

The Minisupercomputer: One User's Experience

*Tom Vollaro
Assistant Director of Computer Technology
Grumman Data Systems
Woodbury, New York*

Introduction

Grumman Data Systems, a division of Grumman Corporation, designs, develops, and supports computer systems for technical and management information. The Data Systems group provides computing services to other Grumman divisions, as well as to a broad range of customers in federal and state government agencies.

Government contracts encompass a broad spectrum of applications. To remain competitive, Grumman requires computer systems that can handle the changing requirements of different jobs. Furthermore, because the majority of these government applications are compute-intensive, Grumman requires tremendous compute power for performing large-scale numeric computations.

Grumman currently runs many of their applications on a Cray XMP/14 supercomputer. However, the introduction of a new class of machine--the minisupercomputer--led the company to consider alternative platforms for both internal use and potential commercial systems integration. In the case of minisupercomputers, it was Grumman's hope to achieve satisfactory levels of performance at dramatically reduced cost. This article discusses Grumman's computing requirements, their criteria for evaluating high performance systems, and their experience with one of minisupercomputer, the Multiflow TRACE 7/200.

Computing Resource Requirements

From participating in the major phases of space station development, to designing fix-based flight simulators, to developing software for encrypting military messages, Grumman runs a broad range of compute-intensive applications. Thus, Grumman's main computing requirement is very high performance on a wide variety of complex applications. The company and its customers cannot accept marginal performance and have no preference for an application-specific platform when acceptable performance can be achieved on a general-purpose system.

Next, because government and DOD applications consist of large amounts of data, Grumman requires a computer system that is capable of processing large-scale programs containing large data sets. Additionally, because development contracts usually run over fixed periods of time, the company requires quick turnaround on a job-by-job basis. Quick turnaround is essential to providing solutions within tight deadlines and evidence of this ability is often necessary to ensure project funding.

Given that sufficient performance levels are possible, many other criteria come into play. One of the most important is the effort required to achieve needed performance from a given system. Ideally, the platform will run existing applications with few, if any, changes to the source code. The cost and delays associated with reworking applications are immediate issues, as are the difficulties of maintaining different versions of an application for different devices. This particular attribute might be considered a subset of a somewhat larger requirement called "user functionality." This includes an attractive environment for program development, good compatibility with industry standards, and a high degree of connectivity with other resources, such as workstations, VAXes, and, of course, the Cray. Finally, each component of cost is an issue: acquisition, integration, maintenance, and program development costs all are addressed.

The Minisupercomputer Alternative

Until recently, users of scientific and engineering computers were forced to accept a difficult tradeoff. Users could choose very high levels of performance, which were usually accompanied by multimillion dollar system costs and cumbersome program development environments. Or they could choose from any number of affordable, easy to use superminicomputers, but only if they could tolerate the many compromises imposed by relatively poor levels of performance. In response to these limitations, a new class of computer has emerged: the minisupercomputer.

Manufacturers of these machines claim to deliver performance levels approaching those of true supercomputers, such as the Cray, but at system costs much closer to those of the traditional supermini. As important, many makers claim their devices are as easy to use as a superminicomputer. Virtually all of them employ some version of the Unix operating system, offer industry-standard languages, connectivity, and communications protocols, and good VAX-compatibility. Grumman found all these attributes well-suited to their computing requirements.

Another feature which characterizes minisupercomputers is some form of parallel architecture. Indeed it is the parallel use of comparatively inexpensive hardware that makes possible the performance and pricing of these systems. Vector, multiprocessor, and very long instruction word (VLIW) architectures were all available to Grumman for possible evaluation. Grumman had some familiarity with vector and multiprocessor designs, as these approaches are incorporated in the Cray, as well as in a number of less costly systems. Despite--and, in part, because of--their familiarity with vector and multiprocessor devices, Grumman was interested in testing a minisuper based on VLIW design. This new system architecture had been claimed to address several limitations of traditional parallel machines.

According to Multiflow Computer, the Branford CT-based maker of the TRACE VLIW computer selected by Grumman, VLIW architecture offers major advantages over other parallel designs. These are that VLIW systems achieve greater speedups than vector or multiprocessor machines (given comparable clock speeds and, by implication, similar hardware componentry and cost); that VLIW systems achieve speedups across a wider range of applications; and, finally, that these gains are delivered without any of the major programming effort that can be required to achieve speedups on other parallel machines. Grumman had seen evidence to support the first two claims, in the form of a variety of benchmark results. In comparison with other minisupercomputers employing vector or multiprocessor architectures, the Multiflow TRACE 7/200 demonstrated not only greater achieved performance, but greater gains across the breadth of applications reflected in benchmarks such as Linpack, Whetstones, and Livermore Loops.

On the better known benchmarks, at least, it appeared that VLIW architecture would satisfy Grumman's most important requirement: very high performance across a range of applications. What was not all clear when they began their evaluation was how it would perform in the real world. For Grumman, the two critical questions were: to what extent could comparable gains be achieved on the company's production codes and how great a programming effort might be required to achieve them?

Real World Performance

To answer these questions, Grumman ported a suite of its own benchmarks to the test system. This suite consists of compute-intensive codes representative of the variety of programs run on a daily basis. All of them heavily exercise CPU, memory, and I/O resources. Some, like the Linpack benchmark, offer good performance gains on vector and multiprocessor machines; others, like the Whetstones benchmark, are inherently sequential and tend not to benefit from older parallel architectures. To test ease of porting as well as performance, five benchmarks of varying length and complexity were used. The five benchmarks were:

- **FLO67**- a 3-D computation of wing analysis in transonic flow using solutions of unsteady euler equations for the simulation; 3,600 lines of code.
- **TNS**- a 3-D fluid dynamics program using an implicit Beam-Warming algorithm; 15,000 lines of code.
- **BEFAP**- a 2-D boundary element structural analysis program; 3,300 lines of code.
- **DYCAST**- a nonlinear, structural dynamic finite element computer code used to simulate automobile and other vehicle crashes; 50,000 lines of code.
- **TOPCAT**- a combat simulation program for the X29 aircraft; 5,000 lines of code.

Benchmark Results

Grumman's results, expressed in cpu seconds, along with the number of days needed to port each benchmark, are shown in the following table:

Grumman Benchmark Summary Report

BENCHMARK problem phase	DAYS TO PORT	CRAY XMP/14	MULTIFLOW TRACE 7/200
FLO67 (92x12x12)	2 Days	20.7	285.5
TNS	2 Days	108.3	663.6
BEFAP	2 Days		
matrix generation		25.2	42.6
decomp & solution		6.1	8.8
domain & stress		16.7	17.9
DYCAST	5 Days	814.9	1003.7
TOPCAT	1 Day	3200.8	2870.1

As shown above, the Multiflow TRACE performed exceedingly well in comparison to the Cray XMP/14, a far more expensive machine. On FLO67, the most highly vectorized benchmark, the Cray was significantly faster than the Multiflow system, though perhaps by not as much as its more than twenty times greater price would suggest. The Cray XMP costs approximately \$9 million, versus \$400,000 for the TRACE 7/200.

What is truly noteworthy is the performance of the TRACE on the other four benchmarks. On TNS, the Multiflow system was able to achieve approximately one-sixth the performance of the Cray; on BEFAP this increased to 50%-90% of the Cray; on DYCAST the TRACE delivered roughly 80% of the Cray's power; and on TOPCAT, the least vectorizable benchmark, the Multiflow 7/200 actually outperformed the XMP/14. From these results, Grumman concludes that minisupercomputers can indeed offer an alternative to supercomputers on many kinds of applications.

The benchmark summary also reveals the ease with which programs were ported to the TRACE. With times ranging from one to five days on programs from 3,300 to 50,000 lines long, Grumman's experience supports the claim made for VLIW ease of use. The performance levels reported above were achieved without modification to the program algorithms and without any special tuning effort.

This finding is most significant for users who depend heavily on proprietary codes. While computer manufacturers and third party software providers can distribute the cost of a difficult port or extensive tuning across a large base of customers, it is the individual user who usually bears the burden of porting and optimizing his particular in-house applications. System architectures vary greatly in this respect and thus remain a primary concern for prospective minisupercomputer users.

In most other respects, the Multiflow TRACE has proven well-suited to Grumman's needs: it employs a Unix operating system (BSD 4.3); runs programs written in standard FORTRAN (with DEC extensions) and C; provides industry-standard communications facilities (Ethernet: TCP/IP, NFS, NCS, X Windows, etc.); and offers extensive DEC compatibility via a DCL command shell, EDT editor, and DECnet networking. Again, it is worth noting that most other minisupercomputers claim to offer similar features.

Finally, Grumman's cost concerns were addressed by their minisupercomputer. The Data Systems group's TRACE was priced at approximately \$400,000, roughly what one would expect to spend for a mid-ranged supermini and not quite 5% the cost of the Cray. Porting and development costs have proven to be comparable to what Grumman has experienced with traditional superminis, as have support and maintenance expenses. Modest operating costs for the machine itself were to be expected given its somewhat ordinary CMOS and Schottky-TTL componentry. Other vendors systems using the same hardware implementation should achieve similar economies.

In Conclusion

Grumman Data Systems' experience with one minisupercomputer has proven quite favorable. In comparison with the Cray XMP/14, performance on a wide variety of jobs has ranged from acceptable to outstanding. The system chosen by Grumman presented no unusual porting challenge, provided a robust development environment, and maintained the portability of company codes. These results would be impressive enough from a far more expensive system. Given the attractive pricing of the machine that produced them, the Multiflow TRACE, they are nothing less than remarkable.