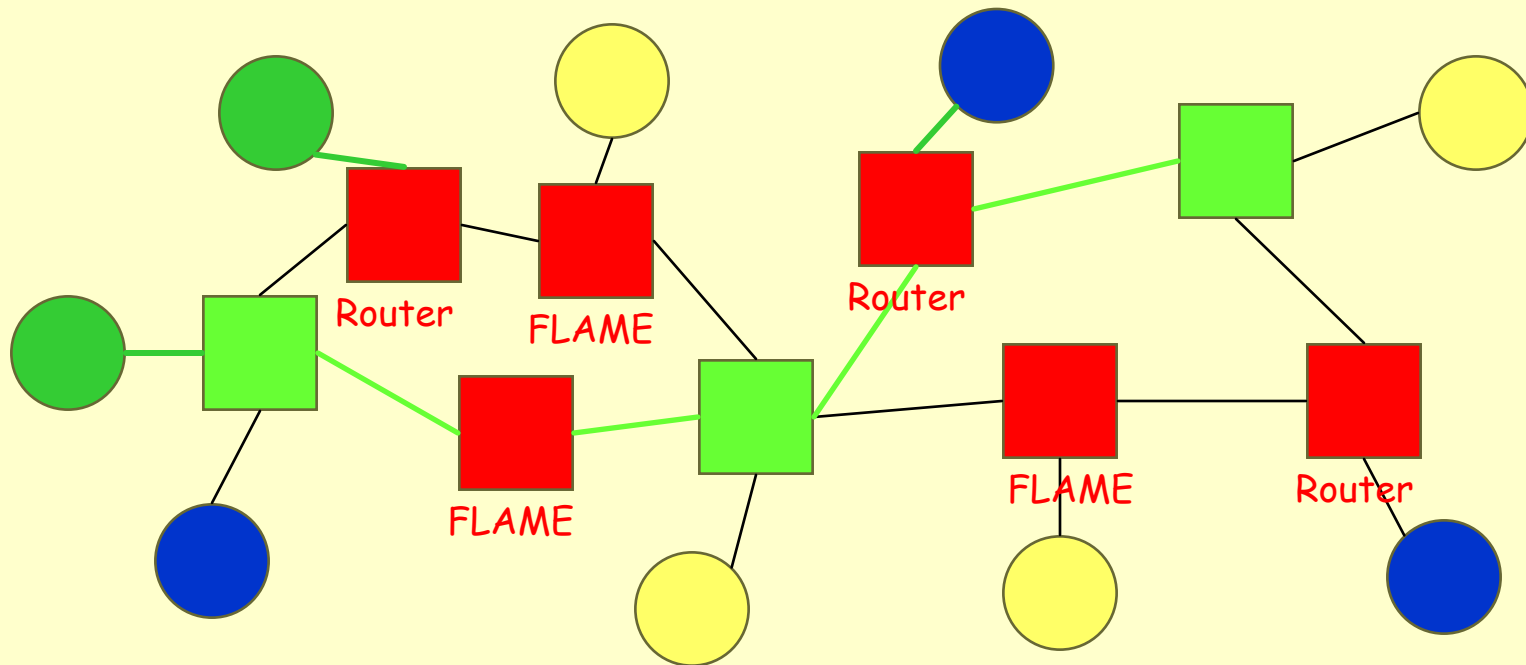
traffic measurement

Accounting

SPYCE platform

FLAME: Safe *in-network* monitoring

- Hosts of various types
- Routers / FLAME nodes
- Diffuse Computing Elements
- P2P participants



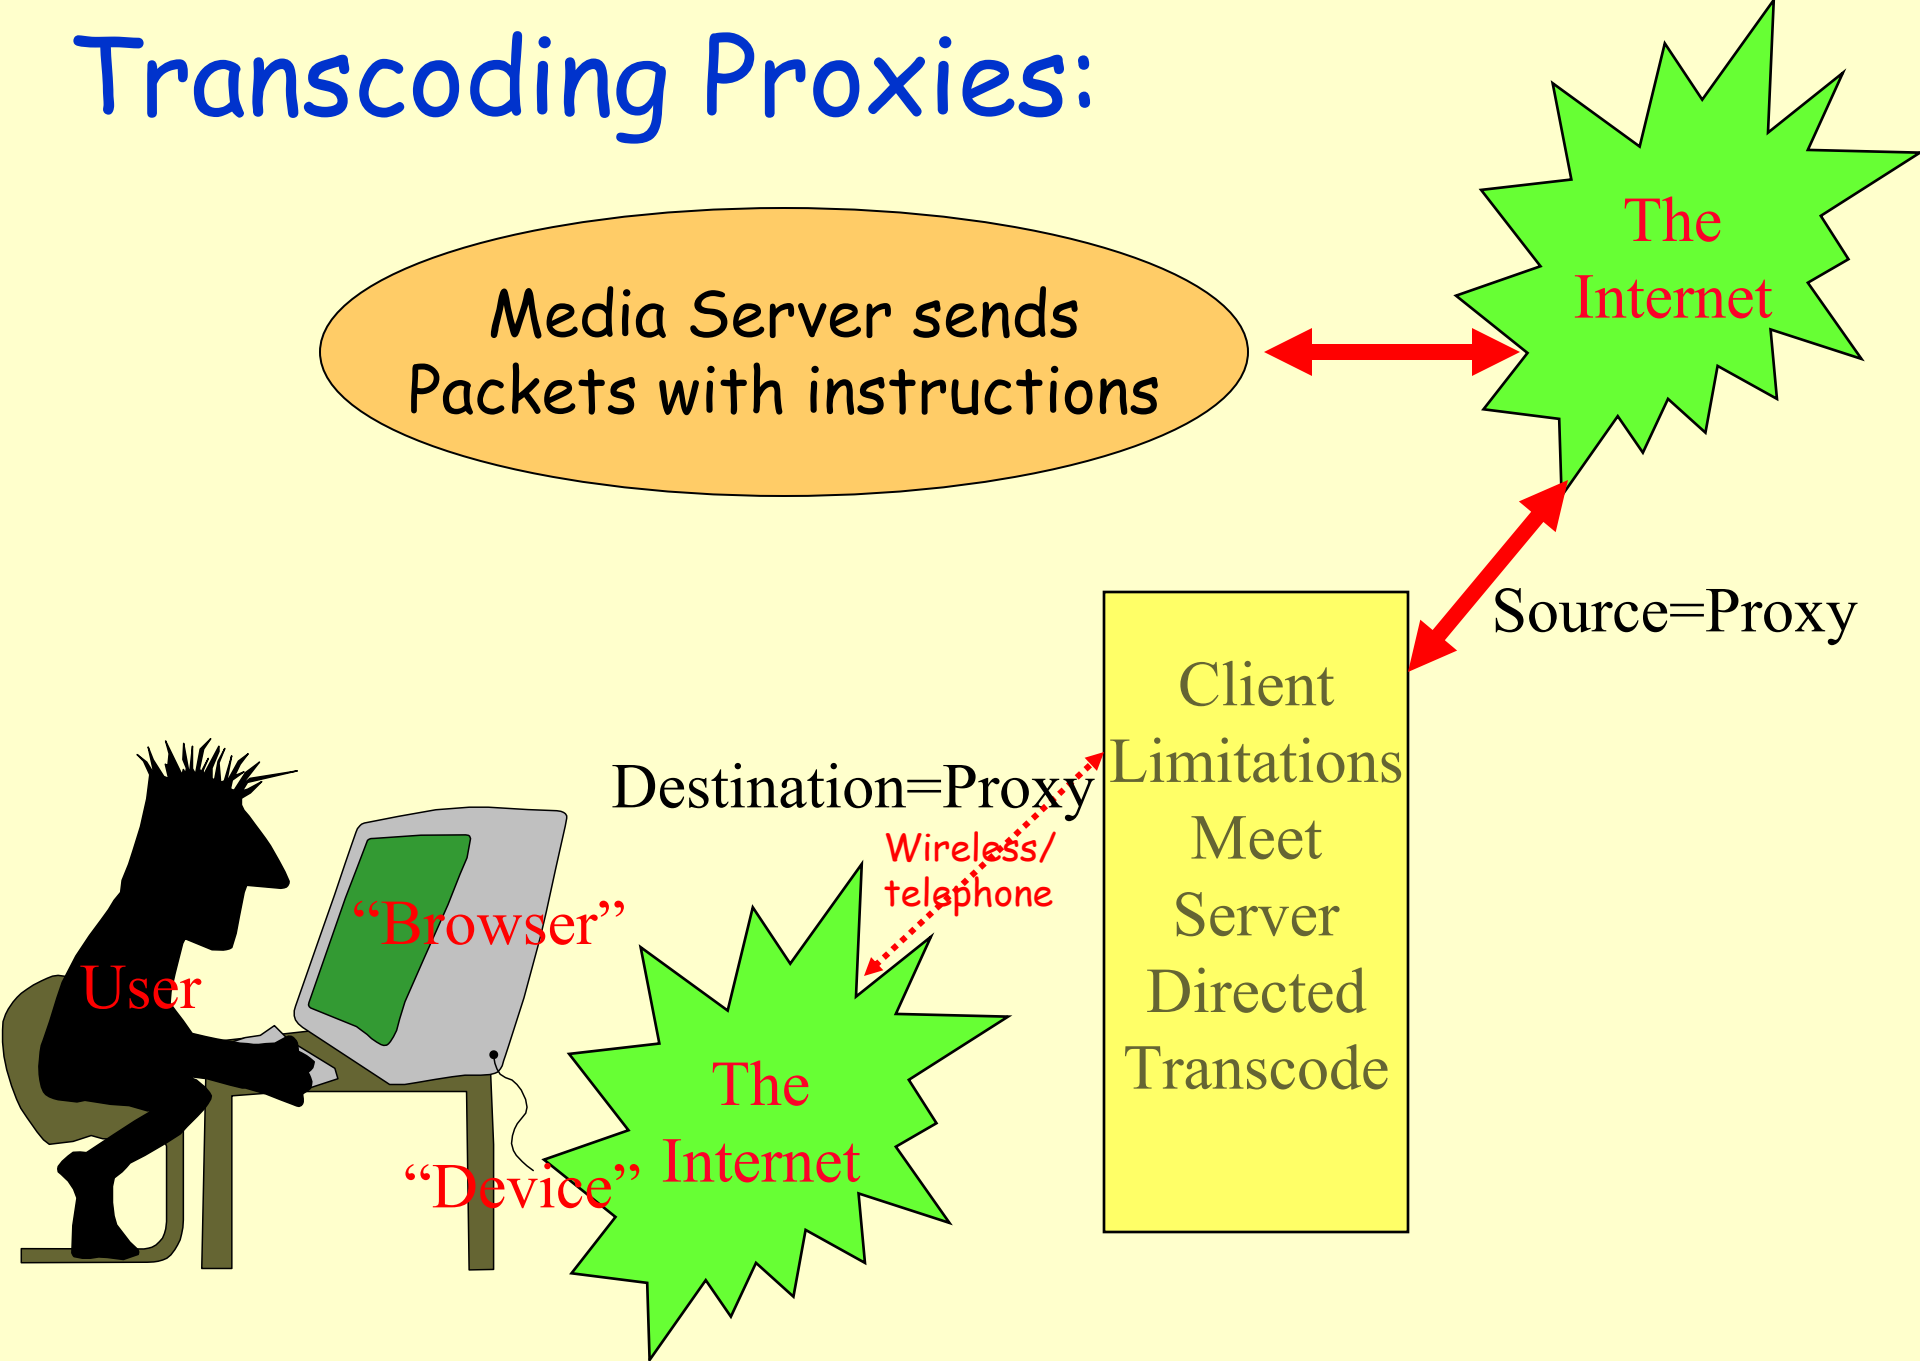
New results

- **FLAME network monitoring:**
 - Role and cost of safety (Cyclone in OpenBSD)
 - Router integration issues (move from OpenBSD software router to 3Com Corebuilder 3500 and other commercial platforms)
- **Market-based resource control with AN:**
 - Further mechanisms - focus on routing
 - Certified resource bounds in packet headers
 - BGP->S-BGP via incentives...

Transcoding proxies

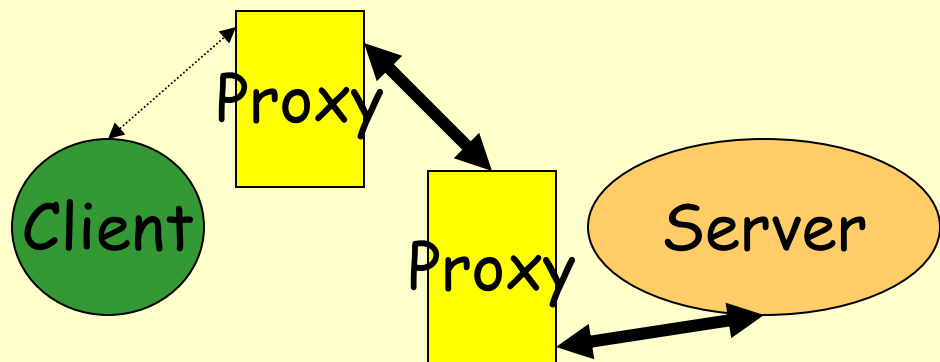

- Client/Server model too simplistic
 - "remote access" style of distributed computing from 70s
 - Convenient model for LANs - Network File System, Remote Procedure Call, *etc.*
- **Diffuse computing** allows a far more general approach
 - WWW can be seen as distributed RPC
 - Diffuse transcoding for network and device heterogeneity

Transcoding Proxies:



Diffuse Computing issues addressed

- Distributes work
 - Easy (in fact desirable) to have multiple proxies
 - Load-balancing with anycast
 - Anycast is path to economics control
- Addresses scale
 - Copes with heterogeneous nets and devices
 - Localizes complexity
 - Diffuse complexity, not centralized at server




Massively Multiplayer Games

- Key area of DoD study (e.g., 2002 ISAT)
 - Commercial systems have outpaced DoD distributed simulations in scale and function
 - Many applications to constructive simulations and hybrid simulations
- **Diffuse Computing** is the path to ultra large scale, robustness, extensibility, *etc.*
 - Natural application of P2P systems
 - What are the systems/networking issues?

P2P for Massively Multiplayer Games

- Hosts - Participate Dynamically as *Peers*
- **Routers**
- Participating Nodes create an *overlay*



Game State

- Highly detailed, becoming more so
- Sets of objects of interest to players
 - Weapons
 - Treasures
 - Buildings, resources, *etc.*
- *Updates* as players alter the world
 - Broadcast does not scale
 - Need consistent game state
 - Issue: partitioning (link/node failure)
 - Issue: scale (200,000+ players in some systems)

Discoveries in MMG research

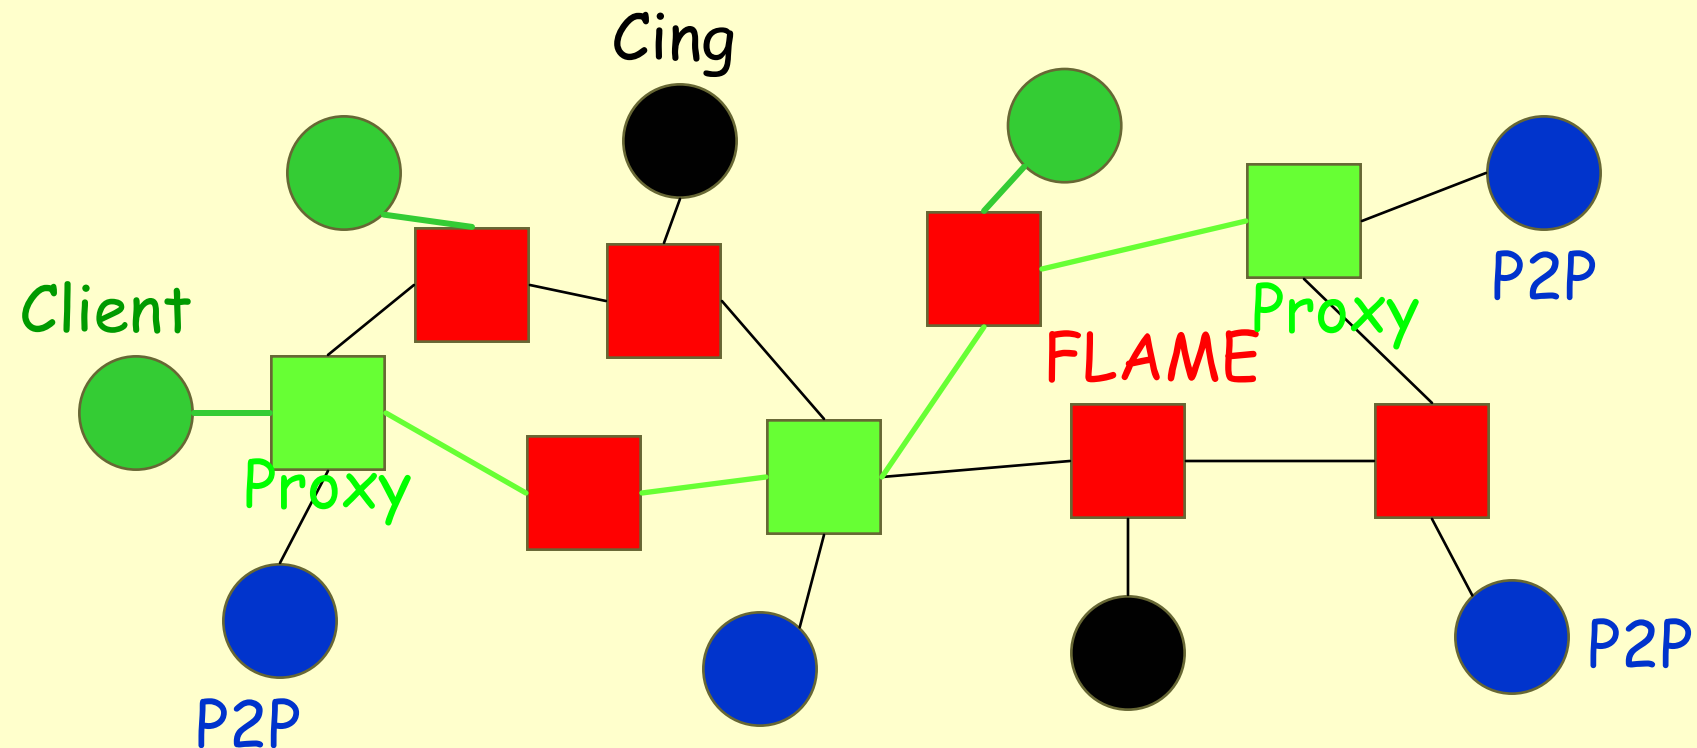
- Interest in game state is localized
 - Players not interested in global picture
 - Also noted in DoD HLA *interest management*
- This can be exploited!
 - Caches/ local replicas / define *regions*
 - New consistency protocol for state updates
 - Reliability with small # of replicas
- Implemented in a simulator for the Pastry P2P system
 - Extensive simulations suggest highly scalable performance, dependent on region structure

Diffuse Computing issues addressed

- Ultra large scale
- Focus on global reliability with local failures
- Tight coupling between networking and distributed computing issues
- Ongoing work, with excellent opportunities to introduce incentives for diffuse control (*e.g.*, for *dynamic* region construction)
- Great opportunity for DoD tech transfer

MANY PLACES for Diffuse Computing!

- Hosts
- Routers
- Network Embedded Diffuse Computing Elements
- Peer-to-peer simulations and multiplayer games



Open Questions, Future Work

- Infuse economic models into proxy model
- Move MMG system from emulator to Pastry (straightforward but *mucho* engineering...)
 - Provide economic incentives to increase locality (adaptive regions) and further accelerate consistency protocols
- Continue to extend and integrate new systems into SPYCE platform (*e.g.*, AN+P2P)
- Embed incentives in algorithms to "coerce" users towards secure systems, *e.g.*, S-BGP

Notes:

- MMG work done by Prof. Honghui Lu, Dr. Bjorn Knutsson, Xu Wai
- Transcoding Proxy work done by Prof. Honghui Lu, Dr. Bjorn Knutsson, Dr. Jeffrey Mogul
- Cing work done by K. Anagnostakis, Prof. Michael Greenwald, R. Ryger
- FLAME work done by K. Anagnostakis, *et al.*
- Papers + Software Distros at:
<http://www.cis.upenn.edu/~spyce>