

**CRAY**  
**RESEARCH**  
*A Silicon Graphics Company*

# A Comparison of Application Performance Across Cray Product Lines

CUG San Jose, May 1997

---

[R. Kent Koeninger](#)

*Software Division of Cray Research, a Silicon Graphics Company  
655 F Lone Oak Drive, Eagan MN 55121*

[kentk@cray.com](mailto:kentk@cray.com)

[www.cray.com](http://www.cray.com)

---

## **ABSTRACT:**

This paper will compare standard benchmark and specific application performance across the CRAY T90, CRAY J90, CRAY T3E, and Origin 2000 product lines.

## **KEYWORDS:**

Application performance, benchmarks, CRAY T90, CRAY J90, CRAY T3E, Origin 2000, LINPACK, NAS Parallel Benchmarks, Streams, STAR-CD, Gaussian94, PAM-CRASH.

---

## **Introduction**

The current product offerings from Cray Research are the CRAY T90, CRAY J90se, CRAY T3E 900, and Cray Origin 2000 systems. Each has advantages that are application dependent. This paper will give comparisons that help classify which applications are best suited for which products.

Most performance comparisons in this paper are measured by Origin 2000 (O2K) processor equivalents. The run-time on the platform in question is divided by the run time on a single processor of an O2K processor. With this technique, larger numbers indicate better performance.

The examples will show that there is no one product best suited for all applications. Some run better with the extremely high memory bandwidth of the CRAY T90 systems, some run better with the high scalability and excellent interprocessor latency of the CRAY T3E systems, some run better with the large caches and SMP


scalability of the Origin systems, and some run better with the good vector price-performance throughput of CRAY J90 systems. By looking at which applications run best on which platforms, one can get a feeling for which platform might best suit other applications.

---

## Product Line Overview

---

### **CRAY T90 Strengths**




The CRAY T90 has extremely high memory bandwidth (880 GB/s peak) with great vector, gather, and random access to memory. It offers fast individual processors (1.8 GFLOPS peak) with very high bandwidth IO (Gbytes/second). Its UNICOS operating system is robust and provides excellent multijob throughput. It has proven high-performance for parallel applications, up to 32 processors.

Cray is designing another generation of CRAY T90 computers that will be faster and will have better price/performance with plug-compatible CPUs.

Note the extremely high memory bandwidth in the following STREAMS benchmark graph.

### **STREAMS Benchmark**




Complex NASTRAN analyses are good examples of applications suited for CRAY T90 performance. For example, the CRAY T90 is the best machine for complex full-automobile acoustic-fluid-vibration analysis (NVH). It can turn around a full analysis (a half-dozen iterations) in one day. This type of analysis is complex in that it couples energy from the car's various metal, plastic, and other parts with the air in the car to determine the noise level at the driver's and passenger's ears.

The models tend to run for 3.5 hours with 1.5 million degrees of freedom and 1000+ modes on CRAY T90 systems. This requires terabytes of IO to about 50 GB of disk. The runs are extremely memory and IO intensive, sustaining 4 GB/s of memory bandwidth over the multiple-hour runs.

This problem will not complete in a single day on systems that do not have the high CRAY T90 memory bandwidth and the fast CRAY T90 IO.

## CRAY J90se Strengths




CRAY J90se systems are compatible with CRAY Y-MP, CRAY C90, and CRAY T90 applications. CRAY J90se systems run the UNICOS operating system, which offers good system throughput and robust production capabilities. CRAY J90 systems have high memory and IO bandwidth and better price/performance than CRAY T90 systems but at a lower per-processor performance.

Single and multiple-PE applications are sometimes faster and sometimes slower on CRAY J90se systems than on Origin 2000 systems.

---

## CRAY T3E 900 Strengths



The CRAY T3E 900 systems are the most scalable solution sold today. They have a very fast interconnect and fast microprocessors which makes them a good choice for 32 to 1024 processor distributed-memory codes. They have fast global IO (gigabytes per second) and are typically programmed with message passing (PVM, MPI, or SHMEM).

Individual CRAY T3E 900 processors are sometimes faster and sometimes slower than Origin 2000 processors. CRAY T3E applications scale to 100s of PEs better than on Origin 2000 systems (today).

The **HIRLAM Weather Code** is a good example of scalable CRAY T3E performance. It is a high resolution limited area model from a Scandinavian Consortium (plus Irish and Spanish contributions). It is a derivative of an ECMWF weather model. It runs at 11.4 GFLOPS on 192 PEs of mixed speed (128 @ 375 MHz + 64 @ 300 MHz). We expect to run this code on over a 1000 processors on a CRAY T3E in Chippewa Falls this summer.

---

## Cray Origin2000 Strengths



Cray Origin2000 individual processor performance and scaling are comparable to CRAY J90se and CRAY T3E 900 speeds (up to 32 PEs).

Cray Origin2000 scales to 32 PEs today. Cray will beta test 64 processor systems in 2QCY97 and will soon release the OS for a 128 PE Single System Image. The Origin product line will eventually scale to 4000+ PEs with Cellular IRIX.

The Origin systems can be programmed with SMP or MPP styles. Codes are scaling very well on Origin2000 using SMP-parallel and message-passing techniques. An example message-passing applications is a **PAM CRASH** analysis of a train crash. This model ran for 40 hours on 32 Origin PEs (using PVM). Compare this

with a 50 hour run on 4 CRAY T90 CPUs (using SMP parallelism). Another example is *MPP-DYNA* which shows excellent speed on 32 Origin PEs (similar to CRAY T3E performance).

### **Cray Origin2000 Scalable Supercomputing**

The *ASCI Blue Mountain* initial system is a prime example of how Cray will scale the Origin 2000 systems to hundreds of PEs in the next two years. Cray delivered an Origin 2000 system to Los Alamos in 1997 and will grow this system to 400+ GigaFLOPS (peak) on 768 processors in 1998, using IRIX Array (cluster) techniques. This initial system will have 2.5 TB of disk storage.

In 1999, Cray will augment this system by delivering a 4 TeraFLOPS peak, 1 TeraFLOPS sustained system to the *ASCI Blue Mountain* and *LANL ACL* programs. This combined system will have over 4000 processors with over 75 TB of disks storage on over 1000 disk channels.

---


## **Standard Benchmarks**

---

### **LINPACK**

LINPACK shows near-peak processor performance characteristics and is a standard benchmark on scientific computers. Note, a 128 PE T3E has equivalent peak performance to a 32 processor CRAY T90 system. Per processor, the CRAY T3E performance is higher than the Origin 2000 performance, which is higher than the CRAY J90 performance. These relationships will vary on other applications.


#### **LINPACK Performance on Cray Systems**



## STREAMS

The following benchmark compares the memory bandwidth of the CRAY T3E, CRAY J90, and Origin 2000 systems. The CRAY T90 bandwidth shown in the earlier STREAMS chart was omitted here to allow an easier comparison among these three systems.


### Memory Bandwidth Measured by STREAMS



## NAS Parallel Benchmarks

The NAS Parallel Benchmarks are a standard set of kernels and pseudo-applications that measure common parallel algorithms. The chart below compares the performance on the pseudo-applications. Note the similar 32 PE performance between the CRAY T3E and Origin 2000 systems. The CRAY T3E scales these applications to the highest numbers of processors of any system available today.

### NPB 1.0 Class A Pseudo-Applications




---

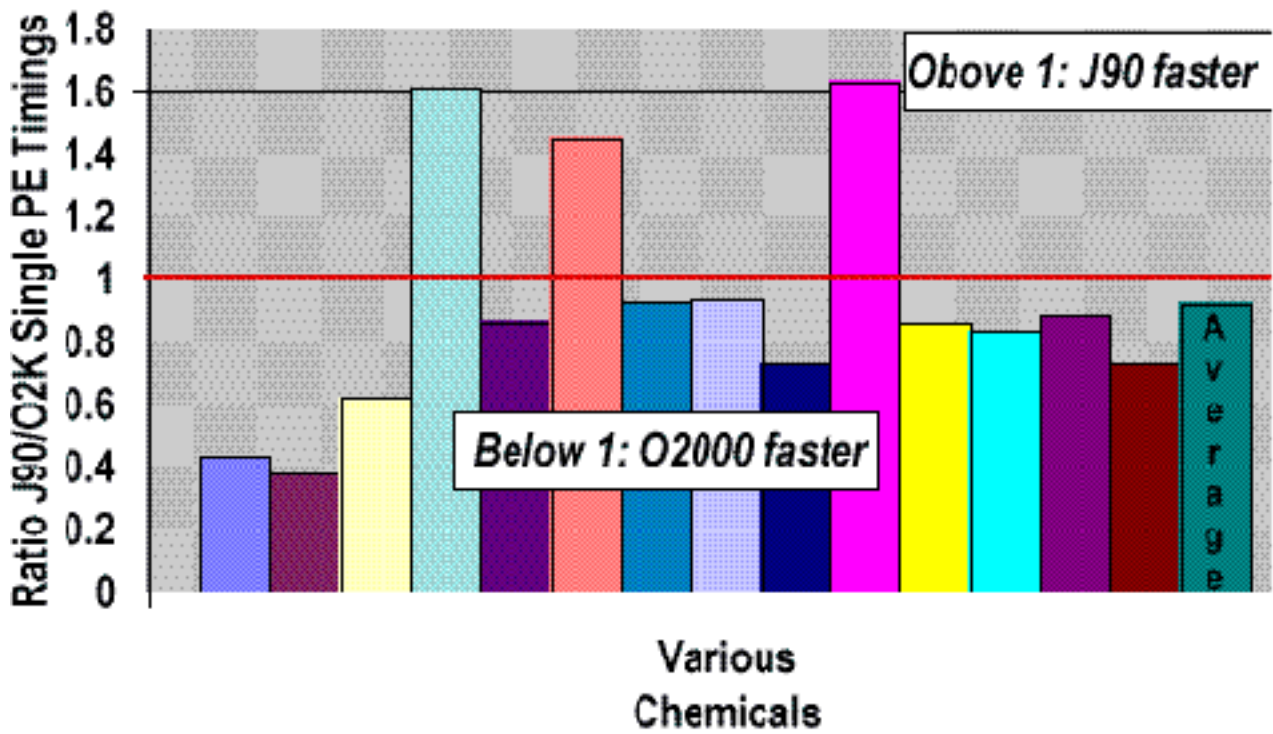
## Application Benchmarks

### Gaussian94

The following chart shows CRAY J90 and Origin 2000 single processor timings for various chemicals. The relative performance varied according to the chemical structure under analysis, with some vectorizing better and some caching better. An average of these of 14 chemicals shows the Origin 2000 was 8% faster than the CRAY J90 on these particular Gaussian94 runs.

#### Gaussian94 on 14 Various Chemicals

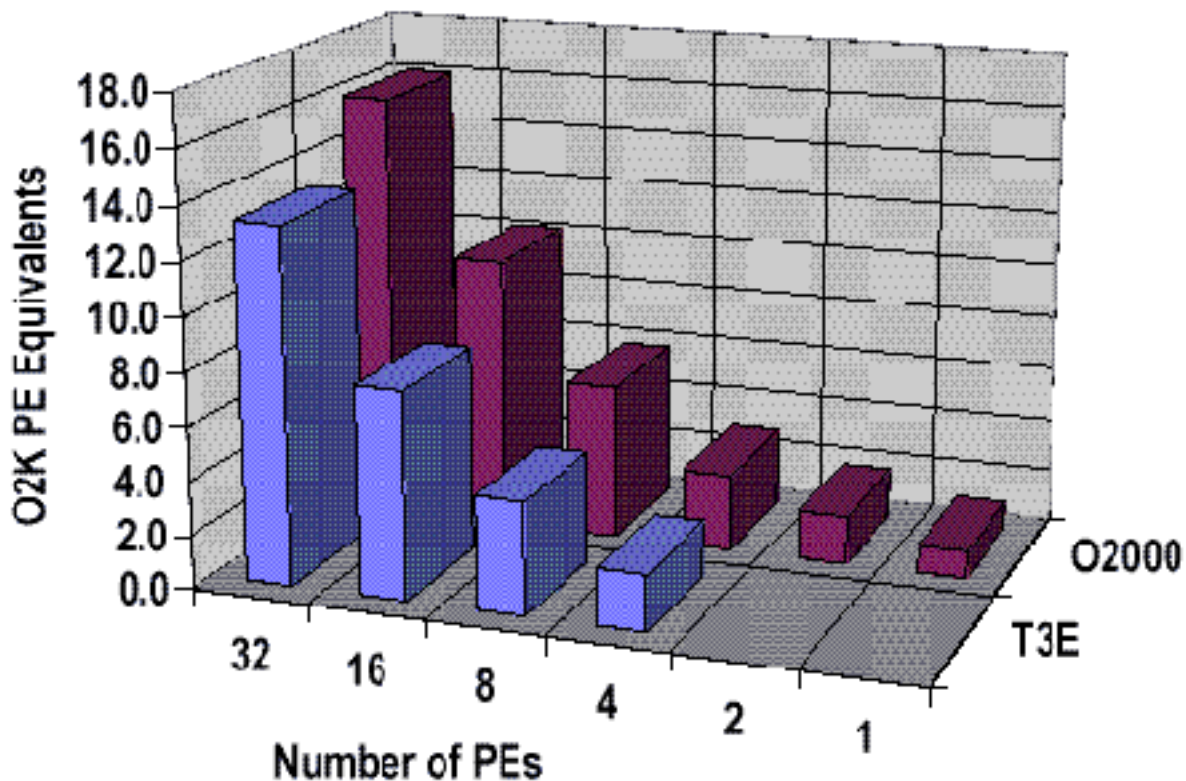




## MPP-DYNA

The following chart shows that performance of the message-passing version of DYNA is similar on CRAY T3E and Origin systems, up to 32 processors.

### MPP-DYNA




---

## NASTRAN Ford "Faster" Model

This section gives more detail on the NASTRAN MHV model mentioned earlier in the CRAY T90 section.

The Ford "Faster" model is a complex MSC/NASTRAN V69 model that runs well on on a T94 system. It uses 1.5 million degrees of freedom for the car body model and 1073 modes (eigenvalues) under 200 Hz. It ran dedicated using 3 CPUs, 90 MW of memory, 450 MW of SSD, and 45 GB disk. It moved 1.2 TB of IO to the I/O cache in the SSD and 520 GB disk IO through the SSD back door. In other words, the 3.5 GB IO cache cut the disk IO in half. The average resource use (over 3.5 hour run time) was 4 GB/sec of memory bandwidth and 100 MB/sec of IO bandwidth. It accumulated 30 minutes IO wait time.

On a CRAY J90 system, "Faster" model runs about 1/7 the speed on a CRAY T90 system. It would not complete a full set of a half-dozen analyses in one day. The complex "Faster" model is too memory and IO intensive for Origin 2000 today, but simpler static NASTRAN models run well on Origin systems today.

---

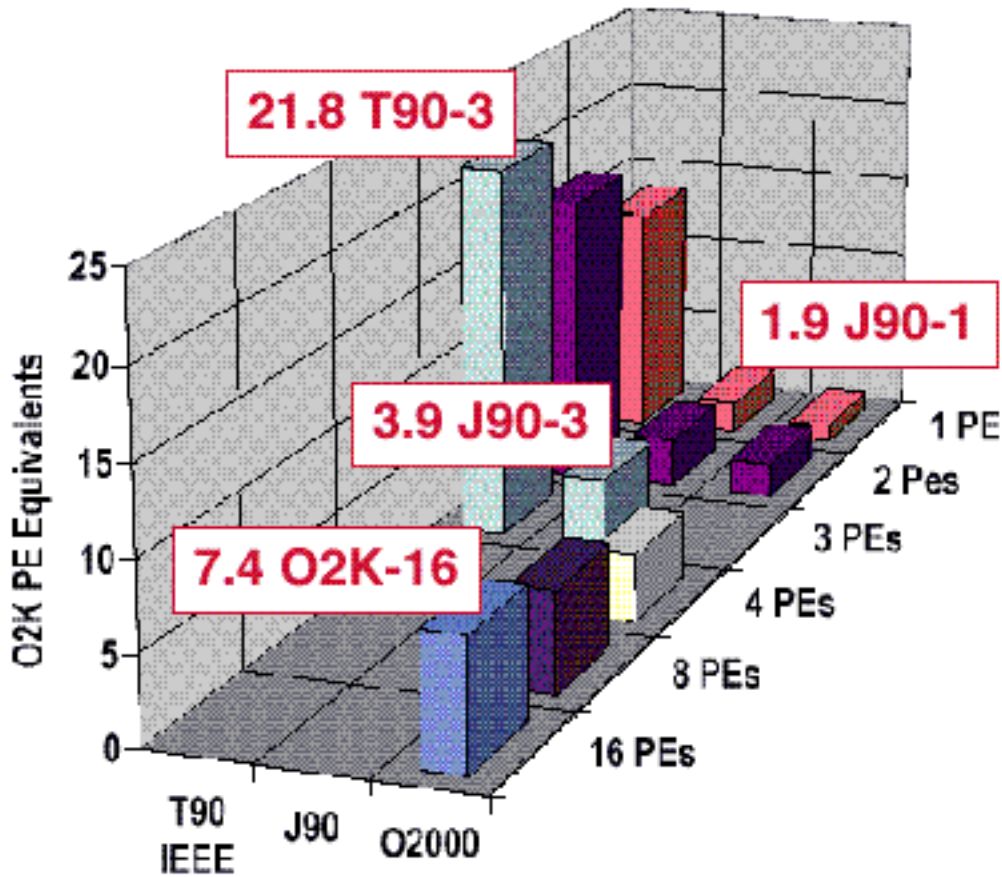
## STAR-CD

The following chart shows an engine case analyses using STAR-CD with 20 iterations on CRAY T90, CRAY J90 and Origin 2000 system. The CRAY T3E is not included because this code is as an SMP-parallel code. This example shows the complex mix of individual and scaled processor performance characteristics.

The one to three processor CRAY T90 timings are two to three times faster than the most parallel Origin runs. The individual CRAY J90 processor speed is about twice as fast as the Origin speed. The scaled Origin

performance is about twice as fast as the scaled CRAY J90 performance.

### *STAR-CD Engine Case Analysis*



---

## Summary

Currently the CRAY T90, CRAY J90, CRAY T3E, and Cray Origin 2000 systems each have superior features which are application dependent. The CRAY T90 systems work best on high-memory bandwidth, high-IO vectorized code. The CRAY J90 systems run similar programs with better price-performance but with lower overall performance. The CRAY T3E systems scale programs to very large aggregate performance. The Cray Origin 2000 systems are just beginning to scale and today provide high performance for SMP and MPP style codes.

The CRAY T90 processors are by far the fastest individual processors of this group but are also the most expensive. The individual CRAY J90, CRAY T3E, and Origin processors offer about the same performance and the ratio of their performance varies according to the application characteristics of vectorization, strided memory access, cache reuse, and other characteristics.

Faster generations of all four of these platforms are in the works so all four of these families will remain competitive for years until they are gradually replaced by future generations of the Origin systems.

---

## Acknowledgments

The author wishes to thank Guangye Li, Steve Behling, Martin Feyereisen, Jeff Brooks, Jeff Zais, C. P. Sosa, John Carpenter, Cheng Liao, Charles Grassl, and all the other people that worked hard to optimize the benchmarks presented in this paper.

---

## Author Biography

R. Kent Koeninger is a Program Manager in the Software Division of Cray Research, a Silicon Graphics Company. Mr. Koeninger has 20 years experience in high-performance computing, including 5 years as a Program Manager at Cray for scalable systems and high-performance IO.

[© Silicon Graphics Inc. and Cray Research Inc. 1997](#)

[Copyright and trademark statement](#)