Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.2 R2 Evaluation

Table of Contents

Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.2 R2 Evaluation

Table of Figures

Albert-Battaglin Consulting Group TAGITT/CATIA 4.2.2 R2 Evaluation

Table of Result Charts

TAGITT/CATIA
4.2.2 R2 Evaluation

abcg

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

First Quarter 2002

Notices

This document was researched and written by The Albert-Battaglin Consulting Group under a contract with IBM. The analysis, views, conclusion and recommendations in this document are those of The Albert-Battaglin Consulting Group and not necessarily those of IBM. This document is provided "AS IS" and may contain errors and inaccuracies. THE ALBERT-BATTAGLIN CONSULTING GROUP AND IBM DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE. It is your responsibility to evaluate the information contained in this document and ensure that your use is in accordance with your company's policies and practices including, in the case of IBM, the IBM Business Conduct Guidelines.

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

Executive Summary

Nine engineering workstations from Hewlett-Packard, IBM, and Silicon Graphics, Inc. (SGI) 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 4.2.2 refresh 02 version of this CAD/CAM application software package. Key results were as follows:

- The IBM IntelliStation POWER 265 (2CPU) GXT6500P was the overall performance leader in the TAGITT/CATIA 4.2.2 R2 test. The winner was closely followed (4% slower) by the single processor version of the same machine. The IBM POWER 265 (1 CPU) GXT4500P, HP c3700 (750MHz) $fx^{10}pro$ and HP c3700 (750MHz) $fx^{5}pro$ machines were tied for third place about 20% slower than the leader.
- For the first time, a dual processor machine, the IBM POWER 265 (2 CPU) GXT6500P was the fastest in overall TAGITT V4 performance. The overhead of managing two processors had only a minor impact on performance in all cases and in several tests, the two processors delivered significant performance benefits.
- For overall graphics performance, the GXT6500P was the clear leader. The IBM POWER 265 (2 CPU) GXT6500P was the fastest overall and IBM POWER 265 (1 CPU) GXT6500P was 18% slower in second place. The IBM 44P-170 (400 MHz) GXT6000P was 43% slower than the winner and the HP c3700 (750MHz) $fx^{10}pro$ was 48% slower in this series of tests.
- The single and dual CPU configurations of IBM's IntelliStation POWER 265 workstations were the leaders in CPU performance. These machines were just over 8% faster than the HP c3700 (750MHz) $fx^{10}pro$ and HP c3700 (750MHz) $fx^{5}pro$ machines.
- In the 4D Navigator test, the value of dual CPU's was clearly seen as the IBM POWER 265 (2 CPU) GXT6500P machine was 37% faster than the second place HP c3700 (750MHz) $fx^{10}pro$. The single CPU IBM POWER 265 (1 CPU) GXT6500P was 48% slower in this test than its two CPU equivalent.
- Machine performance was not the same across all tested applications. Although the IBM POWER 265 (2 CPU) GXT6500P machine was the fastest in nearly all tests, HP c3700 (750MHz) $fx^{10}pro$ also won an individual test (ANSOLID analysis). In addition the relative performance of the machines varied significantly on the different tests.

Summary of Top Three Winners in Each Test:

Throughput Summary Results

Task/Test	First	Second	Third
Overall Throughput	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT4500P** HP c3700 (750MHz) $fx^{10}pro$
Application Throughput	**IBM POWER 265 (2 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P**	HP c3700 (750MHz) $fx^{10}pro$ HP c3700 (750MHz) $fx^{5}pro$	**IBM 44P-170 (400 MHz) GXT6000P**
Graphics Throughput	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P**
System Responsiveness	**IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P** HP c3700 (750MHz) $fx^{5}pro$ HP c3700 (750MHz) $fx^{10}pro$	**IBM 44P-170 (333 MHz) GXT4000P**
CPU Throughput	**IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (2 CPU) GXT6500P**	HP c3700 (750MHz) $fx^{5}pro$ HP c3700 (750MHz) $fx^{10}pro$	**IBM 44P-170 (400 MHz) GXT6000P**
Dynamic Graphics Throughput	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P**
4D Navigator	**IBM POWER 265 (2 CPU) GXT6500P**	HP c3700 (750MHz) $fx^{10}pro$	**IBM POWER 265 (1 CPU) GXT6500P**

Individual Application Test Throughput Results:

Primary Application Test Throughput Results

Task/Test	**First**	**Second**	**Third**
Modeling Solid Model Creation and Modification	**IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P**	HP c3700 (750MHz) *fx⁵pro* HP c3700 (750MHz) *fx¹⁰pro*	**IBM 44P-170 (400 MHz) GXT6000P**
Finite Element Analysis (ANSOLID)	HP c3700 (750MHz) *fx¹⁰pro* HP c3700 (750MHz) *fx⁵pro*	**IBM POWER 265 (2 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (1 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P**
NC Tool Path Generation Throughput	**IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (2 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P**	HP c3700 (750MHz) *fx⁵pro* HP c3700 (750MHz) *fx¹⁰pro*
Detail Drawing Creation	**IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P**	HP c3700 (750MHz) *fx⁵pro* HP c3700 (750MHz) *fx¹⁰pro*	**IBM 44P-170 (400 MHz) GXT6000P**

Secondary Application Test Throughput Results			
Task/Test	**First**	**Second**	**Third**
Solid and Surface Analysis Function	**IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT6500P**	HP c3700 (750MHz) *fx^{5}pro* HP c3700 (750MHz) *fx^{10}pro* Sun Blade 1000 (900MHz) Expert3D	**IBM 44P-170 (400 MHz) GXT6000P**
Read/Write	**IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P**	**IBM 44P-170 (400 MHz) GXT6000P**	Sun Blade 1000 (900MHz) Expert3D
Walk Through	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT4500P** **IBM 44P-170 (400 MHz) GXT6000P**
Bend (Sheet Metal Part Development and Modification)	**IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (2 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P**	HP c3700 (750MHz) *fx^{5}pro* HP c3700 (750MHz) *fx^{10}pro*
Fitting Simulation	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P** HP c3700 (750MHz) *fx^{5}pro* HP c3700 (750MHz) *fx^{10}pro*
Kinematics Simulation	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P**	**IBM 44P-170 (400 MHz) GXT6000P** **IBM POWER 265 (1 CPU) GXT4500P**
Studio	**IBM POWER 265 (2 CPU) GXT6500P**	**IBM POWER 265 (1 CPU) GXT6500P** **IBM POWER 265 (1 CPU) GXT4500P**	Sun Blade 1000 (900MHz) Expert3D HP c3700 (750MHz) *fx^{10}pro*
Image Viewer	**IBM POWER 265 (2 CPU) GXT6500P**	HP c3700 (750MHz) *fx^{5}pro* HP c3700 (750MHz) *fx^{10}pro*	**IBM POWER 265 (1 CPU) GXT4500P** **IBM POWER 265 (1 CPU) 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, 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.

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 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, kinematic and "walk through" analysis, drawing view creation from 3D models, NC tool path generation, 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 range between 0.5 and 5 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.

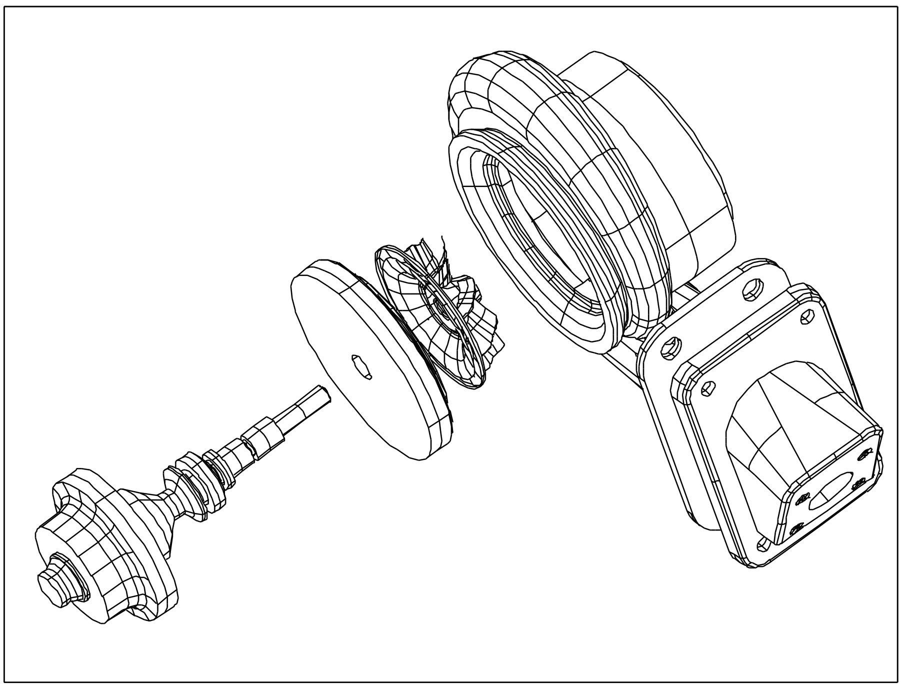

Figure 1 – Albert-Battaglin Consulting Group Turbo Charger Assembly

Models used for TAGITT evaluations consist mostly of real models taken from customer data. In order to provide some level of cross-application comparison, a subset of each TAGITT evaluation consists of one or more parts from the Albert-Battaglin Consulting Group Turbo Charger Assembly depicted in Figure 1. Although conceptually simple, the assembly's geometries still require enough computation so as to make performance differentiation possible. In addition, the simplicity of the models forces operations to be undertaken in a similar fashion on various applications, thus providing better comparative data.

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. Sixty program files are used in the 4.2.2 refresh 2 version of TAGITT/CATIA. About one-fifth were created by Albert-Battaglin Consulting Group to construct, modify and manipulate the Compressor Wheel and Test Fixture parts in a fashion which as closely as possible matches the method used for other TAGITT testing with other leading CAD/CAM systems. The remaining program files were either 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 2. These files cover many areas of the CATIA 4.2.2 R2 product including part modeling, surface intersection, drawing layout, parametric modification and updates, 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.

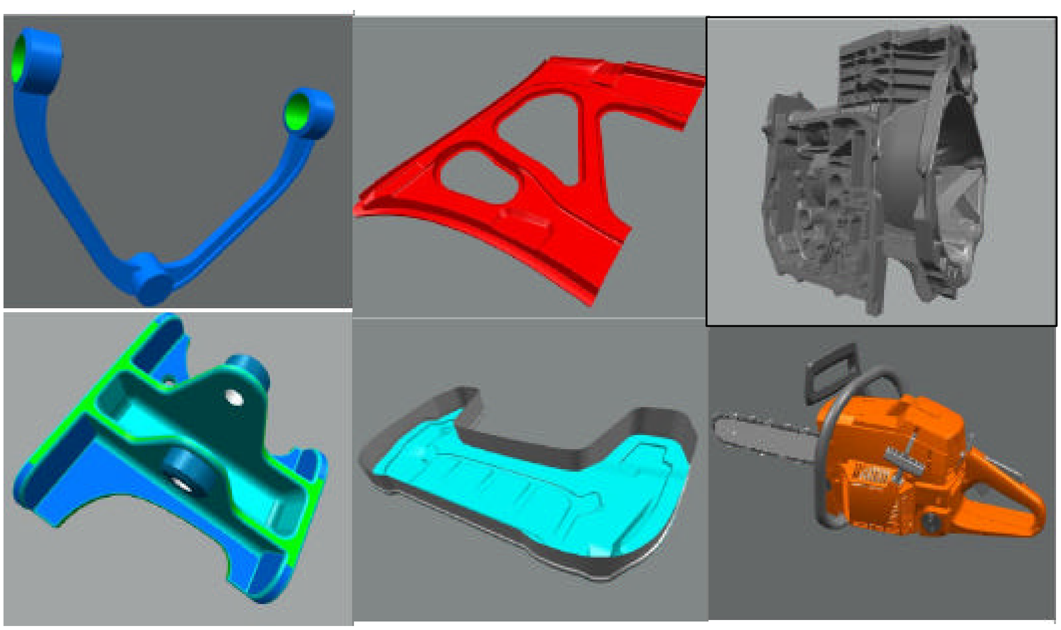

Figure 2 – Some of the CATIA models used for TAGITT/CATIA

This version of TAGITT/CATIA 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 that 0.5 seconds). Some of the models used for the 4D Navigator test are shown in Figure 3.

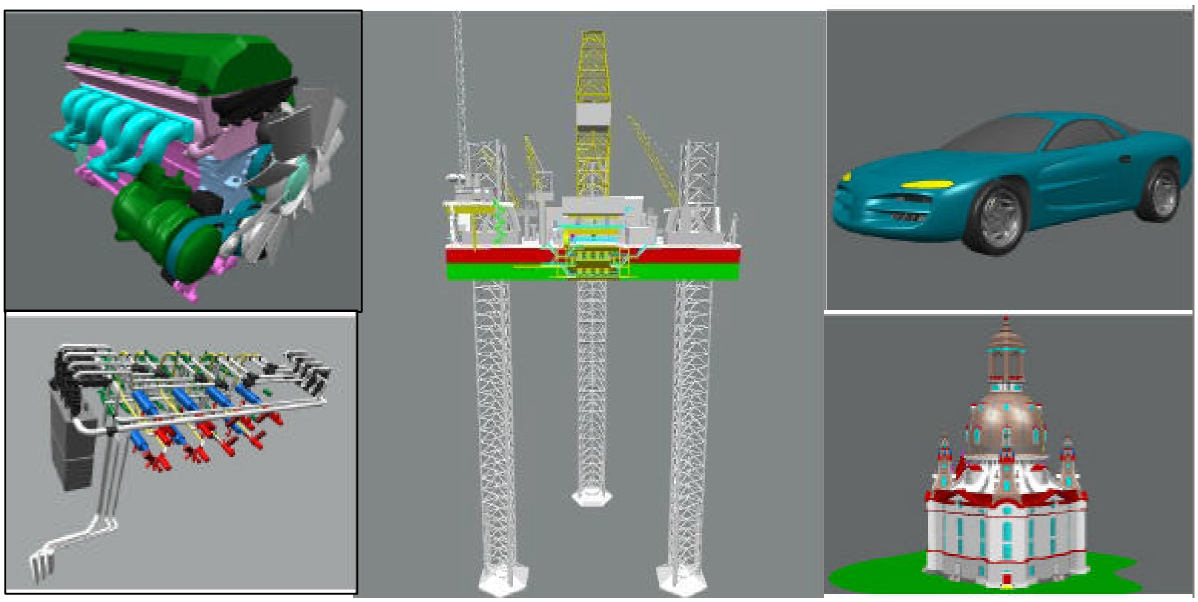

Figure 3 – CATIA 4D Navigator Test Modules

Test Weighting

While running the TAGITT/CATIA 4.2.2 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 4.

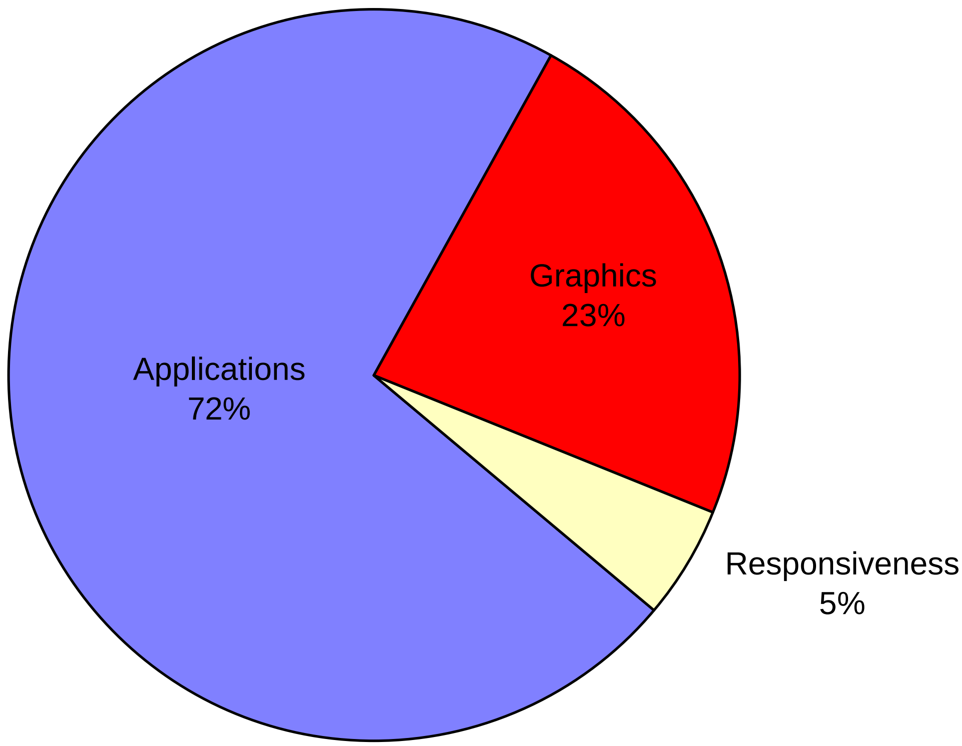

Figure 4 – Overall Throughput Weighting

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 5.

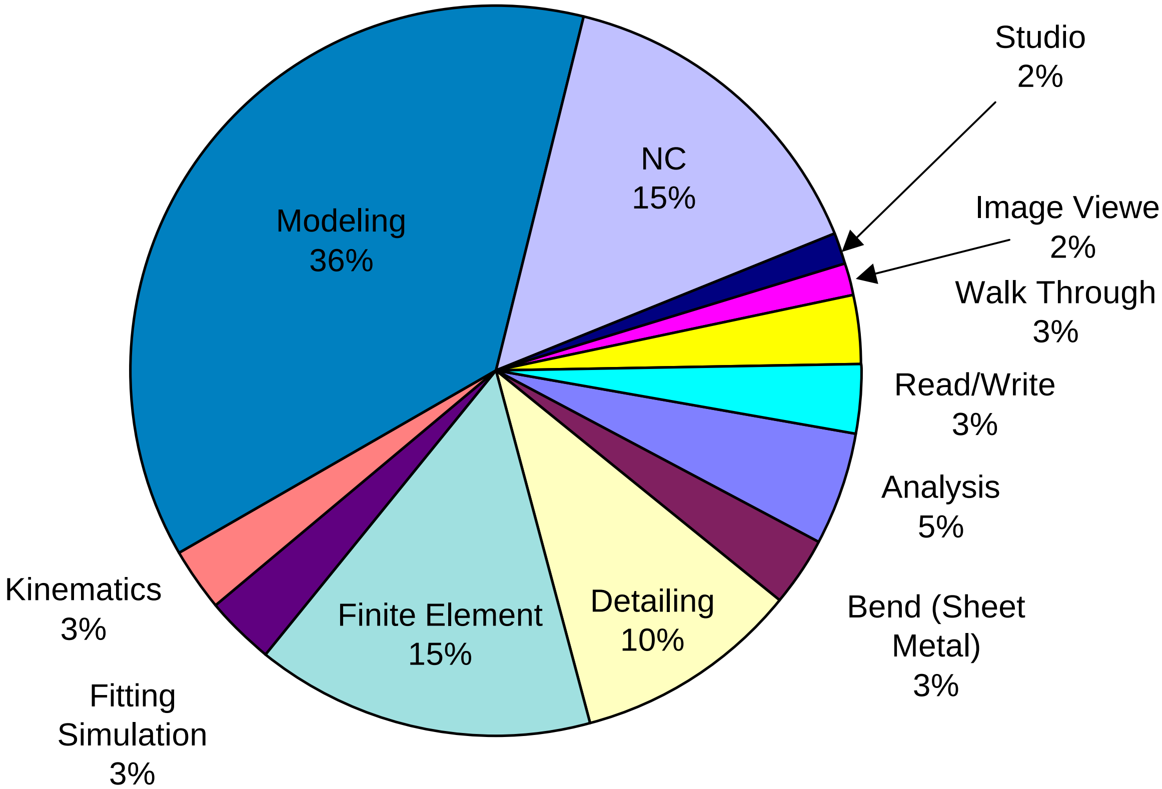

Figure 5 – Application Weighting Factors

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 6.

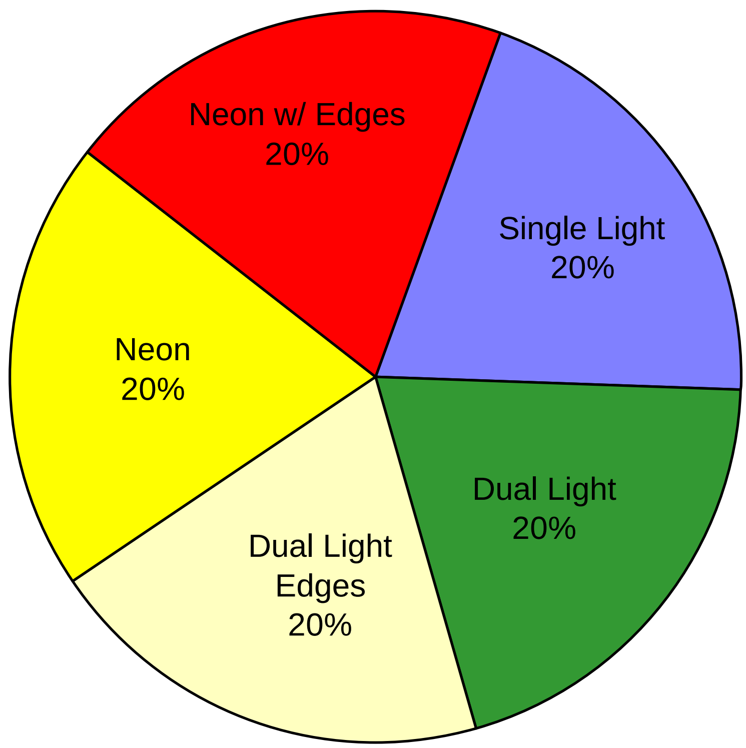

Figure 6 – 4D Navigator Weighting Factors

Graphics throughput is calculated based on the weighted times for the various graphics tests (refer to Figure 7) 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 (15) models used throughout the testing.

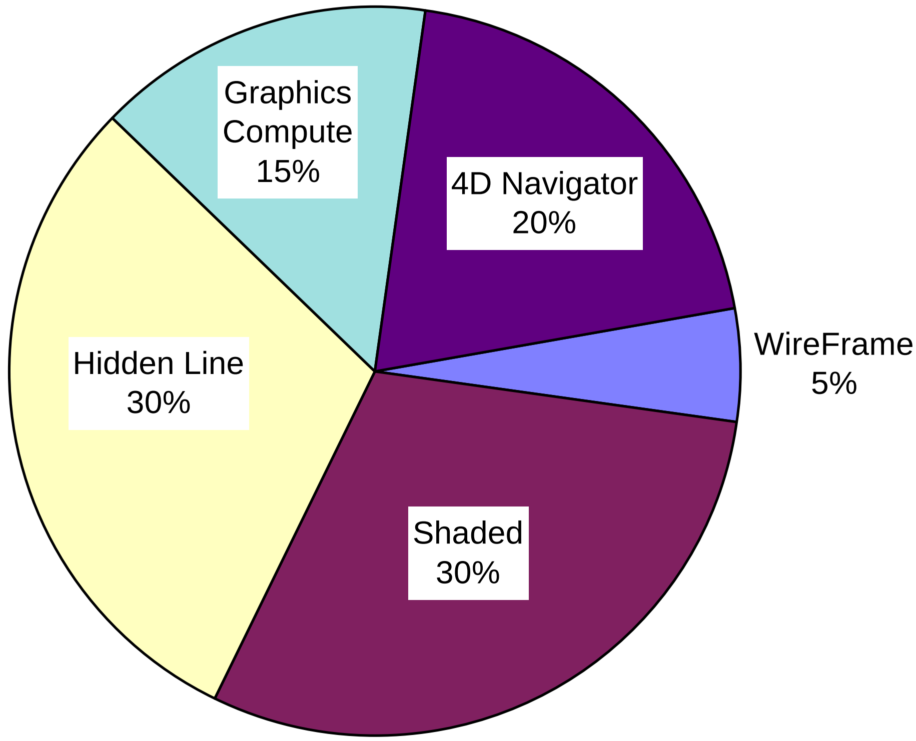

Figure 7 – Graphics Weighting Factors

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 512 MB main memory, 1 GB swap, 20”/21” color monitor, CD 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 SGI machines include the graphics accelerators listed below. IBM IntelliStation POWER 265 machines were available and tested only with 1GB RAM.

Vendor	**Model**	**Graphics**	**CPU**	**L2 Cache Mb**	**Clock MHz**	**OS**	**Abbreviation used in report and on charts**
IBM	IntelliStation POWER 265	GXT6500P	Power3-II (x2)	4	450	AIX4.3.3	IBM POWER 265 (2 CPU) GXT6500P
IBM	IntelliStation POWER265	GXT6500P	Power3-II	4	450	AIX4.3.3	IBM POWER 265 (1 CPU) GXT6500P
IBM	IntelliStation POWER 265	GXT4500P	Power3-II	4	450	AIX4.3.3	IBM POWER 265 (1 CPU) GXT4500P
IBM	44P 170 (400)	GXT6000P	Power3-II	4	400	AIX4.3.3	IBM 44P-170 (400 MHz) GXT6000P
IBM	44P 170 (333)	GXT4000P	Power3-II	1	333	AIX4.3.3	IBM 44P-170 (333 MHz) GXT4000P
HP	c3700	$fx^{10}pro$	PA-RISC 8700	2.25	750	HPUX11.00	HP c3700 (750MHz) $fx^{10}pro$
HP	c3700	$fx^{5}pro$	PA-RISC 8700	2.25	750	HPUX11.00	HP c3700 (750MHz) $fx^{5}pro$
Sun	Blade 1000	Expert3D	UltraSparc-III	4	900	Solaris 8	Sun Blade 1000 (900MHz) Expert3D
SGI	Octane 2	VPro V6	R12000A	2	400	IRIX64 6.5	SGI Octane2 (400MHz) VPro V6

TAGITT/CATIA 4.2.2 R2 Results

Overall Throughput

Chart 1 compares the overall weighted elapsed time to complete the TAGITT/CATIA 4.2.2 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 265 (2 CPU) GXT6500P machine to be the fastest machine overall, 4% faster than the second place IBM POWER 265 (1 CPU) GXT6500P. The IBM POWER 265 (1 CPU) GXT4500P and HP c3700 (750MHz) *fx*^{10}*pro* tied for third place with performance about 20% slower than the fastest machine.

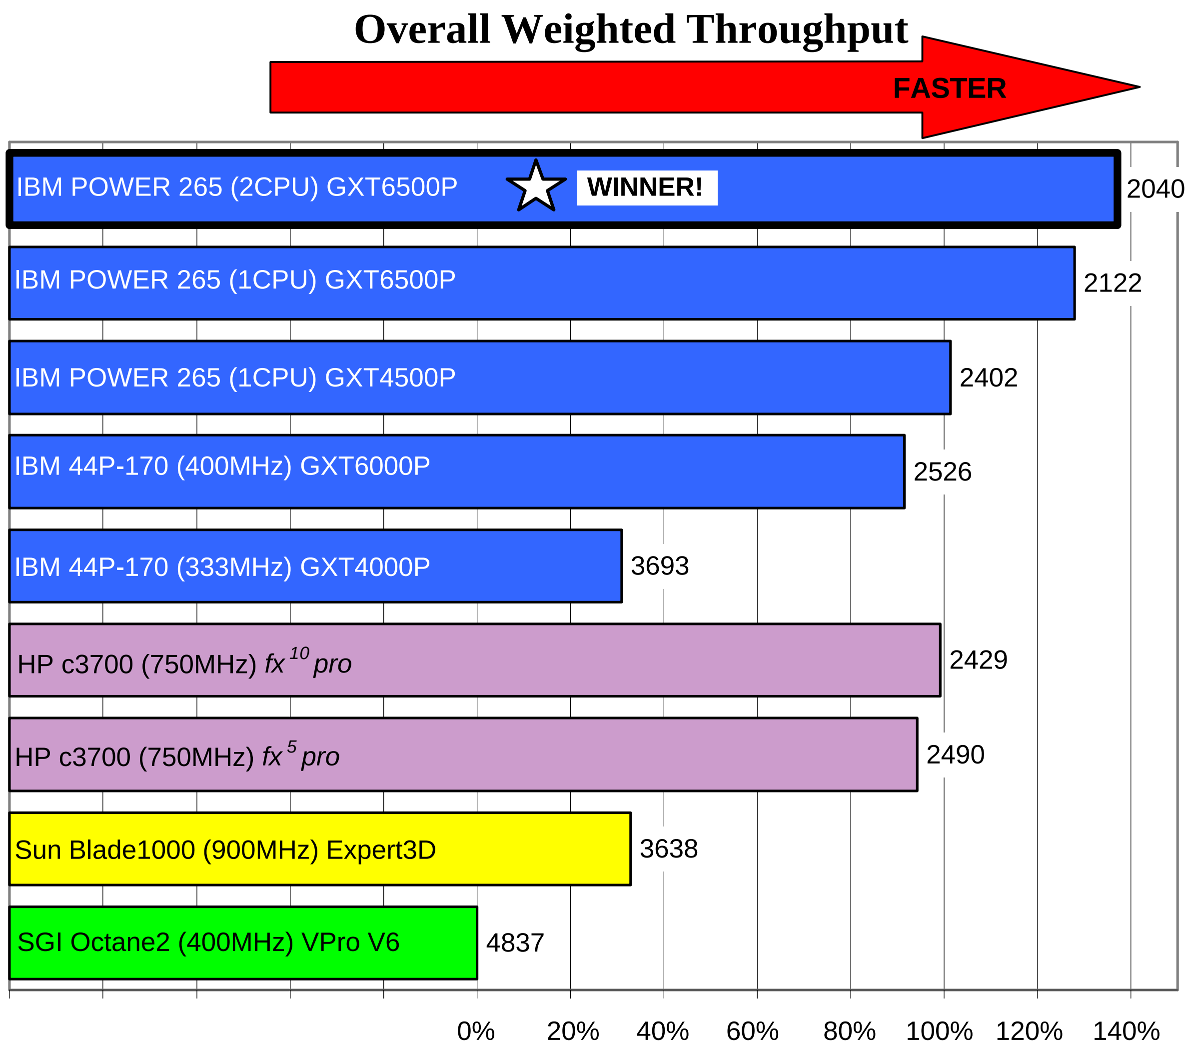

Chart 1 – Overall Weighted Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

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 IBM POWER 265 (2 CPU) GXT6500P was clearly the fastest machine in this test with a time 18% faster than the second place IBM POWER 265 (1 CPU) GXT6500P. The third fastest machine, the IBM 44P-170 (400 MHz) GXT6000P, was 43% slower than the leader. Close behind in forth place, the HP c3700 (750MHz) *fx*[10]*pro* machine was 48% slower than the leader.

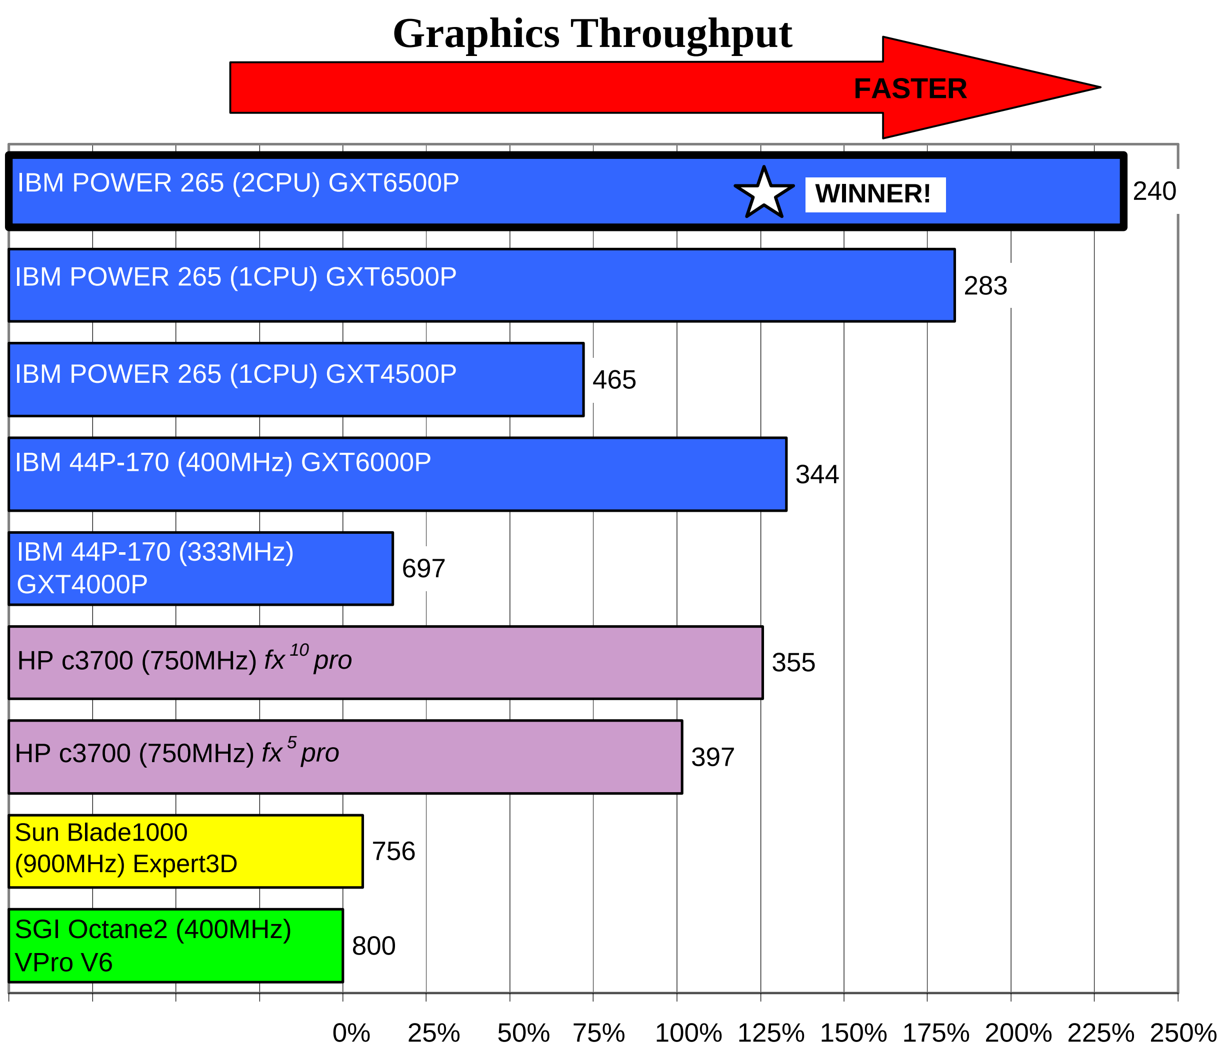

Chart 2 – Graphics Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Application Throughput

Chart 3 shows the weighted cumulative time to complete all of the application-specific tasks in the benchmark such as solid modeling operations, drafting and detailing operations, FEM functions and NC computations. The IBM IntelliStation POWER 265 workstations were the leaders in this test with times within 3% of each other. In a two way tie for second place, the HP c3700 (750MHz) *fx*^{5}*pro* and HP c3700 (750MHz) *fx*^{10}*pro* machines were 12% slower than the leader. The third place IBM 44P-170 (400 MHz) GXT6000P machine finished 19% slower than the leader.

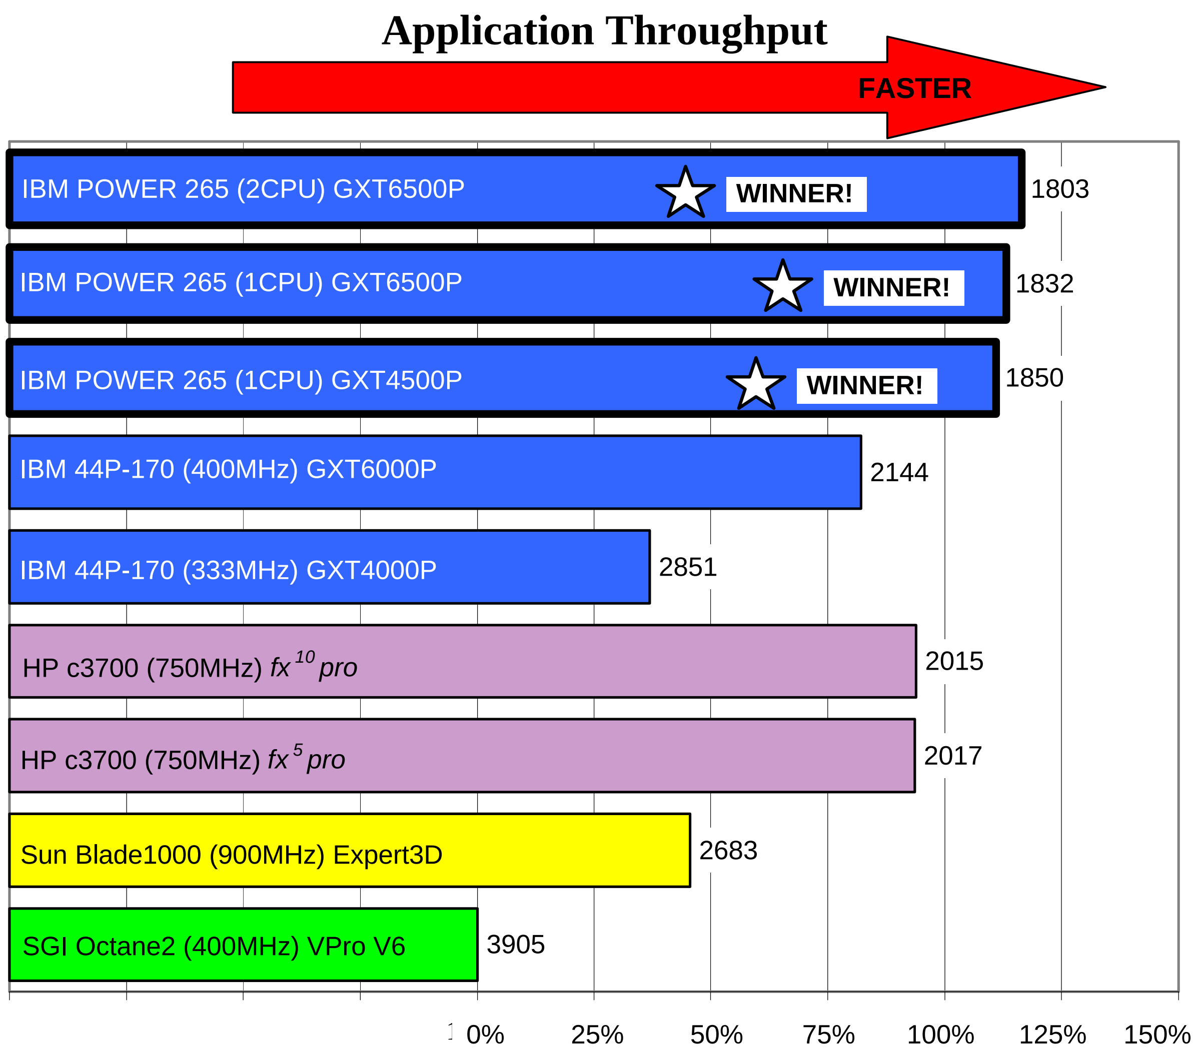

Chart 3 – Application Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

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 three new IBM POWER 265 workstations were the overall winners of this test. Tied for second place, the IBM 44P-170 (400 MHz) GXT6000P, HP c3700 (750MHz) *fx*^{5}*pro* and HP c3700 (750MHz) *fx*^{10}*pro* machines were a significant 40% slower than the leading machines. In third place, the IBM 44P-170 (333 MHz) GXT4000P was about 78% slower than the fastest machine.

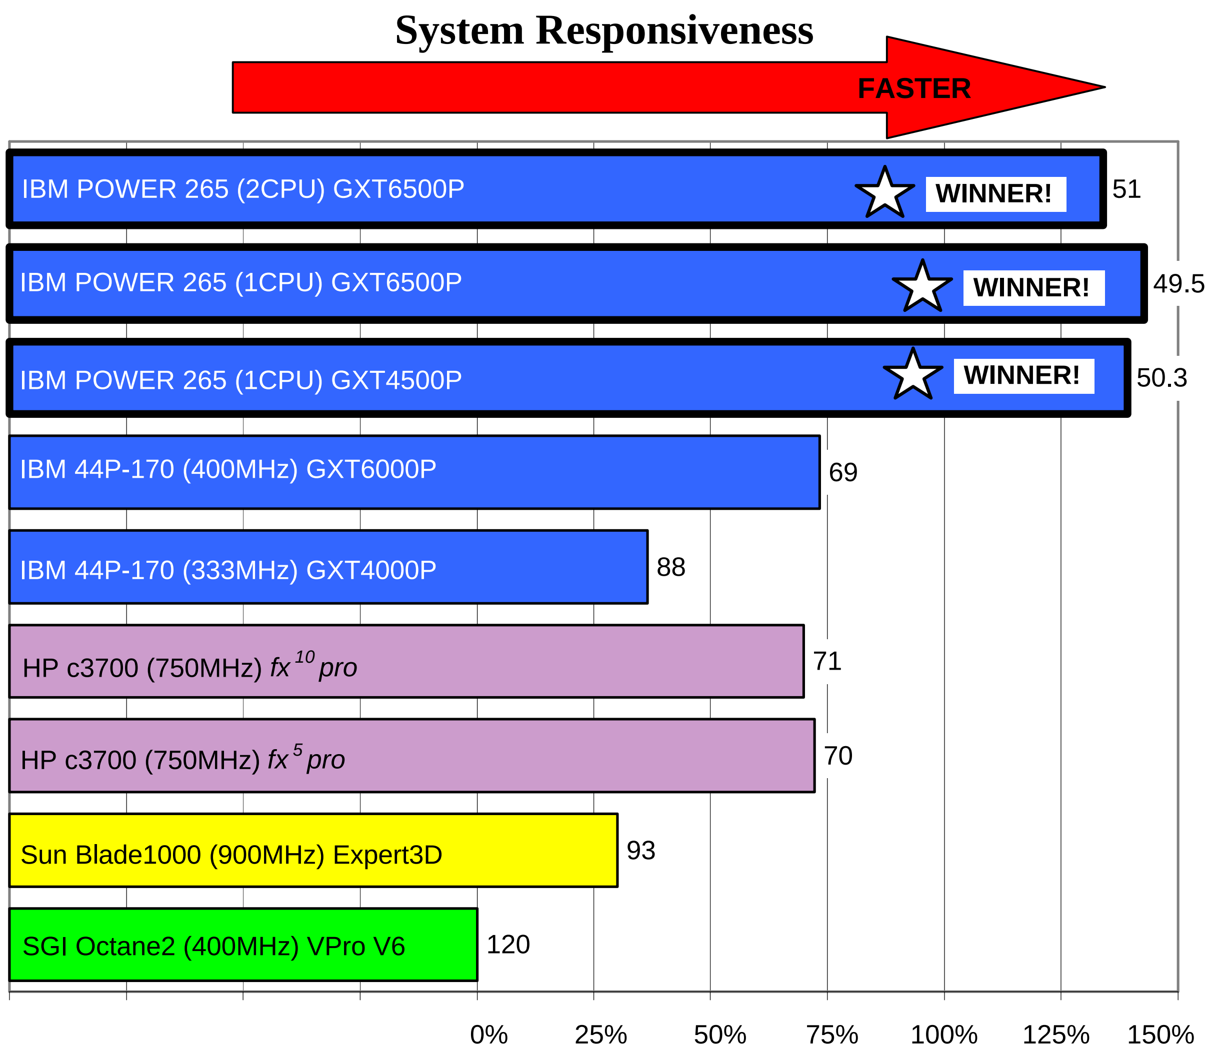

Chart 4 – System Responsiveness Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

CPU Throughput

Chart 5 shows the cumulative time to complete all of the CPU intensive tasks in the benchmark. The overall winners of this test were again the IBM IntelliStation POWER 265 workstations. In second place, the HP c3700 (750MHz) *fx*^{5}*pro* and HP c3700 (750MHz) *fx*^{10}*pro* were 8%-9% slower than the winner. In third place, the IBM 44P-170 (400 MHz) GXT6000P machine finished 17% slower than the leading machine.

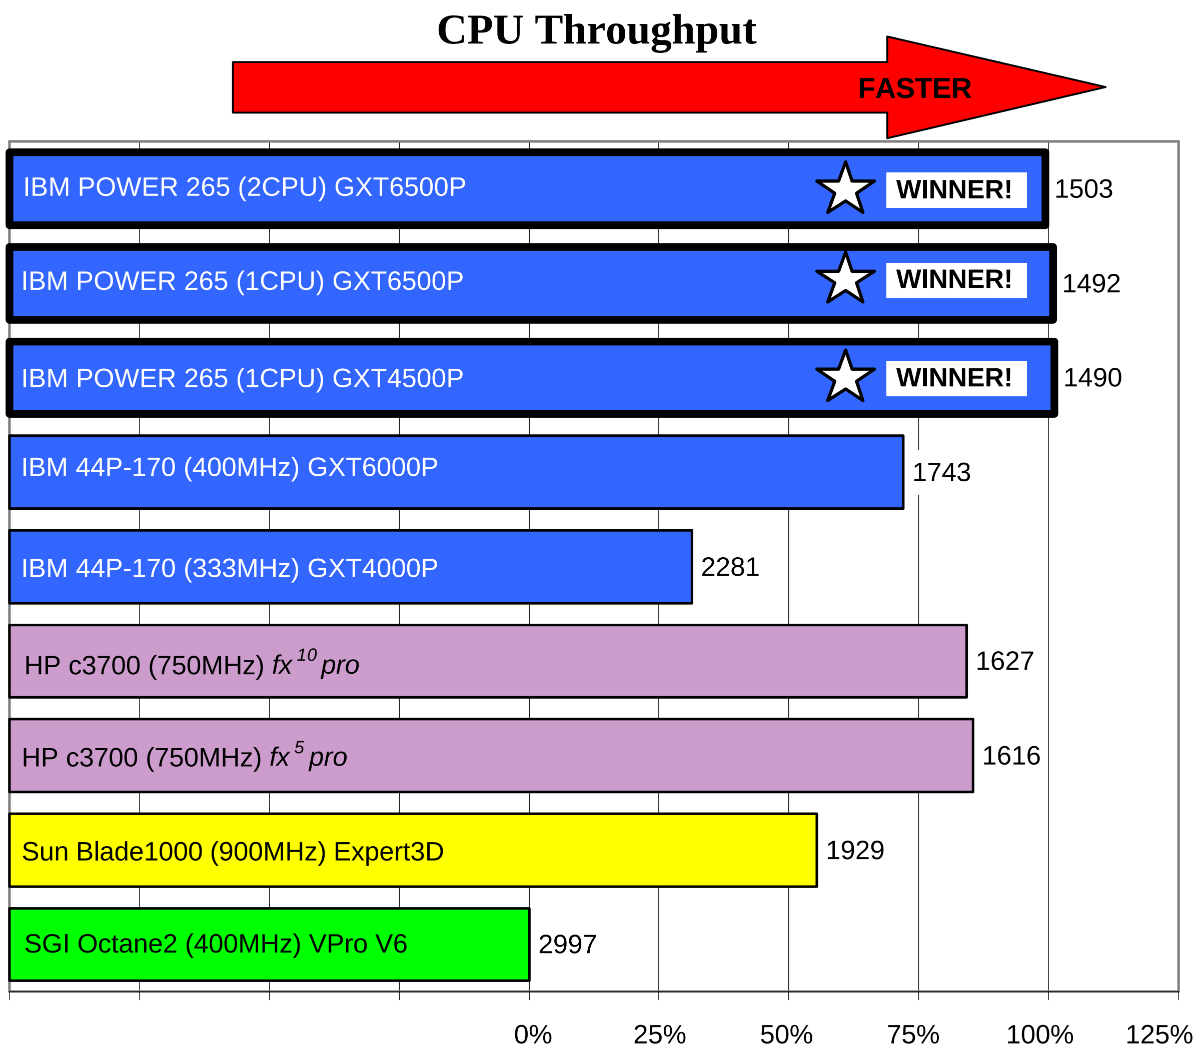

Chart 5 – CPU Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

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 clear winner of this test was the IBM POWER 265 (2 CPU) GXT6500P that was 16% faster than the IBM POWER 265 (1 CPU) GXT6500P in second place. The IBM 44P-170 (400 MHz) GXT6000P machine was third with times 42% slower than the leader. The HP c3700 (750MHz) *fx^{10}pro* machine finished 60% slower than the leader.

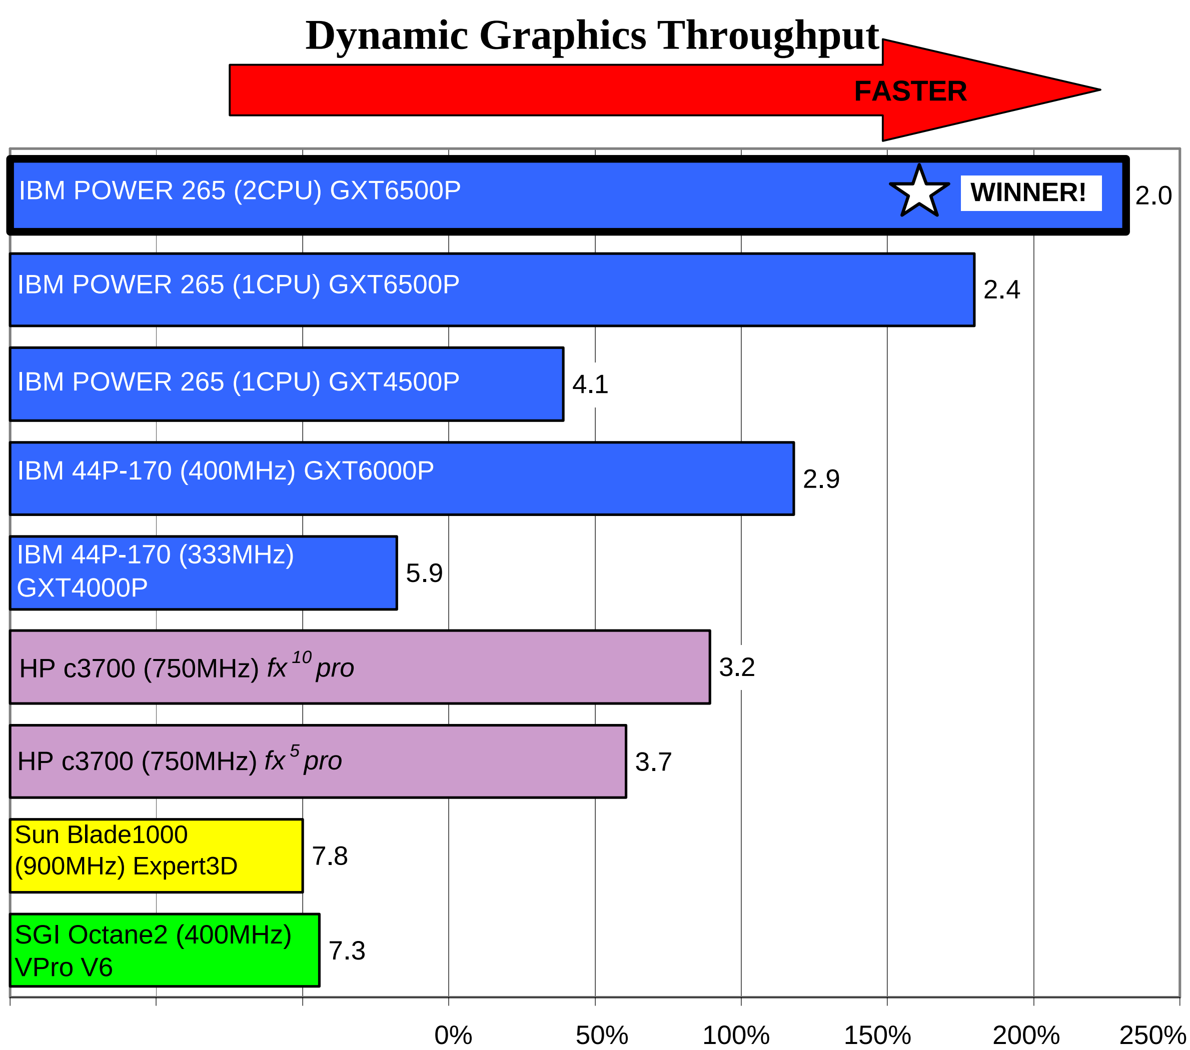

Chart 6 – Dynamic Graphic Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

4D Navigator Throughput

Chart 7 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 265 (2 CPU) GXT6500P by a clear margin. In second place, the HP c3700 (750MHz) *fx[10]pro* machine was 37% slower than the leader. The third place IBM POWER 265 (1 CPU) GXT6500P was 48% slower than the leader showing the benefit of using 2 CPUs for 4D Navigator work.

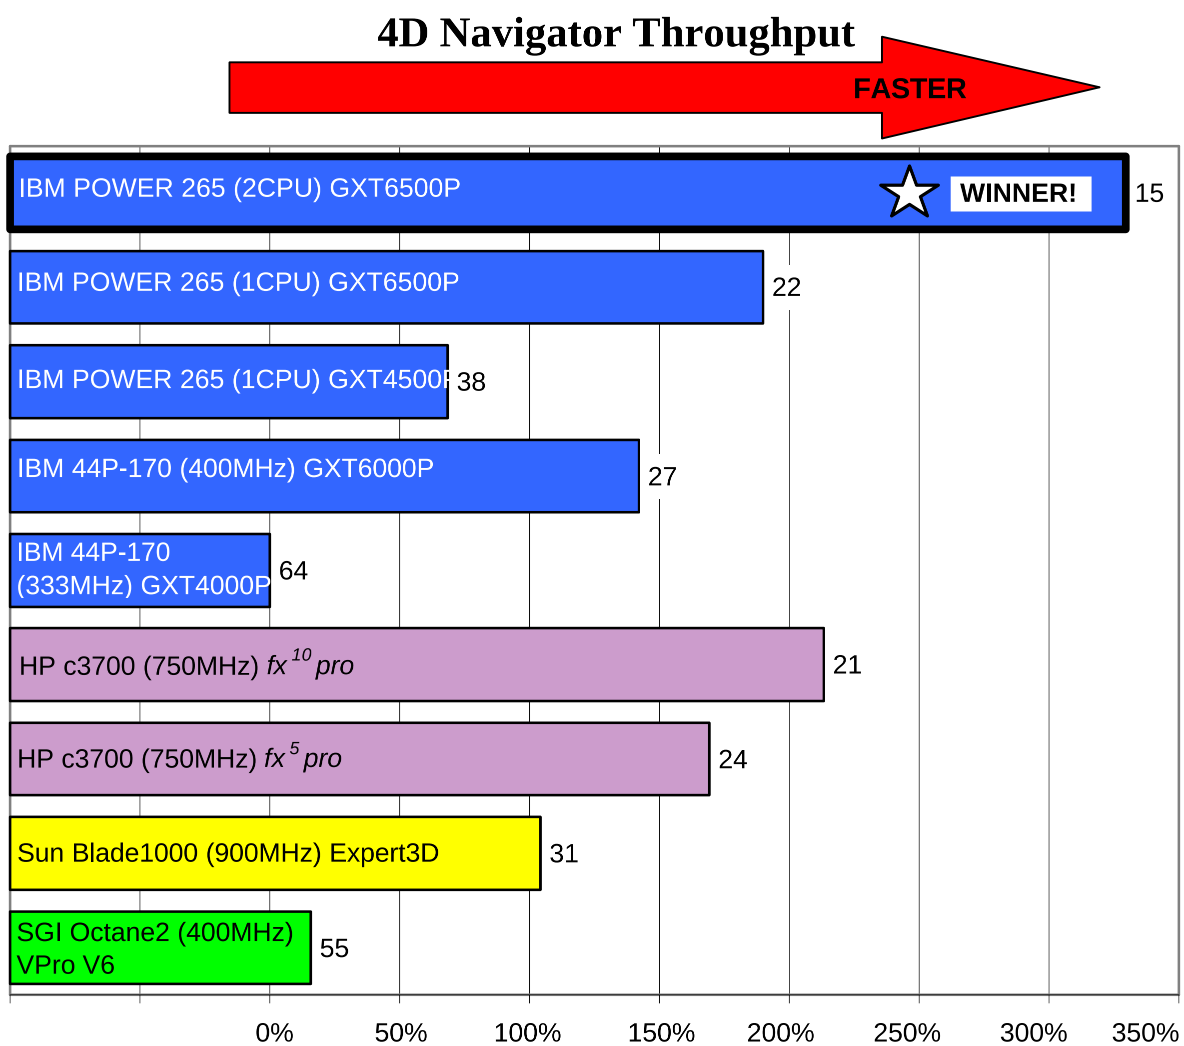

Chart 7 – 4D Navigator Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

Conclusions

The results of the TAGITT/CATIA 4.2.2 refresh 02 evaluation show IBM to have the overall performance lead with its IntelliStation POWER 265 machines based on the Power3-II architecture. The IBM POWER 265 (2 CPU) GXT6500P was a significant 19% faster overall than the nearest competitive machine, the HP c3700 (750MHz) *fx¹⁰pro*. The Sun Blade 1000 (900MHz) Expert3D was also no match for the new IBM machines with overall performance times a disappointing 78% slower than the leader. The single CPU IBM POWER 265 (1 CPU) GXT6500P was only 4% slower than the 2 CPU machine overall because multi-threading is used in only a few places in CATIA V4.

In terms of graphic performance, IBM had the fastest overall graphics performance with its new GXT6500P graphics adapter. The GXT6500P outperformed its nearest competitor (the *fx¹⁰pro*) in all tested configurations. In the IBM POWER 265 (2 CPU) GXT6500P configuration, the machine outperformed the nearest competitor, the HP c3700 (750MHz) *fx¹⁰pro*, by 48% in overall graphics throughput and 60% in dynamic graphics throughput.

Graphics performance is an area where IBM's use of dual processors showed clear benefit. The excellent performance of the GXT6500P graphic adapter appears well matched with the dual Power3-II (450MHz) CPUs in the IBM POWER 265 (2 CPU) GXT6500P for excellent overall results. This was particularly evident in the 4D Navigator test. In the 4D Navigator test, the IBM POWER 265 (2 CPU) GXT6500P outperformed the IBM POWER 265 (1 CPU) GXT6500P that has the same graphics card by a significant 48%.

The wide range of application testing in TAGITT/CATIA 4.2.2 R2 showed performance differences between various CATIA functions. It is impossible to tell from the TAGITT evaluation the reasons for these differences. Although the IBM POWER 265 (2 CPU) GXT6500P machine won nearly all of the individual application tests by a significant margin, there was one notable exception. The HP c3700 (750MHz) *fx¹⁰pro* and HP c3700 (750MHz) *fx⁵pro* showed the best ANSOLID performance by winning the FEM tests by about 10%. In addition, the margin of victory for the IBM machines varied between the various applications. In the Detailing test, for example, the IBM POWER 265 (2 CPU) GXT6500P machine was 47% faster than the Sun Blade 1000 (900MHz) Expert3D while in the model Analysis test the difference was only 8%.

For the first time in TAGITT testing, a dual processor machine was the overall performance winner. In the past, the overhead of managing dual processors has overshadowed the benefits when running CATIA V4. With the IBM IntelliStation POWER 265 workstations this was not the case. Even in non-multi threaded tasks such as modeling, the 2 CPU configuration was nearly (within 1%) as fast as the single processor configuration. The two dual processor machines tested showed clear performance benefits on functions within CATIA V4 that are written to be multi-threaded. In the overall 4D Navigator results, the IBM POWER 265 (2 CPU) GXT6500P outperformed the best single processor system, the HP c3700 (750MHz) *fx*^{10}*pro* by 37% and its single CPU cousin, the IBM POWER 265 (1 CPU) GXT6500P by 48%. The effect was more pronounced in the ray-tracing "Studio" test and the dual application Image Viewer test where the IBM POWER 265 (2 CPU) GXT6500P outperformed the IBM POWER 265 (1 CPU) GXT6500P and GXT4500P by 76% and 97% respectively.

IBM has maintained its CATIA V4 performance leadership with its new IntelliStation POWER 265 workstations. The TAGITT/CATIA 4.2.2 R2 test shows the IBM POWER 265 (2 CPU) GXT6500P and IBM POWER 265 (1 CPU) GXT6500P to be excellent all around performers in the CATIA environment. The IBM POWER 265 (2 CPU) GXT6500P shows significant benefit in performance with specific functions within CATIA such as 4D Navigator, Studio and Image Viewer. The combination of raw CPU speed and excellent graphic performance in the IBM IntelliStation POWER 265 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 4.2.2 R2 tests were completed using released CATIA 4.2.2 refresh 02 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 R4.2.2 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 4.2.2 R2.

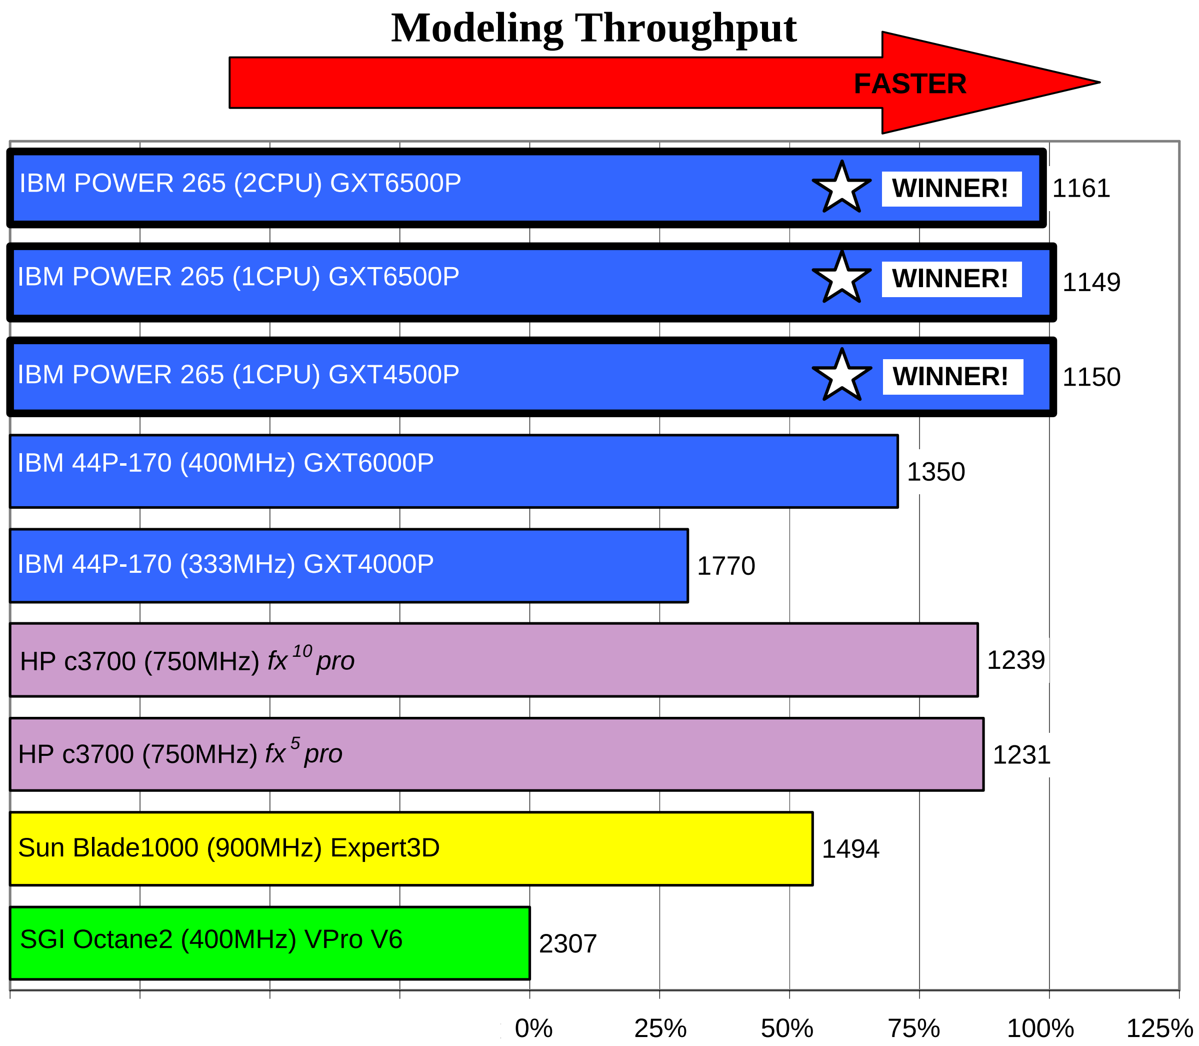

Chart 8 – Solid Model Creation and Modification Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

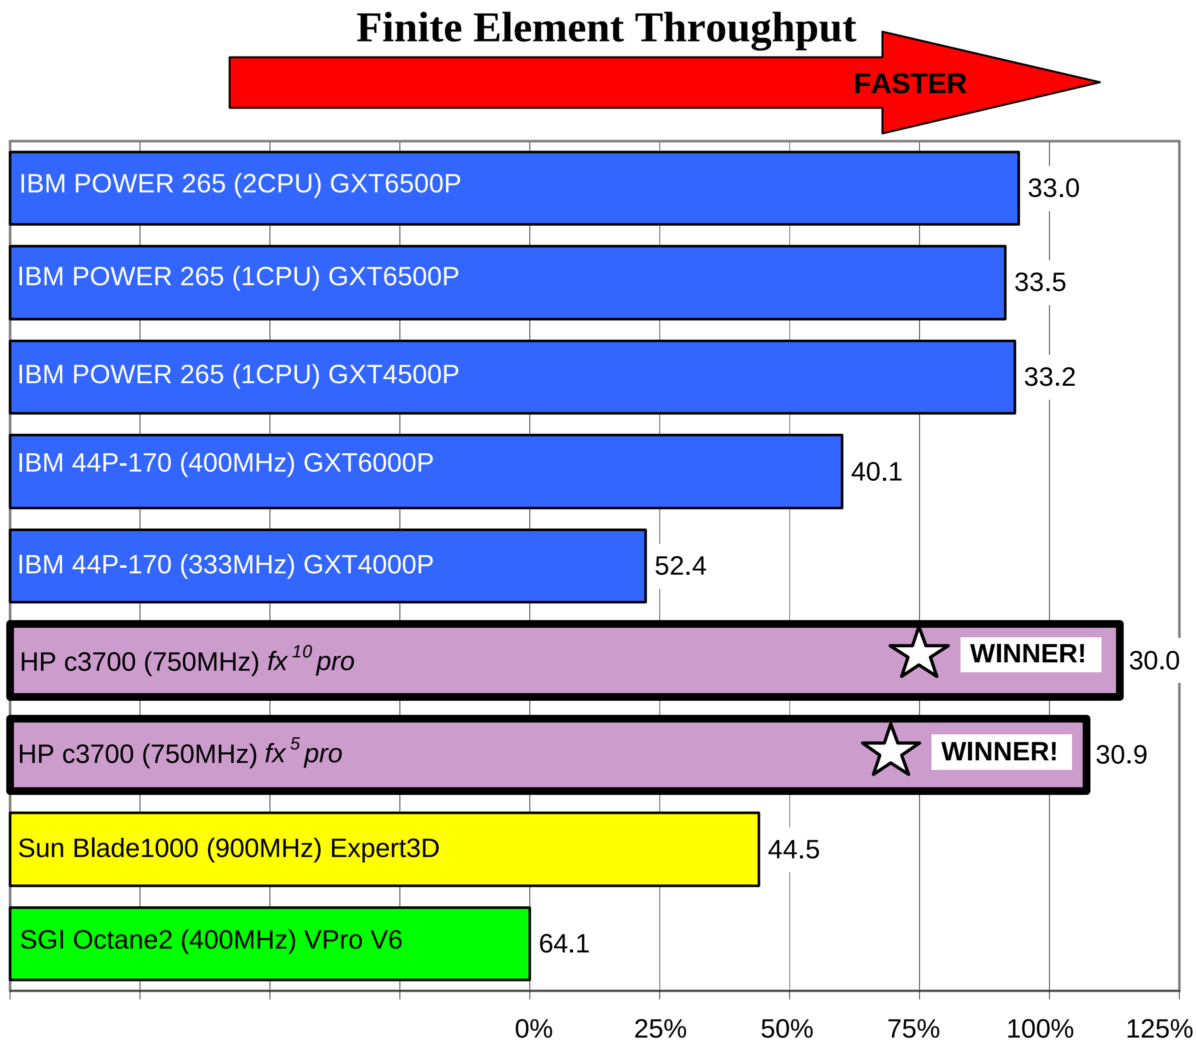

Chart 9 – Finite Element Analysis (ANSOLID) Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

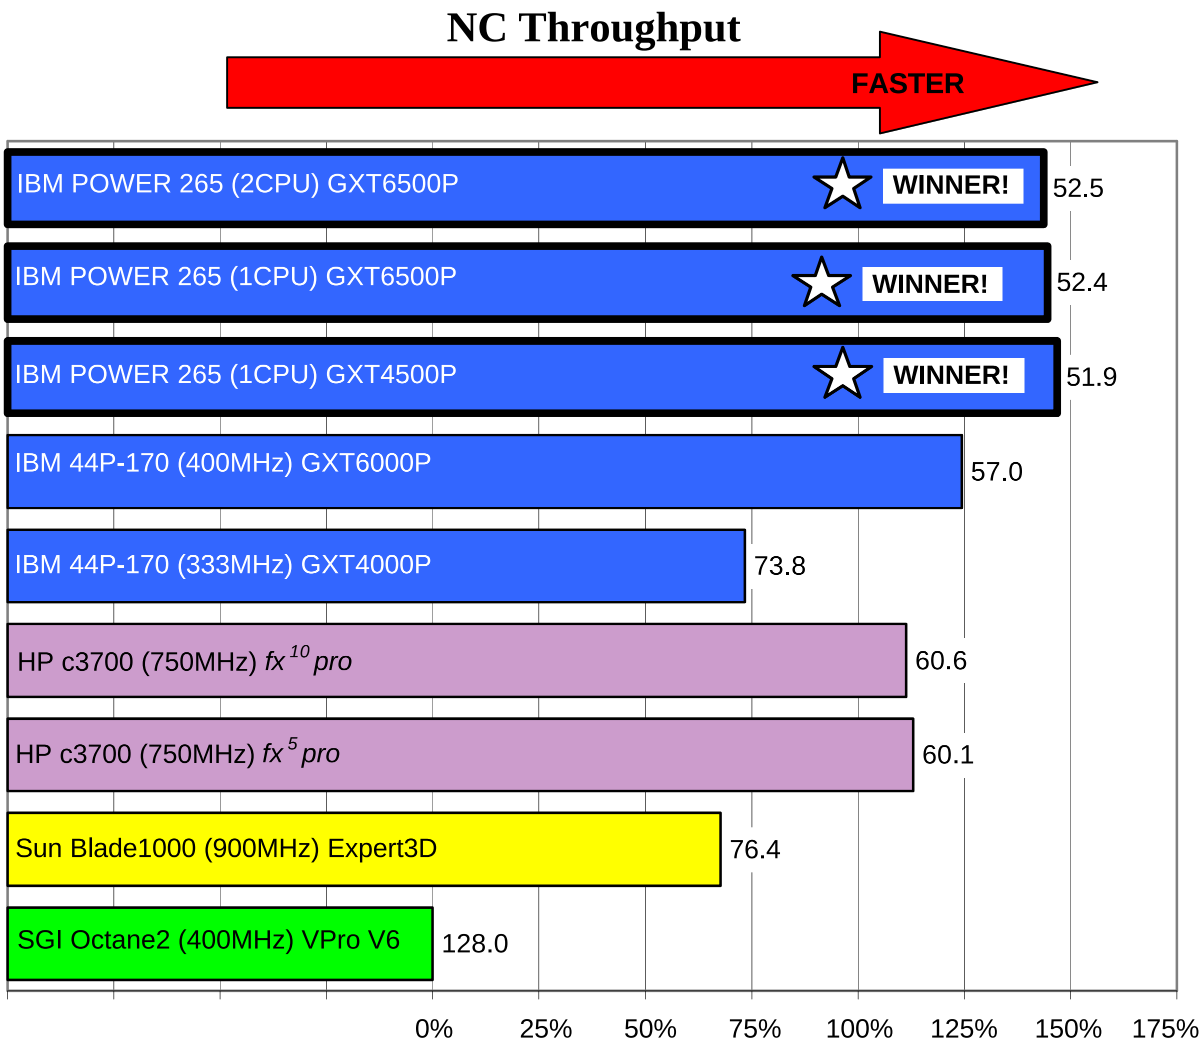

Chart 10 – NC Tool Path Generation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

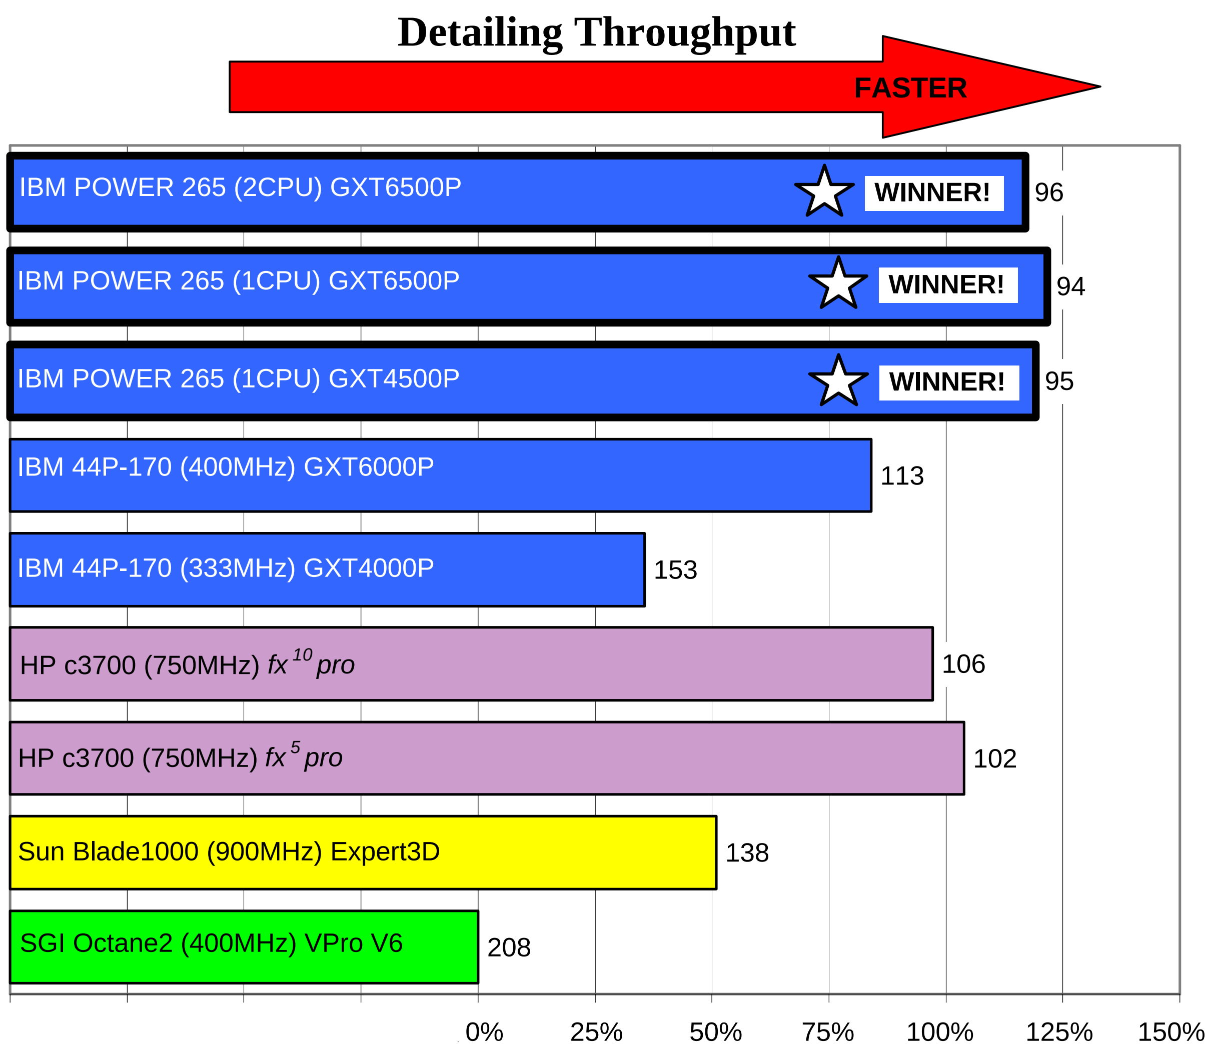

Chart 11 – Detail Drawing Creation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

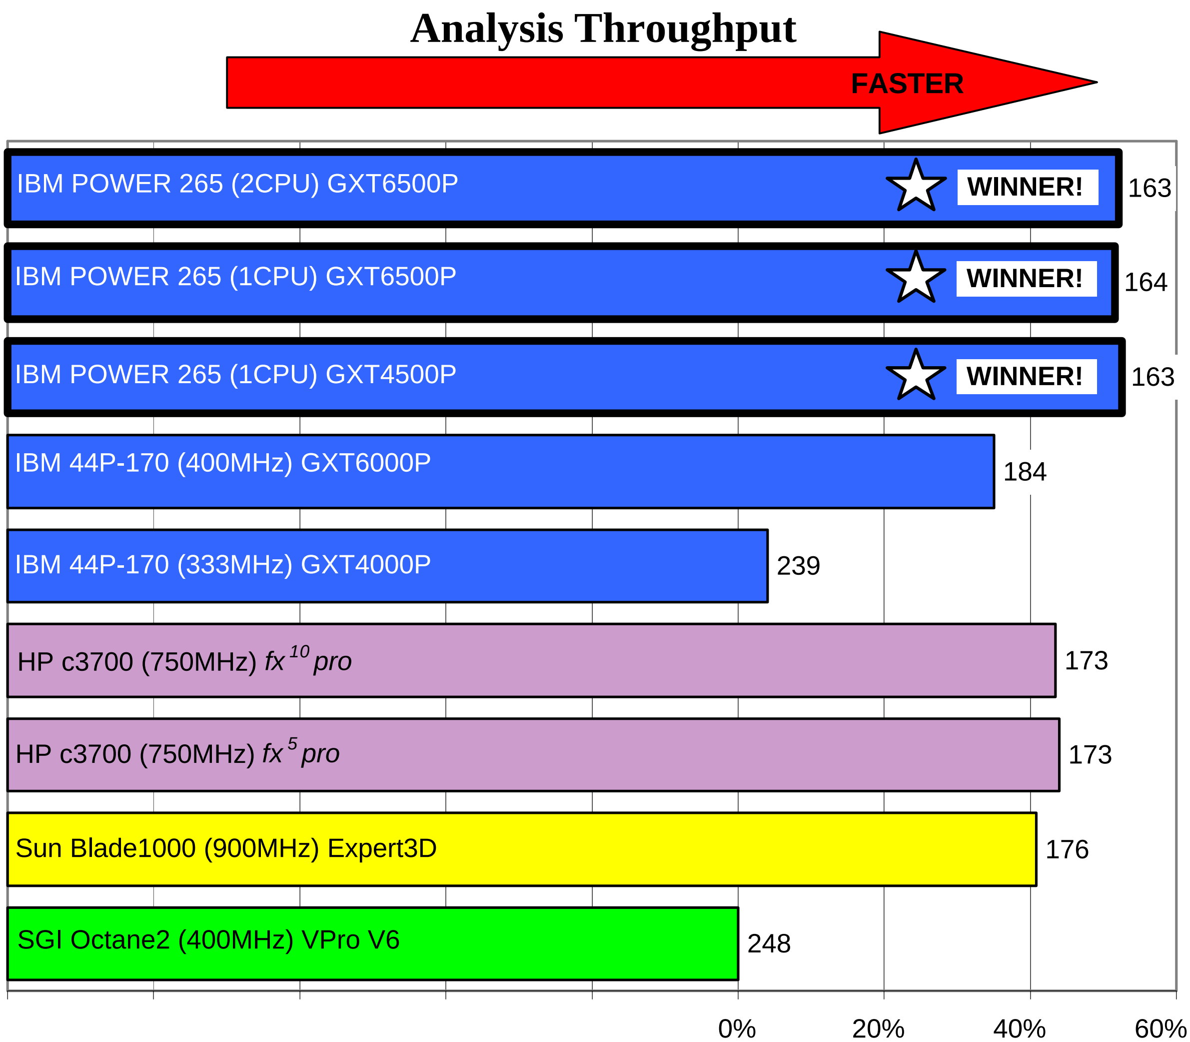

Chart 12 – Solid and Surface Analysis Function Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

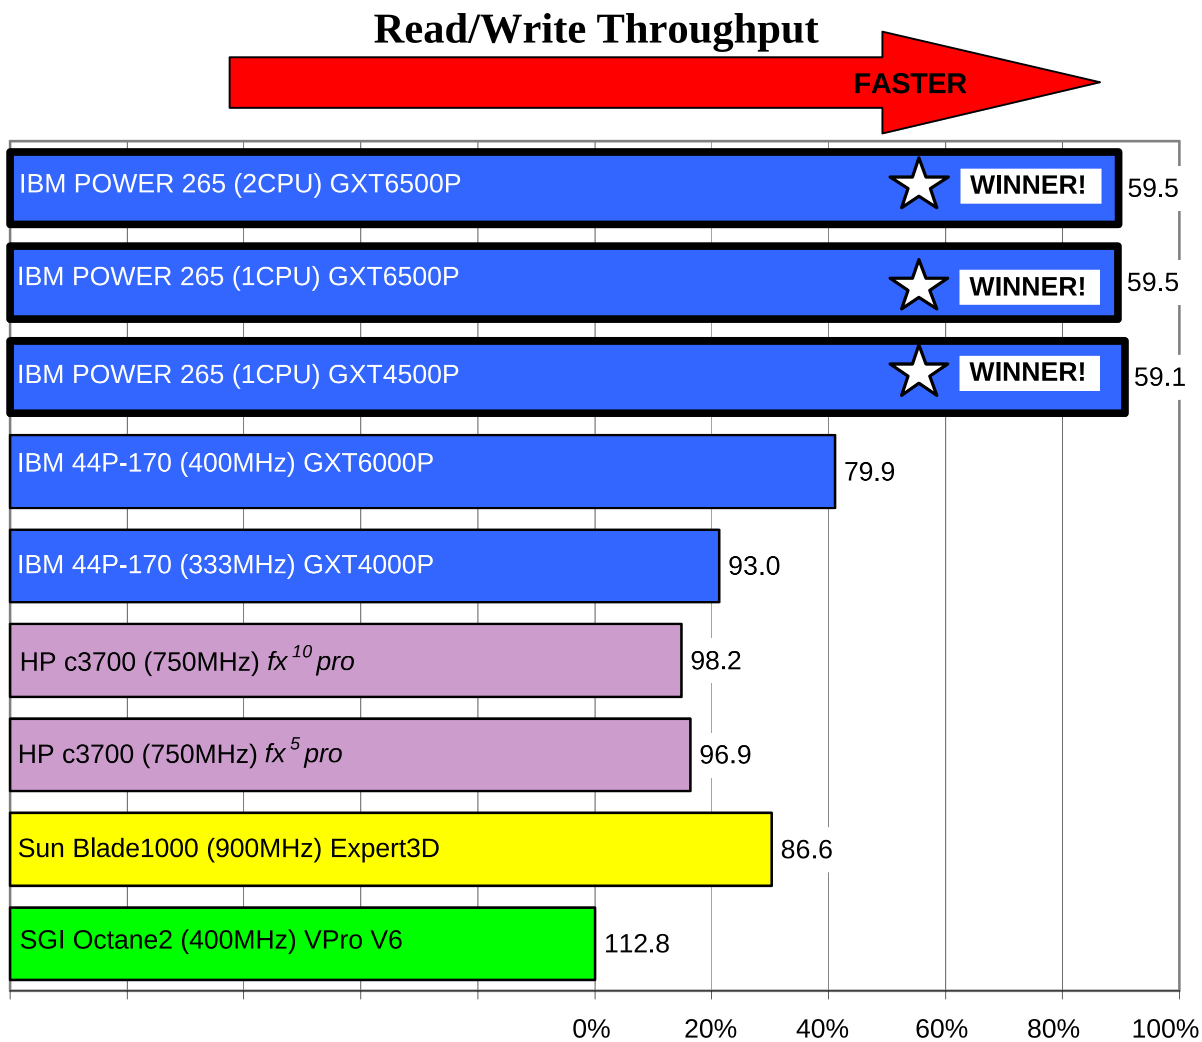

Chart 13 – Read/Write Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

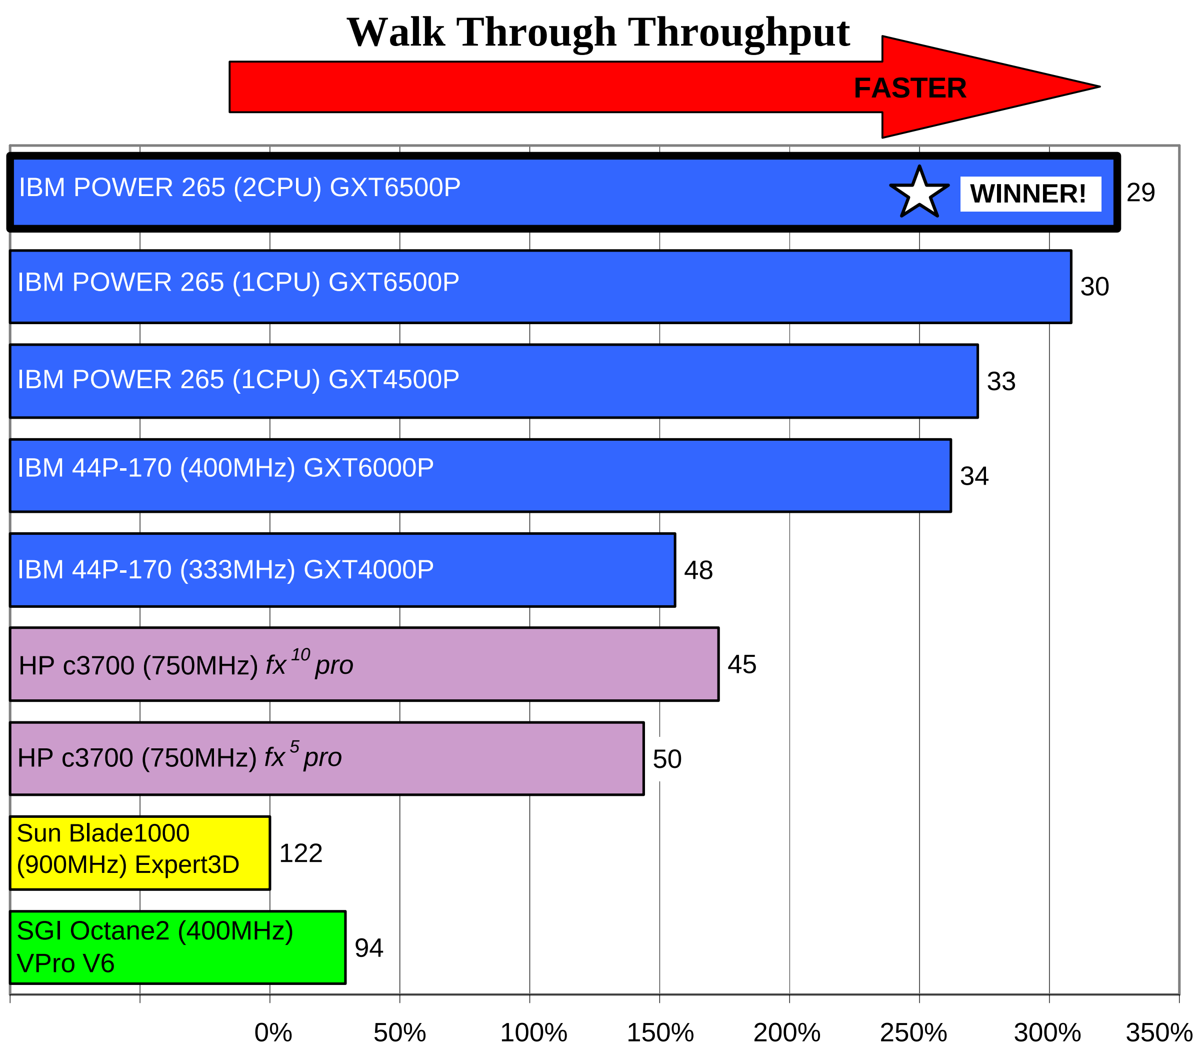

Chart 14 – Walk Through Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

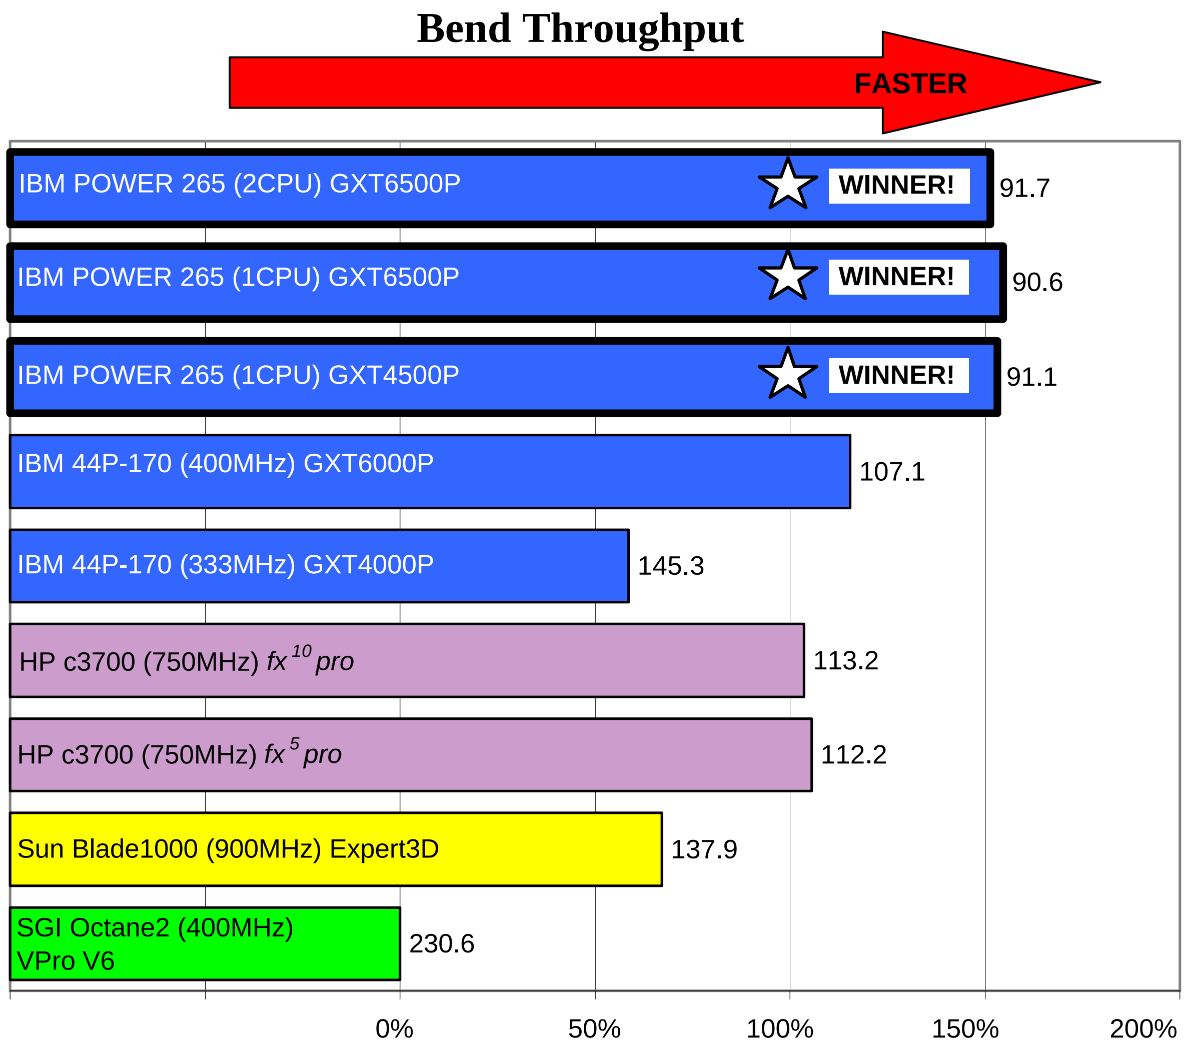

Chart 15 – Bend (Sheet Metal Part Development and Modification) Throughput Relative to Slowest Machine

Test time in seconds shown next to bars (smaller numbers faster)

Longest bar wins! (★)

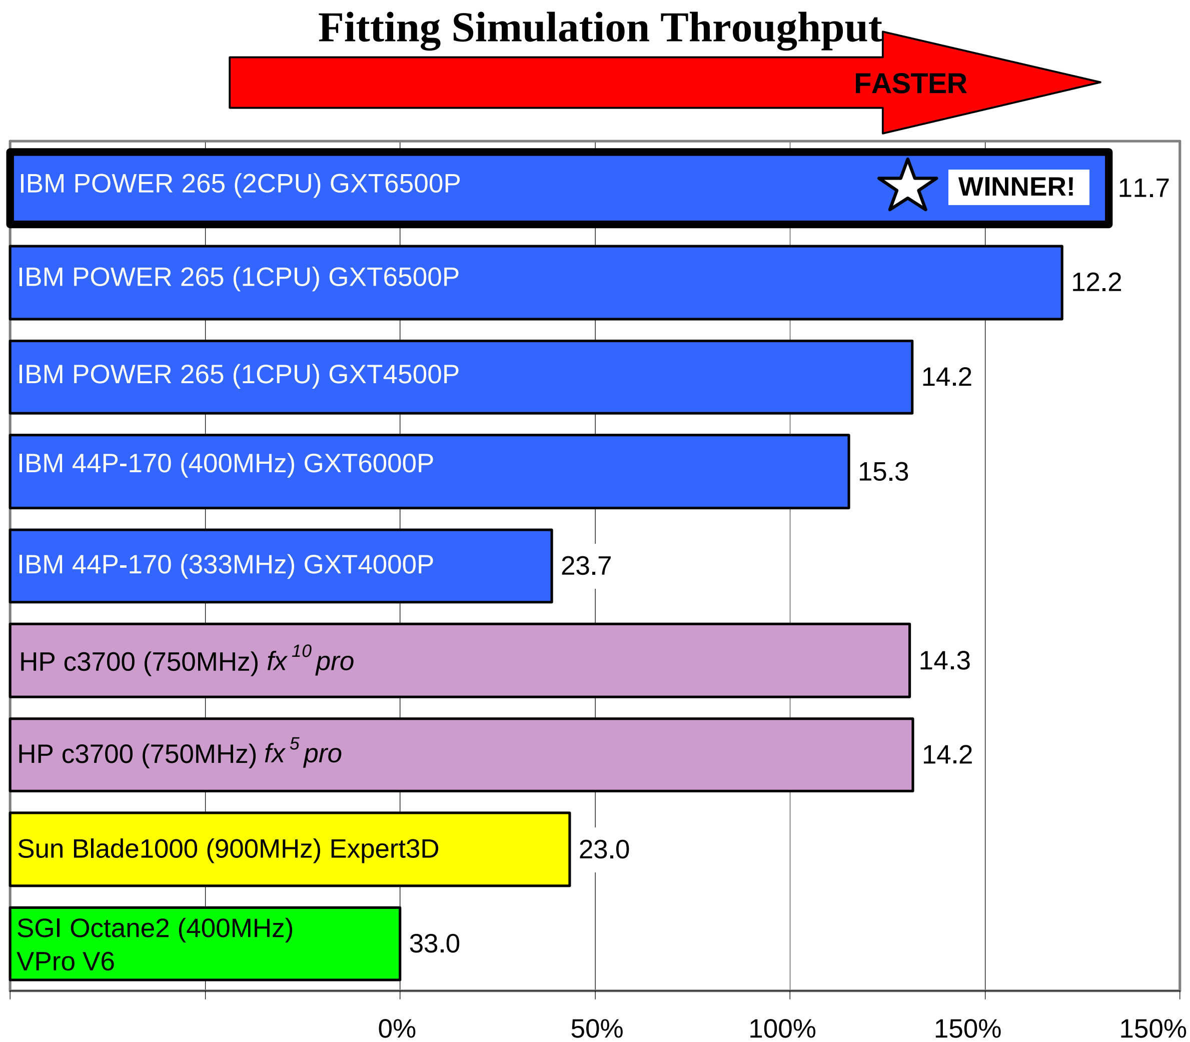

Chart 16 – Fitting Simulation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

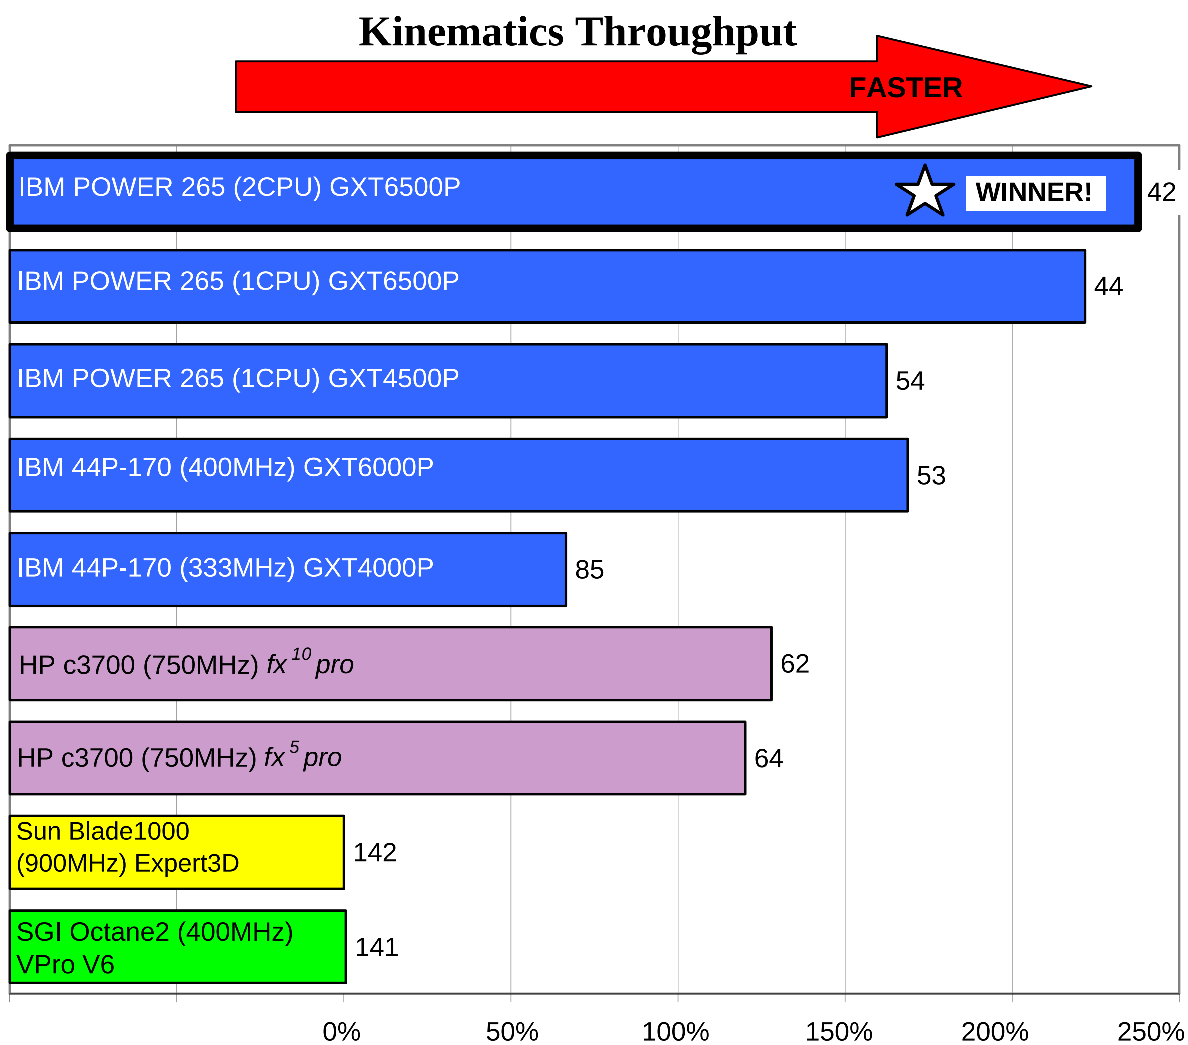

Chart 17 – Kinematic Simulation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

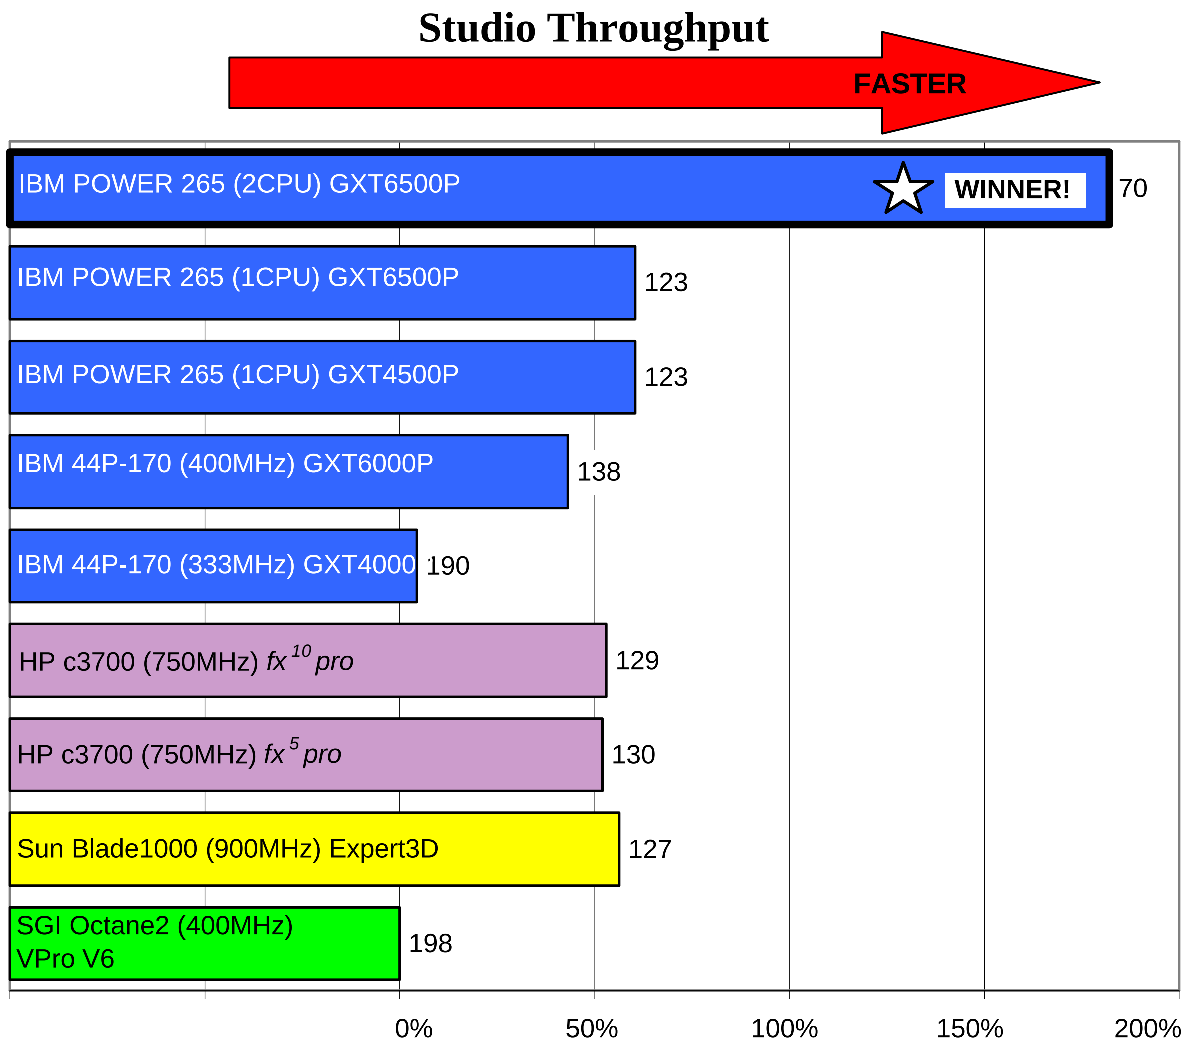

Chart 18 – Studio Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

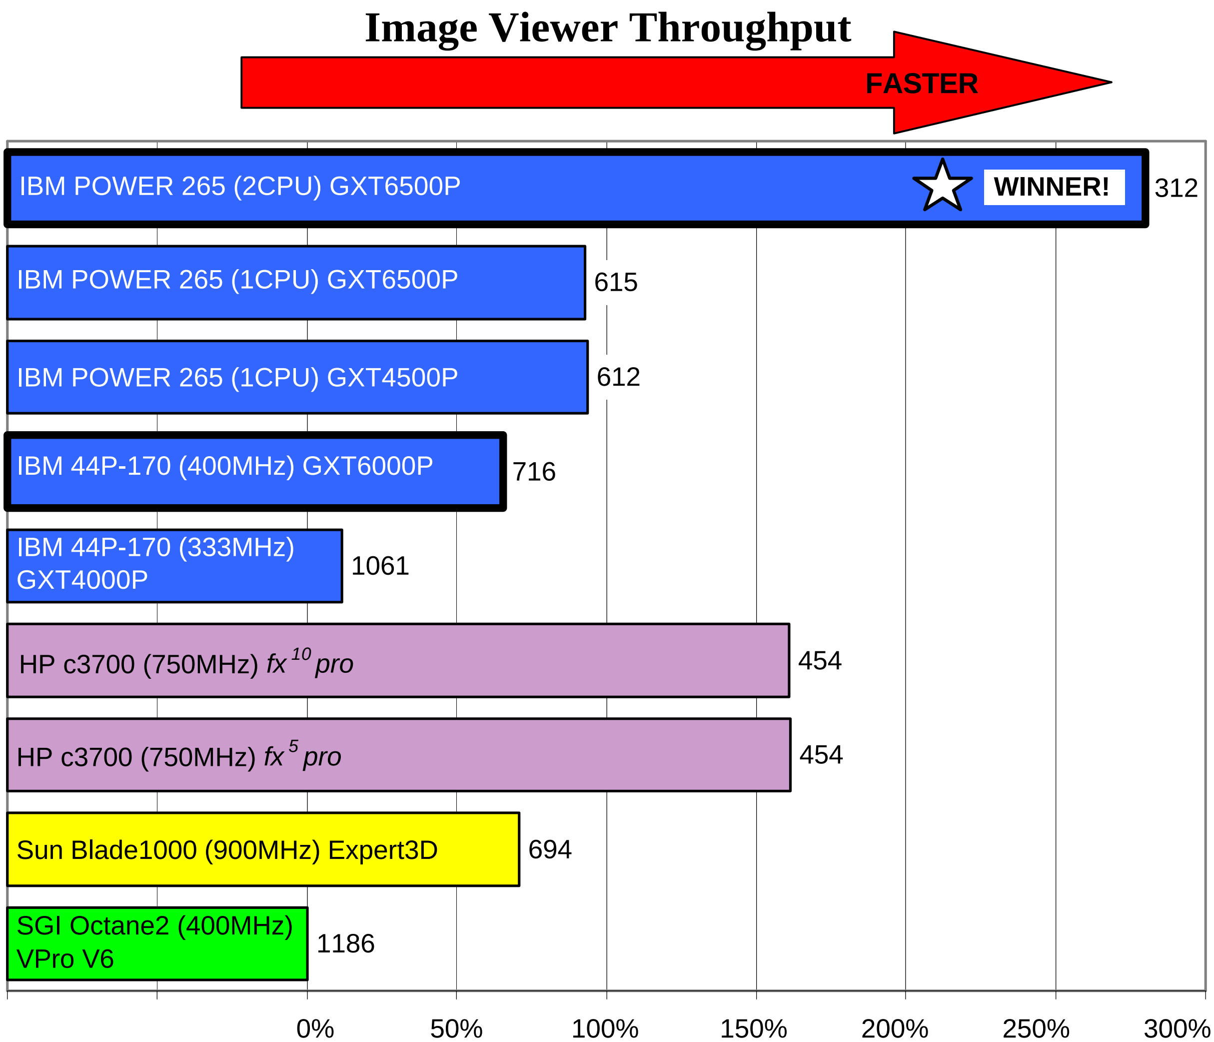

Chart 19 – Image Viewer Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)