Final Report for Period: 08/2006 - 07/2007 **Submitted on:** 10/29/2007
Principal Investigator: Bein, Wolfgang . **Award ID:** 0312093
Organization: U of Nevada Las Vegas
Title:
ITR: Online Algorithms and Information Technology

Project Participants

Senior Personnel

Name: Bein, Wolfgang
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Larmore, Lawrence
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Reischuk, Ruediger
Worked for more than 160 Hours: No
Contribution to Project:
Ruediger Reischuk is the Dekan of Institute f³r Theoretische Informatik Universitõt zu L³beck, Germany. He contributed substantially to the centerpiece of the project: The Knowledge State Approach. Both Principal Investigators visited him in L³beck over the 03/04 winter break. As a result the principal investigators in collaboration with Reischuk finished up the theoretical underpinning of the Knowledge State Approach. The resulting write-up is now pending at the 45th Annual IEEE Symposium on Foundations of Computer Science.
Another paper with Reischuk applied the Knowledge State Approach to the Caching Problem in Shared Memory Multiprocessor Systems. These were published in the Proceedings of the 7th International Symposium on Parallel Architectures, Algorithms and Networks.

Name: Epstein, Leah
Worked for more than 160 Hours: No
Contribution to Project:
Leah Epstein of the Interdisciplinary Center Herzliya, Israel, has contributed to the work on online batching. A joint paper has been accepted atthe 9th Scandinavian Workshop on Algorithm Theory. Work is ongoing. NSF support was for travel only. Epstein's work was also supported by Israel Science Foundation grant 250/01.

Name: Noga, John
Worked for more than 160 Hours: Yes
Contribution to Project:
John Noga of California State University, Northridge has contributed
to the work on online batching. Our problem is related to the TCP acknowledgement problem. Noga and the Principal investigators meet numerous times at Northridge and at UNLV. A joint paper has been accepted at the 9th Scandinavian Workshop on Algorithm Theory. Work is ongoing. The principal investigators are currently also working with Noga on new algorithms for uniform task systems. NSF support was for travel.

Name: Sudborough, Hal
Worked for more than 160 Hours: No
Contribution to Project:
Hal Sudborough of the University of Texas at Dallas has contributed to the investigators' work on a 2-competitive algorithm for block sorting. Block sorting has important applications in Optical Character Recognition. Sudborough visited UNLV for a research visit and the Principal Investigator and Sudborough also had an opportunity to meet during the 7th International Symposium on Parallel Architectures, Algorithms and Networks in Hong Kong. (The next symposium will be held at UNLV.) NSF support was for travel only. Sudborough also used funds from his endowed position at the University of Texas.

Name: Morales, Linda

Worked for more than 160 Hours: No

Contribution to Project:

Linda Morales of the Texas A&M at Commerce has contributed to our work on a 2-competitive algorithm for block sorting. Morales visited UNLV for a research visit and the Principal Investigator and Morales also had an opportunity to meet during the 7th International Symposium on Parallel Architectures, Algorithms and Networks. NSF support was for travel only.

Name: Brucker, Peter

Worked for more than 160 Hours: No

Contribution to Project:

Peter Brucker of the University of Osnabrueck, Germany, came to UNLV for a research visit. Peter Brucker is a leading expert in Scheduling; his Springer Verlag text 'Scheduling' is in its third edition. His visit inspired the principal investigators to conduct the work on online batching problems. The principal investigators also worked on path problems under algebraic Monge properties and produced a paper, which has been accepted for publication in Discrete Applied Mathematics. NSF support was for travel only. Brucker also received an honorarium from UNLV, and his visit was supported by the Deutsche Forschunsgemeinschaft.

Name: Iwama, Kazuo

Worked for more than 160 Hours: Yes

Contribution to Project:

Kazuo Iwama is the head of an internationally recognized research group at Kyoto University in Japan and a preeminent algorithms researcher. Iwama invited the PI for an extended research stay during December 04 - January 05. The funding for this trip was entirely by Kyoto University and added nicely to the NSF project. The research conducted at Kyoto University lead to fundamental results on the delayed server problem.

Name: Golin, Mordecai

Worked for more than 160 Hours: No

Contribution to Project:

Mordecai Golin is an Associate Professor in the Department of Computer Science of the Hong Kong University of Science and Technology. He has visited the PIs in Spring 2005 to work on coding problems.

Post-doc

Graduate Student

Name: Zhang, Qin

Worked for more than 160 Hours: Yes

Contribution to Project:

Ms. Qin Zhang is a Master's student who received her degree in spring 04. The principal investigator as well as the Co-PI spent considerable amounts of time (especially during Summer 03) guiding her Master's thesis. The work with her reviewed various aspects of the Knowledge State Approach. Ms. Zhang was also involved in work on new limited bookmark algorithms for paging. Ms. Zhang is well positioned to pursue a Ph.D. and we anticipate that she will return to the project as a Ph.D. student in 05. Ms. Zhang also received funding from UNLV.

Name: Fuchs, Berhard

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr. Fuchs is a doctoral student at the Zentrum f³r Angewandte Informatik Universitõt zu K÷ln. Mr. Fuchs joined the project for two month on the suggestion of Faigle who is the head of the institute and a recognized researcher in optimization. The investigators studied online matching problems during Mr. Fuchs' stay. The grant paid for room and travel. Mr. Fuchs also had support from the Deutscher Akademischer Austauschdienst. Collaboration with the Zentrum f³r Angewandte Informatik is ongoing.

Name: Kurlinsky, Joshua

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr. Kurlinski is a Master's student who will graduate in Spring 2006. The Co-PI spent considerable amounts of time guiding his Master's project. Mr. Kurlinski's project focuses on implement a specialized convex hull algorithm needed to compute knowledge

states for the 2 server problem.

Name: Oravec, James

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr Oravec is working on his Master's Thesis with the Co-PI and writes programs in Mathematica for simulations regarding the randomized 2-server problem

Undergraduate Student

Name: Agtarap, Randy

Worked for more than 160 Hours: No

Contribution to Project:

Mr. Agtarap is an independent study student.

Technician, Programmer

Other Participant

Research Experience for Undergraduates

Name: Tran, Hung

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr. Tran has worked on a web tutorial:
'Knowledge State Algorithms and the Server Problem.'

Years of schooling completed: Junior

Home Institution: Same as Research Site

Home Institution if Other:

Home Institution Highest Degree Granted(in fields supported by NSF): Associate's Degree

Fiscal year(s) REU Participant supported: 2005

REU Funding: REU supplement

Name: Rider, Jason

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr.Tran helped with php coding on the project

Years of schooling completed: Other

Home Institution: Same as Research Site

Home Institution if Other:

Home Institution Highest Degree Granted(in fields supported by NSF): Bachelor's Degree

Fiscal year(s) REU Participant supported: 2006

REU Funding: REU supplement

Organizational Partners

Kyoto University

Other Collaborators or Contacts

Ulrich Faigle, Head, Zentrum f³r Angewandte Informatik Universitõt zu K÷ln, served as liaison.
Camil Demitrescus, University of Rome, Organizer FOCS 2004, gave talk on online path problems.
Stefan Pickl, Vistor from Zentrum f³r Angewandte Informatik Universitõt zu K÷ln

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)

Findings: (See PDF version submitted by PI at the end of the report)

Training and Development:

Ms. Qin Zhang is a Master's student who received her degree in Spring 04. The work with her reviewed various aspects of the Knowledge State Approach. A second Master's student, Mr. Joshua Kurlinski, has also graduated. His work included implementation of the optimization engine for generating knowledge states automatically. The optimization engine is highly modular with an object oriented design. Java Netbeans was used for development. Kurlinski has focused on well-defined submodules, such as the main module of the method which creates convex hulls. Kurlinski has also developed a solid understanding of the major theoretical issues of the project. Mr. Kurlinski has graduated in December 2005. A third Master's student, James Oravec, has graduated in 2007. He has worked on geometric models and linear programming techniques for generating knowledge state libraries. In this context, Oravec has used Mathematica extensively. Along more theoretical lines Oravec also worked on T-Theory Applications to Online Algorithms for the Server Problem.

Additionally two REU students, Hung Tran and Jason Rider, and an independent study undergraduate student, Randy Agartarp, have worked on a web-based tutorial.
(See: http://www.egr.unlv.edu/~bein/tutorial/).
This effort have focused on the animation of online algorithms on the web using Flash, Java, and PHP.

All students who participated in the project also were taught by the PI how to use open source software in a Linux environment. Operating systems have included use of distributions such as Debian and Ubuntu, Red Hat and Fedora, as well as Suse.

The principal investigators' graduate level courses were close to the theme of the project. Larmore's course page at
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F03/
and
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F04/
have extensive write-ups on the theoretical underpinnings of the knowledge state approach written at a level accessible to graduate students. Some of the materials provided the basis for the paper with Reischuk on knowledge states.

Outreach Activities:

We have posted a tutorial, which is of interest to non-traditional students and the general public.
(See http://www.egr.unlv.edu/~bein/tutorial/)

In December 2004 the Regents of the University of Nevada approved the Center for the Advanced Study of Algorithms (CASA). CASA facilitates national and international collaborations, interdisciplinary work across UNLV's campus and with other academic entities, as well as events for the

general public.

The PI was the General Chair of the 8th International Symposium on Parallel Architectures, Algorithms, and Networks held at the Universityof Nevada. Las Vegas, in December 2005. This made the the project further visible to a larger academic community.

The PI visited Clark County middle school and gave a number of presentations to Clark County high school seniors through the Howard Hughes College of Engineering Minority Engineering Program at the University of Nevada, Las Vegas.

Journal Publications

Wolfgang Bein, Lawrence L. Larmore, Ruediger Reischuk, "Knowledge States for the Caching Problem in Shared Memory Multiprocessor Systems", 7th International Symposium on Parallel, IEEE Computer Society Architectures, Algorithms and Networks, IEEE, p. 307, vol. , (2004). Published,

Wolfgang Bein, Lawrence L. Larmore, "Trackless and Limited Bookmark Algorithms for Paging", SIGACT News, 35, p. 40, vol. 1, (2004). Published,

Wolfgang Bein, Leah Epstein, John Noga, Lawrence Larmore, "Optimally Competitive List Batching", Proceedings of the 9th Scandinavian Workshop on Algorithm Theory, Lecture Notes in Computer Science, Springer Verlag Lecture Notes in Computer Science, Springer Verlag, p. 77, vol. 3111, (2004). Published,

Wolfgang Bein, Lawrence L. Larmore, and Ruediger Reischuk, "Knowledge State Algorithms: Randomization with Limited Information", Arxiv: archive.org/cs/0701142
*Author(s):

*Title:, p. 17 pages, vol. , (2007). Accepted,

Wolfgang Bein, Peter Brucker, Lawrence Larmore, James K. Park, "The Algebraic Monge Property and Path Problems", Discrete Applied Mathematics, p. 455, vol. 145(3), (2005). Published,

Wolfgang Bein, Lawrence Larmore, Ruediger Reischuk, "Knowledge States for the Caching Problem in Shared Memory Multiprocessor Systems. (Journal Version of the Conference Paper.)", International Journal of Foundation of Computer Science, p. , vol. , (2005). Accepted,

Wolfgang Bein, Lawrence Larmore, Charles Shields, Hal Sudborough, "Embedding a Complete Binary Tree into a Three-Dimensional Grid", Journal of Interconnection Networks, p. 111, vol. 5(2), (2004). Published,

Wolfgang Bein, Kazuo Iwama, Lawrence L. Larmore, John Noga, "The Delayed k-Server Problem", 15th Internationa Syposium on Fundamentals of Compuation Theory, Lecture Notes in Computer Science, Springer Verlag, p. 281, vol. 3623, (2005). Published,

Wolfgang Bein, Lawrence Larmore, Linda Morales, Hal Sudborough, "A Faster and Simpler 2-Approximation for Block Sorting", 15th International Symposium on Fundamentals of Computation Theory, Lecture Notes in Computer Science, Springer Verlag, p. 115, vol. 3623, (2005). Published,

Wolfgang Bein, Lawrence L. Larmore, John Noga, "A Better Algorithm for Uniform Metrical Task Systems with Few States", Proceedings of the 8th International Symposium on Parallel Architectures, Algorithms and Networks. IEEE Computer Society, p. 94, vol. , (2005). Published,

Wolfgang Bein, Lawrence L. Larmore, Mordecai Golin, Yan Zhang, "The Knuth-Yao Quadrangle-Inequality Speedup is a Consequence of

Total-Monotonicity", Proceedings of The ACM-SIAM Symposium on Discrete Algorithms (SODA'06), ACM Press, p. 31, vol. , (2006). Published,

Doina Bein, Wolfgang Bein, Srilaxmi Malladi, "Fault Tolerant Coverage Model for Sensor Networks", Proceedings of the 5th International Conference on Computational Science 2005. (ICCS 2005), Atlanta, Georgia, Lecture Notes in Computer Science, Springer Verlag, p. 535, vol. 3515, (2005). Published,

Doina Bein, Wolfgang Bein, Natasa Brajkovska, Shahram Latifi, "An Optimal Embedding of Honeycomb Networks into Hypercubes", Parallel Processing Letters, p. 367, vol. 14, (2004). Published,

T. C. Hu, Lawrence L. Larmore, J. David Morgenthaler, "Optimal Integer Alphabetic Trees in Linear Time", , Proceedings of the 13th Annual European Symposium on Algorithms, Palma de Mallorca, Spain, October 3-6, 2005, Lecture Notes in Computer Science, Springer-Verlag (2005), pp. 226-237., p. 226, vol. 3669, (2005). Published,

Wolfgang Bein, Lawrence L. Larmore, Mordecai Golin, Yan Zhang, "The Knuth-Yao quadrangle-inequality speedup is a consequence of total-monotonicity (Journal Version)", ACM Transactions on Algorithms, p. , vol. , ( ). Accepted,

Wolfgang Bein, Lawrence L. Larmore, John Noga, "Uniform metrical task systems with a limited number of states", Information Processing Letters, p. 123, vol. 104(4), (2007). Published,

Wolfagn Bein, Kazuo Iwama, Jun Kawahara, Lawrence L. Lawrence, James Oravec, "A Randomized Algorithm for Two Servers in Cross Polytope Spaces", Proceedings, 5th Workshop on Approximation and Online Algorithms, Lecture Notes in Computer Science, Springer Verlag, p. , vol. , ( ). Accepted,

Wolfgang Bein, Lawrecne L. Larmore, John Noga, "Equitable Revisited", Proceedings, 15th Annual European Symposium on Algorithms (ESA), Lecture Notes in Computer Science, p. 419, vol. 4698, (2007). Published,

Wolfgang Bein, Lawrence L. Larmore, Ruediger Reischuk, "Knowledge States: A Tool for Randomized Online Algorithms", Proceedings of the Fourty-First Annual Hawaii International Conference on on System Sciences (CD-ROM), January 7-10, 2008, Computer Society Press, p. 10 pages, vol. , (2008). Published,

Lawrence L. Larmore, Lames Oravec, "T-Theory Applications to Online Algorithms for the Server Problem", Arxiv: archive.org/cs/0611088, p. 38 pages, vol. , (2007). Published,

Wolfgang Bein, Leah Epstein, Lawrence L. Larmore, John Noga, "Optimally Competitive List Batching (Journal Version)", Theoretical Computer Science, p. , vol. , ( ). Submitted,

Wolfgang Bein, Lawrence L. Larmore, Linda Morales, I. Hal Sudborough, "A Quadratic Time 2-Approximation Algorithm for Block Sorting", Theoretical Computer Science, p. , vol. , ( ). Submitted,

Books or Other One-time Publications

Bein WW, Chin FYL, Hsu DF, Palis ML (Editors), "Proceedings of the 8th International Symposium on Parallel Architectures, Algorithms and Networks", (2005). Book, Published
Bibliography: EEE Computer Society, 2005, ISBN 0-7695-2509-1.

Web/Internet Site

URL(s):
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F03/;
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F04/

Description:

Larmore's course pages at http://www.egr.unlv.edu/~larmore/Courses/CSC789/F03/
and
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F04/
have extensive write-ups on the theoretical underpinnings of the knowledge state approach written at a level accessible to graduate students.

Other Specific Products

Product Type:

Teaching aids

Product Description:

A substantial quantity of the material from this project is meant to made available on the world wide web in the form of tutorials, where it will be available broadly across the network, especially to non-traditional students. The investigators (with help from two undergraduate students) have created a preliminary version for a tutorial: "Knowledge State Algorithms and the Server Problem."

Sharing Information:

The tutorial is available at http://www.egr.unlv.edu/~bein/tutorial/

Contributions

Contributions within Discipline:

The core of the project was to better understand the true nature of online randomization. Online algorithms are needed for an enormous variety of practical situations; in fact, most real-life problems require online approaches. By developing the theoretical underpinnings of (i.e. rigorous proofs for) the knowledge state approach the investigators have made significant progress to the understanding of this aspect of computer science. The investigators have pointed out that the well known concept of forgiveness is properly contained in the knowledge state approach. They have demonstrated that the distributional model of computation and the proper use of the work function tool can yield competitive algorithms in numerous situations.

The investigators have made considerable progress on the randomized 2-server problem -- a notorious open problem in online algorithms. The new approach has lead to new results for uniform spaces, the line, and cross polytope spaces. These results are promising to finally settle the question of whether there exists an online algorithm with competitive ratio better than 2 for general spaces. The investigations also have shed new light on the issue of the trade-offs between memory and competitiveness. Progress has been made on constructing optimally competitive paging algorithms with a minimal number of bookmarks required. Memory restrictions seem to be less important in today's word of gigabyte storage. Yet, the opposite is true. For real time data streams for example only very limited information can be stored, for web browsing bookmarks are hard to keep coherent due to frequent updates. More specifically, in collaboration with John Noga, the principal investigators have developed an H_k-competitive algorithm with only with $O(k)$ memory; this result had been conjectured for a long time and is now established.

Another example is cache coherent paging in shared

memory multiprocessor systems where the investigators' method gives new insights. Multiprocessor systems with a global shared memory provide logically uniform data access. To hide latencies when accessing global memory each processor makes use of a private cache. Several copies of a data item may exist concurrently in the system. To guarantee consistency when updating an item a processor must invalidate copies of the item in other private caches. This additional complication may degrade the performance of paging algorithms in an online environment. To exclude the effect of classical paging faults, the investigators assumed that each processor knows its own data access sequence, but does not know the sequence of future invalidations requested by other processors. Performance of a processor with this restriction can be measured against the optimal behavior of a theoretical omniscient processor, using competitive analysis.

The investigators mention that their work on block sorting may implications for computational biology as applied areas of Information Technology. Their results on batching are significant for the TCP acknowledgment problem. With TCP there exists a possibility of using a single acknowledgment packet to simultaneously acknowledge multiple arriving packets, thereby reducing the overhead of the acknowledgments.

Work on speeding up standard dynamic programs shows that even offline problems call for online techniques, as dynamic programming follows an online protocol.

Contributions to Other Disciplines:

Society can benefit as the online techniques the proposers develop are applied to many areas, including robitics computer networking, memory management, and databases.

Contributions to Human Resource Development:

Ms. Qin Zhang is a Master's student received her degree in spring 04. The work with her reviewed various aspects of the Knowledge State Approach and given Ms. Zhang insight into theoretical computer science.

The Co-PI has worked with a new Master's student, Mr. Joshua Kurlinski, who has extensively done the implementation work for simulations. His work includes further implementation of the optimization engine for generating knowledge states automatically. The optimization engine is highly modular with an object oriented design. Java Netbeans is used for development. Thus this student project has focused on well-defined submodules, such as the main module of the method that has to do with creating convex hulls. Of course, programming work is meant to raise the curiosity of the student towards the larger issues of the project and the student has developed a solid understanding of the major theoretical issues of the project. He has graduated in 2005.

A third Master's student, James Oravec, has been working on geometric models and linear programming techniques for generating knowledge state libraries. He has used Mathematica extensively.

Additionally two REU students, Hung Tran and Jason Rider, and an independent study undergraduate student, Randy Agartarp, have worked

on a web-based tutorial. This effort has focuses on the animation of online algorithms on the web using Flash, Java, and PHP.

All students also were taught by the PI how to use open source software in a Linux environment. Operating systems has included use of distributions such as Debian and Ubuntu, Red Hat and Fedora, as well as Suse.

The principal investigators' graduate level courses were close to the project. Larmore's course page at
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F03/
and
http://www.egr.unlv.edu/~larmore/Courses/CSC789/F04/
have extensive write-ups on the theoretical underpinnings of the knowledge state approach written at a level accessible to graduate students. Some of the materials provided the basis for the paper with Reischuk on knowledge states.

Contributions to Resources for Research and Education:

In December 2004 the Regents of the University of Nevada approved the Center for the Advanced Study of Algorithms (CASA). CASA facilitates national and international collaborations, interdisciplinary work across UNLV's campus and with other academic entities.

The principal investigator is the Director of the Center and the Co-PI is the Co-Director of the Center. This project was instrumental in getting the Center started.

The investigators have improved instrumentation the at the School of Computer Science within the space limitations of the institution.

The principal investigator has given numerous presentations through the Howard Hughes College of Engineering Minority Engineering Program in which he has described various aspects of the project.

The principal investigator has promoted open source software at the University of Nevada, Las Vegas.

Contributions Beyond Science and Engineering:

The project has strengthened the computer science community in Southern Nevada. Ultimately this is beneficial to the goal of diversifying Southern Nevada economy.

Categories for which nothing is reported:

In Information Technology decisions must typically be made before all inputs are available. Whether it is setting up virtual circuits in order to carry IP traffic over ATM networks, deciding whether to leave a disk spinning in between accesses to data, or keeping cache coherent in a multiprocessor architecture – online algorithms play a crucial role in such diverse areas as machine learning, robotics, operating systems, network routing, distributed systems, databases. Because for many applications (such as data streaming) it is desirable to use little memory, the tradeoff between the quality (competitiveness) of an online algorithm and its memory use is important. Meanwhile, in the context of online algorithms, randomization – a powerful tool in algorithmic design – has not been fully understood.

To address these issues, the investigators have introduced the novel concept of knowledge states, which can be used to construct competitive randomized online algorithms and study the tradeoff between competitiveness and memory. A knowledge state simply states conditional obligations of an adversary, by fixing a work function, and gives a distribution for the algorithm. When a knowledge state algorithm receives a request, it then calculates one or more "subsequent" knowledge states, together with a probability of transition to each. The algorithm then uses randomization to select one of those subsequents to be the new knowledge state. Although before this project the formal definition of knowledge state algorithms did not appear in any publication, many well-known algorithms could be viewed as knowledge state algorithms.

- The investigators have given a formal description of the knowledge state method. They have related the mixed model of online computation to the standard models of online computation, and have described how a behavioral algorithm can be derived from a mixed model description. They also give proofs to show how potentials can be used to derive the competitive ratio of a knowledge state algorithm. (These results in collaboration with Reischuk, Universitaet Luebeck, Germany.)

- Knowledge state algorithms are not always based on hand-designed knowledge states. Instead, investigators sometimes relied on experimental methods to design knowledge state algorithms. The investigators have developed optimization techniques to generate knowledge states automatically. The techniques use convex hull algorithms as well as techniques based on linear programming. (These tools were developed with the help of a number of graduate students.)

- Borodin and El Yaniv list as an open question whether there exists an H_k-competitive randomized algorithm which requires $O(k)$ memory for k-paging. Using the knowledge state approach, the investigators answer this question in the affirmative. (These results in collaboration with Noga, California State University, Northridge.)

- The investigators have given optimally competitive algorithms with minimal memory requirements for the 2-paging and 3 paging problems. For $k = 2$ the number of active pages is never more than three and for $k = 3$ is never more than five.

- The investigators have considered a multiprocessor systems paging problem, where there is a global shared memory to provide logically uniform data access and where to hide latencies when accessing global memory each processor makes use of a private cache. To exclude the effect of classical paging faults and to study the effects of invalidations by updates, the model assumes that each processor knows its own data access sequence, but does not know the sequence of future invalidations requested by other processors. A $\frac{4}{3}$-competitive randomized online algorithm for this problem for cache size 2 was obtained. The investigators have also proved a matching lower bound, thus this online algorithm is best possible. A lower bound of $\frac{3}{2}$ on the competitiveness for larger cache sizes has been shown. (These results in collaboration with Reischuk, Universitaet Luebeck, Germany.)

- It has been a long-standing open problem to determine the exact randomized competitiveness of the 2-server problem, that is, the minimum competitiveness of any randomized online algorithm for the 2-server problem. The investigators have obtained a $\frac{19}{12}$-competitive knowledge state algorithm for the 2-server problem over Cross Polytope Spaces $M_{2,4}$. They have also shown that this ratio is best possible. The space $M_{2,4}$ generalizes uniform spaces, and thus paging. This gives the second non-trivial example of metric spaces with better than 2 competitive ratio. (These results were obtained while the Principal Investigator was a Visiting Professor at Kyoto University. The results are in collaboration with Iwama, Kyoto University, Japan.)

- Although largely unnoticed by the online algorithms community, T-theory, a field of discrete mathematics, has contributed to the devel-

opment of several online algorithms for the k-server problem. The investigators have restated a number of known k-server results using the established terminology of T-theory. A previously unpublished 3-competitiveness proof, using T-theory, for the HARMONIC algorithm for two servers has been presented.

- Regarding randomized algorithms for metrical task systems, the investigators have shown for $k \geq 2$ that there is a $\frac{3}{2}H_k - \frac{1}{2k}$ that competitive randomized stable distribution algorithm for the uniform metrical task system with k states. By refining this technique through linear programming, for $k = 3$, calculations indicate that the competitive ratio C for uniform metrical task systems is bounded by $2.527 \leq C \leq 2.545$. The investigators conjecture that, for fixed k, these refinements will yield better and better bounds which converge to the true competitiveness. (These results in collaboration with Noga, California State University, Northridge.)

- New results regarding the delayed k-server problem were obtained in collaboration with Kazuo Iwama of Kyoto University, Japan. The study of the delayed server problem is important in Information Technology to account for latencies in caching and related problems. Specifically, it was shown that the delayed k-server problem is equivalent to the $k-1$ (regular) server problem in a uniform metric space (i.e. paging.) For general spaces the regular and delayed problem are quite different, and a k-competitive algorithm for trees was given. (The results are in collaboration with Iwama, Kyoto University, Japan.)

In this project, the investigators have also focused on scheduling and arrangement problems.

- In batching the order of jobs to be considered is fixed and one seeks to minimize average flow time of the jobs. The investigators have presented optimally competitive algorithms for s-batch (competitive ratio $C = 2$) and p-batch problems (competitive ratio of $C = 4$). They have derived an optimally competitive $\frac{619}{583}$-competitive for the unit processing case, which is related to the TCP acknowledgment problem. With TCP there exists a possibility of using a single acknowledgment packet to simultaneously acknowledge multiple arriving packets, thereby reducing the overhead of the acknowledgments. Dooly, Goldman, and

Scott introduced the dynamic TCP acknowledgment problem in which the goal is to minimize the number of acknowledgments sent plus the sum of the delays of all data packets which are the time gaps between the packets arrival time and the time at which the acknowledgment is sent. (These results in collaboration with Epstein, University of Haifa, Israel and Noga, California State University, Northridge.)

- The block sorting problem is the problem of minimizing the number of steps to sort a list of distinct items, where a sublist of items which are already in sorted order (called a block) can be moved together in one step. Blocks, once formed, are never broken up, but are only combined to form larger blocks; in the final step, there is just one block. Block sorting is used in connection with Optical Character Recognition (OCR) but might also have implications for computational biology. The investigators have given an approximation algorithm for the block sorting problem with an approximation ratio of 2 and run time $O(n^2)$. (It is interesting to note that Mahajan, Rama and Vijayakumar have studied the problem at the Max-Planck Institute, Germany, and have come up independently with a more complicated and less efficient 2-competitive algorithm.) (These results in collaboration with Hal Sudborough , University of Texas at Dallas and Linda Morales, Texas A&M Commerce.)

The investigators have worked on speeding up naive implementations of dynamic programming. Even though the problems are offline problems this work is related to online algorithms in the sense that the dynamic program follows an online protocol.

- The investigators have shown that the Knuth-Yao technique is a direct consequence of total monotonicity. As well as providing new derivations of the Knuth-Yao result, this also permits showing how to solve the Knuth-Yao problem directly using the SMAWK algorithm. Another consequence of this approach is a method for solving online versions of problems with the Knuth-Yao property. The online algorithms given are asymptotically as fast as the best previously known static ones. For example the Knuth-Yao technique speeds up the standard dynamic program for finding the optimal binary search tree of n elements from $\Theta(n^3)$ down to $O(n^2)$, and the results allow construction of an optimal binary search tree in an online fashion (adding a node to the left or

right of the current nodes at each step) in $O(n)$ time per step. (These results in collaboration with Mordecai Golin, Hong Kong University of Science and Technology, Hong Kong.)

- Regarding path problems, the investigators have obtained new results for algebraic Monge properties. They have shown that the online LARSCH algorithm for finding row minima applies in this case and they show that their bottleneck shortest-path techniques can be used to obtain fast algorithms for a variant of Hirschberg and Larmore's optimal paragraph formation problem, and a special case of the bottleneck traveling-salesman problem. (These results in collaboration with Peter Brucker, University of Osnabrueck, Germany.)

As Information Technology has become more pervasive, the focus on online computing has increased. This project has had a broad focus on online algorithms, especially the effective use of randomization. Research activities have lead to a number of breakthrough results in the field of theoretical computer science – see Section "Findings" of the report – and the activities resulting from the project have strengthened the theoretical computer science community in Southern Nevada.

Partially as a result of the activities of the project, the principal investigators opened the the Center for the Advanced Study of Algorithms (CASA) at the University of Nevada, Las Vegas, during the second year of the grant period. In fact, the center has developed into a hub for the research activities of this project. The principal investigators have deepened ties with research groups headed by Kazuo Iwama of Kyoto University, Japan, Ulrich Faigle, Zentrum für Angewandte Informatik Universität zu Köln, Hal Sudborough at the University of Texas at Dallas, Leah Epstein, Haifa University, Stefan Pickl at the Universität der Bundeswehr München, Germany, John Noga, California State University, Northridge, Mordecai Golin of Hong Kong University of Science & Technology and Camil Demetrescu, University of Rome "La Sapienza", Italy.

- The principal investigator has visited Kyoto University, and specifically the research group of Kazuo Iwama, twice. Iwama is the head of an internationally recognized research group at Kyoto University and a preeminent algorithms researcher. Iwama had invited the pricipal investigator for an extended research stay during December 04 - January 05 and the again for the 06 Fall semester as Kyoto University Visiting Professor. The funding for these research stays were entirely by Kyoto University and added nicely to this NSF project. The principal investigator gave a series of lectures there. The research conducted at Kyoto University has lead to fundamental results on the (deterministic) delayed server problem and the randomized 2-sever problem. The study of the delayed server problem is important in Information Technology to account for latencies in caching. Regarding the randomized 2-server problem, breaking through the "2-competitive barrier" had been ranked by the participants of the 2002 Dagstuhl Workshop to be one of the three most important outstanding online problems.

- The principal investigators have collaborated extensively with Rüdiger

Reischuk, the head of the Institute für Theoretische Informatik Universität zu Lübeck, Germany. Both principal investigators visited him in Lübeck over the 03/04 winter break, and the principle investigator continued visits to Lübeck in 05, 06 and 07. Reischuk visited UNLV for a week in Spring 05. This collaboration lead to the development of the knowledge state method, which is fundamental to this project (cf. Section "Findings"), as well as applications to caching problem in shared memory multiprocessor systems.

- The principal investigators have collaborated extensively with John Noga of California State University, Northridge. A number of visits between California and Nevada took place during the entire project period. New results on memory requirements for paging were obtained. Other work with Noga lead to new randomized algorithms for metrical task systems. Metrical task systems are important as they generalize a host of online problems, including many applications. During these visits the principal investigator also met with Jeff Wiegley, an operating systems expert. Wiegley gave valuable advice on aspects of instrumentation for the project. In related activities, the principal investigator has performed system administration tasks for the project and gained proficiency in numerous aspects of open source software. As part of this interest the pricipal investigator attended Linux Tag in Karlsruhe, Germany in June 2004.

- The principal investigator was invited to visit Hong Kong University of Science & Technology, December 1 - 15, 2006, where he collaborated with Mordecai Golin. Following this, Golin came to UNLV for a week-log visit at the beginning of 2007. Together with the Co-pricipal investigator, these collaborations have lead to significant results in the area of dynamic programming speedup. Even though the problems under investigation are offline this line of research is related to online algorithms in the sense that dynamic programming follows an online protocol.

- The principal investigators have collaborated extensively with Linda Morales (Texas A&M, Commerce) and Hal Sudborough (University of Texas at Dallas). A number of visits between Texas and Nevada took place during the entire project period. During Spring 2007, the principal investigator visited the University of Texas at Dallas for a longer

research stay. The collaboration lead to important results for block sorting. Block sorting is used in connection with Optical Character Recognition (OCR) but may also have implications for computational biology. It is interesting to note that Mahajan, Rama and Vijayakumar have studied the problem independently at the Max-Planck Institute, Germany. The principal investigators together with Morales and Sudborough obtained results superior to those of Mahajan et al.

- The principal investigators have expanded the research partnership with the Zentrum für Angewandte Informatik Universität zu Köln headed by Ulrich Faigle. Both investigators visited Cologne in Summer 2004 and both investigators gave colloquia there. Bernhard Fuchs, a doctoral student there, joined the project at UNLV for two month on the suggestion of Faigle. The principal investigators have worked with Fuchs on online matching problems, and have improved the results in his thesis. The principal investigator has also followed invitations to visit Stefan Pickl, who has moved from Cologne to the Universität der Bundeswehr München, Germany. The principal investigator received a "Lehrauftrag" there during Spring 2007 and has given a series of collaquia in Munich. In turn, Pickl has also visited UNLV.

- The principal investigators have collaborated with Leah Epstein of Haifa University on online batching problems. List Batching is related to TCP acknowledgment. Epstein together with Noga visited UNLV for a week in Spring 2004 and the principal investigators finalized work on their journal paper with Epstein in Eilat, Israel in 2007 during the WAOA work shop.

- The principal investigators are grateful to Peter Brucker of the University of Osnabrück, Germany, who came to UNLV for a research visit. Peter Brucker is a leading expert in Scheduling. His visit inspired the principal investigators to conduct the work on online batching problems. Brucker's visit also presented an opportunity to work on path problems under algebraic Monge properties. The principal investigator gave a colloquium at Universität Osnabrück in Spring 07.

Research Methods

Though a significant portion of the research was by pencil and paper, some of the results used computer calculation. New software for construct-

ing convex hulls as well as specialized linear programming code was written for the project. Results on the randomized 2-server problem and on batching have depended on such software. Specifically for 2-server problem such implementations are necessary, as the approach relies on computational methods to design knowledge state algorithms by computer. The pricipal investigator has continued his interest in open source issues. He has provided stable open source (Debian based) platforms for the participants and has performed various system administration tasks, which benefited the program.

The project involved significant collaboration with students. The pricipal investigator has graduated one Master's student under this project. The student, Ms. Qin Zhang, reviewed various aspects of the knowledge state approach in her thesis. Ms. Zhang was also involved in work on new limited bookmark algorithms for paging. The Co-pricipal investigator has graduated two Master's students, J. Kurlinski, who has worked on geometric aspects of the Knowledge State Method, and J. Oravec, who has done extensive implementation work regarding a knowledge state library for the 2-server problem. These models were mainly implemented in Mathematica.

The pricipal investigator has supervised two REU students, H. Tran and J. Rider, who have worked on a web tutorial and some aspects of computer implementations. Additionally, the pricipal investigator held a number of independent study courses for undergraduate students which taught aspects of the research projects.

Since the research depended on sessions with a number of recognized computer scientists, significant research work took place during various conferences and workshops. The investigators were also invited to a number of prestigious workshops to give invited talks. (Additionally the principal investigator attended a number of STOC and FOCS conferences as SIGACT publicity chair.)

Special Invitations

- The principal investigator was invited to give a semi-plenary talk "Knowledge States: A Tool in Randomized Online Algorithms" ADS 2007 - 3rd Bertinoro Workshop on Algorithms and Data Structures September 30 - October 6, 2007 in Bertinoro, Italy to give a survey talk. The Bertinoro conference center is similar to the Dagstuhl conference center in Germany. The workshop was co-sponsored by BICI, the Bertinoro

Center for Informatics, and DIMACS, the Center for Discrete Mathematics and Theoretical Computer Science, and was organized by Camil Demetrescu and Robert Tarjan. The pricipal investigator received a grant from BICI/DIMACS for room and board.

- The pricipal investigator was invited by Kirk Pruhs to give a semiplenary talk in the online optimization cluster of the 2006 International Symposium on Mathematical Programming (ISMP), Rio de Janeiro, July 30 - August 4 2007. The cluster also featured a number of recognized colleagues such as Albers, Iwama, Erlebach, and others.

- In Fall 2006, the pricipal investigator was part of a three talk mini conference at Kyoto University and gave a talk " Online and Offline List Batching" The other two speakers were Oscar Ibarra and Richard Ladner.

- The pricipal investigator and Co-pricipal investigator were invited to give talks and engaged in research collaborations at the Workshop on On-Line Algorithms 2004 (OLA 2004) at Rungstedgaard in Rungsted Kyst, Denmark, July 5-7, 2004. The workshop was held in the spirit of the series of three workshops on on-line algorithms, which have been held at Dagstuhl in 1996, 1999, and 2002.

Refereed Conferences

- The pricipal investigator gave the co-authored paper "Equitable Revisited" at the 15th Annual European Symposium on Algorithms Eilat, Israel, October 8-10, 2007.

- The pricipal investigator gave the co-authored paper "A Randomized Algorithm for Two Servers in Cross Polytope Spaces" at the 5th Workshop on Approximation and Online Algorithms (WAOA) Eilat, Israel, October 11-12, 2007.

- Golin gave the co-authored paper "The Knuth-Yao Quadrangle-Inequality Speedup is a Consequence of Total-Monotonicity" at the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006.

- The pricipal investigator gave the co-authored paper "The Delayed k-Server Problem" at the 15th International Symposium on Fundamentals of Computation Theory, Lübeck, Germany August 17 - 20, 2005. The Co-pricipal investigator gave the co-authored paper "A Faster and Simpler 2-Approximation Algorithm for Block Sorting", also FCT. Leading up to FCT, the pricipal investigator and Co-pricipal investigator met Sudborough and Morales as well as Iwama during their visits to Europe. On this occasions they continued their work on block sorting with Sudborough and Morales and the delayed server problem with Iwama. Immediately following the FCT conference the principal investigators meet with Reischuk in Lübeck to work on knowledge state algorithms for the CNN problem and the 2-server problem.

- Noga gave the co-authored paper "Optimally Competitive List Batching" at the 9th Scandinavian Workshop on Algorithm Theory, July 2004. Both pricipal investigator and Co-pricipal investigator attended the workshop and they worked at the workshop on the journal version of the paper.

- The pricipal investigator held the 8th International Symposium on Parallel Architectures, Algorithms and Networks - ISPAN 2005 in Las Vegas, NV, December 7-9, 2005, under the auspices of the Center for the Advanced Study of Algorithms. This conference has been held before in Kanazawa, Taipei, Perth, Dallas, Manila, and Hong Kong. The Co-pricipal investigator gave the co-authored paper "A Better Algorithm for Uniform Metrical Task System with Few States."

- The pricipal investigator gave the paper "Knowledge States for the Caching Problem in Shared Memory Multiprocessor Systems" at the 7th International Symposium on Parallel Architectures, Algorithms and Networks - ISPAN 2004, in Hong Kong, May 10- 12, 2004.

- The pricipal investigator gave a talk "Knowledge State Algorithms and the 2-Server Problem" at the Cologne Twente Workshop on Graphs and Combinatorial Optimization (CTW 2004), Como, Italy, May 31 - June 2, 2004.

Colloquia

The pricipal investigator has given numerous colloquia regarding the research of the project:

- Monge Properties. Academia Sinica, Institute of Mathematics, Taipei, Taiwan, May 14, 2004. also at Institut fr Angewandte Systemwissenschaften und Wirtschaftsinformatik, Universität der Bundeswehr München, Germany, June 2005.

- The k-Server Problem -The State of the Art. School of Informatics Kyoto University, December 15, 2004, also at Universität der Bundeswehr München, Germany, June 2005. New Mexico State University, 2004, Mathematisches Institut Zentrum für Angewandte Informatik, Universität Köln, 2003; Universite de Picardie, Jules Verne, Amiens, 2003. Knowledge State Algorithms and Paging, School of Informatics Kyoto University, December 22, 2004. (At Kyoto University the pricipal investigator gave a second talk on randomized paging.)