

SERIES 58 PERFORMANCE

Jim Kramer
Hewlett-Packard Company, San Diego California, USA

Summary

The Series 58 is the new mid-range system in the HP 3000 line of computers. It has a more powerful processor than its predecessor, the Series 48, and its maximum main memory configuration is twice as large.

In this paper I will discuss the hardware characteristics which result in the improved performance of the Series 58, and will report the results of performance tests. For those system owners who might be considering an upgrade, I will present some guidelines to aid in determining whether an upgrade is likely to result in a significant performance improvement.

Series 58 System Description

The Series 58 introduces new processor, memory, and memory controller boards. Otherwise the system hardware is virtually identical to that of the Series 48, except that 1/4 megabyte memory boards will not work on the Series 58.

An upgrade to a Series 58 from a Series 44 or 48 replaces the CPU (two boards replace two or three), Control and Maintenance Processor board, and memory controller board. It also adds two megabytes of memory on a single board. A similar upgrade is available for a Series 39, 40 or 42, and the upgraded system is designated the 42 XP.

Supported memory sizes start at 2 megabytes and extend to 6 megabytes for the Series 42 XP and 8 megabytes for the Series 58. (The 2 megabyte minimums are possible if the system being upgraded has only 1/4 megabyte boards.) A new Series 58 will have 4 megabytes or more, and a "new" 42 XP (Series 42 XP upgrade purchased with a Series 42) will have at least 3 megabytes; note that these minimums are maximums for the 48 and 42 respectively.

The new hardware has been designed to run existing software; the processor identifies itself as a Series 48 processor in order to eliminate the need for changes. (A new machine instruction is available to distinguish the Series 58 processor from its predecessor, if that should be necessary. Since it will cause an

illegal instruction trap on HP 3000's other than the Series 58, it cannot be used until MPE recognizes and emulates it). Supported software has been restricted to T-MIT and its successors to minimize system testing requirements.

The Series 58 and 42 XP, like the Series 42, 48 and 68, are always cached systems.

Performance Considerations

In the following discussion of performance considerations, comments about the Series 58 apply to the 42 XP as well (unless 7 or 8 megabytes of memory is being used).

The performance significance of the Series 58 compared to the 48 is its faster processor and the doubling of the maximum memory size. The processor provides substantially greater throughput than the Series 48 processor, as shown by the test results presented later in this paper.

For upgrade customers, the addition of caching to non-cached systems, and the addition of 2 megabytes of memory are also significant.

One thing to note about the new machines is that they are certain to match or outperform their predecessors. The processor is more powerful, their minimum memory configurations were maximums on the predecessors, and caching is available. The same claim could not be made about some prior upgrades: sometimes a Series 44 (if it only had one disc controller) performed worse than the Series III it replaced, and sometimes a 48 with caching turned on performed worse than the 44 it replaced.

The new systems address the three major performance bottlenecks: processor, memory, and disc I/O. The processor and memory bottlenecks are addressed directly with additional capacity; the disc I/O bottleneck is addressed with caching. Customers who already have caching are almost certainly CPU-limited, and will be helped by the additional processor power.

One can try to imagine a load on a 44 that would not run faster on a 58. It would be disc I/O limited, and the disc I/O would be so nonlocal that caching (and the additional processor and memory to support it) would not help. This situation is unlikely, however. Our experience is that caching does help when the necessary extra CPU and memory capacity are there, because MPE and most applications exhibit significant locality in their disc access patterns.

Please notice how much growth capability has been provided the Series 39, 40 and 42 in the last year. ATP's allow connection of as many as 60 point-to-point terminals. And the 42 XP upgrade

provides a maximum memory size 50% greater than that of the Series 48, and greater processor power.

The Series 58 Processor

As mentioned earlier, the significant performance gains from the new systems are provided by the Series 58 processor and increased memory sizes permitted. Disc caching may also provide performance gains for customers upgrading from uncached systems.

The Series 58 processor contains a 32 Kbyte high-speed memory cache for buffering both instructions and data being read from memory. (Be careful not to confuse the memory cache with MPE's disc caching.) If a required word is in the cache, it is obtained in a single machine cycle. Six cycles are required by the Series 44 processor to retrieve a word from memory.

To take full advantage of the cache speed, hardware was added to speed up functions that in the Series 44 were performed while waiting for data to be read from memory. Microcode of the most frequently used machine instructions was rewritten to make use of the new hardware.

For performance, our main concern is that the processor speedup was achieved mainly through the provision of a memory cache, not through a processor clock speedup.

Initial performance testing concentrated on the new processor. Less emphasis was placed on the effects of additional memory and disc caching.

Performance Test Results

We will begin our summary of performance test results with some tests that do no I/O, and therefore isolate processor performance. It is important to note that these tests are timings of small programs that repetitively execute code within a loop. They therefore make excellent use of the memory cache, and show the series 58 processor to good advantage.

The Sieve of Eratosthenes is a prime number calculation routine that has been coded in many languages and run on many machines. It tests looping, array indexing, and integer arithmetic. The following table shows results on the 48 and 58 processors for four languages. The timings are in seconds required to execute ten loops, each of which calculates 1899 prime numbers.

	Series 48	Series 58	48/58
Fortran	2.71	1.56	1.74
SPL	2.75	1.59	1.73
Basic	2.20	2.00	1.60
COBOL-II	17.30	10.18	1.70

The column labeled 48/58 is the Series 48 timing divided by the Series 58 timing. This is the throughput ratio. The number 1.74 for Fortran indicates that the Series 58 can perform 74% more work of this type in a given time than the 48. (The "response time" improvement is 42% -- $(2.71-1.56)/2.71$. The response time improvement for a stand-alone task is always less than the throughput time improvement.)

The Whetstone benchmark tests floating point computations. It was done for both single and double precision.

	Series 48	Series 58	48/58
Single	9.31	7.02	1.33
Double	26.14	22.57	1.16

The throughput improvement is less than for the Sieve, because floating point instructions do extensive calculations on the data extracted from memory, so that the cache provides less help. The floating point and packed decimal instructions were not re-microcoded. Most applications do not use them as intensively as this benchmark.

An extensive set of tests was done testing various language constructs. The timings were obtained by executing a loop with and without the construct, and taking the difference. Below is a small sampling. The times are given in microseconds.

	Series 48	Series 58	48/58
SPL integer move	2.8	1.2	2.3
SPL real move	4.4	2.5	1.8
SPL long move	8.9	4.9	1.8
COBOL II S9(7) COMP-3 move	41.	34.	1.2
SPL move 80 bytes	62.	46.	1.4
SPL integer add	4.2	2.0	2.1
SPL real add	11.	8.0	1.4
COBOL II S9(7) COMP-3 add	160.	130.	1.2
SPL integer multiply	7.3	5.0	1.5
SPL real multiply	16.0	12.0	1.3
COBOL II S9(7) COMP-3 mult.	220.	190.0	1.2
SPL call/return, no parms	23.	16.	1.4
FWRITE 80 bytes, large BF	2300.	1600.	1.4

The pattern holds that speedups are greatest for the simpler operations which can make good use of the memory cache. The 2.3 speedup (130%) for the SPL integer move (LOAD, STORE) is the current series 58 speed record.

Now we have some tests which are much more representative of actual system usage. Descriptions of the tests come first, followed by a table of the results.

Compilation

1. COBOL Compile -- This test generated 32K lines of output. There were a few include files.
2. Pascal Compile -- A 3500 line program was compiled, no include files, and the output was sent to \$NULL.
3. SPL2 Compile -- SPL2 is not available to customers, but some HP software is written in it. The program source was 5000+ lines, no include files, and the output was routed to \$NULL.

TDP Operations

4. TDP Final 1 -- A 1515 line document was formatted to produce 2281 lines of output. There were 160 includes (each requiring a file open, read and close).
5. TDP Final 2 -- This test is the same as TDP Final 1 without the includes.
6. TDP Text and Keep -- A 10,000 line file was texted, a global change was made which modified and displayed 1136 lines, and the result was kept.
7. ED Text and Keep -- This is the same as the TDP test, but used an unsupported editor (similar to Unix's XED).

Data Base Access

8. ORBIT -- This is a production job from the MIS group of HP's Computer Systems Division. It reads, sorts and processes Image data. The sorting is done by Robelle's SUPRTOOL.
9. NEWJIT MRP -- This is HP application software that runs against an Image data base, and produces output files totalling more than 16K sectors.

All of the tests were run stand-alone on a 4 megabyte system with caching enabled.

Test	Series 48		Series 58		48/58	
	CPU (sec)	Wall (min)	CPU (sec)	Wall (min)	CPU	Wall
1 COBOL	1113	22	860	18	1.29	1.22
2 Pascal	169	4	122	3	1.39	N/A
3 SPL2	297	7	222	5	1.34	N/A
4 TDP Final 1	98	2	72	2	1.36	N/A
5 TDP Final 2	45	1	32	1	1.41	N/A
6 TDP Text/Keep	197	4	152	4	1.30	N/A
7 ED Text/Keep	120	3	84	2	1.43	N/A
8 ORBIT	2197	42	1801	33	1.22	1.27
9 MRP	10206	174	7214	123	1.41	1.41
Totals	14442	259	10559	191	1.37	1.36

The CPU throughput improvement of 1.37 is obtained from the CPU totals, and thus is a weighted average of the CPU throughput improvements for the 9 tests. A straight average yields 1.36.

A TEPE test was run of HPAccess, the new PC Central product which can extract data from the HP 3000 for use on personal computers.

The tests emulated five terminals performing such operations as project, join, select, sort, summarize and output. Two tests were run on both the Series 48 and Series 58. The tests differed only in the amount of data handled.

The following test results were taken from the emulator report. The response times are averages over all transactions for all terminals.

Test	Series 48		Series 58		48/58	
	Wall (min)	Avg Resp (sec)	Wall (min)	Avg Resp (sec)	Wall (min)	Resp Reduct
Test 1	142	70.93	110	49.81	1.29	30%
Test 2	574	392.74	433	290.37	1.33	26%

Since the machines were almost never idle during the testing, especially for Test 2 (about 1% idle), the ratio of wall times can be taken as a good indicator of CPU throughput increases. The last column is the response time improvement.

The venerable EDP benchmark, which has been used to test the

entire product line, was also run. Briefly described, the test comprises multiple terminals doing transaction processing with VPLUS and Image, and additional terminals doing program development. There are also batch COBOL compiles in the background.

For details of the test environment, consult the published performance guide, part # 5954-0401.

The following table shows throughput increases achieved by the Series 58, determined by comparing transaction execution rates at specified response times. The data provided by the tests gives transaction rates and response times for various numbers of terminals. To obtain transaction rates for the specific response times listed below, it was necessary to interpolate (linearly) between the provided points.

Response Time (secs)	Series 48 Trans (/hour)	Series 58/4mb Trans (/hour)	Increase	Series 58/6mb Trans (/hour)	Increase
.5	4486	13428	2.99	13793	3.07
1.0	10879	20753	1.91	22100	2.03
1.5	16057	22147	1.38	24061	1.50
2.0	19253	23541	1.23	26022	1.35
2.5	20447	24935	1.22	26362	1.29

As the systems start to become saturated (2.5 second response), the Series 58 at 4 megabytes provides a 22% throughput increase. Additional memory is shown to help performance throughout, and near saturation the Series 58 with 6 megabytes provides a 29% throughput increase over the Series 48.

The next table shows response time reductions at specified transaction rates.

Trans (/hour)	Series 48 Response (sec)	Series 58/4mb Response (sec)	Reduction	Series 58/6mb Response (sec)	Reduction
10000	.93	.38	58%	.37	60%
14000	1.24	.52	58%	.51	59%
18000	1.80	.76	58%	.66	63%
22000	4.37	1.45	67%	.84	81%

The response time reductions are substantial at all transaction rates. Again the extra memory is shown to be helpful as the system becomes busier.

Who Should Upgrade?

There are two good reasons to upgrade to a Series 58 (or a 42XP): to improve on-line response times, and to improve batch job turnaround times. The case of batch jobs is the easiest to consider: the Series 58 can be depended on to run them significantly faster. If the system being upgraded is uncached and the job is I/O intensive, caching should provide further help. However if the objective of the upgrade is to speed up just one or two critical jobs, it is definitely worthwhile to run them on someone else's Series 58 prior to upgrading to be sure that the hoped-for improvement will be achieved. I say this because a test completed very recently showed only a disappointing 10% speedup for a set of Image reporting jobs running concurrently.

The impact of a Series 58 upgrade on on-line response times is more difficult to predict. Improvements can range from dramatic to barely noticeable.

The dramatic improvements are seen when the poor response is caused by a shortage of a critical performance resource (CPU, disc I/O, memory). When one of these resources is near exhaustion, a small additional demand for it can make response times dramatically worse. Similarly even a small additional supply can make response dramatically better.

The chances for dramatic improvement are even better if the resource in short supply is memory, or if it is disc I/O on an uncached system. Here the benefits of the extra memory and caching add to the benefits of the faster processor.

The barely noticeable improvements, conversely, come when performance resources are adequate but the transactions are inherently long ones. The Series 58 will speed up such transactions, but reducing response time from 10 seconds to 7 or 8 seconds may not be all that was hoped for. The solution in such cases is, unfortunately, reworking the application or upgrading to an even more powerful machine than the Series 58.

How can you tell if poor response is due to a resource shortage or to long transactions? An easy way is to observe the response time for a single user running the problem application. If response for a single user is good, then poor response for many is due to a resource shortage. If response for a single user is poor, then the problem is long transactions. If you can say "Response used to be good, but has gradually gotten worse as we added more and more users," then an upgrade is almost certainly called for.

Another way to tell is by measurement. Tools such as OPT or Surveyor (in the TELESUP account) provide extensive information

about system resource usage. However there are pitfalls to this approach that even experienced performance SE's have fallen into. For instance it is not enough to run OPT, observe that CPU usage is at 100%, and conclude that there is a resource shortage. This is because a significant amount of CPU used by batch activities (including such things as online compiles) really signifies the availability of CPU for short transactions, because of MPE's process priority structure. It is not enough to measure system resource usage; it is also necessary to determine which processes are using the resources.

Conclusion

The Series 58 and Series 42 XP provide help where it is needed in this era of disc caching: CPU and memory capacity. The processor performance increase, coupled with whatever help memory and caching provide, should give a significant boost to the current midrange machines.

Please join me in welcoming these new arrivals to the HP 3000 product line.

Biography

Jim Kramer has been an HP3000 Systems Engineer and Performance Specialist with Hewlett-Packard for eight years, and works in the San Diego, California, field sales office.

