

# Rapid Prototyping of Parallel Robot Vision Systems Using Virtual Reality and Systems Simulation

Progress Report, NSF CDA-9401142

Thomas LeBlanc, Dana Ballard, Christopher Brown, Randal Nelson, and Michael Scott  
Computer Science Department, University of Rochester

July 1999

## 1 Introduction

### 1.1 Institution

The University of Rochester is a small, private University established in 1850, originally as a Baptist-sponsored institution. During the early 20th century, the University grew significantly, in part due to the efforts of George Eastman, the founder of Eastman Kodak. During this period, the Medical School, the Institute of Optics, and the Eastman School of Music, all currently nationally known, were established. Today, the University is home to 4600 undergraduates and 2500 graduate students, and operates with a philosophy of providing the academic opportunities of a renowned research institution in an environment scaled to the individual.

The Department of Computer Science at the University of Rochester offers an intense, research-oriented program leading to the degree of Doctor of Philosophy, with particular emphasis on the areas of computer vision and robotics, knowledge representation and natural language understanding, systems software for parallel computing, and the theory of computation. The focused research interests reflect our desire to achieve excellence in a core of important issues, rather than to try to cover all areas.

### 1.2 Project

The main focus of our research is a laboratory that combines sensory interaction with simulated physical environments (otherwise known as virtual reality), physical and sensory interaction with mixed physical and simulated environments (otherwise known as augmented reality), and the reconstruction of physical environments from sensory inputs (otherwise known as computer vision), with execution-driven simulation of complex parallel systems in the design and control of visually-controlled robotic systems. We also develop systems tools to manage large parallel computing applications and development environments. This RI-supported research leverages Rochester's unique combination of expertise in active vision systems, behavioral robotics, virtual reality, and parallel programming environments and systems. (see the URL <http://www.cs.rochester.edu:80/research/iip/>).

Our laboratory has two parts: one is for building working systems in the *real world*; the other part is for prototyping and experimentation in the *virtual world*. The components of the real-world laboratory are the effectors and sensors for interacting in the real world, and the computational machinery required to run the control algorithms. The virtual-world laboratory includes

the hardware and software for creating the virtual world, (models of sensors, effectors, and their interaction with an environment). System support includes the large parallel computers used for computation-intensive applications and simulations, and research software systems for running them.

## **2 Year's Activities**

### **2.1 Goals, Objectives and Targeted Activities**

The infrastructure constructed over 5 years with this award has been a significant factor in producing an environment that has allowed us to attract top notch faculty and graduate students, and carry out state-of-the-art research. 24 PhDs, and more than 350 publications, over half of them in refereed forums have been produced with direct support of the infrastructure grant. Many of the original visions were achieved; and as is the case in creative endeavors, some lines of research took unanticipated directions. For example, the virtual reality component took off dramatically in ways not originally envisioned, especially regarding study of human performance in virtual environments, and in development of augmented reality technology.

Goals for final year involved continuing use of the infrastructure to address basic questions in vision and robotics and related issues in human perception. Several vision and robotics algorithms developed with simulated imagery were scheduled for transit to real-world environments. We also wanted to continue development of system methodologies to support work of this nature, which involves both massive amounts of computation, and integration of heterogeneous systems.

#### **Specific goals on the real-world side**

1. Use the computational resources provided by the infrastructure to investigate and document the performance of our 3D object recognition system across a number of variables.
2. Investigate the performance of mobile robot exploration and mapping algorithms, developed in simulation, on the physical mobile robots.
3. Apply novel deformable contour tracking algorithms developed in simulated imagery to real image sequences of deforming objects.

#### **Specific goals in the virtual world**

1. Integrate inertial feedback into the driving simulator using a hydraulic sled (acquired with NIH funds, whose leverage has allowed us to improve the capabilities of the virtual reality part of the lab).
2. Study human performance of assembly and manipulation tasks in an integrated visual and haptic virtual world with the goal of discovering strategies relevant to machine implementation of such abilities.
3. Continue work on uncalibrated methods of fusing artificial graphics with real video (augmented reality).
4. Continue work on generating and manipulating photo-realistic, 3D object models from imagery.

#### **Specific Systems support goals:**

1. Investigate fundamental issues in parallel processing including partitioning, software coherence, and interactive client-server computing models.
2. Continue development of an integrated compile-time and run-time system for efficient shared memory parallel computing on distributed memory machines (Cashmere and Treadmarks).
3. Continue studies of particular parallel applications, especially data mining.

## **2.2 Components and Materials Required**

Our budget for the final year was relatively small, with most of the primary infrastructure having been obtained in the first four years of the grant. Two major purchases, made with previous-year funds overlapped the beginning of the current fiscal year.

The first was an upgrade to the vision/robotics controller, to a Sun Multiprocessor with 8 336MHz processors and 3 Gigabytes of main memory. This machine was essential to obtaining the 3D recognition results. The size of main memory permitted experiments to be performed in parallel and in core, representing a speedup of approximately 80X over use of standard desktop, which would require use of virtual memory to handle large databases. Since large-scale performance tests took several days to run, this was critical to obtaining the results that we did.

The second was the upgrade for the Alpha Cluster, which now consists of 32 600MHz processors with 16 GigaBytes of main memory connected by the Memory Channel interface. The new Alpha Cluster has about 4 times the aggregate compute power, 4 times the memory, half the communication latency, and 10 times the aggregate communication bandwidth of the current cluster. Systems simulations on this machine allows us to address the applicability of basic computing assumptions, as different architectural parameters change. This machine was also used in running some of the object recognition experiments.

## **2.3 Indications of Success**

### **Real World Vision/Robot Applications**

- Quantified 3D object recognition system performance on various dimensions (via ROC curves) and identified critical axes for improvement. (e.g., intermediate segmentation). Also developed new low-level contour finder with improved robustness to severe background clutter.
- Developed mechanisms for integrating data acquired from a robot-mounted, real-time stereo range finder into a 3D world model that could be used for navigation and planning. This represents transfer of algorithms developed in simulation into a real-world scenario.
- Developed theory for spatio-temporal tracking of deforming contours. Developed algorithms to implement the theory, and tested them on VR image sequences for which ground truth was known, then transferred algorithms to use real-world image sequences.

### **Virtual World**

- Human studies during driving established utility of virtual routines model, including experiments that showed definitively that information invoking specific routines (such as stopping behavior) is acquired at specific moments in time. Added inertial feedback to driving simulator, and performed experiments as to its effect on human control (it helps).

- Continued experiments in using haptic virtual reality, e.g. we use a grasp and repositioning task to study the relation of weight and density perception.
- The augmented reality project created augmented reality goggles using a Virtual Reality helmet fitted with two cameras. By adding the haptic interface, we were able to touch and manipulate the virtual objects inserted into a real scene.

### **Systems Support**

- Implemented support in Cashmere for very large (out-of-core) data sets, with customized paging policies. Also Completed quantitative comparison of fine-grain (Shasta) and coarse-grant (Cashmere) software coherence on a common platform.
- Continued integration of compiler and run-time support for coherence management. Completed development of the Coign automatic distributed partitioning system. Developed the COP interface proxy system for OS versioning and "remoting".
- Continued algorithmic and behavior studies of parallel applications, with an emphasis on data mining including InterAct, a shared-memory client-server system for remote interactive data-mining.

Altogether there were 35–40 refereed publications and about 20 other publications directly related to the infrastructure. 24 graduate students, three of them women, and seven undergraduates made direct use of the infrastructure. Three PhDs, 6 Masters, and 2 Bachelors were awarded among these students.

## **3 Evaluation**

NSF funding has allowed us to build a laboratory that has unique capabilities to support research into perceptual intelligence through robot vision, psychophysics and cognition, and photo-realistic and touch-realistic simulations for advanced computer control. Almost all our infrastructure funding has been spent on highly-leveraged, unusual hardware that helps us attract and retain the best students and faculty as well as carry out our work. With the infrastructure resources, we have been able to develop a number of successful applications in which the real and simulated worlds were effectively married, realizing improved testing, performance, debugging, and human experimental interfaces.

### **3.1 Outcome**

The grant has helped us retain our position as a leader in active paradigm, biologically-motivated vision and robotics research, and in shared-memory parallel computing. It has helped us recruit and train the very best graduate students (consistently the best in the College on standardized tests), and place them at the best institutions upon their graduation.

At the national and global level, work supported by the grant has its impact through the influence of our papers on other researchers, and through the work of our alumni. Our work is heavily cited. Some of our laboratory techniques have been widely copied (e.g. verging binocular cameras) and others are still on the leading edge (eye trackers in VR helmets, oversize haptic hardware, uncalibrated augmented reality, real-time vision algorithms implemented on special hardware).

### 3.2 Immediate Impact

Undergraduate students supported (and bachelor's degree if awarded in 1998): Aaron Gerega (BS), Peenak Inamdar, Craig Harman, Fred Marcus Henry McCauley, Yasser Mufti (BS), Sean Rodgers.

Graduate students supported (and MS or PhD if awarded): Rajeev Balasubramonian, Jessica Bayliss, Rahul Bhotika (MS), Rodrigo Carceroni, DeQuing Chen, Melissa Dominguez (MS), Christopher Eveland, Isaac Green (MS), Christopher Homan, Bo Hu, Galen Hunt (PhD), Grigorios Magklis, Srinivasan Parthasarathy, Umit Rencuzogullari (MS), Garbis Salgian (PhD), Brandon Sanders, Andrea Selinger, Amit Singhal, Nathan Sprague, Robert Stets, Yiyang Tao (MS), Mike VanWie, Mohammed Zaki (PhD), Zuohua Zhang (MS).

## 4 Representative References

Ballard, D.H., G. Salgian, R.P.N. Rao, and A. McCallum, "On the role of time in brain computation," chapter 5 in L.R. Harris and M. Jenkin (Eds.). *Vision and Action*. Cambridge U. Press, 82-119, 1998.

Carceroni, R.L., C. Harman, C.K. Eveland and C.M. Brown, "Real-time pose estimation and control for convoying applications," in D. Kriegman, G. Hager, and S. Morse (Eds.). *The Confluence of Vision and Control (Lecture Notes in Computer Science Series #237)*. Springer-Verlag, 230-243, 1998.

Carceroni, R.L. and K.N. Kutulakos, "Toward recovering shape and motion of 3d curves from multi-view image sequences," Proc., IEEE Conf. on Computer Vision and Pattern Recognition, Vol. 1, pp. 192-197, Ft. Collins, CO, June 1999.

Dwarkadas, S., K. Gharachorloo, L.I. Kontothanassis, D.J. Scales, M.L. Scott, and R.J. Stets, "Comparative evaluation of fine- and coarse-grain software distributed shared memory," Proc., 5th Int'l. Conf. on High-Performance Computer Architecture (HPCA-5), Orlando, FL, January 1999.

Dwarkadas, S., H. Lu, A.L. Cox, R. Rajamony, and W. Zwaenepoel. "Combining compile-time and run-time support for efficient software distributed shared memory," Proc., IEEE 87, 3, 476-486, March 1999.

Fuentes, O. and R.C. Nelson, "Learning dextrous manipulation strategies for multifingered robot hands using the evolution strategy," *Special Combined Issue of Machine Learning* 31, 223-237 and *Autonomous Robots* 5, 395-405, 1998.

Hunt, G.C. and M.L. Scott, "The Coign automatic distributed partitioning system," MSR-TR-98-40, Microsoft Research Laboratory, August 1998; Proc., 3rd Symp. on Operating System Design and Implementation (OSDI '99), USENIX, 187-200, New Orleans, LA, February 1999.

Hunt, G.C. and M.L. Scott, "Intercepting and instrumenting COM applications," Proc., 5th Conf. on Object-Oriented Technologies and Systems (COOTS'99), USENIX, 45-56, San Diego, CA, May 1999.

Ioannidis, S., U. Rencuzogullari, R. Stets, and S. Dwarkadas, "CRAUL: Compiler and Run-Time Integration for Adaptation Under Load", the *Journal of Scientific Programming*, August 1999.

Kutulakos, K.N. and J.R. Vallino, "Calibration-free augmented reality," *IEEE Trans. Visualization and Computer Graphics* 4, 1, 1-20, 1998.

Nelson, R.C. and A. Selinger, "Large-scale tests of a keyed, appearance-based 3d object recognition system," *Vision Research* 38, 15-16 (Special Issue on Computational Vision), 2469-88, August

1998.

Parthasarathy, S. and S. Dwarkadas, "InterAct: Virtual Sharing for Interactive Client-Server Applications", Fourth Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers, Springer-Verlag, May 1998.

Parthasarathy, S., M. J. Zaki, and W. Li, "Memory Placement Techniques for parallel association mining" In the 4th International Conference on Knowledge Discovery and Data Mining (KDD), August 1998.

Rao, R.P.N. and D.H. Ballard, "Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects," *Nature Neuroscience* 2, 1, 79-87, January 1999.

Vallino, J.R. and C.M. Brown, "Haptics in augmented reality," Proc., IEEE Int'l. Conf. on Multimedia Computing and Systems, IEEE Computer Society, 195-200, Florence, Italy, June 1999.