Greetings

Michael D. O'Dell, Editor-in-Chief

This issue presents a collection of work that covers both a broad sweep of interests and a wide dynamic range—from the long-haul backbone to the processor-memory bus.

Our first paper, "Grasshopper," by Al Dearle and a team of Australian co-workers, discusses operating system machinery for supporting objects as first-class citizens. The long-time operating systems afficionado may notice some interesting new twists on some venerable ideas.

Our second paper, "Delayline," by David B. Ingham and Graham D. Parrington from the University of Newcastle upon Tyne, addresses an interesting problem in the testing and development of large distributed applications. How can one test such a system for robustness in the face of transcontinental transaction delays when the development environment is only hosted on a fast LAN? The authors describe a tool developed for just this purpose. (The image of a coal train trudging across the countryside is just too hard to ignore.)

In their new work on "Swift/RAID," Long, Montague and Cabrera pursue the new holy grail in computing performance tuning, very fast filesystem I/O. They describe their experience with an approach for exploiting the parallelism inherent in fast work stations on fast local networks to provide the horsepower needed for high-speed I/O.

One of the recurring problems in the work station-cluster style of computing is scheduling large and long-running jobs for execution on a collection of servers. In "A Distributed Look-Ahead Workload Assignment Algorithm for Interdependent Tasks," Andreas Winckler describes a system that uses a distributed scheduler for managing these tasks. The challenge is to make such a system responsive without using too many resources for just managing the scheduling.

Microprocessors have gotten faster much more quickly than common memory chips have. The most commonly used emollient for this now-huge disparity is the cache. Most modern microCPUs have some level of on-chip cache(s), but very seldom do they have a large enough cache for the chips to really run well. Hence, most processor-board designers are faced with including an external or secondary cache outside the chip but in front of the main memory system. In their paper, Yen-Jen Oyang and Le-Chun Wu provide an interesting parametric study of secondary cache design and its impact on performance. While not quite a "cookbook," the models and techniques should be of interest, especially to higher-level systems

designers trying to understand the impact of various cache organizations on system performance.

We are especially proud of the geographic range spanned in this issue: Taiwan, Australia, the United Kingdom, Germany, and the United States are all represented; four of the seven continents! We anxiously await our first submission from McMurdo Sound.

Finally, I direct your attention to the announcement which immediately follows this introduction.

That does it for this issue. See you here again soon.