

FRANKENSTEIN JR.

Francis Sullivan, Editor in Chief



THE ABOVE TITLE WAS USED IN NORWAY FOR MEL BROOKS' BRILLIANT SPOOF OF 1930s HORROR MOVIES SUCH AS *FRANKENSTEIN*, *BRIDE OF FRANKENSTEIN*, *RETURN OF FRANKENSTEIN*, AND *FRANKENSTEIN MEETS THE WOLF MAN*; THE ACTUAL TITLE OF THE FILM IS *YOUNG FRANKENSTEIN*. WHY THE

Norwegians chose a different English-language title is a mystery to me and is probably destined to remain one.

But I digress.

Two things became clear to me while watching *Frankenstein Jr.* in Norway: the supreme importance of language and the extreme danger caused by placing the brain of a homicidal maniac into the head of a large body and then reanimating the whole ghastly construct. (I already knew that Mel Brooks was a comic genius because I'd seen *High Anxiety* before leaving the States.) Both lessons—that language really matters and that you should exercise caution when installing a possibly corrupt brain into a giant—are relevant to *CiSE*.

At the time I saw *Young Frankenstein*, English-language films were not dubbed into Norwegian. Although many Norwegians spoke English, and some had even mastered the American accent, it didn't take long to realize that viewers would just read the subtitles if given the option when the dialogue was fast. The most obvious evidence was that I found myself laughing at punch lines a little earlier than the rest of the audience. I soon discovered how to steer the rest of the crowd: if I laughed hard, so did they; if I didn't make a sound, their laughter seemed more subdued. Eventually, I tried laughing when almost nothing was happening on the screen. Sure enough, a ripple of laughter soon bubbled up in the audience. (I only tried this once or twice—I swear!) The lesson learned is that language is important. Or, to be more accurate, latency due to language processing is important.

Frankenstein's monster is a talented chap in at least a few ways—in the Mel Brooks film, he even does a song-and-dance number—but the horrible consequences of combining his evil brain and huge, powerful body cancel

out any enjoyment of his talents. His size and strength turn out to be bad things when judged by the usual standards of civilized behavior.

So what does this have to do with computing? In the next few years, computer firms will produce a new collection of high-end machines—devices that previously went by the name “supercomputer.” In part, this is a reaction to the NEC Earth Simulator's impressive performance, but we shouldn't assign too much weight to desires to hold the world's record in computing speed; for one thing, there is no unambiguous and agreed upon definition of it. More importantly, a persistent set of computational problems exist whose solution requires considerably more computing power than even a large Beowulf cluster can offer or uses algorithms and computational strategies that are not amenable to a Grid-computing approach. This was the case 20 years ago (at the time of the last big surge in supercomputer production), and the same computing issues still hold true today.

We simply can't complete some computations in any reasonable amount of time without extensive use of very large, flat memories combined with huge bisection bandwidth. Such applications can't tolerate the latency that slow node-to-node communication introduces. (I hasten to add, of course, that another approach to the problem in question might not have these requirements, but some problems are just plain hard, no matter what approach you take.)

Some readers will recall that quite a few of the supercomputers we greeted with great enthusiasm 20 years ago are no longer around. Most didn't even leave any descendents. Unless we pay a lot of attention to potential problems due to language latency and bad brain installation, I fear the same thing could happen again. In

Reviewer Thanks

We extend our sincerest thanks to the many people who served as peer reviewers for 2002:

Mary Barth, Univ. Corp. for Atmospheric Research
John Carlis, Univ. of Minnesota
Peter Cheeseman, NASA Ames Research Ctr.
Ed Huai-hsin Chi, Xerox PARC
John Conroy, Ctr. for Computing Sciences
Jean Cote, Queen's Univ.
Roger Curry, Univ. of Calgary
George Djorgovski, California Inst. of Technology
John B. Drake, Oak Ridge Nat'l Lab.
John Dukowicz, Los Alamos Nat'l Lab.
Susan Dumais, Microsoft Research
Francis Giraldo, Naval Research Lab., Monterey
Eui-Hong Han, Univ. of Minnesota
Isaac Held, Geophysical Fluid Dynamics Lab.
Larry Horowitz, Nat'l Oceanic and Atmospheric Administration
Mariano Hortal, European Ctr. for Medium-Range Weather Forecasts
Kuo-lin Hsu, Univ. of Arizona
Roberta Humphreys, Univ. of Minnesota
Christiane Jablonowski, Univ. of Michigan
Ravi Janardan, Univ. of Minnesota
Mahesh Joshi, Univ. of Minnesota
Chandrika Kamath, Lawrence Livermore Nat'l Lab.
Steven Klooster, NASA Ames Research Ctr.
Debby Lanser, Univ. of Twente
Steve Libby, Lawrence Livermore Nat'l Lab.
John Lu, Nat'l Inst. of Standards and Technology
Ivan Marusic, Univ. of Minnesota
Julian Noble, Univ. of Virginia
Stephen Odewahn, Arizona State Univ.
Dale Pace, JHU/APL
Jian Pei, Simon Fraser Univ.
Wanda Pratt, Univ. of California, Irvine
Jim Purser, Nat'l Oceanic and Atmospheric Administration
Naren Ramakrishnan, Virginia Tech.
Vivek Sarin, Texas A&M Univ.
Shashi Shekhar, Univ. of Minnesota
Piotr Smolarkiewicz, Nat'l Ctr. for Atmospheric Research
Michael Steinbach, Univ. of Minnesota
Mark Taylor, Los Alamos Nat'l Lab.
Clive Temperton, Nat'l Oceanic and Atmospheric Administration
Paul Thompson, UCLA Medical Ctr.
Joe Tribbia, Nat'l Ctr. for Atmospheric Research
Timothy Trucano, Sandia Nat'l Labs.
Stacy Walters, Univ. Corp. for Atmospheric Research
Tim Warburton, Univ. of New Mexico
Michael Wehner, Lawrence Berkeley Lab.
David Williamson, Nat'l Ctr. for Atmospheric Research
Beth Wingate, Los Alamos Nat'l Lab.
Patrick Worley, Oak Ridge Nat'l Lab.
Mohammed J. Zaki, Rensselaer Polytechnic Inst.

particular, shared-memory applications should make full use of shared-memory architectures, and it should be easy for users to arrange to do so. In my opinion, tools such as MPI and OpenMP still generate much too much latency overhead. HPF (high-performance Fortran) is a noble research effort that has produced many insights, but I believe it suffers from the attempt to be too general; it is therefore rather complex and less well adapted to specific instances. It's perhaps naive to expect high-performance machines to arrive on the loading dock already equipped with a compiler to make good use of the architecture—or is it?

Finally, there's the question of matching the brain to the body. Some of the new machines will be equipped with special architectures, such as vector processors and PIM chips, which are extremely powerful devices. Much of the Earth Simulator's speed is due to effective use of vector processing. If the rest of the processing suite (including standard, unfancy scalar processing) isn't done well, however, the machine will be unbalanced and socially maladjusted.

The next few years will be an extremely exciting time for those interested in the top end of high-performance computing. Although I've emphasized here some lingering worries, I'm fairly confident that the community has learned well from the hard lessons of the past.



Got research?

Submit your paper to
Francis Sullivan at
fran@super.org.

computing
in SCIENCE & ENGINEERING