Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.4 R2 Evaluation

Table of Contents

Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.4 R2 Evaluation

Table of Figures

Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.4 R2 Evaluation

Table of Result Charts

TAGITT/CATIA
4.2.4 R2 Evaluation

abcg

albert-battaglin consulting group
one ironwood drive
soquel, ca 95073
831) 464-0600

18 October 2004

Notices

Results are shown as times in which smaller numbers represent the faster performance.

Executive Summary

Ten engineering workstations from Hewlett-Packard, IBM and Sun, were evaluated by the Albert-Battaglin Consulting Group using TAGITT: The Albert Group Interactive Throughput Test. TAGITT/CATIA measures workstation performance from a user perspective by simulating interactive work sessions using the 4.2.4 refresh 02 version of this CAD/CAM application software package. Key results were as follows:

- The IBM IntelliStation® POWER™ 275 (1.45 GHz) GXT6500P and IBM IntelliStation POWER 275 (1.45 GHz) GXT4500P were nearly tied as the overall performance leaders in the TAGITT/CATIA 4.2.4 R2 test by a significant margin. The POWER 275 (1.45 GHz) GXT6500P tested 4% better overall than the POWER 275 (1.45 GHz) GXT4500P due to its superior graphics performance. The leading machine performed 16% faster than the second place HP c8000 Fire GL-X1 and HP c8000 Fire GL- T2

- For overall graphics performance, the IBM IntelliStation POWER 275 (1.45 GHz) GXT6500P tied with the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2. The IBM POWER 275 (1.0 GHz) GXT6500P and POWER 275 (1.45 GHz) GXT4500P were just over 20% slower in second place.

- In the dynamic graphics tests, the HP c8000 Fire GL-X1 was the clear winner, followed by the 8% slower HP c8000 Fire GL-T2. The IntelliStation POWER 275 (1.45 GHz) GXT6500P was almost twice as slow as the leading HP machine in these tests.

- Machine performance was not the same across all tested applications. Although the 1.45 GHz IntelliStation POWER 275 machines were the fastest in nearly all tests, the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 also won in a few individual tests (Overall Dynamic Graphics Performance, Studio, and Walk Through). In addition, the relative performance of the machines varied significantly on the different tests.

- The HP c8000 machines tested had 2 processors and 2 GB of main memory. It is not clear from the tests we ran what performance boost was provided by the extra processor and the extra memory. CATIA V4 is mostly single threaded code, so an extra processor typically has very little impact. In fact, the HP c8000 machines did not outperform the single processor IntelliStation POWER 275 machines in the Studio Viewer tests, which in the past have benefited from multiprocessor technology.

- The 1.2 GHz Sun Blade 2000 with XVR-1200 graphics was over two times (112%) slower than the leading IBM machines overall. Finite Element analysis and dynamic graphic test results were notably poor, 234% and 391% slower (or well over three times to almost 5 times slower than the leading machines in these tests).

Summary of Top Three Winners in Each Test:

Throughput Summary Results

Task/Test	First	Second	Third
Overall Throughput	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT6500P**
Application Throughput	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**
Graphics Throughput	HP c8000 Fire GL-X1 **IBM POWER 275 (1.45 GHz) GXT6500P** HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	**IBM POWER 275 (1.0 GHz) GXT4500P**
System Responsiveness	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.0 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2
Dynamic Graphics Throughput	HP c8000 Fire GL-X1	HP c8000 Fire GL-T2	**IBM POWER 275 (1.45 GHz) GXT6500P**
4D Navigator	**IBM POWER 275 (1.45 GHz) GXT6500P**	**IBM POWER 275 (1.0 GHz) GXT6500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2

Individual Application Test Throughput Results:

Primary Application Test Throughput Results			
Task/Test	**First**	**Second**	**Third**
Modeling Solid Model Creation and Modification	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.0 GHz) GXT4500P**
Finite Element Analysis (ANSOLID)	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**
NC Operations STL Generation	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.0 GHz) GXT4500P**
Detail Drawing Creation	**IBM POWER 275 (1.45 GHz) GXT4500P** **IBM POWER 275 (1.45 GHz) GXT6500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**

Secondary Application Test Throughput Results

Task/Test	**First**	**Second**	**Third**
Solid and Surface Analysis Function	**IBM POWER 275 (1.45 GHz) GXT4500P** **IBM POWER 275 (1.45 GHz) GXT6500P**	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**	HP c8000 Fire GL-T2 HP c8000 Fire GL-X1
Read/Write	**IBM POWER 275 (1.45 GHz) GXT4500P** **IBM POWER 275 (1.45 GHz) GXT6500P**	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**	HP c8000 Fire GL-T2 HP c8000 Fire GL-X1
Walk Through	HP c8000 Fire GL-X1	HP c8000 Fire GL-T2	**IBM POWER 275 (1.45 GHz) GXT4500P**
Bend (Sheet Metal Part Development and Modification)	**IBM POWER 275 (1.45 GHz) GXT4500P** **IBM POWER 275 (1.45 GHz) GXT6500P**	HP c8000 Fire GL-T2 HP c8000 Fire GL-X1	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.0 GHz) GXT4500P**
Fitting Simulation	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-T2 HP c8000 Fire GL-X1	**IBM POWER 275 (1.0 GHz) GXT6500P** **IBM POWER 275 (1.0 GHz) GXT4500P**
Studio	HP c8000 Fire GL-T2 HP c8000 Fire GL-X1	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**
Image Viewer	**IBM POWER 275 (1.45 GHz) GXT6500P** **IBM POWER 275 (1.45 GHz) GXT4500P**	HP c8000 Fire GL-X1 HP c8000 Fire GL-T2	**IBM POWER 275 (1.0 GHz) GXT4500P** **IBM POWER 275 (1.0 GHz) GXT6500P**

Why TAGITT?

The Need for Application-Level Testing

When mechanical engineers select and use workstations, performance considerations should be based on the ability of the machine to rapidly complete the users' design task. Users are concerned with throughput: how much faster (or better) could I design my next product if I upgraded to a faster graphics card or a faster CPU? In the workstation industry, MHz, MIPS, MFLOPS, SPECmarks, etc. have become the standards for performance comparison of CPUs. For graphics, 3D vector drawing speed and polygon drawing speed (polygons per second) are often used for comparison. In selecting a workstation for a mechanical design application, the user is faced with a choice between many competitive machines — some with higher MIPS ratings, others with higher vector and/or polygon rates. Without running an actual application benchmark, it is difficult to predict which of the two machines will provide the better performance level for its application.

MHz, MIPS, megaflops, vectors per second, GPC, XPC, OPC and polygons per second all allow users to compare machines, but those specs may be misleading as predictors of engineering task efficiency. Today's CAD/CAM applications are typically very large, complicated programs. The way in which these programs perform in the context of different hardware architectures and with different operating system services and graphic libraries is generally not predictable from the previously mentioned specifications. Although software vendors are striving to make their code highly "portable" so that it runs on a wide variety of machines, the fact is that all applications must be ported and tuned to obtain optimal performance. Each workstation vendor offers unique performance-enhancing capability. Without tuning, application software may or may not take full advantage of the target hardware/operating system platform. Since software developers cannot possibly take advantage of every function in every workstation and/or operating system, performance compromises occur. The user has no way of knowing to what extent his/her application software has been ported and tuned to match capabilities offered by any particular workstation vendor without application level testing.

Test Description

TAGITT, The Albert Group Interactive Throughput Test, was designed to directly measure performance in completing typical engineering design tasks, especially related to solid modeling. For users of solid modeling software, the results provide a comparison of workstations that is more relevant than the typical manufacturer's published specifications of MHz, MIPS, megaflops, vectors per second and polygons per second.

TAGITT testing is typically accomplished by recording and playing back user interaction scenarios. Most CAD/CAM applications include functions to accomplish this task although some are undocumented. Record and playback mechanisms are the preferred method of testing for a number of reasons, including repeatability, accuracy, and user relevance. Although it is often easier to measure times for individual operations or functions, this can be a misleading measure of performance from a user perspective. Users constantly switch between functions and/or modules, which can result in significant performance variation as portions of the software are loaded, unloaded, and accessed from memory. The use of interaction scenarios provides a more realistic measurement of overall system performance. TAGITT tests use built-in timing and data capture mechanisms in order to obtain accurate measurements over a relatively large number of functional tests.

In addition to overall time measurements, TAGITT scenarios normally include interim times for specific functions or operations. These can provide specific performance data for individual functions such as adding a solid feature, shading a model or generating an NC tool path. These times are also used by Albert-Battaglin Consulting Group as a rough method for isolating performance that is compute intensive, graphic intensive or I/O intensive. While it is clear that the interplay between these system aspects is too complicated to be accurately measured at the application level, the measurements can sometimes point to areas for in-depth performance profiling using specialized tools.

TAGITT interaction scenarios consist of a variety of operations, with an emphasis on parametric/variational solid modeling and associated tasks such as part visualization, kinematics and "walk through" analysis, drawing view creation from 3D models, NC tool path generation, STL output, geometric/FE analysis and high quality image rendering and display. Regardless of the task, special emphasis is placed on tasks that are not by nature "interactive." For example, the creation of a line segment in most systems takes place in well under a half second so that performance differences will most likely be unnoticed by users. These times are not considered in TAGITT results. In contrast, updating a solid model or regenerating a drawing layout following a dimensional change can take from many seconds to several minutes and therefore has an impact on a user's productivity. TAGITT evaluations also measure ancillary tasks such as changing functions (through menu picks) or selecting geometry. The time taken for these operations generally ranged between 1 and 15 seconds. Albert-Battaglin Consulting Group feels that the overall responsiveness of the system is reflected in these interaction times. This "responsiveness" is the difference between systems that seem heavy and slow, compared to those that "feel" quick and light. The TAGITT measurements gather data from both of these interaction types and combine them together to create an overall throughput measurement.

Models used for TAGITT evaluations are taken from previously released customer data and no data is ever provided to workstation vendors or third parties by the Albert-Battaglin Consulting Group.

TAGITT/CATIA

TAGITT/CATIA consists almost entirely of a series of CATIA "record files" which are capable of recording and playing back a series of user interactions. The record files capture a majority of user interactions including some simulations of dynamic graphic manipulations performed via the GRAPER utility function. Fifty record files are used in the 4.2.4 refresh 2 version of TAGITT/CATIA. The record files were either created by ABCG, adapted from standard CATIA Operator Exchange test files or were developed around Dassault Systèmes' demo part models. Some of the models used are shown in Figure 1. These files cover many areas of the CATIA 4.2.4 R2 product including part modeling, surface intersection, drawing layout, parametric modification and updates, the skin function, finite element analysis, fitting simulation, kinematics simulation, walk through analysis, NC tool path generation, sheet metal part modeling, studio image rendering and viewing and model storage and retrieval from disk. Combined, these files represent thousands of user interactions and many hours of operator seat time.

Figure 1 – Some of the CATIA models used for TAGITT/CATIA

TAGITT/CATIA also includes CATIA Image Viewer and 4D Navigator tests. For the CATIA Image Viewer test, the time for displaying an image generated by the Studio function as recorded by the Image Information function is used. In contrast to common measurements of 4D

Navigator using a "frame counter" utility, Albert-Battaglin Consulting Group has developed what it feels to be a more representative measurement method. In TAGITT/CATIA, eight different CATIA models are combined and rendered in single light, dual light, dual light with edges, neon and neon with edges modes. These models are each rotated through 360 degrees in 30-degree steps and times are recorded from the 4D Navigator's performance monitor. The combination of models used for the testing of both of these functions is large and complicated enough such that the operations are generally not interactive (i.e. less than 0.5 seconds). Some of the models used for the 4D Navigator test are shown in Figure 2.

Figure 2 – CATIA 4D Navigator Test Models

Test Weighting

While running the TAGITT/CATIA 4.2.4 R2 benchmark, the times for individual operations and scenarios are recorded. To produce Application, Graphics and Overall throughput results, weighted sums of appropriate individual test results are used. Albert-Battaglin Consulting Group sets the weighting factors for each type of operation based on its judgment of the relative importance of each operation. The weighting is Albert-Battaglin Consulting Group's best judgment for a "typical" CATIA user, whether from aerospace, automotive or any other industry. The weights are applied to actual times by averaging the results measured across the various workstations and applying the appropriate factor.

The results presented in this report represent a cross section of different types and sizes of models that can act as a guide for overall workstation performance.

Overall throughput numbers consist of the weighted sum of the Application, Graphics and Responsiveness portions of the test. Application tests measure the times required to complete application related tasks such as changing a solid model, generating a tool path or running a FEM analysis. Graphics portions of the test measure exclusively viewing-related functions such as generating a shaded image or dynamically rotating that image. Responsiveness tests include the times needed to change CATIA functions select options or pick geometry. The weighting used is shown in Figure 3.

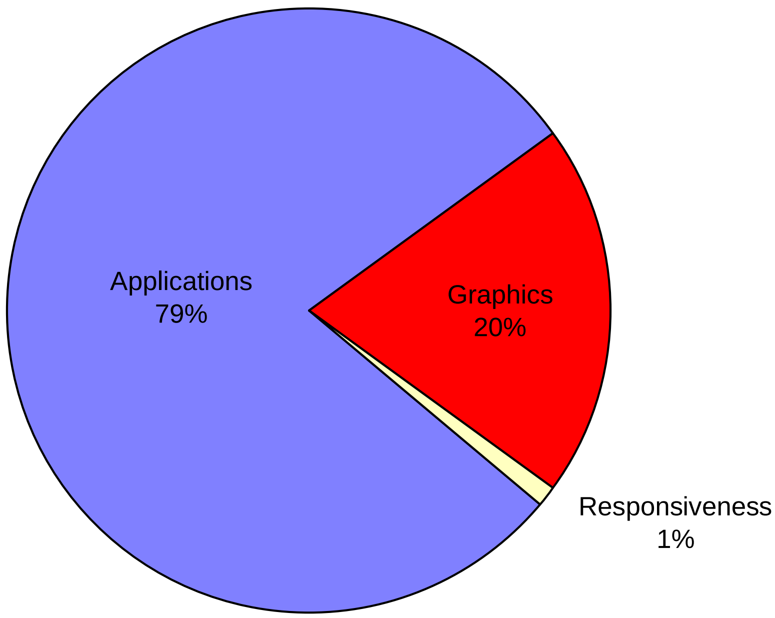

Figure 3 – Overall Throughput Weighting

The Application time for this TAGITT test is the weighted sum of the Analysis, Bend (Sheet Metal), Detailing, Finite Element, Fitting Simulation, Kinematics, Modeling, NC, Studio, Image Viewer, Walk Through, and Read/Write portions of the test. The weightings used are shown in Figure 4.

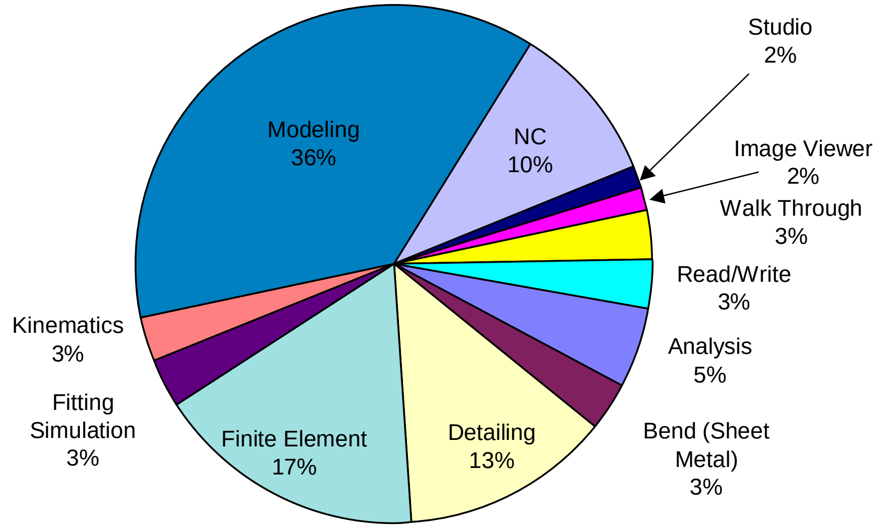

Figure 4 – Application Weighting Factors

4D Navigator throughput is calculated based on the weighted sum for the single light, dual light, neon and neon with edges rendering tests for the five models tested. The weighting factors are shown in Figure 5.

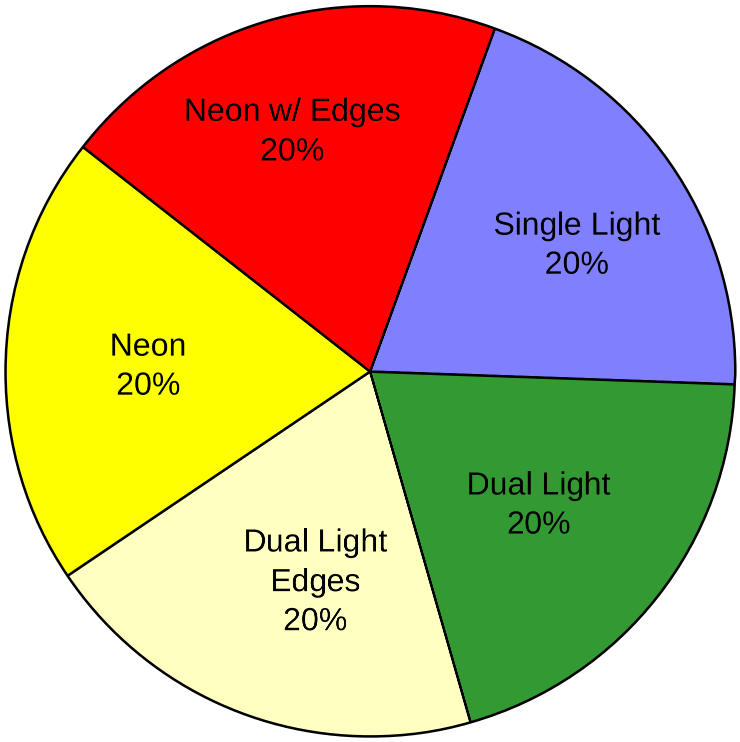

Figure 5 – 4D Navigator Weighting Factors

Graphics throughput is calculated based on the weighted times for the various graphics tests (refer to Figure 6) so as to present one representative number for this aspect of each workstation. These operations include not only the often measured dynamic graphic manipulation (dial turning) functions, but also the “graphics compute” functions which often occur the first time one accesses these operations on a given part. This version of TAGITT also includes 4D Navigator results in the overall Graphics throughput. This combined measurement gives a better overall picture of graphic performance during typical work sessions from a user perspective. Isolated evaluations of dynamic shading or hidden line processing may be good for tuning tasks, but they do not adequately take into account the mix of operations encountered by a user. These graphic compute operations are always much longer than the dynamic manipulations themselves and are also dependent on CPU performance and its interactions with the graphics processor. The Graphics Throughput Time for this TAGITT test is the weighted sum of the 4D Navigator, Graphics Compute, Hidden Line, Shaded Image and Wireframe portions of the test which are each sums from the various (10) models used throughout the testing.

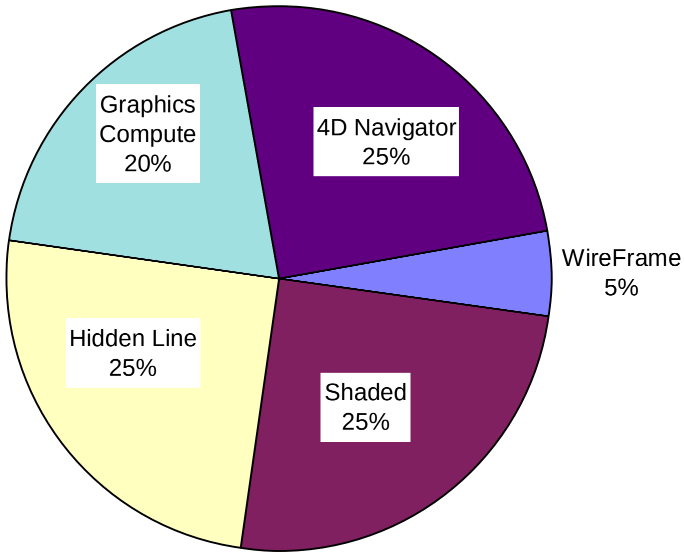

Figure 6 – Graphics Weighting Factors

It should be noted that user perceptible differences in display performance would, in general, only be found when processing large complex models. Overall workstation and graphic performance across the industry has progressed to where “simpler” models can be easily transformed interactively. Unfortunately, we found no easy method for defining a “simple” model. Model size and the number of geometric elements (surfaces, planes, lines, etc.) did not correlate with graphic performance. The performance is based on the complexity of the various geometric components, which is not easy for the user to determine. It is difficult to determine what class of workstation offers “sufficient” performance without examining explicitly the types of models and displays commonly used.

Machines Evaluated

The following table shows the machines tested and their configurations. All machines were configured with 1 GB main memory per processor (i.e. 2 GB for HP c8000 machines), 1 GB swap, 20”/21” color monitor, CD/DVD ROM drive, Ethernet interface, mouse, keyboard and 40 GB or larger hard disk, operating system, and 3D API. Note: throughout the report, references to the HP and Sun machines include the graphics accelerators listed below.

Vendor	Model	Graphics	CPU	L2 Cache MB	CPU MHz	OS	Abbreviation used in report and on charts
IBM	IntelliStation POWER 275	GXT6500P	POWER4+	1.5*	1450	AIX 5L™ V5.1	IBM POWER 275 (1.45 GHz) GXT6500P
IBM	IntelliStation POWER 275	GXT4500P	POWER4+	1.5*	1450	AIX 5L V5.1	IBM POWER 275 (1.45 GHz) GXT4500P
IBM	IntelliStation POWER 275	GXT6500P	POWER4+	1.5*	1000	AIX 5L V5.1	IBM POWER 275 (1.0 GHz) GXT6500P
IBM	IntelliStation POWER 275	GXT4500P	POWER4+	1.5*	1000	AIX 5L V5.1	IBM POWER 275 (1.0 GHz) GXT4500P
IBM	44P 170 (450)	GXT6500P	POWER3-II	8	450	AIX 5L V5.1	IBM 44P-170 (450 MHz) GXT6500P
IBM	44P 170 (450)	GXT4500P	POWER3-II	8	450	AIX 5L V5.1	IBM 44P-170 (450 MHz) GXT4500P
HP	c8000	ATI FireGL X1-256p	PA-8800 (2way)	1.5	1000	HP-UX 11i	HP c8000 (1 GHz) Fire GL-X1
HP	c8000	ATI FireGL T2-128p	PA-8800 (2way)	1.5	1000	HP-UX 11i	HP c8000 (1 GHz) Fire GL-T2
Sun	Blade 2000	XVR-1200	UltraSparc-III cu	8	1200	Solaris 5.9	Sun Blade 2000 (1.2 GHz) XVR1200
Sun	Blade 2000	XVR-1000	UltraSparc-III cu	8	1000	Solaris 5.8	Sun Blade 2000 (1 GHz) XVR1000

*Also has 8MB of L3 cache

TAGITT/CATIA 4.2.4 R2 Results

Overall Throughput

Chart 1 compares the overall weighted elapsed time to complete the TAGITT/CATIA 4.2.4 R2 interactive scenarios including the graphics tests. Albert-Battaglin Consulting Group feels that this number gives the best overall rating of workstation performance. The chart shows the IBM POWER 275 (1.45 GHz) GXT6500P and IBM POWER 275 (1.45 GHz) GXT4500P machines to be the fastest machines overall, a solid 16% faster than the second place HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 machines. The IBM POWER 275 (1.45 GHz) GXT4500P machine was 4% slower that the IBM POWER 275 (1.45 GHz) GXT6500P in overall throughput. The IBM POWER 275 (1.0 GHz) GXT6500P achieved third place with performance that was 32% slower than the fastest machines.

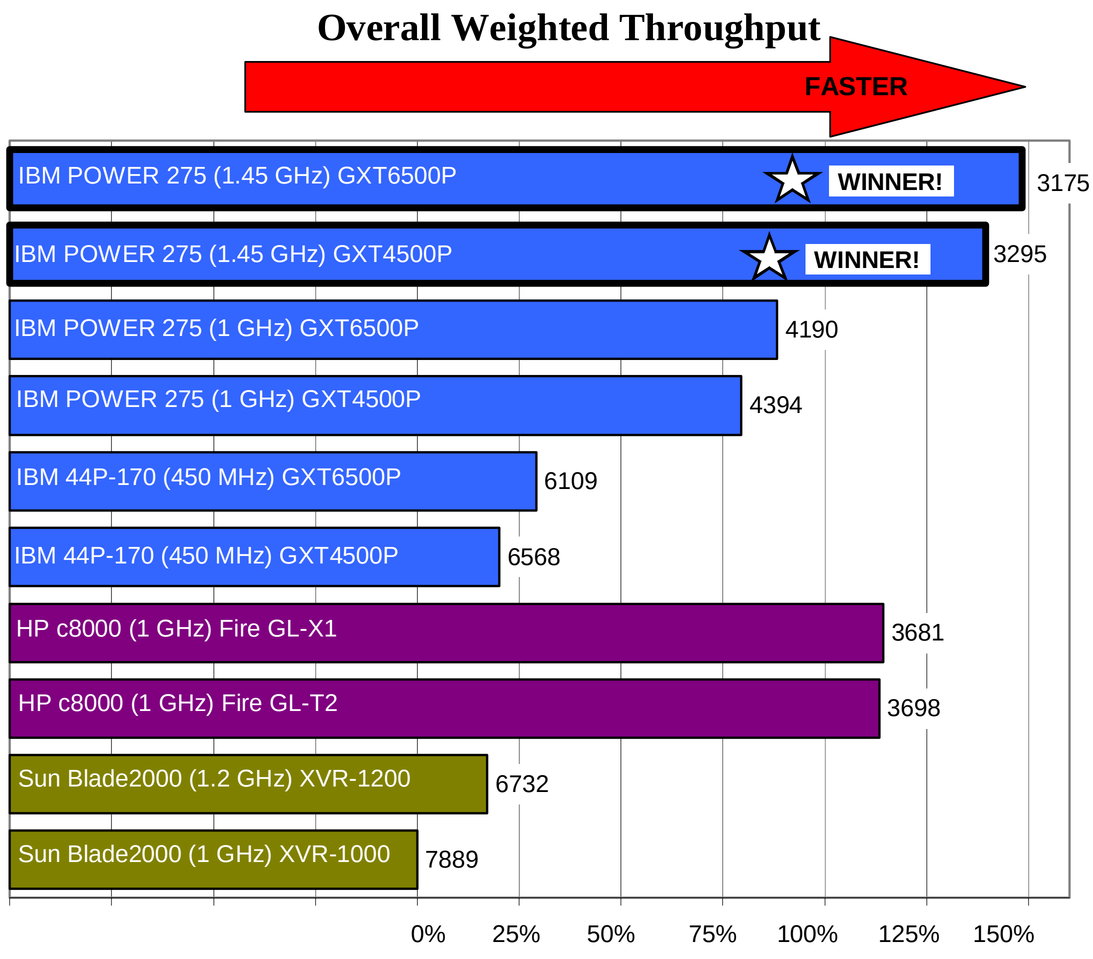

Chart 1 – Overall Weighted Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Graphics Throughput

Chart 2 shows the weighted cumulative time to complete graphic view manipulation operations of parts and assemblies. As described earlier, the Albert-Battaglin Consulting Group Graphics Throughput time includes both initial "loading" of graphics as well as dynamic manipulation times. The HP c8000 Fire GL-X1, IBM POWER 275 (1.45 GHz) GXT6500P and HP c8000 Fire GL-T2 were the fastest machines in this test with a time 21% faster than the second place IBM POWER 275 (1.0 GHz) GXT6500P and IBM POWER 275 (1.45 GHz) GXT4500P. In third place, the IBM POWER 275 (1.0 GHz) GXT4500P was 62% slower than the leader.

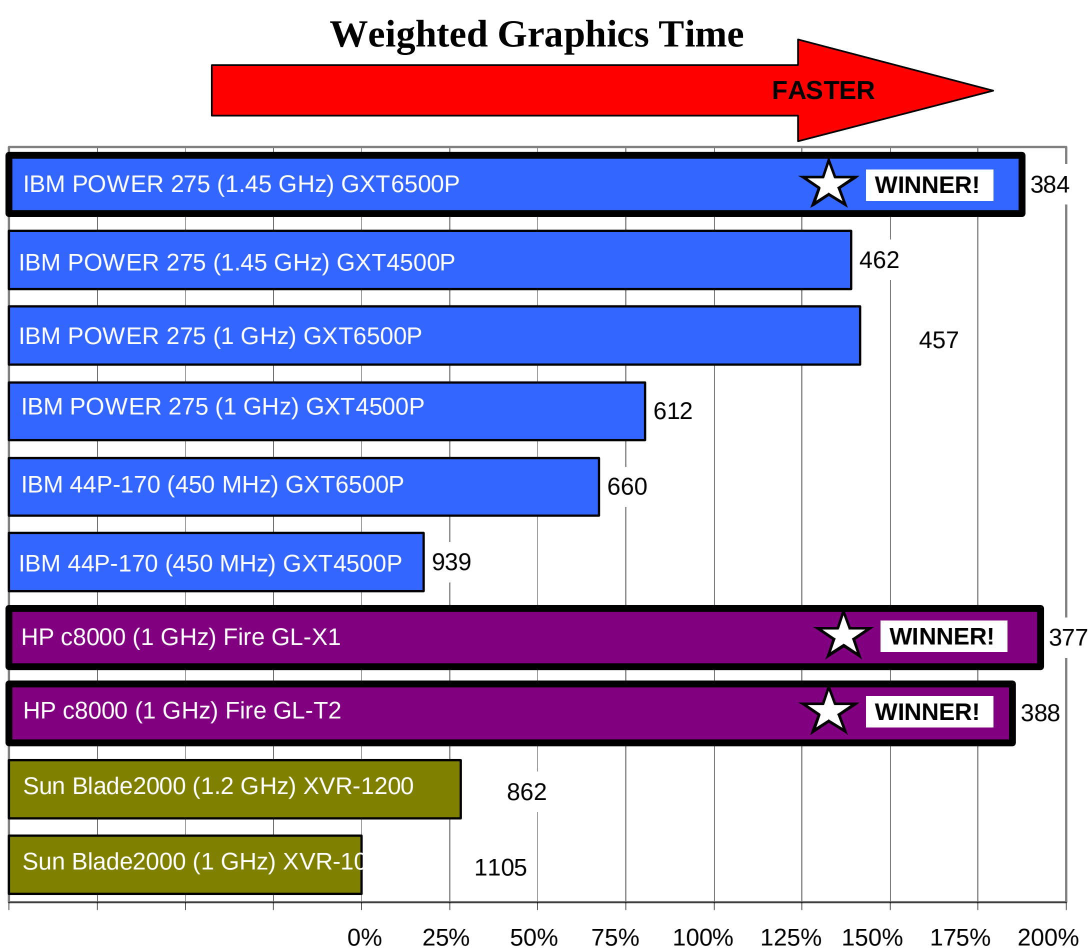

Chart 2 – Graphics Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Application Throughput

Chart 3 shows the weighted cumulative time to complete all of the CPU intensive application-specific tasks in the benchmark such as solid modeling operations, drafting and detailing operations, FEM functions and NC computations. The IBM POWER 275 (1.45 GHz) GXT6500P and IBM POWER 275 (1.45 GHz) GXT4500P workstations were the leaders in this test with nearly identical test times. In a two-way tie for second place, the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 were 19% slower than the leader. The third place IBM POWER 275 (1.0 GHz) GXT4500P and IBM POWER 275 (1.0 GHz) GXT6500P machines finished over 34% slower than the leader.

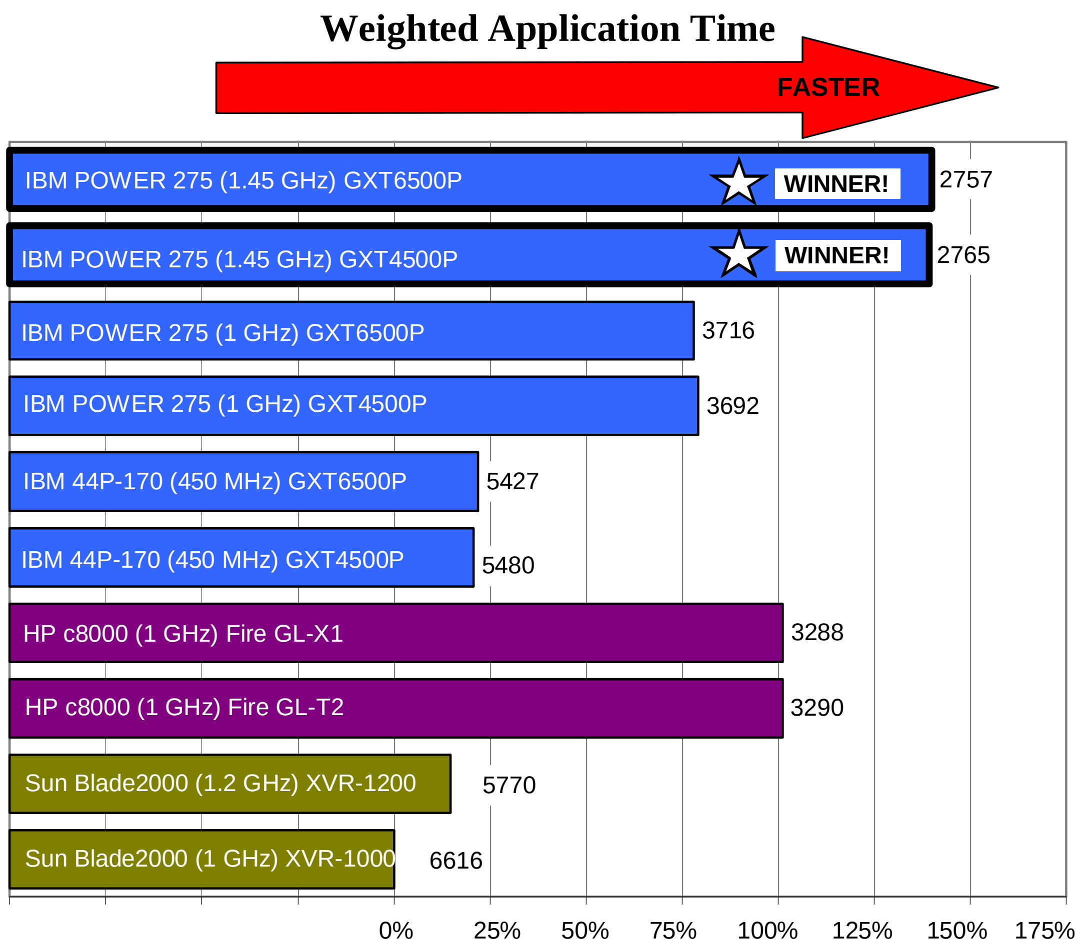

Chart 3 – Application Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

System Responsiveness

Chart 4 shows the cumulative time to complete all of the system responsiveness tests in the benchmark. The tests measure the "quickness" of the system, doing common CATIA V4 tasks such as changing function and selecting elements. Again the IBM POWER 275 (1.45 GHz) workstations were the overall winners of this test. Tied for second place, the two IBM POWER 275 (1.0 GHz) machines were 36% and 39% slower than the leading machines. In third place, the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 were 61% slower than the fastest machines.

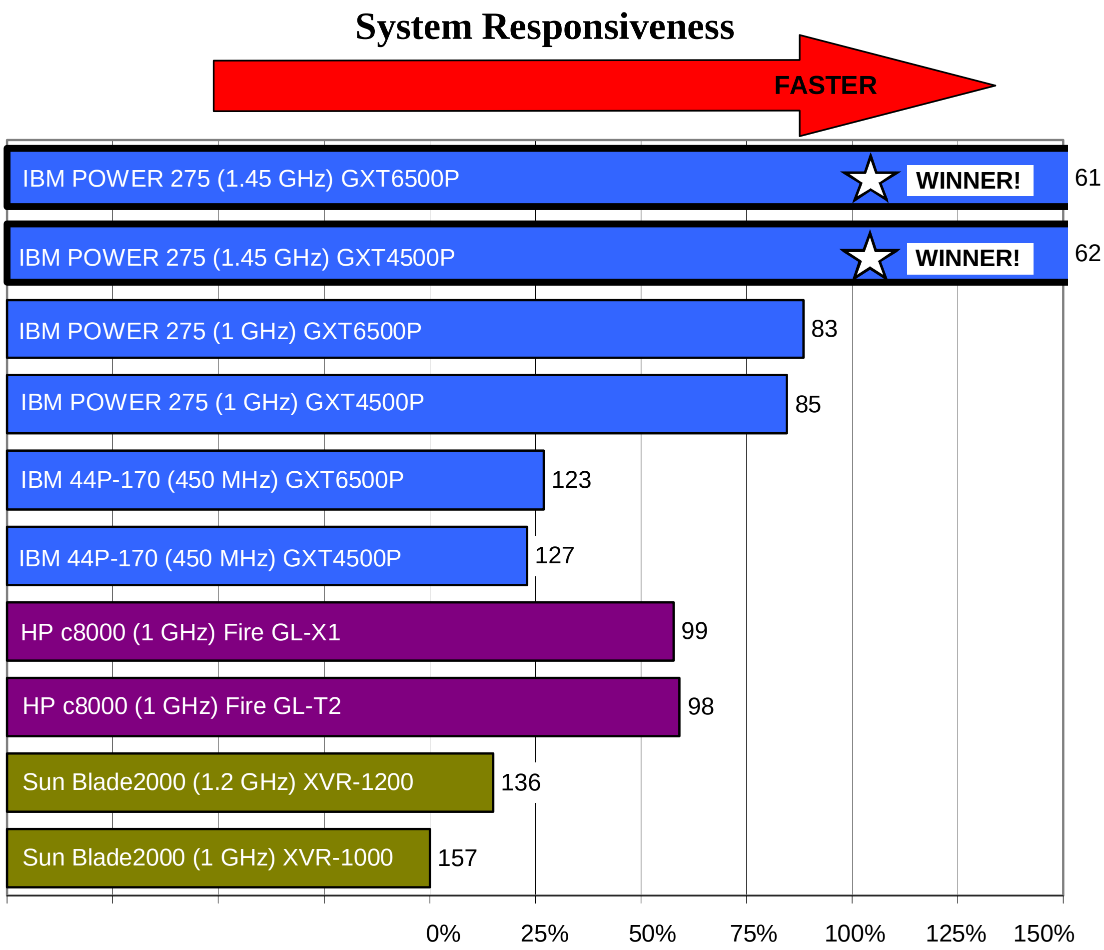

Chart 4 – System Responsiveness Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Dynamic Graphic Throughput

Chart 6 shows the cumulative time to complete all of the dynamic graphics tasks in the benchmark including wireframe, shaded and hidden line dynamic graphic operations. These results do not include the longer graphic computation times that are included in the overall graphics throughput results. The winner of this test was the HP c8000 Fire GL-X1, finishing 8% faster than the HP c8000 Fire GL-T2 in second place. The IBM POWER 275 (1.45 GHz) GXT6500P machine was almost twice as slow as the leading machine in third place. The Sun Blade 2000 (1.2 GHz) XVR1200 finished a disappointing 391% (almost 5 times) slower than the leader, while the Sun Blade 2000 (1.0 GHz) XVR1000 was a dismal 571% slower than the winning HP machine.

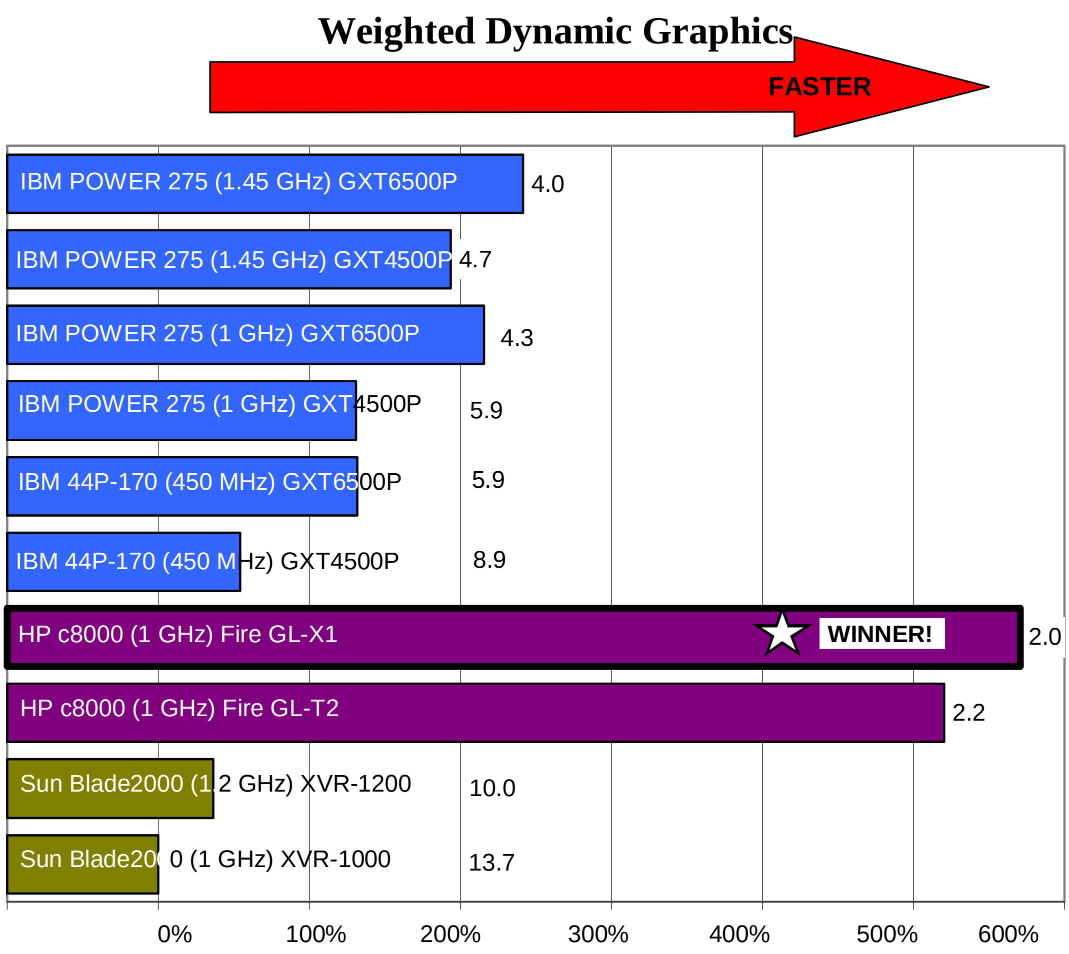

Chart 5 – Dynamic Graphic Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

4D Navigator Throughput

Chart 6 shows the weighted cumulative time to complete all of the 4D Navigator tests in the benchmark including single light, dual light, dual light with edges, neon and neon with edge rendering modes. CATIA's 4D Navigator uses OpenGL graphics so these results will become more important as users begin switching to CATIA V5. The overall winner of this test was the IBM POWER 275 (1.45 GHz) GXT6500P by a clear margin. In second place, the IBM POWER 275 (1.0 GHz) GXT6500P machine was 22% slower than the leader. The third place HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 machines were 33% and 36% slower respectively than the leader.

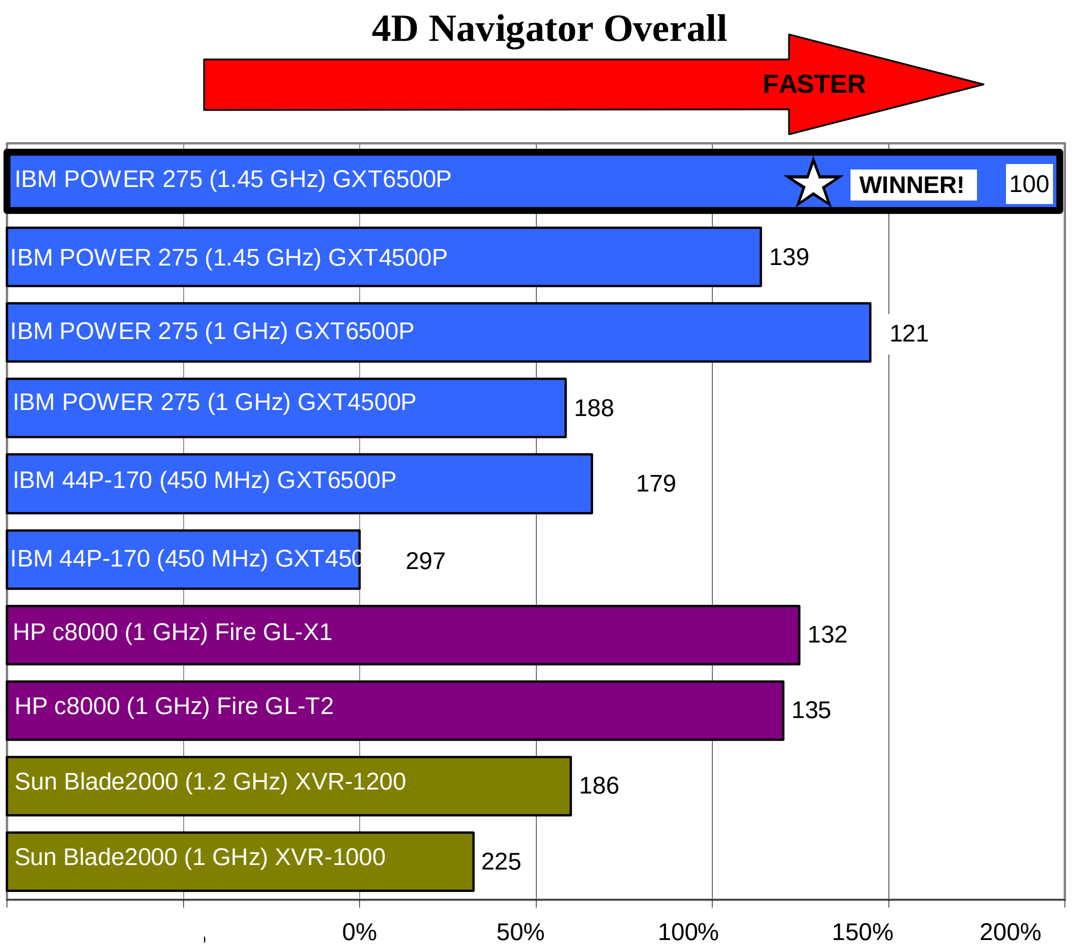

Chart 6 – 4D Navigator Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Conclusions

The results of the TAGITT/CATIA 4.2.4 refresh 02 evaluation show IBM to have a solid overall performance lead with its 1.45 GHz POWER 275 machines based on the Power4+ architecture. The IBM POWER 275 (1.45 GHz) GXT6500P was 16% faster overall than the nearest competitive machine, the HP c8000 Fire GL-X1. The Sun Blade 2000 (1.2 GHz) XVR-1200 was no match for the IBM machines with overall performance times a disappointing 112% slower than the leader. The IBM POWER 275 (1.45 GHz) GXT4500P was only 4% slower than the GXT6500P probably because of the extreme speed of the CPU. IBM's POWER 275 1.0 GHz machines also performed very well. The IBM POWER 275 (1.0 GHz) GXT6500P tested 32% slower than the leading machine while the IBM POWER 275 (1.0 GHz) GXT4500P was 38% slower than the leader.

In terms of overall graphic performance, HP narrowly had the fastest overall graphics performance with its HP c8000 Fire GL-X1 with the IBM POWER 275 (1.45 GHz) GXT6500P machine and HP c8000 Fire GL-T2 machine only 2% slower. This machine outperformed the IBM POWER 275 (1.0 GHz) GXT6500P and IBM POWER 275 (1.45 GHz) GXT4500P machines by just over 20%. It is important to note however that the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 machines were clearly the fastest dynamic graphics machines outperforming the second place IBM POWER 275 (1.45 GHz) GXT6500P by an impressive 96%. The IBM POWER 275 (1.45 GHz) GXT4500P machine tested a further 16% slower than the GXT6500P matched with the same CPU. The overall win of the IBM machine is attributable to superior graphics compute (like buffer regeneration), and better 4DNavigator results. Sun's dynamic graphic results were notably poor, the Sun Blade 2000 (1.2 GHz) XVR-1200 finishing almost 4 times slower than the fastest HP box.

The wide range of application testing in TAGITT/CATIA 4.2.4 R2 showed performance differences between various CATIA functions. It is impossible to tell from the TAGITT evaluation the reasons for these differences. Although the 1.45 GHz IBM POWER 275 machines won nearly all of the individual application tests, there were exceptions. As stated before, the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 won the dynamic graphics test as well as the Studio and Walk Through tests. In addition, the margin of victory for the IBM machines varied between the various applications. In the Modeling test, for example, the IBM

POWER 275 (1.45 GHz) GXT6500P and IBM POWER 275 (1.45 GHz) GXT4500P machines were 14% faster than the HP c8000 Fire GL-X1 and HP c8000 Fire GL-T2 while in the Finite element test the difference was greater, 22%. In the Geometric Analysis tests on the other hand, the Sun Blade 2000 (1.2 GHz) XVR-1200 machine was nearly the same as that of the HP machines and 90% slower than the 1.45 GHz POWER 275 machines.

Although the HP c8000 machines tested had 2 processors and 2 GB of main memory, it was unclear how much, if any, advantage this configuration brings when running CATIA V4. Since CATIA V4 is mostly single threaded code, we have typically found very little advantage to multi-processor configurations. In fact, the c8000 machines did not outperform the single processor IBM POWER 275 machines in the Studio Viewer tests, which in the past have benefited from multiprocessor technology.

IBM has maintained its overall CATIA V4 performance leadership with its 1.45 GHz IntelliStation POWER 275 workstations. HP however is closing the gap with its multi-processor c8000 machines. Combined with the new ATI FireGL X1-256p and ATI FireGL T2-128p graphic cards these machines outperformed IBM's machine in raw dynamic graphic performance and two other minor tests. These wins were not enough however to win in overall performance. The TAGITT/CATIA 4.2.4 R2 test shows the POWER 275 machines to be excellent all around performers in the CATIA environment. The combination of raw CPU speed and good graphic performance in the IBM IntelliStation POWER 275 workstation line, make them excellent choices in CATIA engineering environments.

Methodology

All tests were conducted by Albert-Battaglin Consulting Group personnel. Test conditions were set up to minimize any environmental differences with the various systems. Systems were tested in "lab" environments so that they were isolated from network interference. Albert-Battaglin Consulting Group saw no evidence to suggest that performance was impacted by extraneous network activity. Nearly all timing data was automatically recorded and transferred directly into spreadsheets for analysis. All tests were run at least three times and the average times were used for comparison. For the overall test, the time differences between runs were typically less than 1.5% and in most cases as low as 0.1%.

All TAGITT/CATIA 4.2.4 R2 tests were completed using released CATIA 4.2.4 R2 software. In all cases the software and data and licenses were loaded locally on each workstation. Manufacturers' required and recommended software "patches" or upgrades for CATIA V4 R4.2.4 R2 were installed on all systems prior to testing.

Appendix

The following charts show the results for individual application tests that were combined to form the overall application and overall throughput times for TAGITT/CATIA 4.2.4 R2.

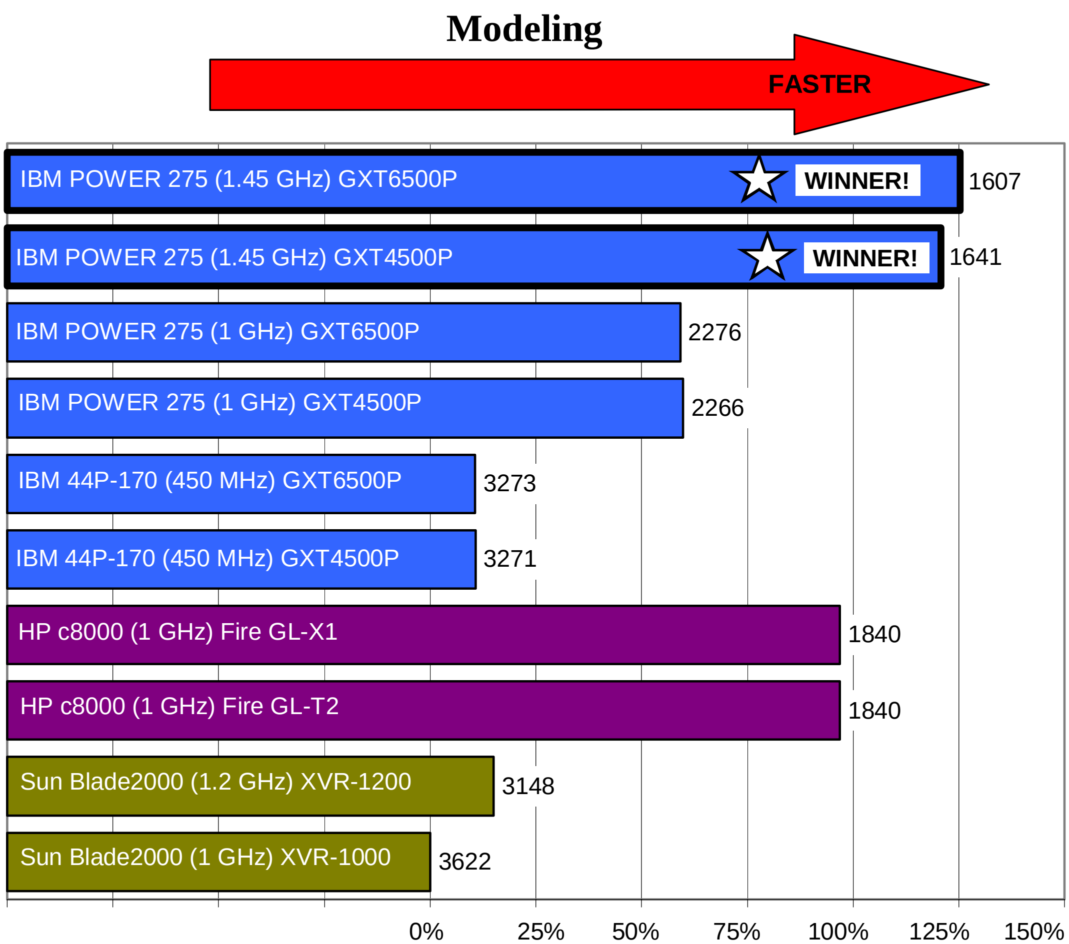

Chart 7 – Solid Model Creation and Modification Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

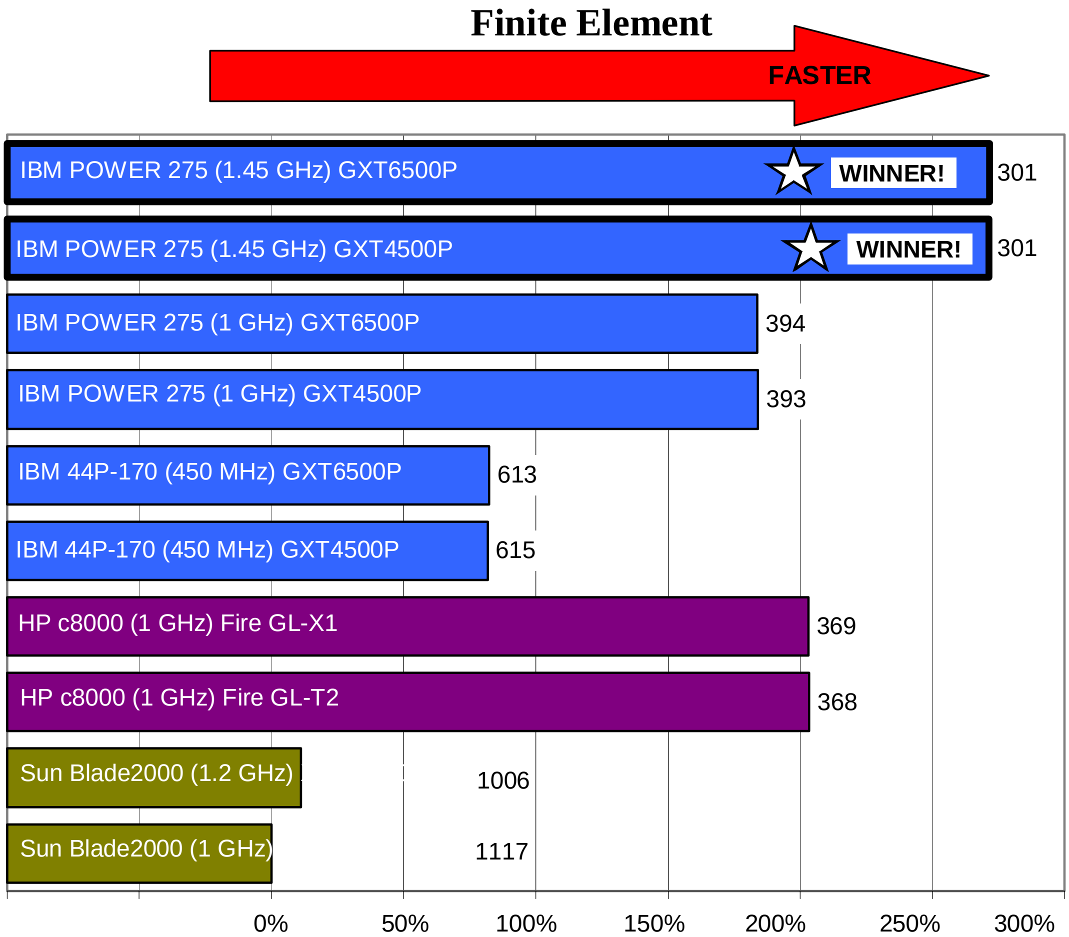

Chart 8 – Finite Element Analysis (ANSOLID) Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

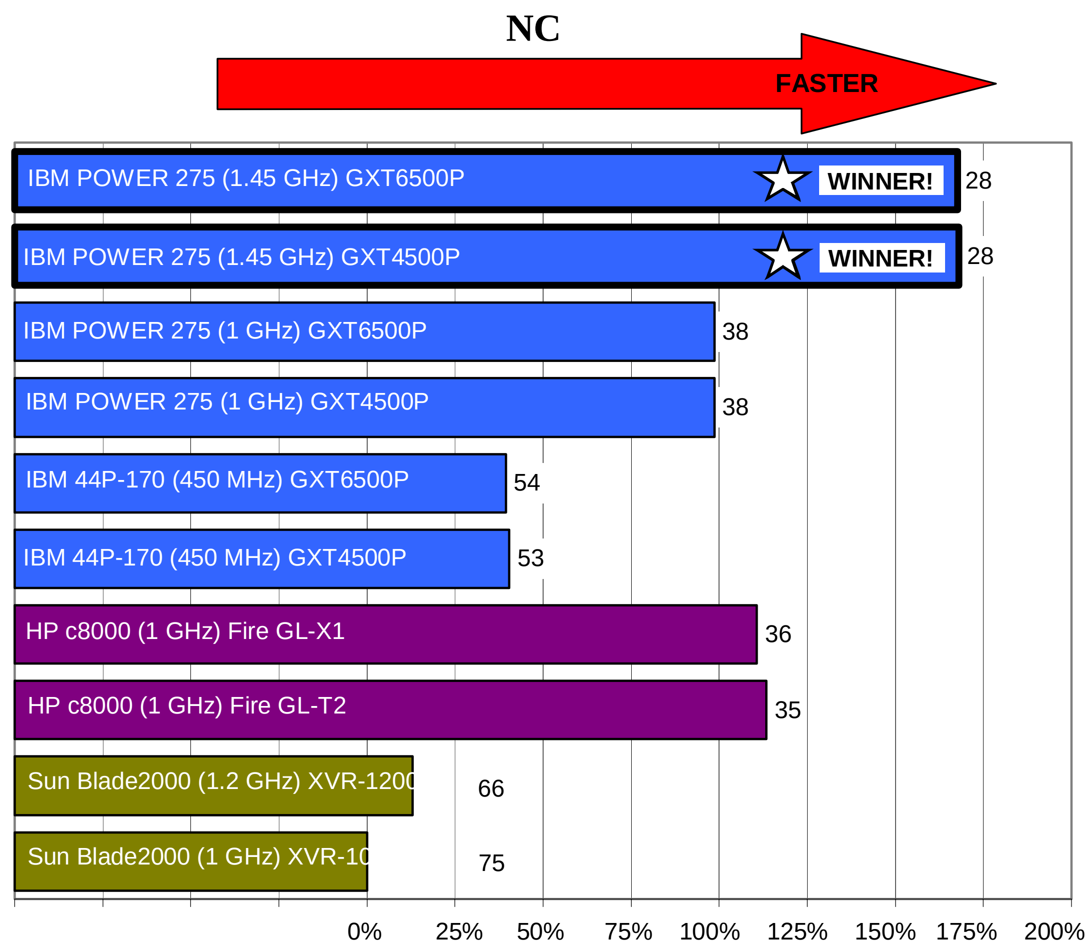

Chart 9 – NC STL Generation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

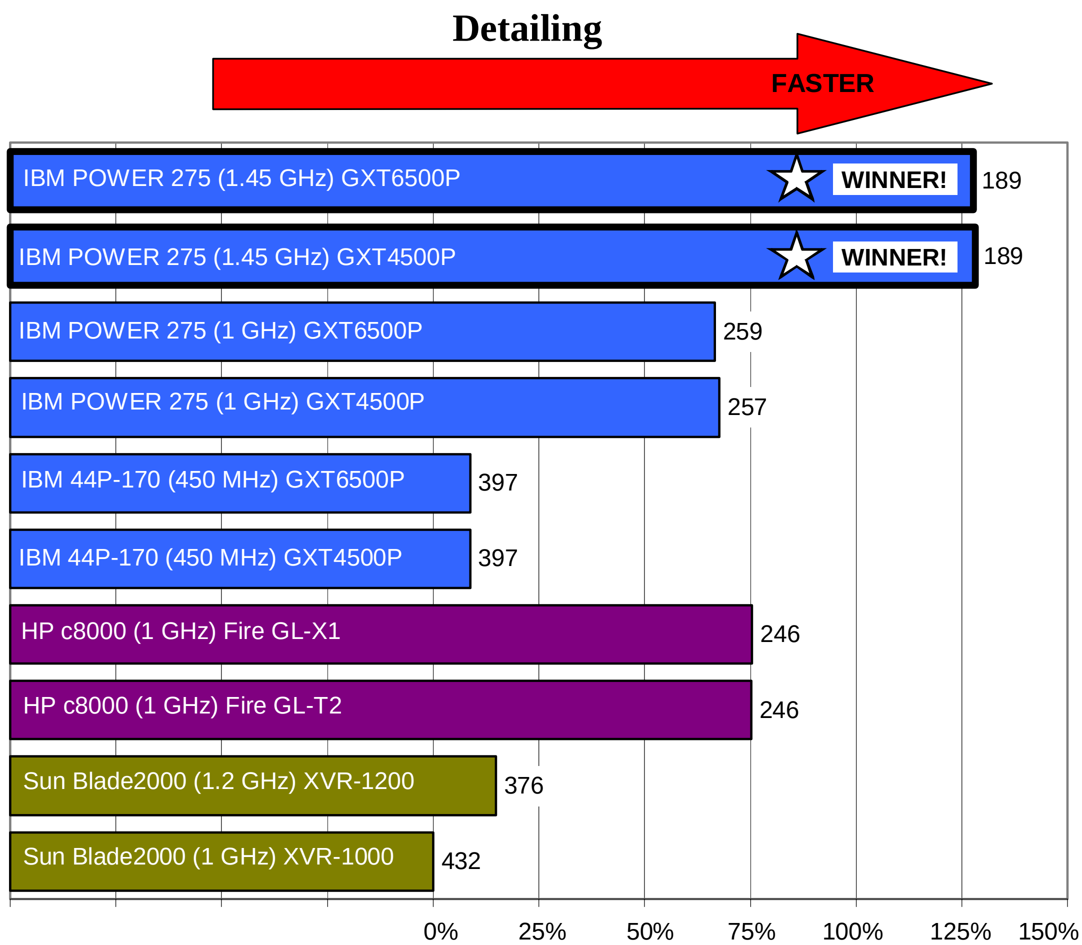

Chart 10 – Detail Drawing Creation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

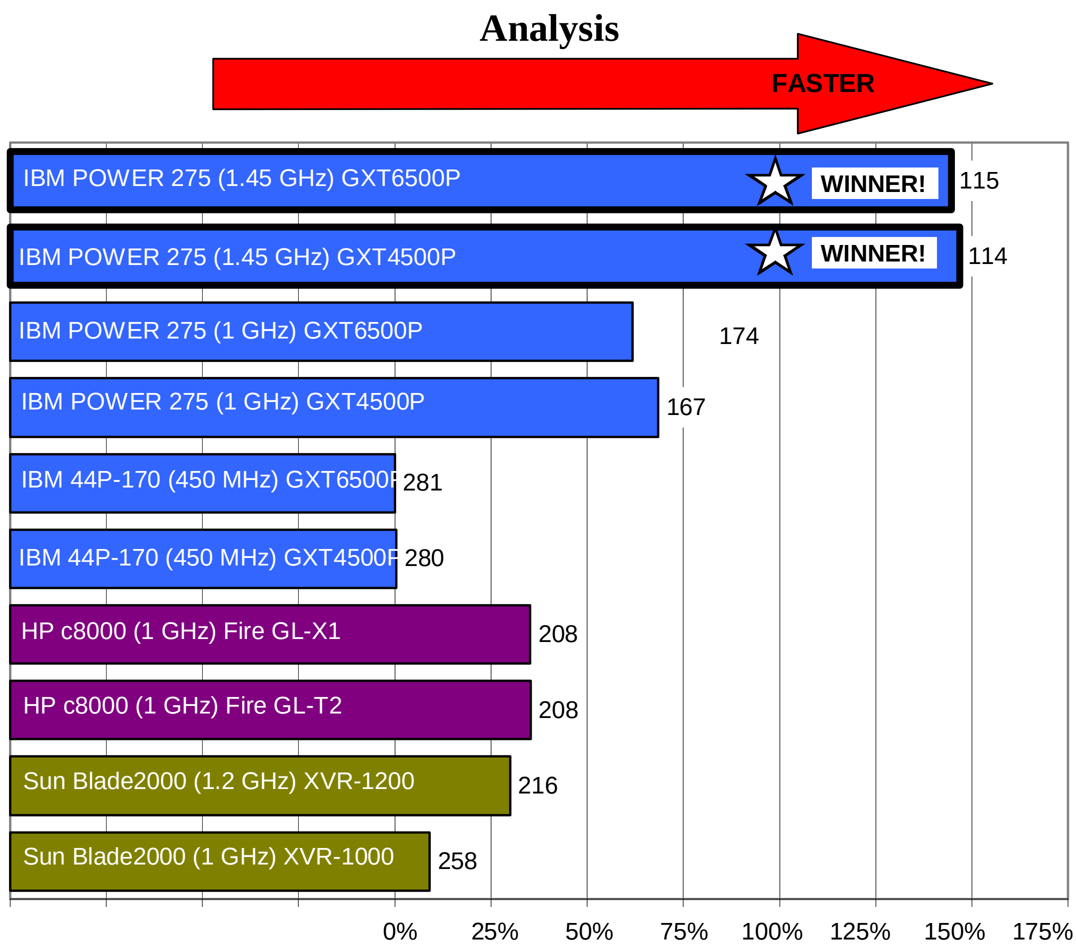

Chart 11 – Solid and Surface Analysis Function Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

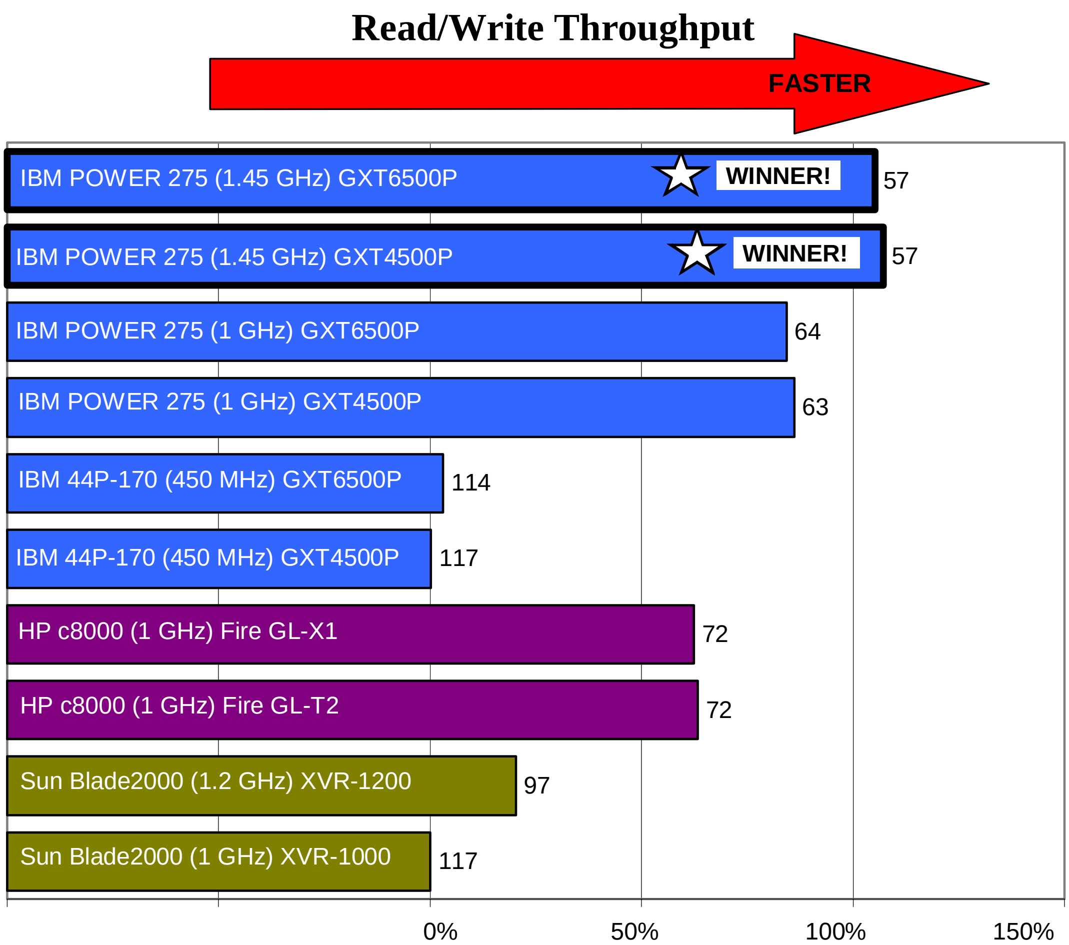

Chart 12 – Read/Write Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

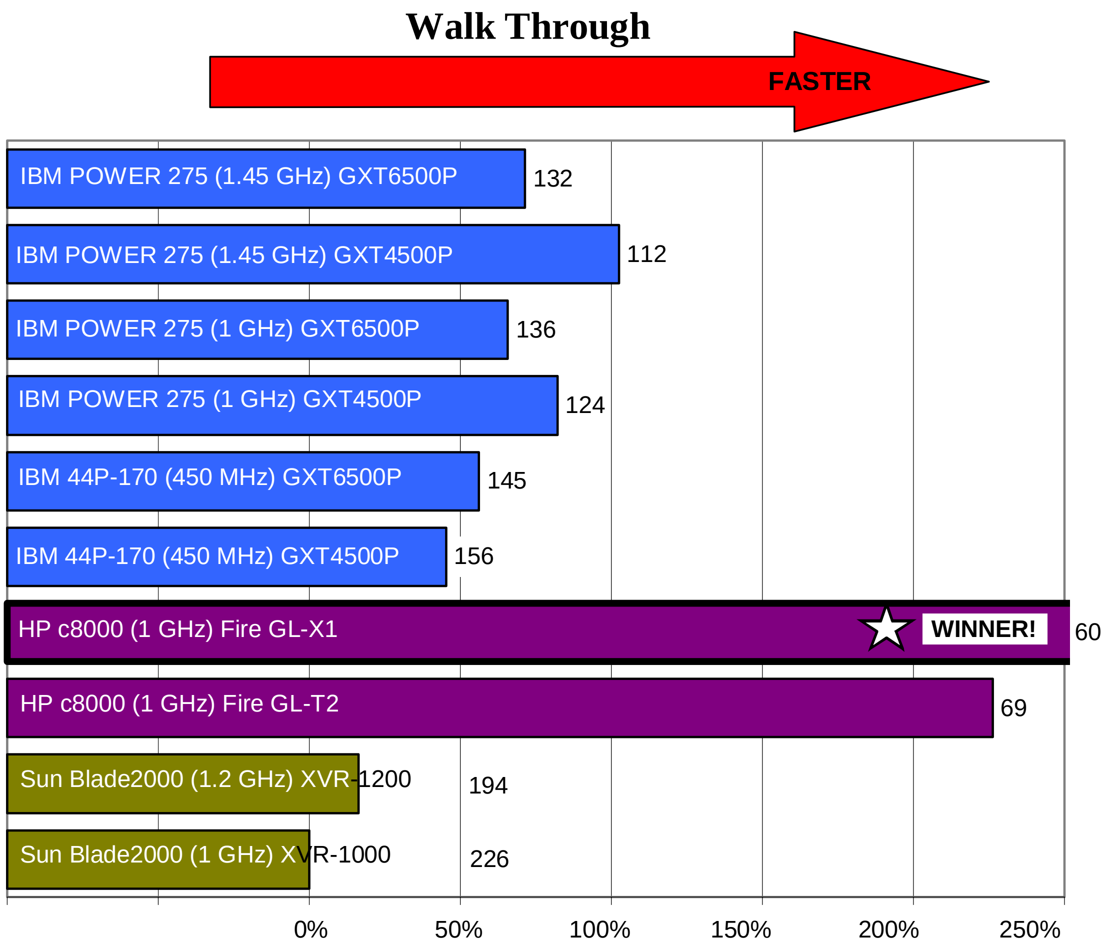

Chart 13 – Walk Through Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

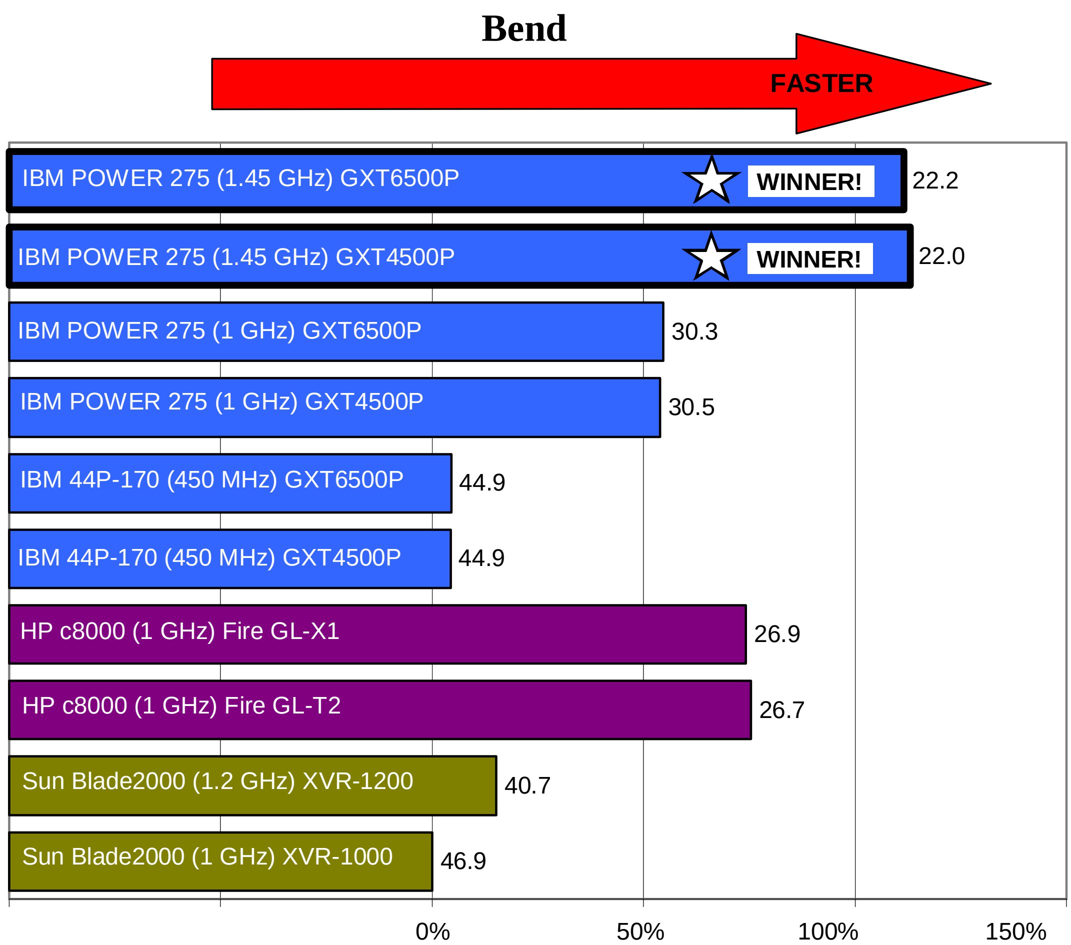

Chart 14 – Bend (Sheet Metal Part Development and Modification) Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

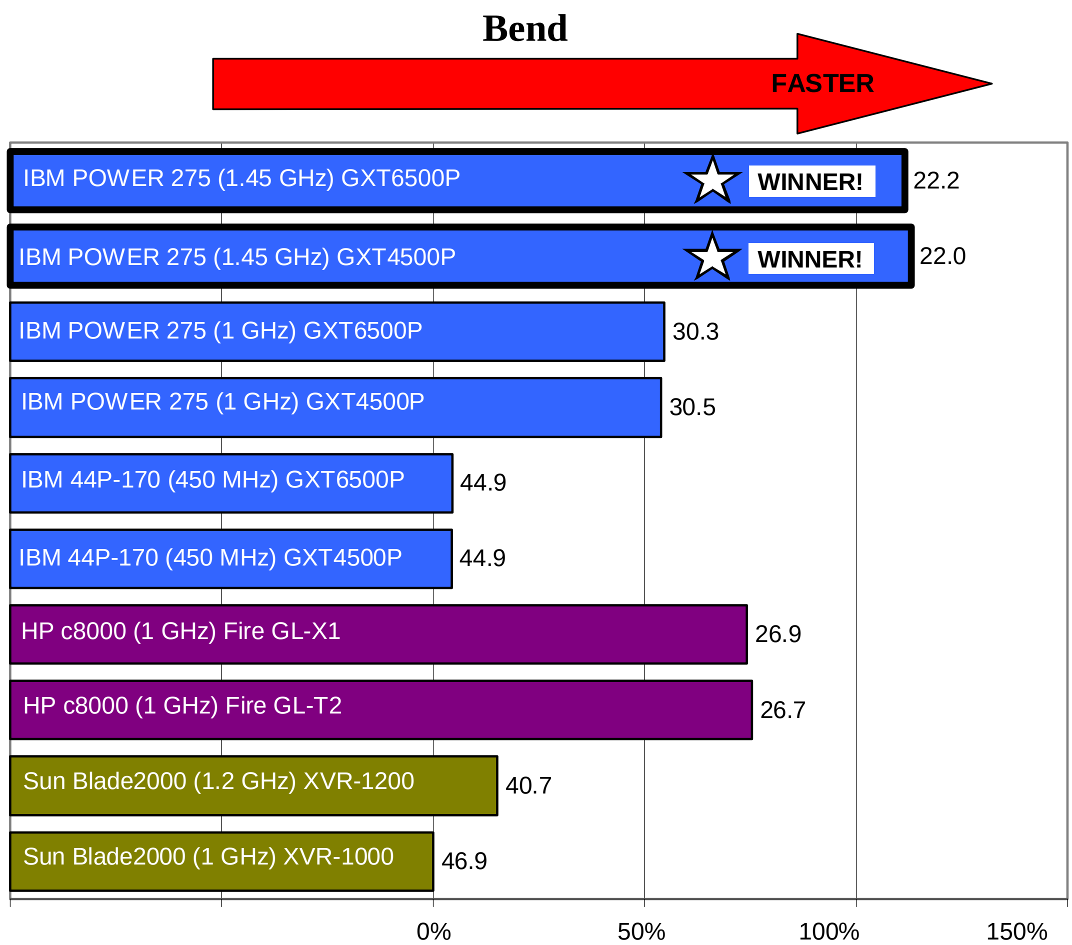

Chart 15 – Fitting Simulation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

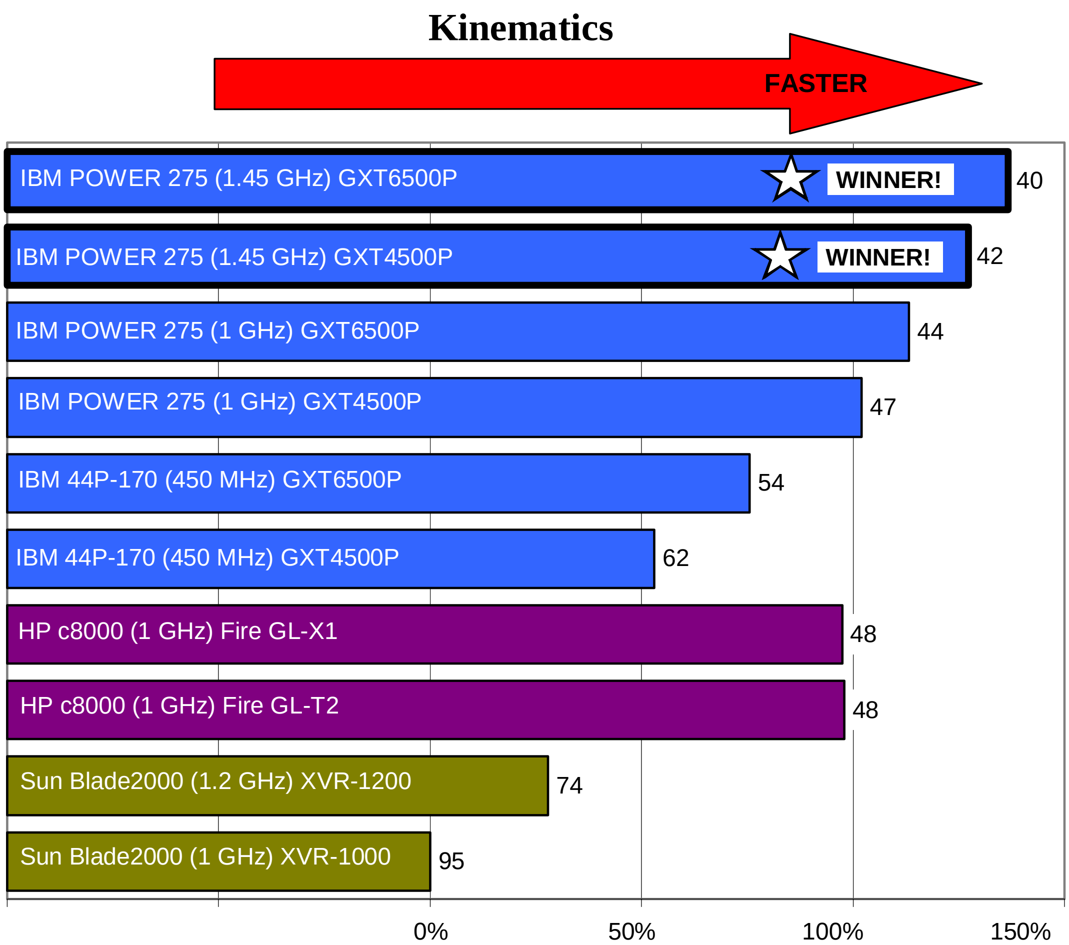

Chart 16 – Kinematics Simulation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

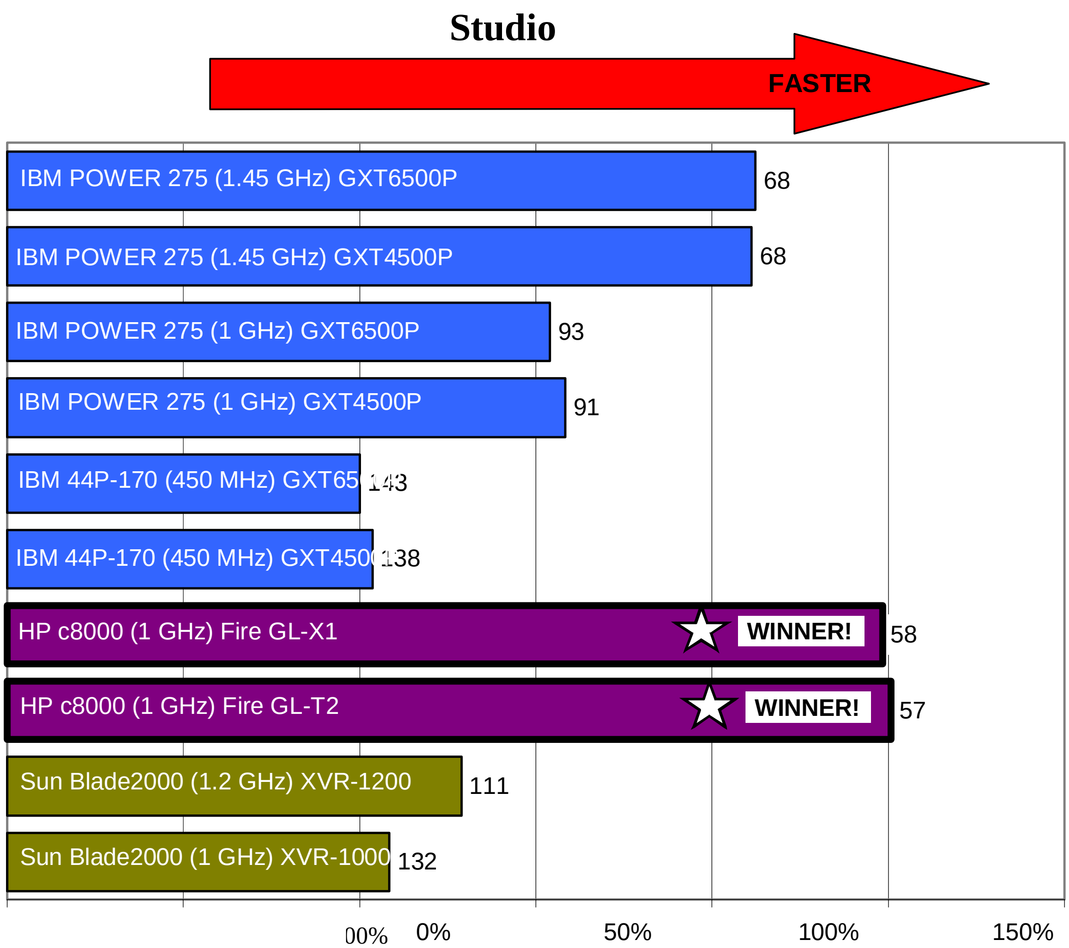

Chart 17 – Studio Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

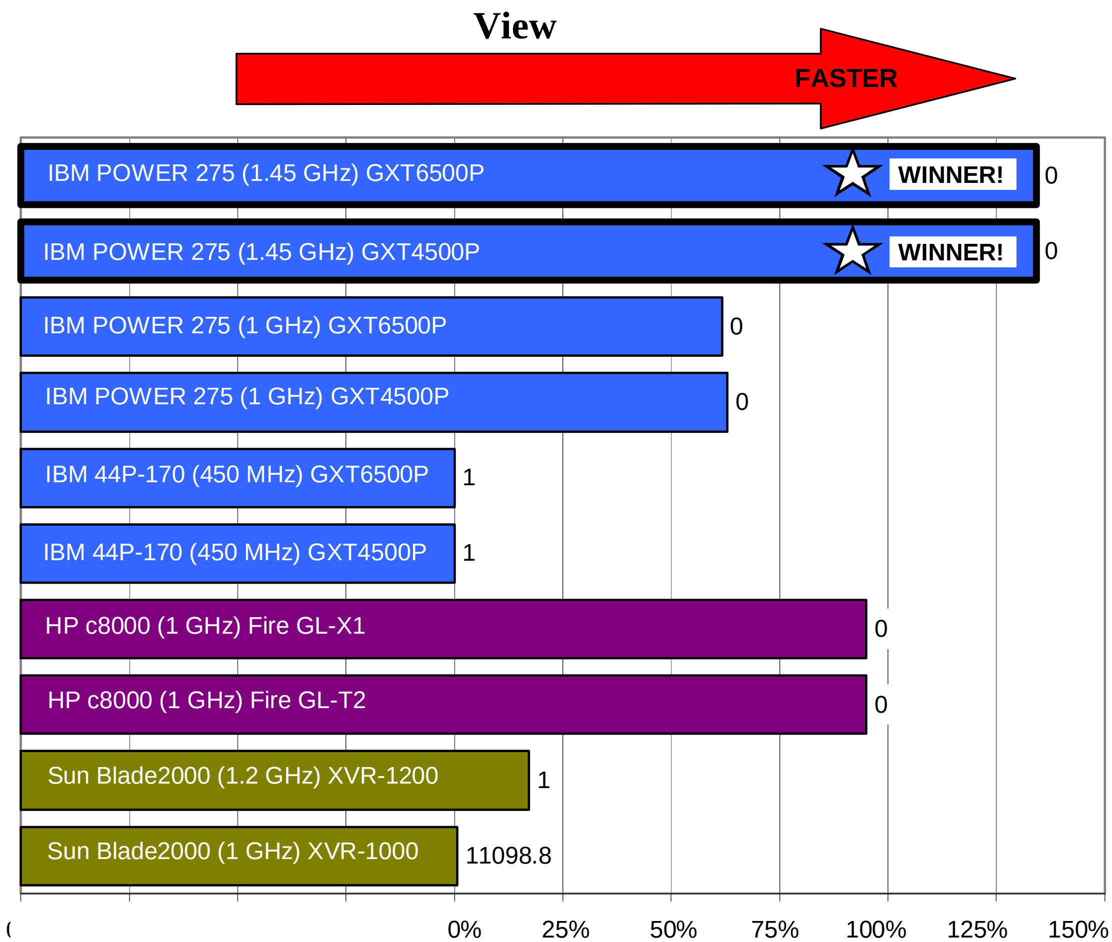

Chart 18 – Image Viewer Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)