



Performance Brief

January 1999
ECG067/0199

Prepared by Workstations

Compaq Computer Corporation

Contents

Benchmarks and Performance3

 Benchmarks3

 Application Benchmarks.....4

 Benchmark Biases4

SPECint and SPECfp.....6

Bench98.....9

ANSYS11

Linpack.....13

Viewperf.....15

Compaq Professional Workstation XP1000 Benchmark Results

Abstract: This paper presents benchmark information for the XP1000. An overview and description of each benchmark is provided, along with the measured performance on that benchmark.

Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, registered United States Patent and Trademark Office.

Microsoft, Windows, and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

Copyright ©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Compaq Professional Workstation XP1000 Benchmark Results
Performance Brief prepared by Workstations

First Edition (January 1999)
Document Number ECG067/0199

Benchmarks and Performance

Computer hardware performance is measured in terms of the maximum rate the system can achieve in executing instructions. The most common measures are Millions of Instructions Per Second (MIPS) and Millions of Floating point Operations Per Second (MFLOPS).

Hardware specifications are of limited value, as they only address maximum theoretical performance and do not measure realistic system performance. Other factors in the workstation design, such as memory bandwidth and latency and I/O performance, often limit a system to a small portion of its theoretical performance.

Additionally, hardware performance measurements and system architectures are not standardized, making it difficult to directly compare vendor provided specifications.

Because of this, there has been a strong migration toward measuring actual performance with *benchmarks*.

Benchmarks

Another approach to performance measurement is to develop a program and to measure the performance of the system running this program. This has the advantage of measuring actual performance that can be achieved on the system. If the program is designed to be portable between systems it can also allow direct comparisons between systems. For this reason, organizations, such as SPEC (Standard Performance Evaluation Corporation) and BAPCo (Business Applications Performance Corp.), and academic/research institutions, such as CERN (European Particle Physics Laboratory), have developed and made tools available to provide cross-platform tests that can help users compare the performance of different platforms. We will refer to these tests as industry-standard benchmarks. Other tests, based on specific applications and test sets, we will refer to as "application-specific" benchmarks. These application-specific tests can be either widely-accepted and virtual industry standards, such as the Bench98 Pro/ENGINEER sponsored by *Pro/E: The Magazine*, or simply a comparison of vendor or OEM supplied test files with a given application. We will briefly discuss the industry-standard benchmarks first, and then cover some of the leading application benchmark tools used.

An objection sometimes raised about industry standard benchmarks is that they do not measure total system performance, but focus on the performance of a single subsystem, such as CPU, memory or graphics. This objection is correct but misleading. While most industry standard benchmarks *are* primarily intended to measure subsystem performance, they can be effectively used *in conjunction* with other measurements to determine overall system performance. Few dispute that industry standard benchmarks do an excellent job at what they are intended – measuring the performance that a system (both hardware and software) can actually achieve.

Several benchmarks exist for measuring CPU performance, including SPECint, SPECfp, and Linpack. Other standard benchmarks are used to measure graphics performance; examples include Viewperf and Indy3D.

Application Benchmarks

Application benchmarks can be at the same time useful, misleading, and a difficult way to measure performance. Since workstations are used to run applications, the only *real* metric for performance is application performance.

Application benchmarks are misleading because they are only valid for that specific application, and sometimes for a specific data set or test file. Even directly competitive applications that do essentially the same things may be dramatically different. An excellent example of this comes from CAD/CAM, using two of the leading applications: Pro/ENGINEER and Unigraphics. Although similar in many ways, these two applications use different graphics approaches, utilize very different parts of OpenGL, and have radically different approaches to implementing user interfaces. Because of this, performance optimization for each application is different. Good performance on one does not necessarily indicate good performance on the other. Likewise, poor performance on one doesn't necessarily indicate poor performance on the other.

Also, different uses (as represented here by data sets or trail files) exercise different parts of the application and may use totally different features and functions.

Application benchmarking requires much work. There are very few broad based, comprehensive benchmarks that span a range of systems and allow easy comparisons. A notable example is Bench98, the benchmark for Pro/ENGINEER reported by *Pro/E: The Magazine*. This benchmark is run across a wide range of systems and allows easy comparison. But even here, graphics is a small part of the benchmark, making it necessary to look further to understand graphics capabilities and differences between systems. Other vendors, such as NewTek (Lightwave), Unigraphics, and Softimage provide test files with their applications that allow some ability to perform system comparisons, but in each instance, the comparison is rather ad-hoc, and requires strict attention to details to ensure consistent test methodologies. As such, interpretation of results is often subject to debate and question over possible bias in the results.

Benchmark Biases

All benchmarks are biased. Understanding this fact is critical to effective use of benchmark data. Biased doesn't mean misleading, it simply means that you need to understand *what* the benchmark is measuring, what systems are being compared, *how* they are being measured, and how the *results* are being used. The bias may be subtle or overt, the benchmark may be well designed or poorly designed, the characteristics being measured may be crucial or irrelevant, and the testing methodology may be valid or flawed. Education on the details of the benchmark is the only solution to guide one through the landmines of potential bias.

Good benchmarks are difficult to design. A good benchmark is one that provides a true indicator of the performance for which the system and the application were designed. Developing a benchmark that provides a true indicator of actual performance is: broad based, portable across different hardware and operating systems, easily run, easily reported, and easily interpreted. This is not a simple task.

An additional challenge arises when a benchmark becomes popular, and vendors begin *optimizing for the benchmark* –changing their system to provide higher performance on the benchmark while not improving application performance. This occurs with all popular benchmarks; in the graphics space, the Viewperf CDRS benchmark is increasingly popular, and many vendors are reporting performance results on this benchmark that do not reflect application performance.

To summarize, benchmarks are a tool –a tool that can be used or misused. Well-designed benchmarks can provide valuable insights into performance. Poorly designed benchmarks may be highly inaccurate and misleading. And no single figure can capture all the information that is needed for a well-chosen system selection.

The following pages provide more information on the most popular workstation benchmarks used today and recent performance figures for the Compaq Professional Workstation XP1000 and competitive systems to help put the benchmark in perspective.

SPECint and SPECfp

Benchmark

SPEC CPU benchmark suite, with results for SPECint95 and SPECfp95.

Source

SPEC, the Standard Performance Evaluation Corporation, is a non-profit corporation formed to "establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers" (quoted from SPEC's bylaws). The founders of this organization believe that the user community will benefit greatly from an objective series of applications-oriented tests, which can serve as common reference points and be considered during the evaluation process. While no one benchmark can fully characterize overall system performance, the results of a variety of realistic benchmarks can give valuable insight into expected real performance.

SPEC basically performs two functions.

- SPEC develops suites of benchmarks intended to measure computer performance. These suites are packaged with source code and tools and are extensively tested for portability before release. They are available to the public for a fee covering development and administrative costs. By license agreement, SPEC members and customers agree to run and report results as specified in each benchmark suite's documentation.
- SPEC publishes news and benchmark results in *The SPEC Newsletter* and *The GPC Quarterly*. Both are available electronically through <http://www.spec.org/> (and available in paper as a quarterly publication if necessary). This provides a centralized source of information for SPEC benchmark results. Both SPEC members and non-SPEC members may publish in the SPEC Newsletter, though there is a fee for non-members. (Note that results may be published elsewhere as long as the format specified in the SPEC Run Rules and Reporting Rules is followed.)

Description

In August 1995, SPEC introduced the [CPU95 benchmarks](#) as a replacement for the older [CPU92 benchmarks](#). These benchmarks measure the performance of CPU, memory system, and compiler code generation. They normally use UNIX as the portability vehicle, but they have been ported to other operating systems as well. The percentage of time spent in operating system and I/O functions is generally negligible.

Although the CPU95 benchmarks are sold in one package, they are internally composed of two collections:

- CINT95, integer programs, representing the CPU-intensive part of system or commercial application programs
- CFP95, floating-point programs, representing the CPU-intensive part of numeric-scientific application programs

Results are reported for CINT95 and for CFP95 on individual report pages, results can be reported for either one or for both.

Integer benchmarks: CINT95

This suite contains eight benchmarks performing integer computations, all of them written in C. The individual programs are:

Number and name	Application
099.go	Artificial intelligence; plays the game of "Go"
124.m88ksim	Moto 88K chip simulator; runs test program
126.gcc	New version of GCC; builds SPARC code
129.compress	Compresses and decompresses file in memory
130.li	LISP interpreter
132.jpeg	Graphic compression and decompression
134.perl	Manipulates strings (anagrams) and prime numbers in Perl
147.vortex	A database program

Floating-point benchmarks: CFP95

This suite contains 10 benchmarks performing floating-point computations. All of them are written in Fortran77. The individual programs are:

Number and name	Application
101.tomcatv	A mesh-generation program
102.swim	Shallow water model with 513 x 513 grid
103.su2cor	Quantum physics; Monte Carlo simulation
104.hydro2d	Astrophysics; Hydrodynamical Navier Stokes equations
107.mgrid	Multi-grid solver in 3D potential field
110.applu	Parabolic/elliptic partial differential equations
125.turb3d	Simulates isotropic, homogeneous turbulence in a cube
141.apsi	Solves problems regarding temperature, wind, velocity and distribution of pollutants
145.fpppp	Quantum chemistry
146.wave5	Plasma physics; electromagnetic particle simulation

What is Measured

Being compute-intensive benchmarks, these benchmarks emphasize the performance of the computer's processor, memory architecture, and compiler. It is important to remember the contribution of the latter two components; performance is more than just the processor.

Units


Larger is better. The results ("SPEC Ratio" for each individual benchmark) are expressed as the ratio of the wall clock time to execute one single copy of the benchmark, compared to a fixed "SPEC reference time." For the CPU95 benchmarks, a Sun SPARCstation 10/40 was chosen as the reference machine. A geometric mean of all benchmarks is used, which means that each benchmark contributes equally to the final result.

Additional Information on Benchmark

More information is available from <http://www.specbench.org>.

Compaq Results:

As of January 1999, the SPECint/SPECfp benchmark results are as follows:



Sources:

Industry standard benchmark, available from the Standard Performance Evaluation Corporation (SPEC). More information available at <http://www.specbench.org>. XP1000 performance, as measured by Compaq and reported to SPEC in accordance with their standard submission process, have not yet been published by SPEC and must therefore be denoted as estimates. Competitive system performance numbers obtained from the SPEC reports contained on their Web site.

Notes

These are industry standard processor benchmarks that are run across a wide range of systems. This test clearly demonstrates the impressive integer instruction and outstanding floating point performance of the Alpha 21264 processor and the added benefit of the high-bandwidth memory architecture of the Compaq Professional Workstation XP1000.

Bench98

Benchmark

Bench98 and Gbench98

Source

Pro/E: The Magazine.

Description

Bench98 is an application benchmark, using the Parametric Technology Corporation (PTC) Pro/ENGINEER (Pro/E) software and platform neutral script files. There are several versions of the benchmark, with Bench98 being the latest.

Pro/ENGINEER is a leading high-end mechanical CAD application. It uses state of the art technology, including solid modeling and parametric design. Pro/ENGINEER is widely used for sophisticated design work and places a strong load on the system.

Bench98 was developed by analyzing actual design engineers use of Pro/E in engineering projects, recording the most frequently used operations, manipulations, and commands. The engineers' actions were then used to develop a script file that reproduced the users' actions. This script file was written using the Pro/E scripting language and can be run on any hardware platform that supports Pro/E.

The script runs at full speed with no user delays. Thus, the system executes the Pro/E operations as quickly as it can.

What is Measured

Bench98 measures the execution of a wide range of Pro/E operations; consequently, it is a good indication of overall Pro/E performance. The results are for a "weighted average." An engineer may use a specific function 100 times in a day. The benchmark runs this function once, and then weighs the results to correspond to actual usage. This benchmark provides two "scores" a summary weighted average total time of operation and a weighted total time spent on graphics operations, called Gbench98.


Units

Units are elapsed time in minutes, where smaller is better.

Additional Information on Benchmark Summary

More information is available from <http://www.proe.com>. The benchmark results are not posted on this Web site; it is necessary to order reprints to get the full set of results. Ordering information for the reprints is included at this site.

Results



Sources

The benchmark is available from *Pro/E: The Magazine* through their Web site at <http://www.proe.com>. XP1000 and DPW600 performance measurements made by Compaq, using the SPIKE post-compilation optimizer on the Pro/ENGINEER software. Competitive performance figures from *Pro/E: The Magazine* 1998 Workstation Benchmark Supplement.

Notes

The Compaq Professional Workstation XP1000 continues the Compaq/DIGITAL legacy of excellent performance with Pro/ENGINEER. Pro/ENGINEER has been optimized for the Alpha platform. The strong floating point performance of Alpha allows the Compaq Professional Workstation XP1000 to excel at operations such as hidden line removal, transparencies, and rotations. Despite the lack of graphics controllers with on-board geometry acceleration available in more costly competitors, Compaq offers the best workstation performance for Pro/E users.

ANSYS

Benchmark

Suite of test cases provided by ANSYS.

Source

ANSYS, Inc. ANSYS provides a suite of applications for Computer Aided Engineering (CAE). ANSYS is best known for their Finite Element Analysis software, which is widely used for solving large, complex analysis problems.

Description

ANSYS provides a suite of test cases representing a variety of engineering problems. Each of these tests are run, and elapsed time (total time required to complete the job) is reported.

Following is a brief description of each of the test cases. These cases involve stress and thermal analysis of mechanical parts. A combination of small and medium size jobs are run to allow reasonable benchmarking times on a range of systems.

Name	Description	Degrees of Freedom
slss-ps	structural linear static shell PCG small	23,658
slss-pm	structural linear static shell PCG medium	71,472
sldb-pm	structural linear static brick PCG medium	73,575
sms-bls	structural modal shell Block Lanczos small	23,658
sms-blm	structural modal shell Block Lanczos medium	73,575
ttns-ps	thermal transient nonlinear shell PCG small	11,912
ttnb-pm	thermal transient nonlinear brick PCG medium	53,405
ttns-fs	thermal transient nonlinear shell frontal small	11,912
ttnb-fm	thermal transient nonlinear brick frontal small	53,405
slss-fs	structural linear static shell frontal small	23,658
sldb-fm	structural linear static brick frontal medium	73,575
ttns-ss	thermal transient nonlinear shell sparse small	11,912
ttnb-sm	thermal transient nonlinear brick sparse medium	53,405
slss-ss	structural linear static shell sparse small	23,658
slss-sm	structural linear static shell sparse medium	71,472
sldb-sm	structural linear static brick sparse medium	73,575

What is Measured

ANSYS stresses computational power, floating point performance, and memory bandwidth. For this report, the elapsed times for the test cases described above are added together to produce a single result covering all of the test cases.

For consistency with other reports from other vendors, 14 of the 16 test cases are used.

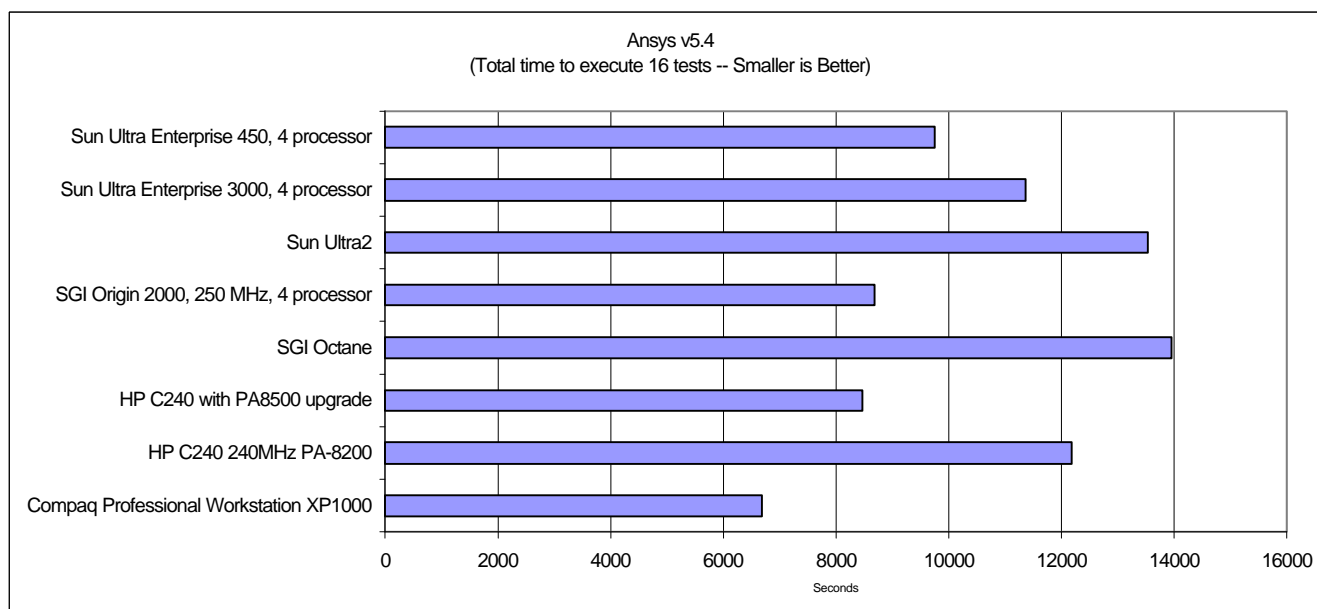
Units

Units are elapsed time in seconds. Lower is better.

Additional Information on Benchmark Summary

Information on ANSYS is available at <http://www.ansys.com>. The benchmark report is at <http://www.ansys.com/ServSupp/Hardware/bm54.html>.

Results



Sources

The benchmark suite is available from ANSYS. XP1000 performance is measured by Compaq. Competitive performance can be obtained from the ANSYS Web site at <http://www.ansys.com/ServSupp/Hardware/bm54.html>.

Notes

The ANSYS Web site does not always provide comprehensive details about system configurations. For this reason, CPU type and/or frequency is not always available to provide background detail for thorough comparison.

Linpack

Benchmark

Linpack

Source

The Linpack software and benchmark are available from <http://www.netlib.org>. Look in the Linpack and Lapack areas. Most of this software was developed under government grants and is either in the public domain or is freely distributed.

Description

Linpack (Linear Algebra Package) is a library of mathematical subroutines focused on linear algebra and matrix manipulations. These functions are at the heart of the solutions to many engineering problems. Since the mathematical functions are standard, many people prefer to use already developed solutions rather than develop their own. One of the most commonly used packages is Linpack. Researchers in universities and government labs developed Linpack and it is available in source form as well as in versions that have been highly optimized for many machine types and architectures.

The Linpack benchmarks are a set of standard problems that utilize the Linpack subroutine library.

What is Measured

The Linpack benchmark measures floating point performance and memory bandwidth. The Linpack benchmark was introduced by Jack Dongarra. A detailed description, as well as a list of performance results on a wide variety of machines, is available in postscript form from *netlib*. To retrieve a copy, send electronic mail to netlib@ornl.gov and by typing the message *send performance from benchmark* or from any machine on the internet type: *rcp anon@netlib2.cs.utk.edu:benchmark/performance performance*.

The benchmark used in the Linpack benchmark is to solve a dense system of linear equations. This performance does not reflect the *overall performance* of a given system, as no single number ever can. It does, however, reflect the *performance of a dedicated system for solving a dense system of linear equations*. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good correction of peak performance.

Two benchmarks are commonly reported: Linpack 100x100 and Linpack 1000x1000. As might be guessed, the 100x100 benchmark uses a 100x100 matrix, and the 1000x1000 uses a matrix of that size.

The Linpack 100x100 benchmark is a very small benchmark by technical computing standards. It is so small that it is difficult to run with high performance or multi-processor systems. Good performance on the 100x100 benchmark requires very low latency through the memory subsystem, floating point unit, and cache subsystem. While it doesn't produce high results, it is a good indicator of performance that can be achieved with many applications.


The Linpack 1000x1000 benchmark is large enough to be "interesting" for high performance systems. It is a large, regular problem, which minimizes the impact of latency and allows good results on multi-processor systems. The 1000x1000 benchmark measures peak floating point performance and memory bandwidth.

While the 100x100 and 1000x1000 benchmarks perform the same kinds of computations and are reported in the same units (MFLOPS), they stress the system differently and report on different aspects of system performance. In most applications of interest to workstation class users, the 100x100 results are the most useful.

Units

Units are MFLOPS (Millions of Floating Point Operations per Second), where larger is better.

Results



Sources

The Linpack benchmark and results for many systems are available from <http://www.netlib.org>. Look in the Linpack and Lapack areas.

Notes

The XP1000 delivers superb performance on both the 100x100 and 1000x1000 benchmarks. The results, especially the 100x100 case, correlate well with application results, such as ANSYS. One might expect as much since the strong floating point performance and high memory bandwidth of the XP1000 that led to impressive ANSYS results also prove beneficial in large scientific problems.

Viewperf

Benchmark

Viewperf OpenGL 3D graphics benchmark suite.

Source

The Viewperf benchmark suite has become an industry standard. It was developed by the Graphics Performance Consortium (GPC) and is currently managed by the SpecBench organization. SpecBench is also responsible for the popular SPECint and SPECfp processor benchmarks as discussed previously. Full information on the benchmarks and on benchmark results for many systems is available on the SPEC Web page at WWW.SPECBENCH.ORG.

Description

Viewperf is a suite of benchmarks measuring OpenGL performance across several different classes of applications. It is platform neutral and runs on any system that has OpenGL, which means that it spans both UNIX and Windows NT operating systems and many different processor and graphics architectures. Viewperf results can be directly compared between systems, greatly enhancing the usefulness of the benchmark.

Viewperf consists of a suite of six different benchmarks. Each benchmark was developed by profiling an actual application, and determining what graphics operations, functions, and primitives were used. The graphics data and graphics calls were extracted from the application and used to create the benchmark. Thus, each benchmark represents the *graphics behavior* of a *specific class of application*.

Within each benchmark, there are sets of specific tests. Each test exercises a specific graphics function—for example, anti-aliased lines, Gouraud shaded triangles, texture mapped triangles, single lights, multiple lights, etc. Performance on each of these tests is reported. Each test is assigned a weighting value, and the weighted values are summed to produce a single *composite result* for that benchmark. Viewperf data is reported in terms of frames per second, so larger values represent better performance.

As a side note, there is some controversy in the Viewperf community over whether to focus on the individual test results or on the composite results. Both results are reported, and the general industry trend is to use the composite results.

There are currently six benchmarks in the Viewperf suite (now in version 6.1), and new benchmarks are likely to be added in the future. The Viewperf 6.1 suite includes CDRS-05, ProCDRS-01, DRV-05, DX-04, Light-02, and AWadv-02. Each benchmark represents a different type of application. It is advisable to choose the benchmark that most closely matches the applications of interest in comparing graphics performance.

CDRS

CDRS is Parametric Technology's modeling and rendering software for computer-aided industrial design (CAID). It is used to create concept models of automobile exteriors and interiors, other vehicles, consumer electronics, appliances, and other products that have challenging free-form shapes. The users of CDRS are typically creative designers with job titles such as automotive designer, products designer, or industrial designer.

CDRS is the most frequently mentioned or promoted subset of the Viewperf benchmark suite. Since the data set is small, vendors have optimized their systems extensively for this benchmark, generating large, easily marketed "scores." There are seven tests specified that represent different types of operations performed within CDRS. Five of the tests use a triangle strip data set from a lawnmower model created using CDRS. The other two tests show the representation of the lawnmower.

ProCDRS

The ProCDRS-01 is the newest Viewperf viewset and is a complete update of the CDRS-03 viewset and models the graphics performance of Parametric Technology Corporation's Pro/DESIGNER industrial design software. The model data was converted from OpenGL command traces taken directly from the running Pro/DESIGNER software, and therefore preserves most of the attributes of the original model.

The viewset consists of ten tests, each of which represents a different mode of operation within Pro/DESIGNER. Two of the tests use a wireframe model, and the other tests use a shaded model. Each test returns a result in frames per second, and a composite score is calculated as a weighted geometric mean of the individual test results. The tests are weighted to represent the typical proportion of time a user would spend in each mode.

The shaded model is a mixture of triangle strips and independent triangles, with approximately 281000 vertices in 4700 OpenGL primitives, giving 131000 triangles total. The average triangle screen area is 4 to 5 pixels. The wire frame model consists of only line strips, with approximately 202000 vertices in 19000 strips, giving 184000 lines total. All tests run in display list mode. The wireframe tests use anti-aliased lines, since these are the default in Pro/DESIGNER. The shaded tests use one infinite light and two-sided lighting. The texture is a 512 by 512 pixel 24-bit color image.

DRV

Developed by Intergraph, DesignReview is a 3D-computer model review package specifically tailored for plant design models consisting of piping, equipment, and structural elements, such as I-beams, HVAC ducting, and electrical raceways. It allows flexible viewing and manipulation of the model for helping the design team visually track progress, identify interferences, locate components, and facilitate project approvals by presenting clear presentations that technical and non-technical audiences can understand.

On the construction site, DesignReview can display construction status and sequencing through vivid graphics that complement blueprints. After construction is complete, DesignReview continues as a valuable tool for planning retrofits and maintenance. DesignReview is a multi-threaded application that is available for both UNIX and Windows NT. Since this is the only multithreaded application in the Viewperf suite, it is important to carefully review system configurations when comparing reported system performance results.

The model in this viewset is a subset of the 3D-plant model made for the GYDA offshore oil production platform located in the North Sea on the southwest coast of Norway. A special thanks goes to British Petroleum, which has given the OPC subcommittee permission to use the geometric data as sample data for this viewset. Use of this data is restricted to this viewset.

DesignReview works from a memory-resident representation of the model that is composed of high-order objects, such as pipes, elbows valves, and I-beams. During a plant walkthrough, each view is rendered by transforming these high-order objects to triangle strips or line strips. Tolerancing of each object is done dynamically and only triangles that are front facing are generated. This is apparent in the viewset model as it is rotated.

Most DesignReview models are greater than 50 megabytes and are stored as high-order objects. For this reason and for the benefit of dynamic tolerancing and face culling, display lists are not used.

DX

The IBM Visualization Data Explorer (DX) is a general-purpose software package for scientific data visualization and analysis. It employs a data-flow driven client-server execution model and is currently available on UNIX workstations from Silicon Graphics, IBM, Sun, Hewlett-Packard, and Digital Equipment. The OpenGL port of Data Explorer was completed with the recent release of DX 2.1.

The tests visualize a set of particle traces through a vector flow field. The width of each tube represents the magnitude of the velocity vector at that location. Data such as this might result from simulations of fluid flow through a constriction. The object represented contains about 1,000 triangle meshes containing approximately 100 vertices each. This is a medium-sized data set for DX.

LIGHT

The Light-02 dataset is based on the Lightscape Visualization System from Lightscape Technologies, Inc. Light contains two integrated visualization components. The primary component utilizes progressive radiosity techniques and generates view-independent simulations of the diffuse light propagation within an environment. Subtle but significant effects are captured, including indirect illumination, soft shadows, and color bleeding between surfaces. A post process using ray tracing techniques adds specular highlights, reflections, and transparency effects to specific views of the radiosity solution.

The Light-02 dataset features four tests using two models, two of the tests are run in wireframe mode and two in fully shaded mode. One model is a "Cornell box" and the other is a model of Parliament Building.

AWADVVS

The Awadvvs dataset is based on Advanced Visualizer from Alias/Wavefront. Advanced Visualizer is an integrated workstation-based 3D-animation system that offers a comprehensive set of tools for 3D modeling, animation, rendering, image composition, and video output.

All operations within AWadvvs are performed in immediate mode (no display lists are used) with double buffered windows. There are four basic modes of operation within Advanced Visualizer, material shading (the most heavily weighted), wireframe, smooth shading and flat shading. All shading operations feature z-buffing, back-face culling, and two local light sources. For this reason, configurations featuring graphics controllers with large frame buffers and on-board geometry accelerators typically offer the strongest performance in this dataset.

What is Measured

Viewperf measures 3D graphics performance using OpenGL.

Units

Weighted results are reported in frames per second with higher scores being better. Each Viewperf benchmark set (such as CDRS or Awadvvs) includes percentage weights for each test. The results for each test are multiplied by their percentage weighting factor, and then all of the tests are added together to produce the weighted composite number for that set. More details on the test cases and their weights are available from the SPECbench Web site.

Additional Information on Benchmark

More information is available from <http://www.specbench.org>.

The Market:

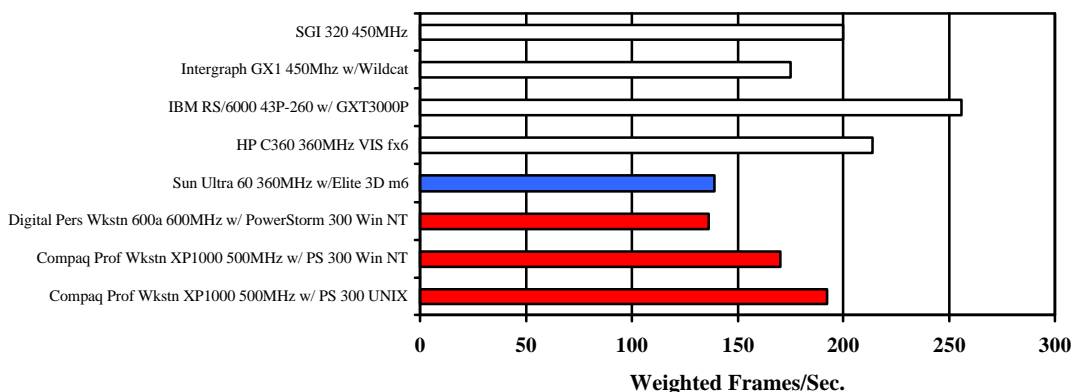
As of January 1999, the Viewperf benchmark results for systems currently on the market are as follows:

Benchmark	Low	High
CDRS-04	25.3	248
ProCDRS-01	2.38	24.45
DRV-05	2.74	20.23
DX-04	3.42	31.6
AWadvvs-02	4.54	61.48
Light-02	.55	3.14

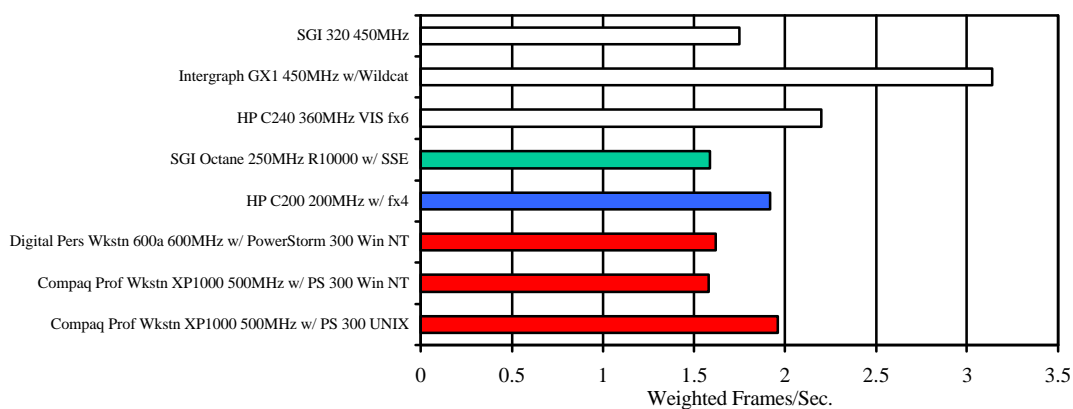
Compaq Results:

As of January 1999, the Viewperf benchmark results for Compaq systems are as follows:


CDRS-04 Results: Larger is Better



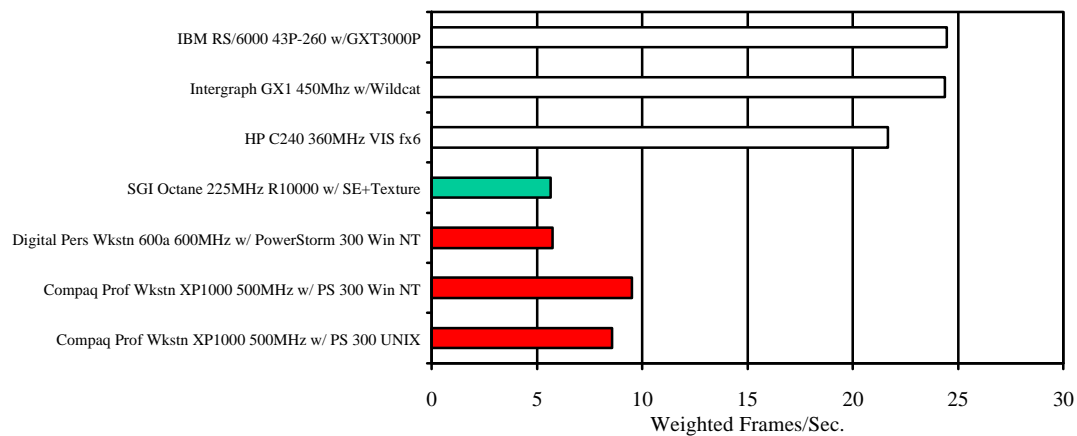


Light-02 Results: Larger is Better



AWadvs-02 Results: Larger is Better



DX-04 Results: Larger is Better**DRV-05 Results: Larger is Better****ProCDRS-01 Results: Larger is Better**

Sources

More information and competitor results are available at
http://www.spec.org/gpc/opc.data/opc_nov98/summary.html.

Compaq Professional Workstation XP1000 results are considered estimates until published by the OPC.

Notes

The results above clearly illustrate a key point. Compaq offers high levels of 3D performance with the PowerStorm 300, compared to graphics controllers in its price and feature class. Essentially this means other graphics controllers in the \$1,500-3,000 range that do not offer onboard geometry acceleration. The PowerStorm 300 does not offer on-board geometry to assist with lighting and transformation operations, and all of these duties are performed by the host Alpha processor. Controllers, such as the Intergraph Wildcat 4000, HP VISUALIZE, and Sun Elite 3D graphics, all offer several dedicated on-board geometry engines and texture acceleration options that can support higher Viewperf scores given the lighting and texturing requirements of many of the Viewperf models. But these controllers also cost as much as \$15,000 (VISUALIZE fx6 on C-class workstations), which allows the Compaq Professional Workstation XP1000 to offer some of the best price/performance ratios in the industry.