

Christos T. Karamanolis

PERSONAL DETAILS

Address: 3305 Hillview Ave, Palo Alto, CA 94034, USA **E-mail:** karamanolis AT yahoo DOT com

EMPLOYMENT

10/2005 – present: Staff Engineer. VMware, Inc. Palo Alto, CA, USA.

Working on disaster recovery and continuous data protection technologies for the next generation of Data Center VMware products (ESX server).

01/2002 – 09/2005: Senior Research Scientist. Hewlett-Packard Labs, Palo Alto, CA, USA.

Technical lead of a number of storage systems research projects (4-6 researchers). Responsible for the design of a planetary-scale distributed file system that supports read/write data access. Developed techniques for automatic, QoS-conscious data placement in large data repositories, such as Content Delivery Networks or large NAS farms. Also worked on approaches to provide differentiated service quality for computing services hosted on a shared infrastructure. Consultant to HP's Storage Division and the National Labs (LLNL) for the design of a cluster file system for use in large NAS clusters. Member of the leadership team that sets HP's technology directions for enterprise storage.

05/2000 – 12/2001: Research Scientist. Hewlett-Packard Labs, Palo Alto, CA, USA.

Designed and developed a distributed file system that performs storage virtualization at the file level; files are stored on different, potentially distributed file servers according to content type, size or load balancing requirements, but they are accessed through a single, uniform namespace. Developed a file-access protocol (extended NFS) that splits the data and metadata access paths, when clients are directly attached to a SAN.

01/2000: Offered tenured Lectureship. Department of Computing, Imperial College, UK.

10/1998 – 04/2000: Post-doctoral Research Fellow. Department of Computing, Imperial College.

Member of the team that developed the Darwin/Regis middleware platform for the development and execution of distributed applications. Responsible for the design of configurable protocol stacks that are composed out of lightweight sub-protocols; stacks are configured according to application needs and system parameters; they allow the development of verifiable protocols and extensive code reuse. Investigated the use of automated software verification methods for the design of provably correct protocols. Project partners: U. of Newcastle, INRIA research labs, Bull, Philips. Also, the manager (part-time) of an IT group. Taught a graduate course on Distributed Algorithms.

03/1997 - 09/1998: Computer Engineer. 4th Army Corps, Greek Army (compulsory national service).

Member of technical staff in the IT Department, C⁴I Systems group. Responsible for the design, deployment and management of classified computer networks.

09/1992 - 03/1997: PhD student & Research Associate. Department of Computing, Imperial College.

Designed and implemented novel replication protocols for the provision of critical computing services in large, open distributed systems. Proposed algorithms for group communication (reliable, ordered multicast) and membership protocols. Communication protocols for client applications that access replicated services in a range of system models and with different application requirements.

07/1991 – 08/1992: Software developer / consultant.

Designed and developed engineering applications for customers in the energy and mineral extraction business (C, C++). IT consultant to the Commission of the European union (DG-23).

01/1989 - 09/1990: Systems Administrator (part-time). Computer Technology Institute, Patras, Greece.

Responsible for a cluster of DEC VAX/VMS systems used by 200 researchers and students.

06/1987 - 08/1987: Applications Programmer. North Aegean Petroleum Company, Kavala, Greece.

Summer internship. Developed engineering applications (numerical methods) in Pascal and Fortran.

EDUCATION

PhD in Distributed Computing (1992 - 96).

Department of Computing, Imperial College, University of London, UK.

Thesis title: “Configurable Highly Available Distributed Services”.

Awards/scholarships: EPSRC, The British Council, NATO (Science Fellowship), Greek Government.

Diploma of Engineering in Computing (1986 - 91).

Department of Computer Engineering and Informatics, University of Patras, Greece.

Diploma thesis title: “An X.400 compliant message handling system”.

First Class Honors (top 10% of class). Top 2% in the National University Entrance Examinations (1986).

RELEVANT EXPERIENCE

Technical skills:

- **Operating Systems.** Extensive experience with kernel development (Linux, VMware ESX kernel). Developed a distributed file system in the Linux kernel that was based on the existing NFS code base. Also, a systems administrator for several years both at the Computer Technology Institute and at Imperial College; have extensive experience with Unix variants, *VMS* and *Windows variants*.
- **Programming languages.** Developed a compiler front-end in *Java*, for an Architectural Description Language called *Darwin*. Participated in the design and development of a run-time system for the deployment of Darwin programs on a Java/RMI environment. Implemented communication protocols and application case studies (a highly available name service) in *C++*. Used *C* for developing a distributed file system (Linux kernel), for application programming and for the implementation of the memory paging module of an experimental OS. Used *Pascal* and *Fortran* for developing engineering applications.
- **Communication protocols.** *TCP/IP*, *Multicast IP*. Proposed and published novel algorithms for group communication (reliable, ordered multicast and membership service) and replication protocols, on top of IP multicast.
- **Concurrent programming.** Developer of *Regis*, a middleware platform aimed at the construction and execution of distributed programs. *Regis* follows a constructive approach to the development of distributed programs by separating program structure from computation and communication. It incorporates a package of non-preemptive threads and a flexible communication system based on configurable protocol stacks composed out of light-weight sub-protocol layers.

Awards / recognition:

- **O-1** special status visa and later Green Card by the US Government, as an “exceptional scientist”.
- **Best Paper in Distributed Systems**, awarded by the British Computer Society and sponsored by AT&T Research Labs. Paper: “Client-Access Protocols for Replicated Services” (*IEEE Transactions on Software Engineering*, Vol. 25, No. 1, January/February 1999).
- **Post-graduate scholarships**: NATO, British Council, EPSRC, Greek government.

Professional activities:

- Program committee, Int. Workshop on Applications & Economics of P2P Systems (AEPP), Nov. 2005.
- Program committee, International Conference on Parallel Processing (ICPP), June 2005.
- Program committee, 3rd USENIX Conference on File and Storage Technologies (FAST), March 2004.
- Program Committee, International Workshop on Large-Scale Group Communication, October 2003.
- Acted as reviewer for international scientific journals (IEEE Trans. on Computers, Trans on Software Engineering, ACM TOSEM, Distributed Computing) and conferences (SOSP, OSDI, ASPLOS, ICDCS, Sigmetrics, Infocom, DSN, SRDS, IPDPS). Manuscript reviewer (Distributed Systems, Distributed Algorithms) for McGraw-Hill.

Teaching and communication skills:

- 1999 and 2000 spring terms: *taught* a graduate course on Distributed Algorithms, Dept. of Computing, Imperial College. Advisor of two graduate (Masters) students.
- 1992 – 97: *teaching assistant* for: Data Communication & Networks, Databases, Operating Systems, Distributed Systems, Distributed Programming.
- Have given numerous talks at *conferences, seminars, industry fora* and *corporate customers*. Have established contacts and regular *collaboration* with many researchers world-wide, both in industry (HP, Microsoft Research, Intel, Bull) and in academia (UC Berkeley, CMU, Duke, UC Santa Cruz, Newcastle, Toronto, Bologna).

PUBLICATIONS

Journal Publications:

1. Magnus Karlsson, Christos Karamanolis and Xiaoyun Zhu. Triage: Performance Differentiation for Storage Systems using Adaptive Control. *ACM Transactions on Storage*, Vol. 1, No. 4, November 2005.
2. Christos Karamanolis and Jeff Magee. Client-Access Protocols for Replicated Services. *IEEE Transactions on Software Engineering*, Vol. 25, No. 1, January/February 1999.
Winner of the 1999 "Brendan Murphy Memorial Prize" for the Best Paper in Distributed Systems, awarded by the British Computer Society and sponsored by AT&T Research Labs.
3. Christos Karamanolis and Jeff Magee. Construction and Management of Highly Available Services in Open Distributed Systems. *Distributed Systems Engineering Journal*. Vol. 5, No. 1, March 1998.

Conference Publications:

4. Magnus Karlsson and Christos Karamanolis. Non-intrusive Performance Management for Computer Services. In proc. of the *7th ACM/IFIP/Usenix Middleware Conference (Middleware)*, Melbourne, Australia, November 2006.
5. Magnus Karlsson, Christos Karamanolis and Jeff Chase. Controllable Fair Queuing for Meeting Performance Goals. In proc. of the *IFIP International Symposium on Computer Performance Modeling, Measurement and Evaluation (PERFORMANCE)*, Juan-les-Pins, France, October 2005.
6. Christos Karamanolis, Magnus Karlsson and Xiaoyun Zhu. Designing Controllable Computer Systems. In proc. of the *10th USENIX Workshop on Hot Topics in Operating Systems (HotOS)*, Santa Fe, NM, USA, June 2005.
7. Magnus Karlsson, Christos Karamanolis and Xiaoyun Zhu. An Adaptive Optimal Controller for Non-Intrusive Performance Differentiation in Computing Services. In proc. of the *5th International Conference on Control and Automation (ICCA)*, Budapest, Hungary, June 2005.
8. Magnus Karlsson, Christos Karamanolis and Xiaoyun Zhu. Triage: Performance Isolation and Differentiation for Storage Systems. In proc. of the *12th IEEE International Workshop on Quality of Service (IWQoS)*, Montreal, Canada, June 2004.
9. Lawrence You, Christos Karamanolis. Evaluation of Efficient Archival Storage Techniques. In proc. of the *21st IEEE Symposium on Mass Storage Systems and Technologies (MSST)*, College Park, MD, USA, April 2004.
10. Magnus Karlsson, Christos Karamanolis. Choosing Replica Placement Heuristics for Wide-Area Systems. In proc. of the *24th International Conference on Distributed Computing Systems (ICDCS)*, Tokyo, Japan, March 2004.
11. Zhichen Xu, Magnus Karlsson, Chunqiang Tang and Christos Karamanolis. Towards a Semantic-Aware File Store. In proc. of *HotOS-IX*, Kauai, Hawaii, USA, May 2003.
12. Y. Saito, C. Karamanolis, M. Karlsson, M. Mahalingam. Taming aggressive replication in the Pangaea wide-area file system. In proc. of the *5th Symposium on Operating Systems Design and Implementation (OSDI)*, Boston, MA, USA, December 2002.

13. Y. Saito, C. Karamanolis. Pangaea: a symbiotic wide-area file system. In proc. of the *10th ACM-SIGOPS European Workshop*, Saint-Emilion, France, September 2002.
14. M. Mahalingam, C. Karamanolis, M. Karlsson, Z. Xu. Locating Logical Volumes in Large-Scale Networks. In proc. of the *19th IEEE Symposium on Mass Storage Systems (MSST)*, Adelphi, MD, USA, April 2002.
15. A. Bhide, C. Karamanolis, et al. File Virtualization with DirectNFS. In proc. of the *19th IEEE Symposium on Mass Storage Systems (MSST)*, Adelphi, MD, USA, April 2002.
16. Zheng Zhang and Christos Karamanolis. Designing a Robust Namespace for Distributed File Services. In proc. of the *20th IEEE Symposium on Reliable Distributed Systems (SRDS)*, New Orleans, LA, USA, October 2001.
17. C. Karamanolis, D. Giannakopoulou, J. Magee, S. Wheeler. Model Checking of Workflow Schemas. In proc. of the *4th International Enterprise Distributed Object Computing Conference (EDOC)*, Makuhari Japan, September 2000.
18. Christos Karamanolis and Jeff Magee. Client-Access Protocols for Replicated Services. In proc. of the *3rd IEEE International Conference on Engineering of Complex Computer Systems (ICECCS)*, IEEE Computer Society Press. Como, Italy, September 1997.
19. Christos Karamanolis and Jeff Magee. A Replication Protocol to Support Dynamically Configurable Groups of Servers. In proc. of the *3rd International Conference on Configurable Distributed Systems*, IEEE Computer Society Press. Annapolis MD, USA, May 1996.
20. Christos Karamanolis and Jeff Magee. Configurable Highly Available Distributed Services. In proc. of the *14th IEEE Symposium on Reliable Distributed Systems (SRDS)*, IEEE Computer Society Press. Bad Neuenhar, Germany, September 1995.
21. Christos Karamanolis and Jeff Magee. Highly Available Distributed Services. In S. Krakowiak and M. Shapiro, editors, proc. of the *European Research Seminar on Advances in Distributed Systems (ERSADS)*, INRIA. L'Alpe d'Huez, France, April 1995.

Unrefreed Reports:

22. M. Karlsson, C. Karamanolis. Bounds on the Replication Cost for QoS. Hewlett-Packard Labs, Technical Report No.HPL-2003-156, July 2003.
23. M. Karlsson, C. Karamanolis, M. Mahalingam. A Framework for Evaluating Replica Placement Algorithms. Hewlett-Packard Labs, Technical Report No.HPL-2002-219, August 2002.
24. C. Karamanolis, L. Liu, M. Mahalingam, D. Muntz, Z. Zhang. An Architecture for Scalable and Manageable File Services. Hewlett-Packard Labs, Technical Report No. HPL-2001-173, July 2001.
25. C. Karamanolis and Z. Zhang. Cross-Partition Protocols in a Distributed File Service. Hewlett-Packard Labs, Technical Report No. HPL-2001-129, May 2001.
26. M. Mahalingam, C. Karamanolis, L. Liu, D. Muntz, Z. Zhang. Data Migration in a Distributed File Service. Hewlett-Packard Labs, Technical Report No. HPL-2001-128, May 2001.
27. C. Karamanolis, M. Mahalingam, D. Muntz, Z. Zhang. DiFFS: a Scalable Distributed File System. Hewlett-Packard Labs, Technical Report No. HPL-2001-19, January 2001.
28. C. Karamanolis, D. Giannakopoulou, J. Magee, S. Wheeler. Modelling and Analysis of Workflow Processes. Department of Computing, Imperial College. Technical Report No. 99/2, September 1999.
29. Christos Karamanolis and Vassos Hadzilacos. A Modular Approach to Fault-Tolerant Multicasts. Imperial College and University of Toronto. Technical Report, March 1997.

PATENTS

1. "Separate Read And Write Servers In A Distributed File System", US patent number 6883029, issued 19-Apr-2005.
2. "Logical volume-level migration in a partition-based distributed file system", US patent number 6775673, issued 10-Aug-2004.

3. "Updating references to a migrated object in a partition-based distributed file system", US patent number 6775672, issued 10-Aug-2004.
4. "Object-level migration in a partition-based distributed file system", US patent number 6772161, issued 3-Aug-2004.
5. "Namespace Management in a Distributed File System", US patent number 6687701, issued 3-Feb-2004.
6. "Namespace Service in a Distributed File System Using a Database Management System", US patent number 6625604, issued 23-Sep-03.
7. "Distribution of Physical File Systems", US patent number 6601070, issued 29-Jul-03.

18 more patents pending with the United States Patent and Trademark Office.

SELECTED TALKS

- "Designing controllable computer systems", 10th Workshop on Hot Topics in Operating Systems (HotOS-X), Santa Fe, NM, USA, June 2005.
- "Taming aggressive replication in the Pangaea wide-area file system", graduate seminar, UC Santa Cruz, May 2003.
- "Archival Storage Systems", Intel Research, Pittsburgh, May 2003.
- "Cluster File Systems", CERN, Switzerland, March 2003.
- "Designing a Robust Namespace for Distributed File Services", 20th IEEE Symposium on Reliable Distributed Systems (SRDS-2001), New Orleans, LA, USA, October 2001.
- "Client-service state consistency in the presence of replication", AT&T Labs, Cambridge, UK, (09/99), Bell Labs, NJ, USA (12/99), HP Labs, CA, USA (12/99).
- "Client-access protocols for replicated services", 3rd IEEE International Conference on Engineering of Complex Computer Systems (ICECCS'97), Como, Italy, Sept. 1997.
- "A replication protocol to support dynamically configurable groups of servers", Third International Conference on Configurable Distributed Systems, Annapolis, May 1996.
- "Management of availability for distributed services", poster demonstrating the results of the MADS EPSRC project, ITeC'96, Leeds University, April 1996.
- "Highly available distributed services", *invited talk* (while on a research visit supported by *Cabernet*), University of Bologna, Italy, October 1995.
- "Configurable highly available distributed services", 14th IEEE Symposium on Reliable Distributed Systems, Bad Neuenhar, Germany, September 1995.
- "Highly available distributed services", European Research Seminar on Advances in Distributed Systems, INRIA, L'Alpe d'Huez, France, April 1995.