

EXHIBIT C

Part 3



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide



THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [categorizing network page](#)

Found 2 of 185,942

Sort results by Save results to a Binder
 Display results Search Tips
 Open results in a new window

Try an [Advanced Search](#)
 Try this search in [The ACM Guide](#)

Results 1 - 2 of 2

Relevance scale

1 [Performance Workload Char. and Adaptation: Improving web performance by client](#)

[characterization driven server adaptation](#)

Balachander Krishnamurthy, Craig E. Wills

May 2002 **Proceedings of the 11th international conference on World Wide Web**

Publisher: ACM Press

Full text available: pdf(241.76 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We categorize the set of clients communicating with a server on the Web based on information that can be determined by the server. The Web server uses the information to direct tailored actions. Users with poor connectivity may choose not to stay at a Web site if it takes a long time to receive a page, even if the Web server at the site is not the bottleneck. Retaining such clients may be of interest to a Web site. Better connected clients can receive enhanced representations of Web pages, such ...

Keywords: client characterization, client connectivity, server adaptation

2 [Intelligent web information access: An intelligent search agent system for semantic](#)
 [information retrieval on the internet](#)

Carmine Cesarano, Antonio d'Acierno, Antonio Picariello

November 2003 **Proceedings of the 5th ACM international workshop on Web information and data management**

Publisher: ACM Press

Full text available: pdf(332.66 KB)





Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we describe a prototype system for information retrieval on the Internet. Our idea is that the Web has to be searched both *semantically* and *syntactically*. In order to automatically categorize the web pages *on the fly* we propose a novel approach based on ontology and semantic networks and we describe a prototype system based on the Intelligent Agent Paradigm. Preliminary experiments are shown and discussed while describing open problems and on-going research.

Keywords: information retrieval, ontology, semantic network, web agents

Results 1 - 2 of 2

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

+abstract:network +abstract:page abstract:user-defined abstr



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used network page user defined header

Found 186 of 185,942

Sort results by: Save results to a Binder
Display results: Search Tips
 Open results in a new window

Try an [Advanced Search](#)
Try this search in [The ACM Guide](#)

Results 1 - 20 of 186

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Relevance scale

1 [Device-aware desktop web page transformation for rendering on handhelds](#)

A. Artail, Mackram Raydan
November 2005 **Personal and Ubiquitous Computing**, Volume 9 Issue 6

Publisher: Springer-Verlag

Full text available: [pdf\(752.77 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper illustrates a new approach to automatic re-authoring of web pages for rendering on small-screen devices. The approach is based on automatic detection of the device type and screen size from the HTTP request header to render a desktop web page or a transformed one for display on small screen devices, for example, PDAs. Known algorithms (transforms) are employed to reduce the size of page elements, to hide parts of the text, and to transform tables into text while preserving the structu ...

Keywords: Context awareness, Mobile device types, Small screen devices, Transcoding, Web browsing, Web page rendering, Wireless devices

2 [IEEE/ACM Transactions on Networking \(TON\) Volume 1 Issue 1](#)

February 1993 issue , Volume 1 Issue 1

Publisher: IEEE Press

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Vol. 1, No. 1 (February 1993). Published bimonthly by The Institute of Electrical and Electronics Engineers, Inc. and the Association for Computing Machinery, Inc. Annual subscription: \$235 (nonmembers), \$22 (IEEE and ACM members). Subscriptions should be sent to the IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331. ACM order number 11509300. ISSN 1063-6692. The explosive growth in the research, development, ...

3 [Miscellaneous: Dynamic page migration with stochastic requests](#)

Marcin Bienkowski

July 2005 **Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures SPAA '05**

Publisher: ACM Press


Full text available: [pdf\(231.03 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The page migration problem is one of subproblems of data management in networks. It occurs in a distributed network of processors sharing one indivisible memory page of size

D. During runtime, the processors access a unit of data from the page, and the system is allowed to migrate the page between the processors. The problem is to compute (on-line) a schedule of page movements to minimize the total communication cost. The Dynamic Page Migration problem is an extension to the page migration ...


Keywords: data management, dynamic networks, mobile networks, online algorithms, page migration

4 Reducing network latency using subpages in a global memory environment

 Hervé A. Jamrozik, Michael J. Feeley, Geoffrey M. Voelker, James Evans, Anna R. Karlin, Henry M. Levy, Mary K. Vernon


September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5

Publisher: ACM Press

Full text available:  pdf(1.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


New high-speed networks greatly encourage the use of network memory as a cache for virtual memory and file pages, thereby reducing the need for disk access. Because pages are the fundamental transfer and access units in remote memory systems, page size is a key performance factor. Recently, page sizes of modern processors have been increasing in order to provide more TLB coverage and amortize disk access costs. Unfortunately, for high-speed networks, *small* transfers are needed to provide ...

5 Page table management in local/remote architectures

 M. A. Holliday


June 1988 **Proceedings of the 2nd international conference on Supercomputing**

Publisher: ACM Press

Full text available:  pdf(783.30 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


We conjecture that a paged memory with page migration by the operating system may be an effective system environment for a local/remote shared memory architecture executing a single parallel computation. Implementing a paged memory in such an architecture raises several issues with respect to page table management. These issues include page table placement, page table replication level, and page table storage overhead. We discuss these issues, propose alternative solutions, and present an e ...

6 Caching in the Sprite network file system

 M. Nelson, B. Welch, J. Ousterhout


November 1987 **ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5

Publisher: ACM Press

Full text available:  pdf(212.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper describes a simple distributed mechanism for caching files among a networked collection of workstations. We have implemented it as part of Sprite, a new operating system being implemented at the University of California at Berkeley. A preliminary version of Sprite is currently running on Sun-2 and Sun-3 workstations, which have about 1-2 MIPS processing power and 4-16 Mbytes of main memory. The system is targeted for workstations like these and newer models likely to become avail ...

7 [An improved network clustering method for I/O-efficient query processing](#) 


 Sung-Ho Woo, Sung-Bong Yang

November 2000 **Proceedings of the 8th ACM international symposium on Advances in geographic information systems**

Publisher: ACM Press

Full text available:  pdf(651.16 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Efficient network query processing is extremely important in Geographical Information Systems (GIS) and Intelligent Transportation Systems (ITS) which include various applications of transportation, utility and communication networks, etc. In order to reduce the I/O cost in network query processing a given network should be stored with high disk-space utilization and a low edge-cut ratio. To do so the nodes in the network should be clustered in such a way that each cluster fits in a disk page ...

8 [Campus computing newsletter editors & the network: a new role for the 1990s](#) 


 Wendy Rickard Bollentin


August 1990 **Proceedings of the 18th annual ACM SIGUCCS conference on User services**

Publisher: ACM Press

Additional Information: [full citation](#), [abstract](#), [index terms](#)


Campus computing newsletter editors have traditionally been responsible for disseminating relevant technical and general information to a specific user community; most often members of the computing department. As information technology and computer networks become more readily available to a wider range of individuals — and classrooms — editors must now also address a more diverse user community. How do campus computing publications specialists keep up with these and ...

9 [Caching I: Fighting against two adversaries: page migration in dynamic networks](#) 

 Marcin Bienkowski, Mirosław Korzeniowski, Friedhelm Meyer auf der Heide


June 2004 **Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures**


Publisher: ACM Press

Full text available:  pdf(236.15 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Page migration is one of the fundamental subproblems in the framework of data management in networks. It occurs in a distributed network of processors sharing one indivisible memory page of size D , which is stored in one of the processors. During runtime, processors access unit size data items from the page, and the system is allowed to move the page from one processor to another in order to minimize the total communication cost. This problem was considered in the online setting numerous t ...


Keywords: data management, online algorithms, page migration

10 [Shared memory parallelism: Hardware profile-guided automatic page placement for ccNUMA systems](#) 

 Jaydeep Marathe, Frank Mueller

March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '06**


Publisher: ACM Press


Full text available:  pdf(167.64 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Cache coherent non-uniform memory architectures (ccNUMA) constitute an important class of high-performance computing plat-forms. Contemporary ccNUMA systems, such as the SGI Altix, have a large number of nodes, where each node consists of a small number

of processors and a fixed amount of physical memory. All processors in the system access the same global virtual address space but the physical memory is distributed across nodes, and coherence is maintained using hardware mechanisms. Accesses to ...


Keywords: NUMA, hardware performance monitoring, page placement, profile-guided optimization

11 Windows scheduling as a paradigm for pushing information in wireless networks 


 Amotz Bar-Noy

September 2003 **Proceedings of the 2003 joint workshop on Foundations of mobile computing**

Publisher: ACM Press

Full text available:  pdf(100.69 KB) Additional Information: [full citation](#), [abstract](#)


Given are n positive integers w_1, w_2, \dots, w_n called windows. The windows are associated with n equal length information pages. In the windows scheduling problem, the goal is to schedule all the pages on minimum number of identical broadcasting channels such that the gap between two consecutive appearances of page i on any of the channels is at most w_i time slots where a time slot is the broadcasting time of one page. Our main application for windows scheduling is ...

12 Reliability for NSWAP 


America Holloway, Jennifer Barry, Heather Jones


May 2005 **Journal of Computing Sciences in Colleges**, Volume 20 Issue 5

Publisher: Consortium for Computing Sciences in Colleges

Full text available:  pdf(120.31 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)


We present reliability schemes for network swapping enabled Linux clusters. As processor speeds have continued to increase, disk speeds have failed to keep pace, resulting in application performance being restricted by slow disk access. This problem can be alleviated in cluster systems by using idle RAM of other nodes as swap space, a strategy known as network swapping. Nswap, a network swapping system, allows each node in the cluster to either swap-out pages or swap-in pages, according to its c ...

13 Evolution of Xerox's network systems architecture 


 Lawrence Garlick


April 1983 **ACM SIGCOMM Computer Communication Review, Proceedings of the symposium on Communications Architectures & Protocols COMM '83**, Volume 13 Issue 2

Publisher: ACM Press

Full text available:  pdf(106.18 KB) Additional Information: [full citation](#), [abstract](#)


Research and development at Xerox's Palo Alto Research Center (PARC) and Office Systems Division (OSD) have led to many advances in an evolving network architecture. As early experimenters in local area networking, researchers at PARC built a prototype of what is becoming the basis for an international standard. Their early work with networked personal workstations, distributed networked services, and local network interconnection provided valuable insight into requirements for higher level p ...

14 Tradeoffs between false sharing and aggregation in software distributed shared memory 


 Cristiana Amza, Alan Cox, Karthick Rajamani, Willy Zwaenepoel

June 1997 **ACM SIGPLAN Notices, Proceedings of the sixth ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '97**, Volume 32 Issue 7

Publisher: ACM Press

Full text available:  [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Software Distributed Shared Memory (DSM) systems based on virtual memory techniques traditionally use the hardware page as the consistency unit. The large size of the hardware page is considered to be a performance bottleneck because of the implied false sharing overheads. Instead, we show that in the presence of a relaxed consistency model and a multiple writer protocol, a large consistency unit is generally not detrimental to performance. We study the tradeoffs between false sharing and aggreg ...

15 [Managing server load in global memory systems](#) 


 Geoffrey M. Voelker, Hervé A. Jamrozik, Mary K. Vernon, Henry M. Levy, Edward D. Lazowska


June 1997 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '97**, Volume 25 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(2.26 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

New high-speed switched networks have reduced the latency of network page transfers significantly below that of local disk. This trend has led to the development of systems that use network-wide memory, or *global* memory, as a cache for virtual memory pages or file blocks. A crucial issue in the implementation of these global memory systems is the selection of the target nodes to receive replaced pages. Current systems use various forms of an approximate global LRU algorithm for making the ...

16 [Proceedings of 3rd ACM SIGCOMM workshop on Network and system support for games](#) 


 Wu-chang Feng


August 2004 proceeding

Publisher: ACM Press

Additional Information: [full citation](#), [abstract](#)

It is a great pleasure to welcome you all to the *ACM SIGCOMM 2004 Workshops!* We are pleased to present an outstanding program consisting of four workshops: (1) Future Directions in Network Architecture (FDNA); (2) Network and System Support for Games (NetGames); (3) Practice and Theory of Incentives in Networked Systems (PINS); and (4) Network Troubleshooting: Research, Theory, and Operations Practice Meet Malfunctioning Reality (NetTs). Workshops were first introduced as a part of the ACM ...

17 [Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture](#) 


 Kevin Fall, S. Keshav


August 2004 proceeding

Publisher: ACM Press

Additional Information: [full citation](#), [abstract](#)


It is a great pleasure to welcome you all to the *ACM SIGCOMM 2004 Workshops!* We are pleased to present an outstanding program consisting of four workshops: (1) Future Directions in Network Architecture (FDNA); (2) Network and System Support for Games (NetGames); (3) Practice and Theory of Incentives in Networked Systems (PINS); and (4) Network Troubleshooting: Research, Theory, and Operations Practice Meet Malfunctioning Reality (NetTs). Workshops were first introduced as a part of the ACM ...

18 [Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations practice meet malfunctioning reality](#) 

 Jon C.R. Bennett, Mark Allman
September 2004 proceeding
Publisher: ACM Press
Additional Information: [full citation](#), [abstract](#)



It is a great pleasure to welcome you all to the *ACM SIGCOMM 2004 Workshops!*We are pleased to present an outstanding program consisting of four workshops: (1) Future Directions in Network Architecture (FDNA); (2) Network and System Support for Games (NetGames); (3) Practice and Theory of Incentives in Networked Systems (PINS); and (4) Network Troubleshooting: Research, Theory, and Operations Practice Meet Malfunctioning Reality (NetTs).Workshops were first introduced as a part of the ACM ...

19 [Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems](#) 

 Dina Katabi, Rahul Sami
September 2004 proceeding
Publisher: ACM Press
Additional Information: [full citation](#), [abstract](#)

It is a great pleasure to welcome you all to the *ACM SIGCOMM 2004 Workshops!*We are pleased to present an outstanding program consisting of four workshops: (1) Future Directions in Network Architecture (FDNA); (2) Network and System Support for Games (NetGames); (3) Practice and Theory of Incentives in Networked Systems (PINS); and (4) Network Troubleshooting: Research, Theory, and Operations Practice Meet Malfunctioning Reality (NetTs).Workshops were first introduced as a part of the ACM ...

20 [Link mining: a survey](#) 

 Lise Getoor, Christopher P. Diehl
December 2005 **ACM SIGKDD Explorations Newsletter**, Volume 7 Issue 2
Publisher: ACM Press
Full text available:  pdf(209.81 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Many datasets of interest today are best described as a linked collection of interrelated objects. These may represent homogeneous networks, in which there is a single-object type and link type, or richer, heterogeneous networks, in which there may be multiple object and link types (and possibly other semantic information). Examples of homogeneous networks include single mode social networks, such as people connected by friendship links, or the WWW, a collection of linked web pages. Examples of ...

Results 1 - 20 of 186

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)