

**Response to Senate Armed Services Committee
Request for Information on The National Security
Implications of Decontrolling Export to Tier III
countries of High Performance Computers between
2,000 and 40,000 MTOPS**

D. Nielsen, Jr., R.B. Christensen

February 10, 2000

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
<http://apollo.osti.gov/bridge/>

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Response to Senate Armed Services Committee Request for Information on The National Security Implications of Decontrolling Export to Tier III countries of High Performance Computers between 2,000 and 40,000 MTOPS

Dale Nielsen Jr., Randy B. Christensen

Lawrence Livermore National Laboratory

10, Feb, 2000

Introduction

The Senate Armed Services Committee has asked Lawrence Livermore National Laboratory (LLNL) to assist the Committee in assessing the national security implications of any relaxation of export controls on High Performance Computers (HPCs). In particular, LLNL was tasked to assess: 1) How these computers could be used militarily, 2) Whether they could be effectively controlled and for how long, 3) The significance and quality of available foreign computers projected to be available on the open market in the next two years, 4) Alternative control mechanisms based on technology or capability that might be more effective to protect national security, and 5) Issues related to other critical technologies.

Abstract

The SASC requested information about the national security implications of decontrolling the export of computers with performance ratings between 2,000 and 40,000 MTOPS (Millions of Theoretical Operations Per Second). We respond to this request by: 1) Discussing computer performance thresholds for nuclear weapons design application software, 2) Examining currently available hardware (including desktop PCs, workstations, Commercial-Off-The-Shelf (COTS) parallel computers, high-end-supercomputers, and highest-end supercomputers) with respect to their utility for nuclear weapon design and their controllability, 3) Briefly commenting on examples of computers available from foreign countries, 4) Discussing the implications of newly emerging COTS parallel clusters for various computer export control strategies. In addition, we are making available to the SASC the DOE sensitive technologies list.

Note: This document is largely based on material developed by LLNL for DOE in response to a request from the National Security Council to provide information on the extent to which operations directly relating to national security issues can be performed on various configurations of computers.

Conclusions

- Advanced computing hardware *by itself* contributes little to the development or maintenance of nuclear weapons. However, in conjunction with nuclear and non-nuclear test data, computing may be very beneficial.
- Computer performance at the level of 2,000 MTOPS is useful and sufficient for many two-dimensional nuclear weapon design simulations if relevant nuclear test data is available.
 - All of the weapons in the enduring US stockpile were designed on computers with a performance lower than 2,000 MTOPS, **in conjunction with a robust nuclear testing program.**

- Computer performance at the level of 50,000 MTOPS is useful for nuclear weapon design simulations in support of safety, stockpile maintenance and minor (i.e., small) modifications of existing, nuclear-tested weapon designs when relevant nuclear test data is available. It is not sufficient for major (i.e., large) modifications or new designs.
- Meaningful control of computers available to foreign nuclear states at the level of 2,000, or even 40,000 MTOPS (and soon at even higher levels) is difficult if not impossible.
 - It is possible to build effective parallel computers today for nuclear weapon design simulations up to about 50,000 MTOPS out of widely available, commercial-off-the-shelf (COTS) computers and networking technology. The level of performance offered via this approach will climb rapidly.
 - Finished high-performance computers (HPCs) still offer a small (but declining advantage) relative to COTS technologies.
- Computer performance at the level of 30 Million MTOPS (i.e., at the scale being developed in the DOE's Accelerated Strategic Computing Initiative - ASCI) is required as an essential element of the US Stockpile Stewardship Plan.
 - 1) It takes many tens to hundreds of the most powerful HPCs connected by the highest performance, vendor-proprietary networking technology to create ASCI-scale computers. **These systems can and should be subject to export control.**
- The scientific skills and understanding, software technology, and computing hardware necessary to support nuclear weapon simulations are widely available throughout the world.
- The data from nuclear tests necessary to validate nuclear weapon simulation codes and otherwise support the design activities of advanced nuclear weapons is the only remaining requisite which is not widely available to the world.

1) How High Performance Computers could be used militarily

We will address this issue in the context of nuclear weapon design.

Methodology / Analysis Framework

The effect on a foreign nuclear weapons program of acquiring US computers may be framed in terms of a matrix of four categories of nuclear stockpile activities versus two testing scenarios. The stockpile activities are Safety, Reliability and Maintenance, Modification, and New Designs. The testing scenarios considered here are No Testing and Unconstrained Testing. The effect of computers on the activities in each of these eight regimes will be examined.

Note: The response to the NSC request referred to above includes a discussion of a third testing scenario - nuclear testing with yields below accepted detection thresholds. The conclusions reached in this document are unaffected by this discussion.

	Safety	Reliability & Maintenance	Modification	New Designs
No Testing				
Unconstrained Testing				

Table 1

Definitions

The four stockpile activities are defined as follows:

- Safety refers to all of the assessment activities focusing on accidents (crashes, fires, etc.), radiation shielding, and the abnormal initiation of high explosives in a nuclear warhead.
- Reliability and Maintenance activities are limited to those efforts which aim to restore aged or damaged weapons to their original design specifications, i.e., to a design which has been extensively tested.
- Modification includes changes to a weapon's design which maintain the intended performance or modestly enhance performance in some operationally significant way. Such performance enhancements would include, for example, modest increases in nuclear yield (by special nuclear materials for example), increased radiation shielding, increased handling durability, increased environmental operating range, etc. Modifications represent deviations from nuclear tested designs.
- A "new design" is a warhead that has never been in the stockpile and undergoes development using an untested nuclear physics package. Present US nuclear design laboratory judgment is that a new design would require a full-system nuclear test before it could be certified for the stockpile.

The two testing scenarios are defined as follows:

- No testing is here defined to mean no supercritical assembly. This represents the current US interpretation of the limits imposed by the CTBT.
- Unconstrained testing includes tests which fully exercise both the primary and secondary (where applicable).

Thresholds

The question to be asked in each of the eight regimes is: "What is possible as a function of the computer power and simulation technology obtained by the foreign nuclear weapon

programs?" In other words, would giving existing foreign nuclear weapon programs access to more capable computers, *by itself*, lead to significant new capabilities for the foreign country's stockpile? More precisely, are there clearly defined thresholds of computer capability at which significant new weapons stockpile activities become possible?

The notion of well-delineated thresholds of computer power which trigger major new computational capabilities is important for establishing controls, *but it is important to keep in mind that the realities of how nuclear weapon design simulation is done do not provide very many sharp transitions.*

The issue of computer capability thresholds may also be cast in terms of specific performance levels, such as,

- 1) What types of nuclear weapons calculations can a foreign nuclear program perform with the computers available at current and anticipated export control thresholds, both military end use and civilian end use?
- 2) What types of nuclear weapon calculations can the foreign nuclear program perform with Commercial Off The Shelf (COTS) parallel clusters?
- 3) What types of nuclear weapon calculations are enabled by access to even the most capable of today's commercially available computers?

This paper will lay groundwork for answering such questions by presenting three ways of characterizing the simulation activities in each of the eight regimes. First, a brief description will be given of the numerical simulations which are performed in the US to support each of the four categories of stockpile activity. Second, the computational performance levels necessary to properly address issues in each of the eight regimes will be given. Finally, a categorization of what can and cannot be done in each regime will be given at a particular threshold level, in this case, the highest levels available in standard commercial computers from US computer vendors (about 50,000 to 70,000 MTOPS).

Computational Characteristics of Nuclear Stockpile Activities

SAFETY – Currently the US uses 3D calculations as the standard for Safety assessments, although some special cases can be done in 2D. As done traditionally, these simulations require physical models for only hydrodynamics and criticality. Supercomputers of the past, offering performances roughly equivalent to today's fastest workstations, were sufficient for these kinds of safety assessments when coupled with nuclear testing. Computers capable of delivering 30,000 MTOPS to the application would provide significant improvements for these traditional safety assessments. Modern safety assessments, covering a wide range of scenarios, must treat a much larger set of chemical and physical models and require more spatial resolution. These modern safety assessments will tax the computational abilities of our largest supercomputers for many years.

RELIABILITY AND MAINTENANCE – Issues associated with Reliability and Maintenance require "full-physics" simulations. This means, that in addition to the physical models necessary for safety, models must be provided for the full range of physical processes present in a nuclear explosion. These additional processes require much greater computational power than does hydrodynamics. Under-resolved 2D "full-physics" calculations, which were adequate when performed in conjunction with nuclear testing, are comparable in computational difficulty to traditional safety calculations. Modern, well-resolved "full-physics" calculations in both 2D and 3D will be challenging for the largest supercomputers we can expect to see for the next decade. Well-resolved 3D simulations using the best-quality physics models require computational capabilities 1,000 times greater than the most capable supercomputers commercially available today. Because of this, the US DOE ASCI Program is now attempting to work directly with US computer vendors to provide new supercomputers for the US nuclear weapons design laboratories that can provide these levels of performance.

MODIFICATION AND NEW DESIGN – These categories are more difficult to characterize simply. They both depend *very* strongly on the size or extent of a proposed

“modification,” or the newness of a “new” design. For a large modification or a new design which lacks a solid “nuclear test pedigree,” there is no computer commercially available today which provides the confidence obtained from nuclear testing.

Absolute Performance Levels Required in Each Regime

Below is a table that shows the approximate level of computing capability needed to adequately address design challenges in each regime. Color is used in this table to separate three levels of computing power.

• **Green**

This indicates performance levels in the “multi-megaflop” range. This range of performance is readily available in the desktop PC and open, commercial workstation market, and extends up to 2,000 MTOPS. This is the range of computer power used (along with active nuclear testing) during the design of the entire US stockpile.

• **Yellow**

This indicates performance levels in the “multi-gigaflop” range. This range extends from 2,000 MTOPS (low-end commercial computers) up to 100,000 MTOPS (about the limit of today’s commercially available supercomputers from US and Japanese manufacturers). The previous export control limit for shipments to military end-users in Tier 3 countries is 6,500 MTOPS, near the bottom of the range. Note that the historic rate of improvement for computer performance predicts that in ten years this entire range will be widely available at desktop PC prices.

• **Red**

This indicates performance levels in the “multi-teraflop” range. This is the performance range that the DOE ASCI program will be providing in support of the US stockpile stewardship program. It extends from 300,000 MTOPS to about 25-30 Million MTOPS. At the high end this is about 4,000 times greater than the previous export control limit for military end-users in Tier 3 countries.

	Safety	Reliability & Maintenance	Modification	New Designs
No Testing (but with access to archived nuclear test data)	Traditional: From 300 to 30,000 MTOPS Modern: 25 Million MTOPS	Traditional: From 300 to 15,000 MTOPS 3D: 25 Million MTOPS	Small: 2,000-15,000 MTOPS Other mods: 25 Million MTOPS	Not Possible
Unconstrained Testing	300 MTOPS	300 MTOPS	300 MTOPS	300 MTOPS

Table 2

Stockpile Activities Which Are Possible With Currently Available Computers From Commercial Vendors

The following table summarizes the influence of computation and simulation on US (and presumably Russian) nuclear stockpile activities under the two testing scenarios, using

the largest computers currently available as standard commercial products in the US today, *but not including the greatly increased capabilities yet to be delivered by the DOE ASCI Program*. The focus on US and Russian capabilities means that this table refers to boosted primaries and thermonuclear secondaries and assumes access to archived nuclear test data. Our experience with proliferators has shown that unboosted single stage weapons can be added to a stockpile without nuclear testing. While it is possible to design boosted primaries and thermonuclear secondaries without a nuclear test, the decrease in confidence relative to proven systems is so great that we consider it extremely unlikely that the Russians would do so.

	Safety	Reliability & Maintenance	Modification	New Designs
No Testing (but with access to archived nuclear test data)	Doable:* • Traditional Not doable: • Modern	Doable:* • Issues within test history Not doable: • 3D issues (like many aging problems)	Doable:* • Small mods Not doable: • Other mods	Not doable
Unconstrained Testing	Doable	Doable	Doable	Doable

* The word doable is qualified here because, even though these calculations can be done, historically the calculations have always been validated with nuclear tests.

Table 3

Gaining Confidence From Simulations¹

The confidence one can place in the results of a nuclear weapon simulation is determined by how well the simulation code has been validated. The validation of such a code depends on several factors.

- The degree to which the underlying physics models are known. Today, there is high confidence that these are correct in many areas of nuclear investigation. These models accurately reflect the nature of the world and are quite stable.
- The quality of the numerical methods and algorithms. The numerical methods necessary for nuclear weapons simulation have much in common with simulations of many other application areas. The experience gained in other application areas is readily transferable to nuclear weapon simulation.
- The resolution of all important processes and spatial scales. The nature of the problem dictates the spatial and temporal resolution required, which in turn dictates the overall computational requirements. If the machine being used is not able to offer the necessary memory and CPU cycles then one will have little confidence in the results, even when the physics models, numerical methods and algorithms are sound.
- The fidelity of the computation. The code's prediction of results must correspond to the real world as represented by empirical data from nuclear and other relevant non-nuclear tests.

Of these four factors, the first two are widely available around the world. Only the third

¹ The following discussion and illustration are liberally adapted from [GWH98]

factor, the issue of resolution of physical processes and spatial scales, depends on the memory size and CPU speed of computers and therefore is directly affected by export control of HPC hardware. The fourth factor requires high-quality data from nuclear tests. **Advanced computing hardware by itself contributes little to the development or maintenance of nuclear weapons. However, in conjunction with nuclear and non-nuclear test data, computing may be very beneficial.**

When a validated code is used to model small variations from a tested design (e.g., replacing one small component with another one substantially similar to it), weapon designers will have a great deal of confidence in the computed results. The greater the deviation from validating test data, the lower the confidence in simulated results. There will be little confidence in computed results alone for a simulation of a new weapons design concept, or a phenomenon that has not occurred during actual testing.

When test data are available, incremental increases in computing power provide incremental benefit. Designers are able to run bigger problems, more problems, improve the resolution of a problem, refine approximations, use improved physics models, reduce the turnaround time, or do a combination of all of these. **When test data are not available, an increase in computing performance has little or no effect.** This relationship is shown in Figure 1. The figure is *not* drawn to scale.

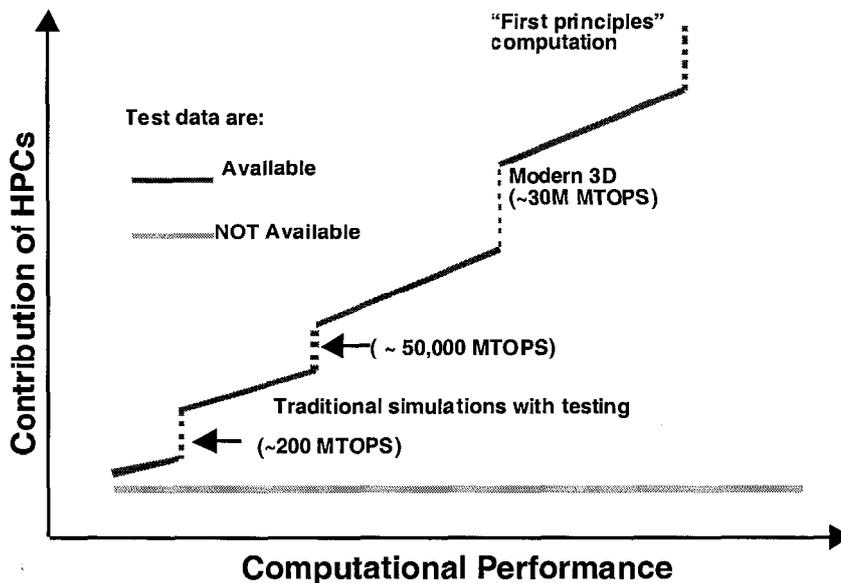


Figure 1 – Important transitions in nuclear weapon design simulation

Note that the dotted lines in figure one are not thresholds at which sudden increases of capability are realized. Rather these are meant to indicate general levels of computing at which nuclear design code developers begin to seriously view the indicated activities as being credible. Also, the fact that the green line is drawn perfectly flat should also not be interpreted too strictly. Increasing computer power does in fact provide *some* increased contribution. Such an increase, however, pales into insignificance when compared to the confidence obtained from validating a code against full-scale nuclear tests.

2) Can HPCs between 2,000 and 40,000 MTOPS be effectively controlled and for how long?

Commodity microprocessors exist today which have MTOPS ratings of ~3000 (PowerPC G4). By the end of 2000 this is expected to move up to the ~6000 MTOPS level with the introduction of the Intel Itanium microprocessor. Small servers based on the Itanium could be expected to perform at the ~20,000 to 40,000 MTOPS level. Individual personal computers, workstations, and small servers (4-8 cpus) based on these commodity microprocessors are so inexpensive, plentiful, and widely available as to be essentially uncontrollable. Moreover, the rapidly evolving technology of COTS cluster computing raises the level of easily accessible HPC significantly above the level of individual small servers.

“Clustering” Computers

The concept of clustered computers is a relatively new concept in the forum of export control discussions, and refers to the strategy of using computer networking technology to connect (or integrate) individual computers together for the purpose of creating a single “parallel” computer that has a combined performance distinctly higher than its individual component computers. Although new to the export control discussion, the concept of clustering has been studied for many years in academia, and has been recently popularized by the Beowulf Project at the California Institute of Technology [BW]. In the 1995 report by Goodman, Wolcott, and Burkhart “Building on the Basics: An Examination of High-Performance Computing Export Control Policy in the 1990s” [GWB95] the concept of clustering computers was investigated and judged to be infeasible due to the low performance of the available networking technology at that time. Only three years later, an updated review by the same researchers [GWH98] reported that COTS networking technology had improved to the point that it had become a viable method for the creation of effective parallel computers, albeit at a performance level which still trailed the performance achievable by vendor-specific, proprietary integration technology. Figure 2 is borrowed (with permission) from the 1998 report and shows this newly emerging class of computers positioned between low-performance traditional LAN (local area network) technology and the highest-performance proprietary hardware based interconnects. Note that this figure places higher performing regimes to the bottom right of the chart.

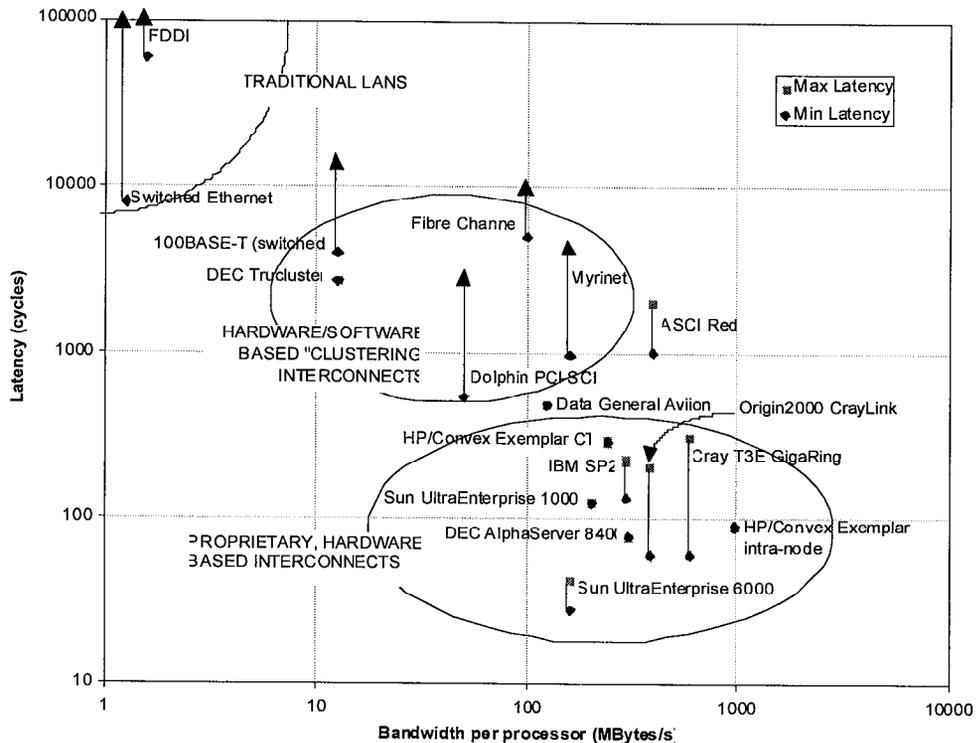


Figure 2: Three regimes of computer performance

COTS Networking Performance Increases One Hundred-Fold

For many years the only widely available COTS networking technology was Ethernet operating at 10 Megabits/sec. Later this was followed by “Fast” Ethernet and FDDI at 100 Megabits/sec. Today several COTS networking options are available with performance from 800 Megabits/sec up to 1,000 Mbits/sec (i.e., 1 Gigabit/sec), for example Gigabit Ethernet, Myrinet, Hippi, and FiberChannel. This dramatic network performance increase is driven by the same advances in semiconductor technology that are driving CPU performance to follow Moore’s law (doubling performance every 18 months). *The highest performance COTS networking now equals the networking used by some commercial parallel computer manufacturers and is fully 38% of the highest-performance vendor-proprietary networking technology.*

Networking Performance Determines Peak Cluster Capability

When COTS networking technology is used to create COTS parallel computers out of either PCs or small servers, the key technical hardware factor, which determines the quality of the resulting cluster, is the balance between network performance and computational power. A proper balance will lead to effective parallel machines. For a given networking performance and computational power there will exist some upper limit to the number of processors which can be effectively connected together. Said another way, for any given networking performance, the more powerful the computational nodes, the fewer of them that can be joined together *effectively*. The result is that, to a rough approximation, the overall power of the final cluster created with a given networking performance level is independent of the power of the computers that are being clustered.

Making COTS Clusters Work

Claims for the value of COTS clusters have been met with some skepticism. In determining just how useful COTS clusters are, there are four aspects of computing which must be addressed: 1) hardware performance, 2) system software, 3) application software and 4) “production support” for nuclear design computations.

Hardware Performance

To be a suitable platform for serious computational simulation, COTS clustering must result in effective hardware for parallel computation. Because of the performance of today’s COTS networking hardware and its flexibility in supporting a wide range of interconnection topologies used by traditional parallel computer manufacturers (and studied in the academic literature), COTS clusters perform up to about a third as well as today’s best distributed memory parallel computers. The cost of COTS networking adds only between 10% and 30% to the cost of the computers, which are being clustered. Finally the ease of integrating the individual computers of a cluster together has been greatly improved by the introduction of self-configuring network switch and interface hardware. Performance, flexibility, low cost, and the ease of integration, taken together, make the development of effective parallel computers using COTS networking hardware dramatically easier than it was just three years ago.

System Software

The system software necessary to integrate a cluster of computers into an effective single computer with a correspondingly higher capability than its individual constituent processors is now widely available. In the earliest days of parallel computing, government laboratories, academic research projects, and a few very brave parallel computer vendors created custom hardware and system software. The critical gap between application level software and the underlying inter-processor communication hardware was filled by what was called a communication software library, and each computer vendor, and research project provided their own version of this software. This made the application software developer’s job of porting his application between different parallel computers time consuming and tedious. Beginning in 1993 a broadly based group of parallel computer vendors, software writers, and application scientists collaborated on the development of a standard portable communication library definition called MPI (which stands for Message Passing Interface)[GLS]. During the past six years there have been several “reference” implementations of this software developed and posted on the Internet for all to study and use.

Application Software

Because the system software used on COTS clusters is the same as that used on vendor assembled parallel computers, programming a COTS cluster is no different than programming most other commercial (distributed memory) parallel computers. The number of skilled computer scientists required to create parallel versions of nuclear design applications is relatively small: a few tens of B.S. level computer scientists. The number of skilled people to create the non-parallel versions of the nuclear design applications in the first place is somewhat larger: several tens to a couple hundred of B.S. and Ph.D. level computer scientists, physicists and mathematicians. All computer science departments at universities throughout the world teach the concepts of simulation and parallel processing necessary to support the development of parallel applications for COTS clusters using the standard MPI interface. Examples of parallel implementations of many of the physics packages involved in nuclear design simulation are individually available on the Internet and have been for many years.

Production Support

Finally, there is the little-discussed aspect of actually providing for COTS clusters the sort of production computing support necessary to enable an advanced nuclear weapons design effort. Production support includes determining an efficient schedule for running hundreds of separate problems, as well as input and output of data between a computer, its local disks, networks and archival storage. This is the unglamorous part of high performance computing, but it is absolutely essential to making what is initially a prototype or technology demonstration into a robust production computing platform. U.S. practice currently is to expect this infrastructure from computer vendors. Such software is moderately well developed for commercial use, but is much less well developed for large parallel scientific applications. However, not having vendor software provide this production computing infrastructure does not make production computing impossible. It simply means that without software to automate production activities, the scheduling, processing of input and output and archiving must be done manually – making production computing into a very labor-intensive operation. While this is not a cost effective strategy in the U.S., in foreign countries such as India and China where skilled labor is plentiful and low cost, this is much less of a problem.

The highest performing HPCs use proprietary networking and/or memory system technology and can provide performance levels that are much greater than COTS clusters

This class of HPCs consists of computers which have been integrated by commercial computer vendors using *proprietary* networking and/or memory system technology to achieve the highest possible levels of integration and represents a configuration which has a potential peak performance level which is truly distinct from, and higher than any other class of HPC or cluster.

The best way to understand both the potentials and the limitations of COTS clusters with respect to vendor-assembled HPCs is to place COTS technology into a context of both computer power and parallel integration technology. Doing so requires the use of four new category designations. These new designations make the large-scale trends in computing easier to understand and forecast, as shown in the section below entitled “Four Distinct Classes of Computers in Today’s Market”.

Four Distinct Classes of Computers in Today’s Market

The best way to understand the relationship between COTS clusters and finished HPCs is to evaluate their overall effectiveness in relation to peak computational performance and the quality of integration, taking into account the hardware technology, which is used. Figure 3 graphically shows just such an evaluation. The first class presented is desktop PCs and workstations. Typically such computers use one (or sometimes two) commercial microprocessors. These same commercial microprocessors are the foundation technology out of which all of the other three classes are built as well. When computers (or individual processors) are connected together using COTS networking, the result is a COTS parallel cluster. If instead of COTS networking, a computer vendor connects multiple (typically from 2 to 128) processors together using proprietary shared memory system technology the result is an HPC, or high-end supercomputer. It is this high-performance, high-quality integration which allows the HPC to be effective in running a larger class of scientific simulation algorithms than clusters can run. Finally, when a computer vendor connects finished HPCs together using proprietary networking technology the result is a “Highest-End Supercomputer,” or ASCI-scale Supercomputer.

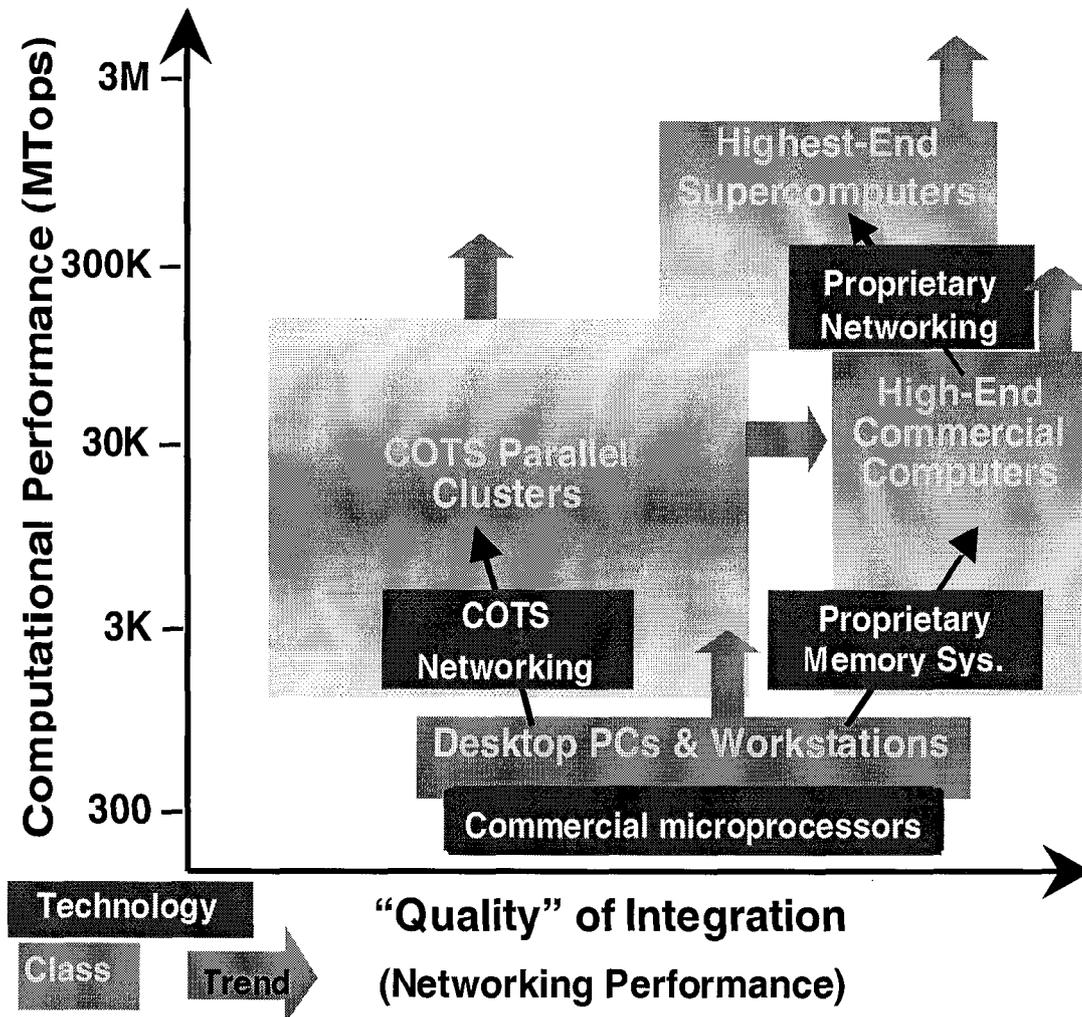


Figure 3: Four Classes of Computers in Today's Commercial Market

A few words about Figure 3 are in order. When the shading in a box is graduated, (tapering gradually to white) this is meant to indicate that there is not really a clearly defined limit. This usually means that even though computers of a given size can physically be connected, the effectiveness of the resulting computer is decreasing with increasing size. In the cases of high-end commercial computers and highest-end supercomputers, the upper limit is not graduated, meaning there is currently a well-defined highest performance level. In the case of graduated shading on the bottom of a box, this is meant to indicate that smaller configurations are possible, usually clear down to the level of desktop PCs and workstations, but such cases are not of interest to this discussion.

The vertical scale in Figure 3 is quantitative and very precise. The unit of measure is Millions of Theoretical Operations Per Second, or MTOPS. The upper limits of COTS clusters and commercial high-end computers are roughly the same. It is this overlapping range of performance, which contributes to much confusion when comparing COTS clusters and high-end supercomputers (or HPCs).

Figure 4 below indicates what each level of performance on the vertical axis of Figure 3 means in terms of historical, current and future DOE nuclear weapon simulation capabilities. Also superimposed on this scale is the previous military threshold for shipment to Tier 3 countries under a general license (meaning no prior U.S. governmental approval required) and an indication of the upper limit of both commercial HPCs and

COTS clusters (about 50,000 MTOPS). Note that Figure 4 extends to a 10-fold higher level than Figure 3.

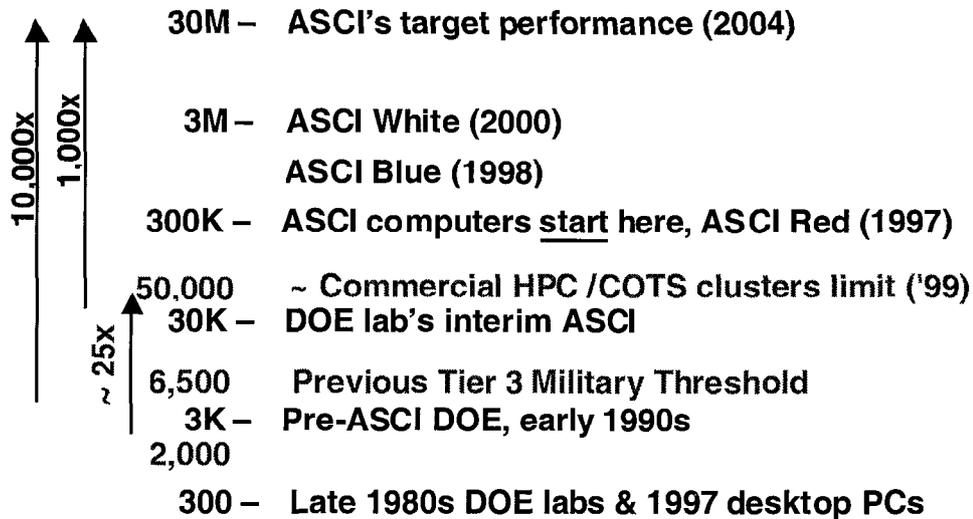


Figure 4: A Guide to Important Ranges of Computer Performance

Computing Performance Trends

None of the boundaries for the above four classes are static because the performances of the four technologies used are constantly being improved. Increasing microprocessor performance lifts the peak performance available in all four classes. The same semiconductor technology advances which allow microprocessor performance to double every eighteen months (Moore's Law) also drive the development of network technology. As the performance of COTS networking improves over time, the "COTS Parallel Cluster" box will grow to the right and approach the integration quality available in today's vendor assembled HPCs. Higher performance COTS network technology also allows a greater number, and more powerful microprocessors to be effectively integrated. The number of processors which vendors have been able to successfully integrate using proprietary shared memory technology also increases with time. The DOE ASCI program relies on the multiplicative effect of simultaneous advances in networking, shared memory supercomputers, and individual microprocessors to maintain its aggressive push to 30 Million MTOPS by 2004.

How Long Can Computer Exports Be Controlled?

The advances in computer power over the last several decades have made inexpensive computers which far outstrip the power of the computers used to develop the nuclear weapons in the enduring U.S. stockpile globally available. During the period when Cray 1s, XMPs and YMPs were considered supercomputers, there were at most a few hundred computers that needed to be "tracked". Today, computers more powerful than Cray YMPs are ubiquitous. Figure 5 shows a history of the performance levels for computers which were used during the development of the current U.S. stockpile along with a display of the number of nuclear tests conducted by the U.S. The time periods for the "peak computational activity" for each enduring stockpile system are also displayed. The rapid increase in the power of computers which is being sought by the ASCI program as a response to the end of nuclear testing represents a dramatic increase over the levels of computing power used when nuclear testing was available. The increase in computer power

represented in Figure 5 covers nine orders of magnitude, i.e., a billion-fold increase of computer power.

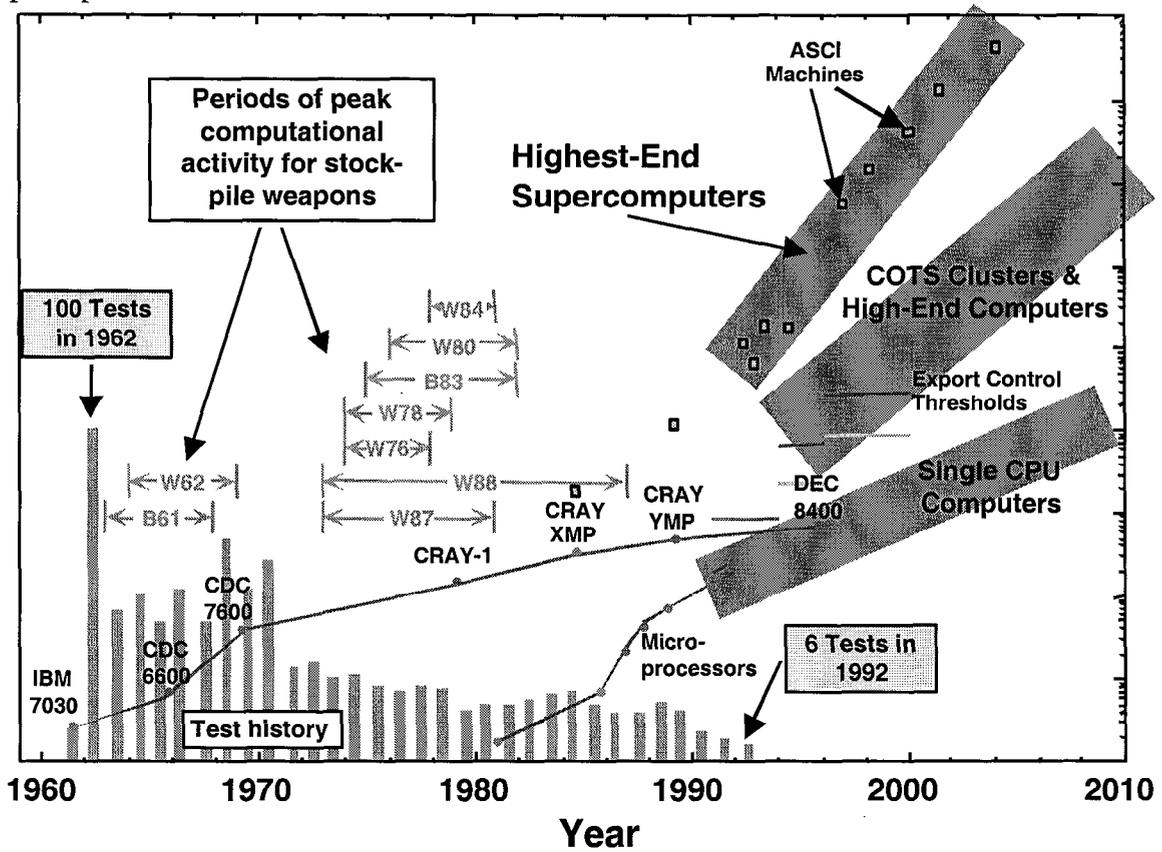


Figure 5
History and Projected Future of Computers Used for Nuclear Weapon Design

Moore's law is expected to continue in effect as a valid predictor of computer power and cost for at least the next decade. With this assumption, computers will continue to double in power at constant cost every eighteen months, or stated in another way, the cost of a given level of computing power will decrease by half every eighteen months. Using the cost/performance of the DOE ASCI White (10 TeraOps) machine as a starting point, the cost of a one TeraFlop (300,000 MTop) computer will fall to one million dollars in mid 2005, and by January 2010, the cost will drop further to only \$134K. COTS clusters today are being constructed for about one third of what IBM charges per MTop, which makes the cost of a COTS TeraFlop even lower. Long before 2010, entry-level ASCI-scale performance will be within the reach of any country in the world. At such a time, the benefits of controlling computer exports will (in the context of nuclear weapon design) be minimal or non-existent. It is not a matter of if, but when - and the when is the current decade.

3) Significance and quality of available foreign computers projected to be available on the open market in the next two years

While LLNL has not carried out an extensive survey of foreign computers available now or in the next two years, we will present two brief examples of relevant systems from Japan, which has the strongest national computer industry after the United States.

The first example is the recent announcement [HIT99] by Hitachi Ltd., a major Japanese computer vendor, of a new line of machines, available in the second half of this year, ranging from "workstations to 32-way servers and cluster systems featuring mainframe technologies" based on Intel's new Itanium chips. The 32-way servers should individually have performance ratings in excess of 100,000 MTOPS. Of particular note is Hitachi's mention of "cluster systems" based on these individually powerful computers. Clustering even a modest number of these 32-way servers together would provide a formidable computing capability. This is a direct result of the effect of Moore's law on computer performance and networking and shows how even systems based on commodity parts are reaching very impressive levels of capability.

The second example is the SX-5 series of High-End supercomputers from NEC [NEC99], another major Japanese computer manufacturer. The members of the SX-5 series are custom-designed vector supercomputers in the same mold as the Cray C-90, but with even greater peak performance. Each individual cpu in the highest performing member of the series (the SX-5/A) has a peak speed of 8 GigaFlops (vs 2GigaFlops for the C-90), with up to 16 CPUs sharing a common, high-speed memory for a total peak speed of 128 GigaFlops. In addition, NEC offers a proprietary, high-speed network to connect up to 32 of these computers for a total peak speed of $32 * 128 \text{ GFlops} = 4 \text{ TeraFlops}$ to form a Highest End supercomputer. This is in the same range as the fastest current ASCI machines (ASCI Blue Pacific at 3.9 TeraFlops). While there are no known actual purchases of the SX-5 configured at or near this capability level, it is clear that NEC is capable of providing near ASCI-class systems.

What these two examples demonstrate is that the US has no monopoly on manufacturers capable of providing high-performance systems based on commodity parts or very high-performance systems based on custom designs.

4) Alternative control mechanisms based on technology or capability that may be more effective to protect national security

The current state of the computer market as represented in Figure 3 taken together with information about the utility of nuclear weapon design simulation activities as detailed in Tables 1 and 2 has strong implications for the availability of computing useful to other nuclear capable countries. Controlling computers has always been a surrogate for controlling sensitive applications, which means that raising export control levels over time reduces the number of important applications which can be effectively controlled. The recent emergence of clustering as a viable method for creating parallel computers from ubiquitous computing and networking technology changes past conclusions in the analysis of the tradeoffs for each of several different strategies for controlling computers.

The nature of the increase in the performance levels of computing which are widely available is due to the large and sudden jump in the number of computers, which can be effectively clustered. Formerly, the performance levels available from computers has increased at a rate that was very predictable (Moore's law), and which was relatively constant when looked at on a time scale extending over ten or twenty years. The sudden arrival of COTS networking coupled with the ability to connect large number of computers means that the peak performance levels which are widely available has taken a large and discontinuous jump higher. After this quantum leap, the rate of future computer performance increases will actually be faster than Moore's law because the individual processors will continue to follow Moore's law at the same time that COTS networking performance advances will increase the number of computers which can be effectively

coupled together. The following sections will each analyze the effect of this COTS networking-induced discontinuity on several possible strategies for controlling computers.

Strategy #1 – Retain export control focus on “finished” vendor-integrated HPCs

This strategy represents the status quo. It continues current procedures for licensing, enforcement, periodically updating assignment of countries to tiers² and updating the performance levels at which controls are applied.

PRO – Would lead to a minimal increase in the performance level of HPCs, which could be legally acquired directly by foreign end-users.

CON – Would not resolve the increasing inconsistency between controlling HPCs but not clusters. Ignoring the existence of clusters leads to a false sense of security. Manufacturers of finished HPCs will be at a disadvantage in their attempts to sell to foreign users with respect to manufacturers of clusters. Controlling the size of the computers available to act as “nodes” of a COTS parallel cluster only marginally controls the size of the parallel cluster that can be created.

Strategy #2 – Seek to control integration technology (i.e., networking)

This strategy seeks to control the number/capability of COTS parallel clusters that can be created by controlling networking hardware (in the form of switches and interface cards). In order to be effective, it implies that networking technology must originate solely in the U.S. or in countries with binding agreements to enforce U.S. export control objectives. Today’s export control regulations do require an individual license when shipping the technology to military end-users in Tier 3 countries, but no license is required when selling to civilian end-users.

PRO – If this were possible, it would be the single most effective means of controlling the number/capability of COTS parallel clusters that could be created. It would slow the increase of peak computational performance widely available in foreign countries back down to the “Moore’s law” rate. The discontinuous increase in the rate of peak performance that is widely available could be undone, and the genie put back into the bottle.

CON – The hardware used in high-speed COTS networking is already widely available and relatively inexpensive throughout the world. Establishing controls at this point would be extremely difficult if not impossible. For example, the Gigabit Ethernet networking hardware specification is the result of an international standardization effort, and is projected to be a three billion-dollar market by the year 2000. Control of such technology will require both the ability to track millions of small, relatively inexpensive objects and monitor a manufacturing capability that is globally distributed. As a second example, the VIA (virtual interface architecture) networking standard was developed and announced in December 1997 as an open standard by Intel, Compaq and Microsoft. The object of the VIA standard was to make clustering of Intel-based servers commonplace in competition with more traditional HPCs in the high-volume commercial computing marketplace. The scientific computing market has always been a small appendage to commercial computing market, which means that restricting the availability of COTS networking to control scientific computing performance levels will directly affect the much larger commercial computing market.

Strategy #3 – Acknowledge clusters and raise limits

This strategy seeks to reset the scale at which computers are controlled to roughly match that which can be effectively provided by COTS clusters to important applications. Today

² The word “tiers” refers to a classification scheme in which all of the countries of the world are placed into four different groups for the purpose of exercising different levels of control on the export of sensitive technology. Tier 1 has the highest privileges, and tier 4 has the lowest. Russia, China, India, Pakistan and Israel are all classified as tier 3 countries.

this limit is around 50,000 to 60,000 MTOPS. Periodic review of the export control levels would not simply track Moore's law, but rather the combined effects of Moore's law for individual computers combined with the increasing ability of COTS networking to link larger numbers of computers together.

PRO – Today's large inconsistency between the levels at which HPCs are controlled and the performance levels achievable by COTS parallel clusters would be eliminated. The true state of U.S. exposure to changing national security threats will be acknowledged and addressed. Applying export controls at the current peak level of single COTS parallel clusters would still provide control far below the level of ASCI-scale computers which the U.S. DOE Stockpile Stewardship Program has identified as essential for the quality of simulation needed during an era of no nuclear testing. Export control resources would be focused on computers that truly matter. (It takes on the order of 100 High-End supercomputers to create an ASCI-scale supercomputer. The COTS networking which is effective at connecting COTS desktop PCs and workstations into COTS clusters is very much less effective in connecting together computers which each have a CTP of 50,000 MTOPS.)

CON – Appropriately responding to the emerging realization of the many implications of the actual availability of high-performance computing will require extensive review and discussion in the national security community. Political accusations of giving away high-performance computing are virtually certain.

Comment on split military/commercial limits within a country

The computer export control performance levels for Tier 3 countries are set separately for military and civilian end users. Specifically, computer manufacturers may sell computers up to 12,300 MTOPS under a general license (i.e., no prior governmental review) to civilian end users, while the limit for military end users is only 6,500 MTOPS.

PRO – Such a dual threshold offers the promise of more stringent control of military activities while simultaneously allowing computer manufacturers greater opportunities to sell commercial computers.

CON – The actual effectiveness of a dual threshold in limiting the availability of computers to military end users to the lower level is doubtful, and must be carefully reviewed. Eliminating the separate military threshold and using the commercial threshold for military end users may more closely correspond to the true availability for military end users in tier 3 countries.

5) Issues related to other critical technologies

The DOE has developed a list [DOE99] of areas of technical subject matter or technologies containing "sensitive" information. The list identifies subjects related to the development and production of weapons of mass destruction (nuclear, chemical, and biological) and their delivery systems (including missiles), conventional weapons, and other technologies deemed significant to the national security of the United States. This list will be provided to the SASC.

References

- [APF96] Accelerated Strategic Computing Initiative: PathForward Project Description, December 27, 1996, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Sandia National Laboratory, 1996.
- [BW] www.beowulf.org
- [DOE99] Department of Energy (DOE) Notice 142.1, July 14, 1999
- [GLS] “Using MPI – Portable Parallel Programming with the Message-Passing Interface”, W. Gropp, E. Lusk, and A. Skjellum, MIT Press, 1994.
- [GWB95] Goodman, S. E., P. Wolcott, and G. Burkhart, *Building on the Basics: An Examination of High-Performance Computing Export Control Policy in the 1990s*, Center for International Security and Arms Control, Stanford University, Stanford, CA, 1995.
- [GWH98] “High Performance Computing, National Security Applications, and Export Control Policy at the Close of the 20th Century”, Seymour E. Goodman, Peter Wolcott, and Patrick Homer, published by the Institute for International Studies, Stanford University, 1998.
- [HIT99] <http://www.idg.net/idgns/2000/01/14/HitachiDetailsRolloutPlansForIA64.shtml>
- [NEC99] <http://www.tgc.com/sponsors/nec/930.html>
- [SM88] “Seismic Verification of Nuclear Testing Treaties”, OTA ISC 361, 1988.
- [SS98] Unclassified extracts from a classified report “Chinese Computational Capabilities for Supporting Stockpile Maintenance(U)”, Steven J. Sackett, COVD-99-0214, S/NF/Sigma 2, November 1998.