

# Guest Editors' Introduction

# Parallel Rendering

## Chuck Hansen

Advanced Computing  
Laboratory  
MS B-287  
Los Alamos National Laboratory  
Los Alamos, NM 87545  
Internet hansen@acl.lanl.gov

## Tom Crockett

Institute for Computer  
Applications in Science and  
Engineering  
Mail Stop 132C  
NASA Langley Research Center  
Hampton, VA 23061-0001  
Internet tom@icase.edu

## Scott Whitman

David Sarnoff Research Center  
CN 5300  
Princeton, NJ 08543  
Internet slim@sarnoff.com

Massively parallel computers have emerged as valuable tools for performing scientific and engineering computations, far outstripping the capabilities of independent workstations in both sheer floating-point performance and memory capacity. As the resolution of simulation models increases, graphics algorithms that take advantage of the large memory and parallelism of these architectures are becoming increasingly important. This issue of *IEEE Parallel & Distributed Technology* highlights some recent work in parallel computer graphics, specifically parallel rendering.

It is well known that some computer graphics problems are embarrassingly parallel; ray-tracing algorithms with replicated databases and fractal geometry are classic examples. By replicating data on each processor, these problems require little if any communication, thus removing a major impediment to linear speedup.

Generally, however, parallel rendering algorithms operate on data sets that are distributed among the processors. This requires massive communication due to the mapping from 3D "object space" to 2D "image space." Hence the algorithm designers must carefully balance computation and communication. Add to this the nuances of parallel architectures and a large algorithmic design space, and the task of achieving efficient, scalable performance becomes very demanding. The problems in parallel rendering algorithms are similar to those in all problems on parallel computers: How do we partition a large data set among the memories of the computer to minimize communication and network contention, maximize processor use, and achieve good speedup relative to a serial approach?

To help address these issues, we organized the 1993 Parallel Rendering Symposium, which brought together researchers from around the world to survey the state of the art in this growing field. The three parallel rendering articles in this issue are updated and revised from papers that appeared at the symposium. They emphasize a systems view of parallel rendering, providing insight into the parallel aspects of computer graphics research without delving into the more esoteric graphics issues.

Thomas Crockett and Tobias Orloff look at polygon rendering on a MIMD distributed memory system, the Intel iPSC/860. They present an asynchronous algorithm and examine its behavior both analytically and experimentally, with an emphasis on the trade-offs between message size and performance.

Michael Cox and Pat Hanrahan look at the problem of polygon rendering by pixel merging. They analyze different pixel merging algorithms based on network bandwidth and develop an algorithm for

shared-memory bus architectures. Their work has many similarities with cache snooping algorithms.

Paul Mackerras and Brian Corrie describe a method for exploiting data coherence in image-space volume rendering. They use a caching scheme to implement a distributed virtual memory system that effectively minimizes communication by taking advantage of data locality during rendering.

Although its roots go back many years, the field of parallel rendering is by no means mature. As you will see in this issue, efficient parallel rendering is an elusive goal, with more questions than answers. As parallel computers become more powerful and affordable, the integration of graphics with parallel applications will be a central issue for the visualization community. Efficient parallel rendering algorithms are a prerequisite to this process.

## ACKNOWLEDGMENTS

The 1993 Parallel Rendering Symposium was held in conjunction with the Visualization '93 conference, and was sponsored by the IEEE Computer Society Technical Committee on Computer Graphics in cooperation with ACM SIG Graph. The July 1994 issue of *IEEE Computer Graphics and Applications* features five articles — four from the symposium — on the graphics aspects of parallel rendering.



**Chuck Hansen** is a technical staff member at the Advanced Computing Laboratory at Los Alamos National Laboratory, and an adjunct faculty member at the University of New Mexico and New Mexico Institute of Technology. He received his BS in computer science from Memphis State University in 1981 and a PhD in computer science from the University of Utah in 1987.



**Thomas W. Crockett** is a staff scientist at the Institute for Computer Applications in Science and Engineering. He received his BS in mathematics from the College of William and Mary in 1977.



**Scott Whitman** is a member of the technical staff at David Sarnoff Research Center. He received his PhD and MS in computer science from Ohio State University and his BS in applied mathematics from Carnegie Mellon University.