

## Indice Approfondimenti

- **Workload per il Web (pp. 1-5, inglese)**
- **Proxy caching (pp. 6-11)**
- **Bibliografia (pp. 12-19)**

## Web performance is different

- **Enormous variations**
  - geographical location
  - day of the week
  - hour of the day (understanding peak periods)
- **Workload is heavy-tailed distributed**
  - Very large values are possible with non-negligible probability
- **Dynamic nature of Web transactions**
- **Unpredictable nature of information retrieval and service request**
  - It is difficult to size server capacity to support demand created by load spikes
- **Traffic is bursty in several time scales**
  - The maximum throughput decreases as the burstiness factors increase

## Workload characterization

- **Main components**
  - Client, server, network, protocol
    - ◆ Characterization at different levels
- **Focus on**
  - *arrivals*
    - ◆ session, client/user times, protocol characteristics
  - *object characteristics*
    - ◆ size, popularity, type
  - *service characteristics*
    - ◆ static, volatile, dynamic, and secure

## Workload: arrivals

- **Session**
  - Session length: heavy tailed distribution [Hub98]
  - Session arrival: Poisson process [Wil98, Liu00]
  - User request patterns [Pir99a, Pit99b]
- **User/client times**
  - User think time: heavy tailed distribution [Cro97a, Bar98, Arl00, Mor00]
  - Client parsing time [Bar98, Bar99b]
- **HTTP protocol characteristics**
  - HTTP/1.0 vs. HTTP/1.1 [Hei97, Bar98, Bar99b, Kri99]

## Workload: object characteristics

- **Size**
  - Unique objects, transferred objects [Cro97a, Arl00]
  - Heavy tailed distribution
    - ♦ Most transfers are small
- **Popularity**
  - Reference frequency follows a Zipf-like behavior [Cro97a, Arl00, Jin00]
- **Type**
  - Page composition [Arl00, Bar99a]
  - Analysis at different granularity level:
    - ♦ coarse grain level: no distinction among object type [Arl97, Bar98]
    - ♦ medium grain level: base, embedded, single object [Bar99b]
    - ♦ fine grain level: HTML, image, audio, video, application, dynamic, ... objects [Arl00, Mah00]
  - Most transfers are still for HTML and image objects [Arl00]

## Some workload references

- **Significant amount of research on different Web-server environments [Arl97, Cro97a, Bar98, Arl00, Pit99a, Mah00]**
- **Some recent studies focused on characterization of heavily accessed and dynamic Web-server environments [Iye99, Arl00, Squ00]**

## Potenziamenti del Web

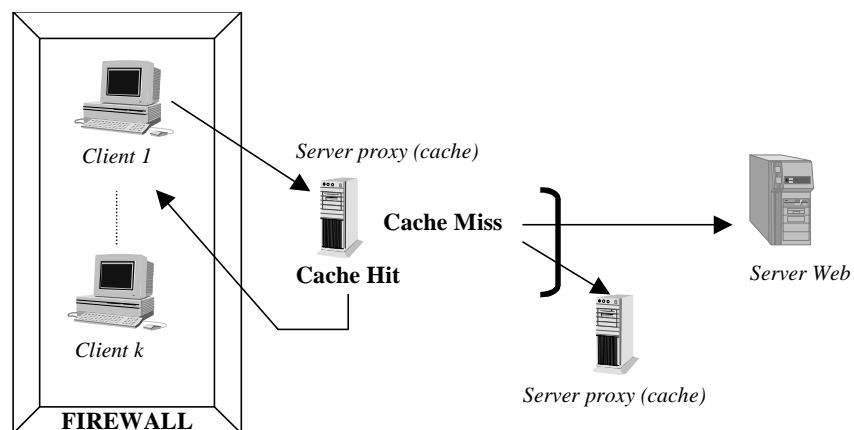
- Azioni a livello di RETE
- Azioni a livello di SISTEMA
- Azioni a livello di INFRASTRUTTURA
  - DNS
  - Caching



© Michele Colajanni, 2001

6

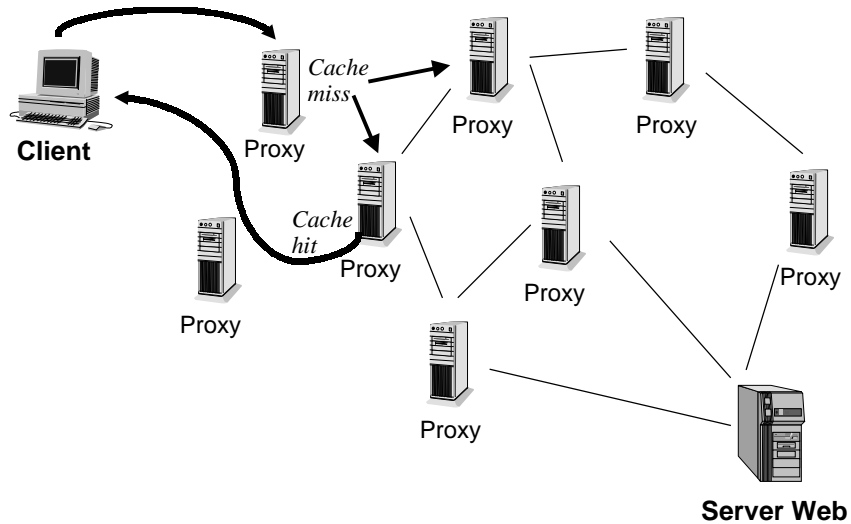
## Caching proxy (singolo)



© Michele Colajanni, 2001

7

## Caching proxy (cooperativo)



© Michele Colajanni, 2001

8

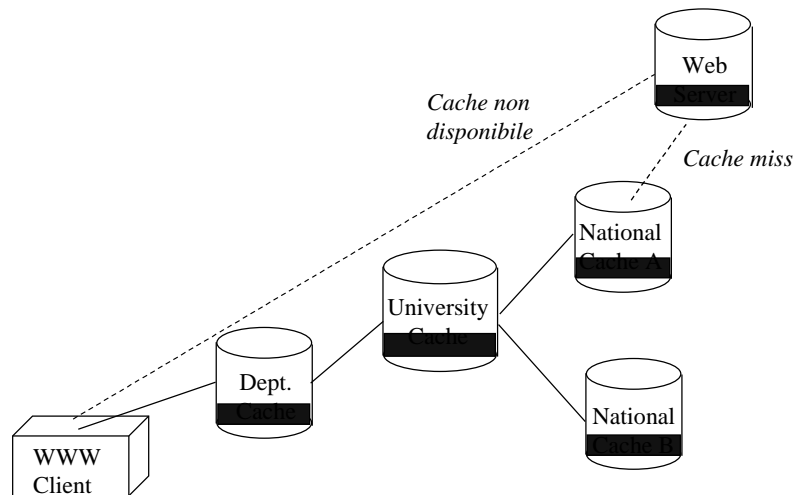
## Sistemi di caching cooperativi

- **Organizzazione delle cache**
  - **Gerarchica**
    - ♦ Harvest, Squid
  - **Orizzontale**
  - **Ibrida**
- **Cooperazione**
  - **On demand.** Es., **Internet Cache Protocol (ICP)** effettua una query multicast ai proxy cooperativi nel caso di cache miss
  - **Periodica**
- **Architettura**
  - **Centralizzata** (con *directory server*)
  - **Distribuita**

© Michele Colajanni, 2001

9

## Es. sistema di caching gerarchico



© Michele Colajanni, 2001

10

## Problemi del proxy caching

- Tradeoff tra contattare Web server e altri proxy
  - non c'è garanzia di migliori *prestazioni*
  - non c'è garanzia di migliore *prossimità Internet*
- Consistenza delle informazioni in cache
- Aumento delle richieste per *pagine dinamiche*
- Altri problemi “legali” e “economici”
  - privacy (verso autorità di controllo sul Web...?)
  - copyright
  - acquisto delle informazioni
  - siti con banner pubblicitari

© Michele Colajanni, 2001

11

## Bibliografia (1)

- [Ade] Adero Inc., <http://www.adero.com>
- [Aka] Akamai Technologies, <http://www.akamai.com>
- [Alt] Alteon Web Systems, <http://www.alteonwebsystems.com>
- [And96] E. Anderson, D. Patterson, E. Brewer, "The Magicrouter, an application of fast packet interposing", University of California, Berkeley, May 1996.  
<http://www.cs.berkeley.edu/~eanders/projects/magicrouter/osdi96-mr-submission.ps>
- [And97] D. Andresen, T. Yang, O.H. Ibarra, "Toward a scalable distributed WWW server on workstation clusters", *J. Parallel and Distributed Computing*, Vol. 42, pp. 91-100, 1997.
- [Arl97] M.F. Arlitt, C.L. Williamson, "Web-server workload characterization: The search for invariants", *IEEE/ACM Trans. on Networking*, Vol. 5, No. 5, pp. 631-645, Oct. 1997.
- [Arl00] M. Arlitt, T. Jin, "A workload characterization study of the 1998 World Cup Web site", *IEEE Network*, May/June 2000.
- [Aro99] M. Aron, P. Druschel, W. Zwaenepoel, "Efficient support for P-HTTP in cluster-based Web servers", *Proc. USENIX 1999*, Monterey, CA, June 1999.
- [Arr] ArrowPoint Communications, <http://www.arrowpoint.com>
- [Bar98] P. Barford, M.E. Crovella, "Generating representative Web workloads for network and server performance evaluations", *Proc. ACM Sigmetrics 1998*, pp. 151-160, July 1998.
- [Bar99a] P. Barford, A. Bestavros, A. Bradley, M.E. Crovella, "Changes in Web client access patterns: Characteristics and caching implications", *World Wide Web*, Vol. 2, No. 1-2, Mar. 1999.
- [Bar99b] P. Barford, M.E. Crovella, "A performance evaluation of HyperText Transfer Protocols", *Proc. ACM Sigmetrics 1999*, Atlanta, May 1999.
- [Bar99c] P. Barford, M.E. Crovella, "Measuring Web performance in the wide area", *ACM Performance Evaluation Review*, Sept. 1999.

## Bibliografia (2)

- [Bar00] G. Barish, K. Obraczka, "World Wide Web caching: Trends and techniques", *IEEE Communications*, Vol. 38, No. 5, May 2000.
- [Bec98] M. Beck, T. Moore, "The Internet2 Distributed Storage Infrastructure project: An architecture for Internet content channels", *Proc. 3th Int'l Web Caching Workshop*, Manchester, UK, June 1998.
- [Bes98] A. Bestavros, M. E. Crovella, J. Liu, D. Martin, "Distributed Packet Rewriting and its application to scalable server architectures", *Proc. 6th IEEE Int'l Conf. Network Protocols*, 1998.
- [Bri95] T. Brisco, "DNS support for load balancing", *RFC 1794*, Apr. 1995.
- [Bun99] R.B. Bunt, D.L. Eager, G.M. Oster, C.L. Williamson, "Achieving load balance and effective caching in clustered Web servers", *Proc. 4th Int'l Web Caching Workshop*, San Diego, pp. 159-169, Apr. 1999.
- [Car97] R.L. Carter, M.E. Crovella, "Server selection using dynamic path characterization in wide-area networks", *Proc. IEEE Infocom 97*, Kobe, Japan, Apr. 1997.
- [Car99a] V. Cardellini, M. Colajanni, P.S. Yu, "Dynamic load balancing on Web-server systems", *IEEE Internet Computing*, Vol. 3, No. 3, pp. 28-39, May/June 1999.
- [Car99b] V. Cardellini, M. Colajanni, P.S. Yu, "Redirection algorithms for load sharing in distributed Web-server systems", *Proc. IEEE 19th Int'l Conf. on Distributed Computing Systems*, Austin, TX, pp. 528-535, June 1999.
- [Car99c] V. Cardellini, M. Colajanni, P.S. Yu, "DNS dispatching algorithms with state estimators for scalable Web-server clusters", *World Wide Web*, Baltzer Science, Vol. 2, No. 3, pp. 101-113, July 1999.
- [Car00a] V. Cardellini, M. Colajanni, P.S. Yu, "Impact of workload models in evaluating the performance of distributed Web-server systems", in *System Performance Evaluation: Methodologies and Applications*, E. Gelenbe ed., CRC Press, Mar. 2000.
- [Car00b] V. Cardellini, M. Colajanni, P.S. Yu, "Geographic load balancing for scalable distributed Web systems", submitted to *IEEE Mascots 2000*, April 2000.

## Bibliografia (3)

- [Cas00] E. Casalicchio, M. Colajanni, "Global scheduling algorithms for high performance Web server clusters", *Proc. of Cluster 2000*, IEEE Computer Society, Nov. 2000.
- [Cho00] C.-F. Chou, L. Golubchik, J. Lui, "Striping doesn't scale: How to achieve scalability for continuous media servers with replication", *Proc. 20th ICDCS*, pp. 64-71, Apr. 2000.
- [CisDD] Cisco's DistributedDirector, <http://www.cisco.com/warp/public/cc/cisco/mkt/scale/distr/index.shtml>
- [CisLD] Cisco's LocalDirector, <http://www.cisco.com/warp/public/cc/cisco/mkt/scale/local/>
- [CisMN] Cisco's MultiNode Load Balancing, <http://www.cisco.com/warp/public/cc/cisco/mkt/iworks/data/mnlb/>
- [Coh99] A. Cohen, S. Rangarajan, H. Slye, "On the performance of TCP splicing for URL-aware redirection", *Proc. 2nd USENIX Symp. On Internet Technologies and Systems*, Oct. 1999.
- [Col98a] M. Colajanni, P.S. Yu, V. Cardellini, "Dynamic load balancing in geographically distributed heterogeneous Web-servers", *Proc. IEEE 18th Int'l Conf. on Distributed Computing Systems*, Amsterdam, pp. 295-302, May 1998.
- [Col98b] M. Colajanni, P.S. Yu, D.M. Dias, "Analysis of task assignment policies in scalable distributed Web-server system", *IEEE Trans. on Par. and Distr. Systems*, Vol. 9, No. 6, June 1998.
- [Con99] M. Conti, E. Gregori, F. Panzieri, "Load distribution among replicated Web servers: A QoS-based approach", *Proc. 2nd Workshop on Internet Server Performance*, Atlanta, May 1999.
- [Cro95] M.E. Crovella, R.L. Carter, "Dynamic server selection in the Internet", *Proc. 3rd IEEE Workshop on Arch. and Implem. of High Performance Comm. Subsystems*, New York, Aug. 1995.
- [Cro97a] M. E. Crovella, A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes", *IEEE/ACM Trans. on Networking*, Vol. 5, No. 6, Dec. 1997, pp. 835-846.
- [Cro97b] M.E. Crovella, L. Lipsky, "Long-lasting transient conditions in simulations with heavy-tailed workloads", *Proc. 1997 Winter Simulation Conf.*, Atlanta, GA, 1997.

## Bibliografia (4)

- [Dam97] O.P. Damani, P.E. Chung, Y.Huang, C. Kintala, Y.-M. Wang, "ONE-IP: Techniques for hosting a service on a cluster of machines", *J. Computer Networks*, Elsevier, Vol. 30, 1997.
- [Dia96] D.M. Dias, W. Kish, R. Mukherjee, R. Tewari, "A scalable and highly available Web-server", *Proc. 41st IEEE Computer Society Int'l Conf.*, pp. 85-92, Feb. 1996.
- [Dik00] S.G. Dikes, C.L. Jeffery, K.A. Robbins, "An empirical evaluation of client-side server selection algorithms", *Proc. IEEE Infocom 2000*, Mar. 2000.
- [Edd] Eddieware, <http://www.eddieware.org>
- [Ege94] K. Egevang, P.Francis, "The IP Network Address Translator (NAT)", *RFC 1631*, May 1994.
- [F5] F5 Networks, <http://www.f5labs.com>
- [Fou] Foundry Networks, <http://www.foundrynetworks.com>
- [Fox97] A. Fox, S.D. Gribble, Y. Chawathe, E.A. Brewer, P. Gauthier, "Extensible cluster-based scalable network services", *Proc. 16th ACM Symp. On Oper. System Princ.*, France, 1997.
- [Gan00] X. Gan, B. Ramamurthy, "LSMAC: An improved load sharing network service dispatcher", *World Wide Web*, Baltzer Science, Vol. 3, No. 1, 2000.
- [Gar95] M. Garland, S. Grassia, R. Monroe, S. Puri, "Implementing Distributed Server Groups for the World Wide Web", Tech. Rep. CMU-CS-95-114, Carnegie Mellon Univ., 1995.
- [Hab98] M.A. Habib, G. Abdulla, E.A. Fox, "Web traffic characterization with time zones", *Proc. ISAVIIA 1998*, Germany, Aug. 1998.
- [Har99] M. Harchol-Balter, M.E. Crovella, C.D. Murta, "On choosing a task assignment policy for a distributed server system", *J. of Parallel and Distributed Computing*, Vol. 59, pp. 204-228, 1999.
- [Hei97] J. Heidemann, K. Obraczka, J. Touch, "Modeling the performance of HTTP over several transport protocols", *IEEE/ACM Trans. on Networking*, Vol. 5, No. 5, pp. 616-630, Oct. 1997.



## Bibliografia (5)

- [Hun98] G.D.H. Hunt, G.S. Goldszmidt, R.P. King, R. Mukherjee, "Network Dispatcher: A connection router for scalable Internet services", *J. Computer Networks*, Elsevier, Vol. 30, 1998.
- [Hyd] HydraWeb Technologies, <http://www.hydraweb.com>
- [IBMND] IBM Network Dispatcher, <http://www.ibm.com/software/network/dispatcher/>
- [Iye99] A.K. Iyengar, M.S. Squillante, L. Zhang, "Analysis and characterization of large-scale Web server access patterns and performance", *World Wide Web*, Vol. 2, No. 1-2, Mar. 1999.
- [Iye00a] A.K. Iyengar, D. Rosu, D. Dias, "Hint-based acceleration of Web proxy caches", *Proc. 19th IEEE Int. Performance, Computing, and Communications Conference*, Phoenix, Feb. 2000.
- [Iye00b] A.K. Iyengar, J. Challenger, D. Dias, P. Dantzig, "High-performance Web site design techniques", *IEEE Internet Computing*, pp. 17-26, Mar./Apr. 2000.
- [Jin00] S. Jin, A. Bestavros, "Temporal locality in Web request streams", *Proc. ACM Sigmetrics 2000*, Santa Clara, CA, June 2000.
- [Kar98] M. Karaul, Y.A. Korilis, A. Orda, "A market-based architecture for management of geographically dispersed, replicated Web servers", *Proc. 1st Int'l Conf. Information and Computation Economics*, pp. 158-165, Charleston, SC, 1998.
- [Kri99] B. Krishnamurthy, J.C. Mogul, D.M. Kristol, "Key differences between HTTP/1.0 and HTTP/1.1", *Proc. 8th Int'l World Wide Web Conf.*, Toronto, May 1998.
- [Kwa95] T.T. Kwan, R.E. McGrath, D.A. Reed, "NCSA's World Wide Web server: Design and performance", *IEEE Computer*, No. 1, pp. 68-74, Nov. 1995.
- [Lie98] P.W.K. Lie, J.C.S. Lui, L. Golubchik, "Threshold-based dynamic replication in large-scale video-on-demand systems", *Proc. 8th Int. Workshop on Continuous-Media Databases and Applications*, pp. 52-59, 1998.
- [Lin] Linux Virtual Server, <http://www.linuxvirtualserver.org>

## Bibliografia (6)

- [Liu00] Z. Liu, N. Niclause, C. Jalpa-Villanueva, "Web traffic modeling and performance comparison between HTTP1.0 and HTTP1.1", *System performance evaluation: Methodologies and applications*, E. Gelenbe ed., CRC Press, Mar. 2000.
- [Luo98] A. Luotonen, *Web Proxy Servers*, Prentice Hall, 1998.
- [Mah00] A. Mahanti, C. Williamson, D. Eager, "Traffic analysis of a Web proxy caching hierarchy", *IEEE Network*, May/June 2000.
- [McM99] P. McManus, "A passive system for server selection within mirrored resource environments using as path length heuristics", <http://proximate.appliedtheory.com>, Apr. 1999.
- [Mir] Mirror Image Internet, <http://www.mirror-image.com>
- [Mog95] J.C. Mogul, "Network behavior of a busy Web server and its clients", Research Report 95/5, DEC Western Research Laboratory, Oct. 1995.
- [Mol00] M. Molina, P. Castelli, G. Faddis, "Web traffic modeling exploiting TCP connections' temporal clustering through HTML-REDUCE", *IEEE Network*, May/June 2000.
- [Mor00] R. Morris, "Variance of aggregated Web traffic", *Proc. IEEE Infocom 2000*, 2000.
- [Mos97] D. Mosedale, W. Foss, R. McCool, "Lesson learned administering Netscape's Internet site", *IEEE Internet Computing*, Vol. 1, No. 2, pp. 28-35, Mar.-Apr. 1997.
- [Mou97] A. Mourad, H. Liu, "Scalable Web server architectures", *Proc. IEEE Int'l Symp. on Computers and Communications*, Alexandria, Egypt, pp. 12-16, July 1997.
- [Obr99] K. Obraczka, F. Silva, "Looking at network latency for server proximity", Tech. Rep. USC-99-714, Information Science Institute, Oct. 1999.
- [Pai98] V.S. Pai, M. Aron, G. Banga, M. Svendsen, P. Druschel, W. Zwaenepoel, E. Nahum, "Locality-aware request distribution in cluster-based network servers", *Proc. 8th Int'l Conf. on Arch. Support for Programming Languages and Operating Systems*, San Jose, CA, Oct. 1998.

## Bibliografia (7)

- [Pan98] R. Pandey, J.F. Barnes, R. Olsson, "Supporting Quality of Service in HTTP servers", *Proc. ACM PODC'98*, pp. 247-256, Puerto Vallarta, Mexico, 1998.
- [Pax97a] V. Paxson, "End-to-end routing behavior in the Internet", *IEEE/ACM Trans. on Networking*, Vol. 5, No. 5, pp. 601-615, Oct. 1997.
- [Pax97b] V. Paxson, S. Floyd, "Why we don't know how to simulate the Internet", *Proc. 1997 Winter Simulation Conf.*, Atlanta, GA, 1997.
- [Pir99] P. Piroli, J.E. Pitkow, "Distributions of surfers' paths through the World Wide Web: Empirical characterization", *World Wide Web*, Baltzer Science, Vol. 2, No. 1-2, Mar. 1999.
- [Pit99a] J.E. Pitkow, "Summary of WWW characterization", *World Wide Web*, Baltzer Science, Vol. 2, No. 1-2, pp. 3-13, Mar. 1999.
- [Pit99b] J.E. Pitkow, P. Piroli, "Mining longest repeating subsequences to predict World Wide Web surfing", *Proc. USENIX 1999*, Monterey, CA, June 1999.
- [Rad] Radware Networks, <http://www.radware.com>
- [ResCD] Resonate Central Dispatcher, [http://www.resonate.com/products/central\\_dispatch](http://www.resonate.com/products/central_dispatch)
- [ResGD] Resonate Global Dispatcher, [http://www.resonate.com/products/global\\_dispatch](http://www.resonate.com/products/global_dispatch)
- [Rnd] RND Networks, <http://www.rndnetworks.com>
- [Sch95] R.J. Schemers, "Ibmnamed: A load balancing name server in Perl", *Proc. 9th Systems Administration Conference*, Monterey, Sep. 1995.
- [Sch00] T. Schroeder, S. Goddard, B. Ramamurthy, "Scalable Web server clustering technologies", *IEEE Network*, May/June 2000.
- [Sin98] A. Singhai, S.-B. Lim, S.R. Radia, "The SunSCALR Framework for Internet servers", *Proc. FCTS 1998*, Munich, Germany, June 1998.

## Bibliografia (8)

- [Son00] J. Song, E. Levy, A. Iyengar, D. Dias, "Design alternatives for scalable Web server accelerators", *Proc. IEEE Int. Symp. on Performance Analysis of Systems and Software*, Austin, Texas, April 2000.
- [Squ00] M.S. Squillante, D.D. Yao, L. Zhang, "Internet traffic: Periodicity, tail behavior, and performance implications", *System performance evaluation: Methodologies and applications*, E. Gelenbe ed., CRC Press, pp. 23-37, Mar. 2000.
- [Sri98] P. Srisuresh, D. Gan, "Load sharing using IP Network Address Translation (LSNAT)", *RFC 2931*, 1998.
- [Vas00] N. Vasiliou, H.L. Lutfiyya, "Providing a differentiated Quality of Service in World Wide Web server", *Proc. Performance and Architecture of Web Servers Workshop*, Santa Clara, California, June 2000.
- [Vin00] R. Vingralek, M. Sayal, Y. Breitbart, P. Scheuermann, "Web++: Architecture, design, and performance", *World Wide Web*, Baltzer Science, Vol. 3, No. 2, 2000.
- [Wil98] W. Willinger, V. Paxson, "Where Mathematics meets the Internet", *Notices of the American Mathematical Society*, Vol. 45, No. 8, Aug. 1998.
- [Yos97] C. Yoshikawa, B. Chun, P. Eastham, A. Vahdat, T. Anderson, D. Culler, "Using Smart Clients to build scalable services", *Proc. USENIX 1997*, Jan. 1997.
- [Zhu99] H. Zhu, B. Smith, T. Yang, "Scheduling optimization for resource-intensive Web requests on server clusters", *Proc. 11th ACM Symp. On Parallel Algorithms and Architectures*, June 1999.