Technical Report
CMU/SEI-90-TR-7
ESD-90-TR-208

Carnegie-Mellon University
Software Engineering Institute

Hartstone Benchmark
Results and Analysis

Patrick Donohoe
Ruth Shapiro
Nelson Weiderman
June 1990

Hartstone Benchmark
Results and Analysis

Patrick Donohoe
Ruth Shapiro
Nelson Weiderman

Real-Time Embedded Systems Testbed Project

Distribution/
Availability Codes

Dist | Avail and/or
Special

A-1

This technical report was prepared for the

SEI Joint Program Office
ESD/AVS
Hanscom AFB, MA 01731

The ideas and findings in this report should not be construed as an official
DoD position. It is published in the interest of scientific and technical
information exchange.

Review and Approval

This report has been reviewed and is approved for publication.

FOR THE COMMANDER

Karl H. Shingler
SEI Joint Program Office

Table of Contents

List of Figures

List of Tables

Hartstone Benchmark Results and Analysis

Abstract: Hartstone is a series of timing requirements for testing a system's ability to handle hard real-time applications. It is specified as a set of processes with well-defined workloads and timing constraints. The name *Hartstone* derives from *HArd Real Time* and the fact that the workloads are based on the well-known Whetstone benchmark. This report describes the results obtained by running Version 1.0 of the Hartstone benchmark, an Ada implementation of one of the requirements, on a number of compiler/target processor combinations. The characteristics and expected behavior of the benchmark are described, actual results are presented and analyzed, and the lessons learned about the compilers and processors, and the benchmark itself, are discussed. Nothing in this report should be taken as an endorsement of, or an indictment of, a particular product. Users of Ada technology are encouraged to experiment with the Hartstone benchmark relative to their own particular application requirements.

1. Introduction

Hartstone comprises a series of requirements to be used for testing the ability of a system to handle hard real-time applications. Its name derives from HArd Real Time and the fact that the computational workload of the benchmark is provided by a variant of the Whetstone benchmark [Curnow 76], [Harbaugh 84], [Wichmann 88]. "Hard" real-time applications *must* meet their deadlines to satisfy system requirements; this contrasts with "soft" real-time applications where a statistical distribution of response times is acceptable [Liu 73]. The rationale and operational concept of the Hartstone are described in [Weiderman 89]; in particular, Hartstone is seen as an evaluation tool:

> The currently available evaluation technology for Ada does not adequately address deadline-driven computing. For the most part the existing test suites and checklists address the questions of functionality (does Ada provide a function), capacity (how big can various aspects of programs be), or performance (how fast can certain operations or programs be performed). None adequately addresses the question of whether a set of activities (when there may be rapid switching of attention between those activities) can meet preestablished deadlines. The answer to this question has more to do with the ability to keep accurate time and perform in a predictable, deterministic fashion than it has to do with the throughput of the system.

Five test series of increasing complexity are defined in [Weiderman 89] and one of these, the Periodic Harmonic (PH) Test Series, is described in detail. The PH test series has been implemented in Ada at the Software Engineering Institute (SEI) by the Real-Time Embedded Systems Testbed (REST) project, and a user's guide has been published [Donohoe 90]. The user's guide describes the design of Version 1.0 of the benchmark and the four experiments defined for the test series. It also provides guidelines for performing the experiments and interpreting the results obtained, and presents results for Version 1.0 of the XD Ada VAX/VMS - MC68020 cross-compiler. This report continues and extends the work begun in the user's guide. Those readers needing a brief description of the characteristics of the PH Series are referred to the properties of the task set as described in Appendix A.

One goal of the Hartstone work is to provide compiler vendors, application developers, and compiler evaluators with a tool to gauge the suitability of Ada compilers for hard real-time systems development. For vendors, Hartstone is a tool to measure the runtime system's handling of multi-tasking and time management. Hartstone does not test how the scheduling algorithm handles tasks of equal priority or the ability to do time-slicing, but it *does* test the ability of a runtime system to do preemptive scheduling efficiently. For application developers, the benchmark could be modified or adapted as an application modelling tool that is representative of a particular hard real-time problem domain. For compiler evaluators, it provides a composite, synthetic benchmark that exercises a combination of Ada language features and provides a measure of the efficiency of the runtime system and the generated machine code. For all these groups, Hartstone provides a standard for communicating application performance requirements.

The Real-Time Systems Testbed (REST) Project at the Software Engineering Institute (SEI) has compiled and run the Hartstone benchmark on several Ada compilation systems and processors; this report is both a summary of that effort and an analysis of the benchmark itself. Chapter 2 describes the testing environment in which the Hartstone results reported here were obtained. It lists the compilation systems under which Hartstone was compiled and linked, and the processors on which it was run. It also summarizes the preparatory work required and the experiences in making Hartstone compile and run under the various Ada compilation systems. Chapter 3 analyzes the expected behavior of Hartstone's periodic tasks, the nature of the Hartstone experiments, and the results obtained from the various compilation systems and target machines. The final chapter, Chapter 4, summarizes some of the lessons learned about the Hartstone benchmark and about the compilation systems tested and draws some conclusions about the state of current Ada compiler technology with respect to timing characteristics. Appendix A presents a rather detailed argument for why the flaws associated with using the Ada delay statement to implement periodicity do not cause a problem in the case of a harmonic task set. Appendix B presents the results of some supporting benchmark tests. These results illustrate the wide variations in the performance of certain Ada features that are critical to the overall performance of the Hartstone benchmark. Appendix C through Appendix I present the actual Hartstone results. A final appendix, Appendix J, explains how to get copies of Hartstone documentation and source code.

The authors wish to caution readers in the strongest possible terms not to draw inappropriate conclusions from the data provided in this report. These data are not provided for comparison purposes, but rather to show a wide range of outcomes and behaviors. There are a large number of independent variables involved. The results do not represent controlled experiments across product lines. Among the uncontrolled variables are the speed and type of the underlying processors, the existence of floating point coprocessors, the existence and type of underlying operating system, and the math libraries which may be neither optimized nor provided by the vendor. Also, we used the compiler and runtime system defaults and did not take any measures to improve performance. We deviated from the defaults only when the program did not work with the defaults (e.g., because of insufficient stack/heap sizes). Our results are indicative of what a new user might see from a newly installed compiler, not what might be obtainable after some serious tuning. All the above factors can be expected to have significant impacts on the results.

The SEI owns licenses for all the compilers used in the testing. We attempt to be an ordinary customer with no special privileges or obligations. The compilers may not be representative of current compiler technology and some may not be the latest vendor releases. They do, however, demonstrate clearly the variations that are possible, and the need for careful evaluation of the timing characteristics of a compiler for real-time applications.

2. The SEI Hartstone Testing Experience

The purpose of this chapter is to describe the testing environment and the experience gained in the benchmarking process. It includes preliminary PIWG testing as well as the process of compiling, linking, and running the Hartstone benchmark. It also describes some of the changes that were required and some surprising operational behavior.

2.1. The Real-Time Embedded Systems Testbed

The REST Project has a host-target configuration of machines for evaluating Ada compilers. The primary host is a cluster of DEC MicroVAX II and 3200 machines, running version 5.3-1 of the VMS operating system. Ada code may be compiled and linked on any of the host machines and downloaded to a number of different target machines. Output from the target machines is routed to a window or file on a host machine. The results documented in this report were obtained from both self-targeted and cross-development systems. The target machines were both embedded "bare-board" systems and non-embedded systems. The embedded targets were:

- An Intel iSBC 386/116-M02 board, comprising a 16 MHz i80386 CPU, a 16 MHz i80387 math co-processor, two megabytes of RAM, and a 64-kilobyte instruction/data cache.

- A Fairchild SBC-50 board (MIL-STD-1750A Instruction Set Architecture), comprising a 15 MHz CPU and 128 kilobytes of RAM.

- A Motorola MVME133A board, comprising a 20 MHz MC68020 CPU, a 20 MHz MC68881 floating-point co-processor, one megabyte of RAM, and a 256-byte instruction cache.

Additionally, Hartstone was run on the following machines:

- A DEC MicroVAX 3200 running VMS 5.3-1.

- A Rational R1000, Series 200, Model 20, running the Delta 1 release of the Rational Environment.

The compilers for the embedded targets were:

- DDC-I DACS-80386PM VAX/VMS - i80386, release 4.3

- Systems Designers XD Ada VAX/VMS - MC68020, Version 1.0

- Tartan Ada VMS/1750A, Version 2.14

- TeleSoft TeleGen2 VAX/VMS - MC68020, release 3.22a

- Verdix VADS VAX/VMS - MC68020, Version 5.7

For the non-embedded systems, the compilers were, respectively:

- DEC VAX Ada 2.1

- Rational Environment, Delta 1 release

2.2. Preliminary PIWG Benchmarking

Readers should consult the Hartstone User's Guide [Donohoe 90] for a more detailed discussion of a number of factors that affect the performance of the benchmark. These are:

- The implementation of task periodicity
- The resolution of the **delay** statement
- The resolution of Calendar.Clock
- The accuracy of the fixed-point type Duration
- The implementation of mathematical library functions
- Floating-point precision
- Miscellaneous overhead factors

The first three items listed above are the most important, in the sense that the performance of Hartstone is critically dependent on the performance of these features. The paradigm for task periodicity used in Hartstone depends heavily on the resolution of Calendar.Clock and the **delay** statement. Resolution is defined as the rate at which the clock ticks (in the case of Calendar.Clock) or the step size with which delays are processed (in the case of the **delay** statement. This paradigm is discussed in more detail in Appendix A, while specific measurements of Clock and **delay** statement resolution are presented in Appendix B. The results in Appendix B were obtained by running selected tests from the ACM Performance Issues Working Group (PIWG) benchmark suite (12/12/87 release).[1] In general, there was a wide range of resolutions supported by the compilation systems tested, with a corresponding wide variation in the results obtained from the Hartstone benchmark.[2] The value of Duration'Small was 61 microseconds for all except the Verdix compiler, which had a Duration'Small equal to 1 millisecond. For all compilers except the Verdix compiler, the type Float was represented in 32 bits and Float'Digits was 6. For Verdix, the type Float was represented in 64 bits; type Short_Float was 32 bits with a corresponding Float'Digits value of 6 (Short_Float was used for testing). Specific measurements of the other factors affecting Hartstone's performance are also included in Appendix B. The results for the "large" Whetstone benchmark (as distinct from Hartstone's Small_Whetstone procedure) are also included as an indication of the overall efficiency with which the various compilers handle a composite mathematical benchmark.

---

[1]The name, address, and telephone number of the current PIWG chairperson and other officers may be found in Ada Letters, a bimonthly publication of the ACM Special Interest Group on Ada (SIGAda).

[2]For a discussion of the issues involved in the measurement of Calendar.Clock and delay statement resolution, see [Clapp 86], [Pollack 90], and [Roy 90].

2.3. Compiling and Linking Hartstone

Some preliminary work was needed before commencing Hartstone compilations. For each compiler and target it was first necessary to modify the character strings describing the compiler and target system, in the Experiment package of Hartstone, to reflect the system under test. These descriptions are printed in the Hartstone output to identify the system under test. Secondly, Hartstone requires some mathematical functions that are typically provided in a vendor-supplied mathematical library. The names of these libraries are non-standard, so the correct name had to be used in the context clause of Hartstone's Workload package to ensure that compilations would succeed. Mathematical libraries caused some minor problems, for two reasons: not all vendors supply a mathematical library as part of the Ada compilation system, and there is no consistency in the naming of mathematical functions. The Numerics Working Group of SigAda (NUMWG) has proposed a standard for an elementary math function library. When Ada vendors provide an NUMWG-compliant math library, the above problems will be avoided.

For compilers that did not provide a mathematical library (e.g., the Tartan 1750A compiler and the Rational Environment) the solution was to use the one provided in the PIWG benchmark suite. The PIWG suite contains two versions of the Whetstone benchmark, one of them containing an all-Ada mathematical library (the A000093 test). Libraries such as this are less efficient than vendor-supplied libraries that are coded in assembly language and optimized for the target processor (or co-processor); however, the PIWG library at least provides an accepted "standard" library for compilers that don't have their own. Using the PIWG library in Hartstone also means that comparisons of the performance of Hartstone's Small_Whetstone and PIWG's "full" Whetstone procedures are possible. Details on what mathematical libraries were used for each system appear in the appendices.

Whether supplied by the vendor or not, most of the libraries differed in their naming of the functions provided, notably in the names of the base 10 and natural logarithm functions. In an attempt to make Hartstone more portable, the natural logarithm function, needed by the Small_Whetstone procedure in the Workload package, is named "Log" within the procedure; this name is actually a renaming of the real function that is effected by a renaming declaration at the beginning of the package. Such renaming was not always successful due to different numbers of parameters in the function call (Verdix and DDC-I have a parameter giving the base of the log). Verdix rejected the renaming immediately, while DDC-I appeared to accept the renaming, but then generated a compile-time error message when the statement with a reference to "Log" was compiled. In those cases, the renaming idea was abandoned and the actual function name was used within the procedure.

Apart from the changes required to accommodate the mathematical functions, only the Verdix and Tartan compilers required changes to the code of the Hartstone benchmark to ensure successful compilation. The Verdix compiler reported an error when compiling the body of the Periodic_Tasks package. This package names the Workload package in a context clause and also uses the name "Workload" for a local variable. The Verdix compiler improperly rejected this (a violation of the 1815A standard), so the variable was renamed.

The Tartan compiler reported an internal compiler error while compiling the body of Periodic_Tasks. The text of this error is given in Appendix G. The error was traced back to the pragma inline statement for the Small_Whetstone procedure. It was necessary to remove this pragma and recompile package Workload in order to successfully compile the body of Periodic_Tasks.

In all cases, Hartstone was compiled with full optimization in effect and all runtime checks enabled. Run-time checks were left enabled because of the possibility of some compilers not observing suppression of checking. We believe that these options will be the most likely options to be chosen during the initial testing of compilers.

Once Hartstone was compiled, linking proved to be straightforward. During actual execution of Hartstone code, it was discovered that, for some of the embedded cross-compilers, the default target memory layout (size and placement of code, stack, and heap areas) defined at link time was insufficient to allow Hartstone to execute. In most cases this was easily remedied, either by modifying a file of link-time options or by modifying qualifiers of the actual linking command. The exception was the Verdix compiler, which required a user to re-build the runtime kernel after changing the memory layout.

It was observed that the size of executable modules generated by different compilers differed considerably, as shown in the table below. The sizes are given in VMS blocks, where a block is half a kilobyte. There was a difference of, at most, one block (if at all) between executable modules for the four Hartstone experiments built by the same compiler and linker; the sizes shown are for Experiment 1. These figures represent storage space requirements and not necessarily final executable sizes because some may not include runtime system code. No attempts were made to optimize the space required by an executing program.

Compiler VMS Blocks

DDC-I DACS-80386PM 4.3(3.1a) 523 DEC VAX Ada 2.1 59 Rational Environment N/A Systems Designers XD Ada 1.0 225 Tartan Ada VMS/1750A 2.14 112 TeleSoft TeleGen2 3.22a 280 Verdix VADS 5.7 339

2.4. Running Hartstone and Logging Results

All of the cross-compilers supported cross I/O so that output from the Hartstone program running on a target could be displayed in a window on the host system. In general, such a window was "mapped" to a serial port on the host, which in turn was directly connected to a serial port on the target. By using a command qualifier parameter when allocating the host serial port, it was also possible to capture Hartstone output in a file on the host system. For some of the cross-compilation systems, this routing of output to a host file was performed automatically by the download program. Logging of Hartstone output was easily accomplished using either scheme, with two exceptions: the DDC-I compiler tended to scramble the Hartstone summary output, and the Verdix output was transferred at a speed on the order of several characters per second. The workaround for the DDC-I problem was to insert delay statements that slowed down the volume

of summary output produced at the end of a Hartstone experiment; this was particularly necessary for experiment 4, where the output from all the extra tasks had to be dealt with. (Increasing the size of the ALTYPEAHD buffer on the VMS host did not solve the problem.) For Verdix, the workaround was to run Hartstone experiments with the full output option disabled, so that only a summary of each experiment was displayed. It should be noted that output does not interfere with running the tests. All output is processed when the test is completed.

3. Analysis of Results

The purpose of this chapter is to present some examples of Hartstone benchmark results and explain them in the context of an Ada runtime system.

3.1. Hartstone Periodic Tasks

The Periodic Harmonic (PH) Test Series is the simplest of the five test series defined in [Weiderman 89] for the Hartstone benchmark. The Ada implementation (the "Delay/ND" design discussed in [Weiderman 89]) consists of a set of five periodic Ada tasks that are independent in the sense that their execution need not be synchronized; they do not communicate with each other. Each periodic task has a *frequency*, a *workload*, and a *priority* which is a function of the frequency. Task frequencies are harmonic: the frequency of a task is an integral multiple of the frequency of any lower-frequency task. Frequencies are expressed in Hertz; the reciprocal of the frequency is a task's period, in seconds.

A task workload is a fixed amount of work, which must be completed within a task's period. The workload of a Hartstone periodic task is provided by a variant of the well-known composite synthetic Whetstone benchmark [Curnow 76] called Small_Whetstone [Wichmann 88]. Small_Whetstone has a main loop which executes one thousand Whetstone instructions (a Whetstone instruction is roughly equivalent to one floating point operation), or one *Kilo-Whetstone*. A Hartstone task is required to execute a specific number of Kilo-Whetstones within its period. The rate at which it does this amount of work is measured in Kilo-Whetstone instructions per second, or *KWIPS*. This *workload rate* of a task is equal to its workload multiplied by the task's frequency. The *deadline* for completion of the workload is the beginning of the task's next period.

Consider a set of 5 periodic harmonic tasks with the following frequencies, workloads, and workload rates:

Task	Frequency	Workload	Workload Rate
1	2 Hertz	32 Kilo-Whets	64 KWIPS
2	4 Hertz	16 Kilo-Whets	64 KWIPS
3	8 Hertz	8 Kilo-Whets	64 KWIPS
4	16 Hertz	4 Kilo-Whets	64 KWIPS
5	32 Hertz	2 Kilo-Whets	64 KWIPS

In the Hartstone benchmark, prior to the beginning of the test, all tasks are provided with a parameter giving the starting time (we say that all tasks are *in phase* with one another). For the case given above, all of the tasks will be ready to run when Task 1 is ready, at the rate of 2 Hertz. In the case of a harmonic task set, the sequence of task executions repeats every time the lowest frequency task starts. The period of the lowest-frequency task is known as the major cycle. In the case of a 10 second test, there will be 20 major cycles each of which is identical. In order to complete without missing any deadlines, it is important that each task be given the opportunity to finish its work before the end of its period. If all of the tasks are at the same priority, this cannot

be guaranteed since it would then be possible for the task which must meet 32 deadlines per second to wait behind the other four tasks. Assume that 1 Kilo-Whet is equal to 10% of Task 5's period. Even if Task 5 must only wait behind Task 1, it will not meet its deadline because Task 1 runs for 32 Kilo-Whets or 320% of Task 5's period. Therefore, distinct priorities are assigned to the tasks based on their frequencies with higher priority given to tasks which execute more often. This is consistent with the *rate-monotonic* scheduling discipline [Liu 73, Sha 89] which is an optimal scheduling strategy for periodic tasks that have statically assigned priorities. It provides analytical proof that all deadlines will be met if the total workload is below a specified level.

3.2. Runtime Overhead

The execution time of a task consists of overhead and useful work (in this case the Small_Whetstone procedure). Interrupt processing, task switching, delay expiration, delay queue management, and time management are all examples of overhead in an Ada runtime system. In order for a task to meet its deadline, it must be able to wait for all higher priority tasks to finish their work as well as have sufficient time to execute its own workload. Since Task 5 is of the highest priority, when its delay expires it only needs to wait for any lower priority task to be preempted, and for the resultant queue management, before it is able to perform its workload. It may also experience delay expirations from other tasks as well as interrupts from the hardware clock. Task 5 should meet its deadlines, then, as long as its period is greater than the sum of all overhead components and execution time of its workload. The overhead will be longer or shorter depending upon a number of factors including whether another task is running when its delay expires.

Task 4 is of a lower priority than Task 5 but higher than all other tasks. Therefore, in order for Task 4 to meet its deadline, it must be able not only to perform its own workload and experience various overhead components, but must suffer preemption by Task 5 as well. Depending on the period of Task 5, Task 4 may be preempted multiple times before it can meet its deadline. Figure 3-1 shows two timelines that illustrate the effect of Task 4 being preempted twice during its execution by Task 5. The periods of Task 5 and Task 4 are shown below their respective timelines.

The same concepts hold true for Tasks 3 through 1. Task 3 must be able to handle preemption by both Tasks 4 and 5 in addition to its own workload and various overhead components in order to meet its deadline.

The overhead that occurs when one or more delays expire can be considerable. The resulting clock and queue management operations can add significant additional overhead during which time even the highest priority task, Task 5, may be affected. Delay expirations are of particular interest at the beginning of the major cycle. Task 5 must wait to begin execution until the expiring delays for Tasks 1-4 have been processed. Evidence of this was seen while using a logic analyzer to measure execution times of portions of an Ada runtime system. In one case, this blocking factor (processing the delay expirations of lower priority tasks) amounted to almost 50% of the overhead for Task 5 during the first period at the beginning of the major cycle in Exper-
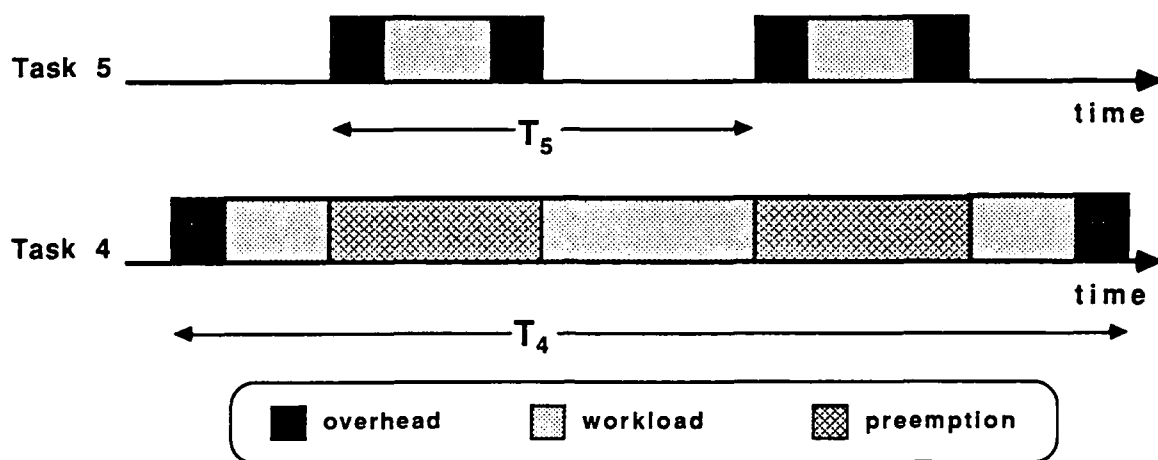
Figure 3-1: Effects of Preemption

iment 1. In experiment 4 (where more tasks are added and all have delays expiring at the beginning of the major cycle), this factor may become more noticeable depending on how efficiently an Ada runtime system handles a large number of queued tasks.

For one of the compilation systems tested, detailed measurements of the runtime system were obtained using a logic analyzer. The following sequence of events was observed at the beginning of a Hartstone major cycle. With the ready queue empty, and a null loop within the runtime system executing, an interrupt was received signaling the expiration of a timer. The runtime system responded by removing 3 tasks from the delay queue and placing them in the ready queue. There was a fixed amount of overhead involved in performing this for one task and a small additional amount for the 2 other tasks. The timer was then reset for the 2 additional delays remaining. The total time elapsed from the start of the interrupt was 139 microseconds. Task 5 was among those placed in the ready queue and was chosen to run based on its priority. Task 5 began executing and was interrupted 136 microseconds later when the timer expired again. The runtime system then saved the state for Task 5, removed the 2 tasks from the delay queue and added them to the ready queue, decided Task 5 was the highest priority task and restored its state. This took 184 microseconds. Task 5 was then able to complete its execution. The runtime system then spent 57 microseconds performing Task 5's delay statement and switching to Task 4 which was ready to run. Task 5 had a compute time of 1.6 milliseconds. Thus, the associated overhead discussed above represents 24% of its execution time. What is notable here is the performance penalty of a number of delay expirations one tick apart as opposed to a number of delay expirations that happen simultaneously. This example illustrates how micro-level analysis may be used to explain some of the macro-level behavior exhibited by the Hartstone benchmark.

Hartstone users should also be aware of interrupts from events external to the Hartstone program as potential sources of overhead. On a system with an underlying operating system these events might include daemons for servicing a clock or for performing checkpoint services or for multiprogramming. It may be difficult or even impossible to prevent some of these sources of interrupts with the result that deadlines are often missed by large amounts. Bare target machines are

more controllable, but the Ada runtime system may service interrupts from a hardware clock in order to update the software clock used in package Calendar. These interruptions are usually short (several machine instructions), but they do contribute to the overhead at unpredictable times. The frequency of such interrupts may vary widely from one runtime system to another. Any interruptions that require processing time that is a significant proportion of the period of Task 5 are likely to cause deadlines to be missed prematurely.

3.3. Hartstone Metrics

The *raw speed* of the benchmark is the number of Kilo-Whetstone instructions per second (KWIPS) achieved by the Small_Whetstone procedure. This calibration test is performed by the Experiment package when an experiment is initialized. The resultant non-tasking workload rate can be no worse than that achievable by splitting the same workload among the five Hartstone tasks (due to the lack of context switching overhead); it provides a metric against which the performance of the Hartstone task set can be measured. Both the raw speed calibration test and a Hartstone task include the overhead of calling the Small_Whetstone procedure. The performance requested of Hartstone tasks is expressed as a percentage workload *utilization*, which is computed as the ratio of the requested task speed (in KWIPS) to the raw benchmark speed. The raw speed is assumed to represent 100% utilization. The utilization required of the entire task set is the sum of the individual task utilizations. Successive tests in an experiment increase the requested utilization to the point where deadlines are not met. The *step size* of an experiment is an indication of the extra work required of the task set when the next test in an experiment is derived from the current test. Like the workload utilization, it is expressed as a percentage of the raw speed. The step size is the granularity, or resolution, of an experiment.

In addition to utilization, there are other important metrics that are included with the output of the Hartstone benchmark. They are the number of deadlines that are processed per second (for the entire task set and for Task 5), and the total number of tasks. Deadlines per second is an important metric because it is a direct measure of the amount of tasking overhead being incurred. For every deadline, there is a delay expiration and for every delay expiration there is a task switch. There should be no task switches that do not correspond to delay expirations because a task either completes its workload or is interrupted because a higher priority task becomes ready due to a delay expiration. Task 5 is important because it has the highest relative overhead, a result of the highest frequency and the shortest period. The total number of tasks is important because of the extra overhead incurred during queue management when delays expire for many tasks at the same time. Good Hartstone results can thus be equated to high breakdown utilization (i.e., utilization at which deadlines are first missed) while processing many deadlines per second, with many tasks and with at least one high frequency task. Different experiments attempt to increase these metrics in different ways.

In Experiment 3, workload is increased without increasing overhead. In Experiments 1, 2, and 4 of the Hartstone, both overhead and workload increase. In Experiment 1 and 4 overhead increases faster than workload, while in Experiment 2 workload and overhead are increased in roughly the same proportion. Details of how workload is increased appear in the following sec-

tions. It is important to consider, therefore, the amount of overhead being performed by the benchmark relative to the maximum achievable utilization. To illustrate this point, consider the following two baseline task sets. In both cases the raw speed is 1795 KWIPS and the step size is 1.78%. The step size is derived from the fact that only the frequency of Task 5 is being increased and that it is increasing by 16 Hz. Thus the step size is 32 KWIPS (16 Hz times 2 KWI per period) divided by the raw speed (1795).

Baseline Set 1:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	32.00	2	64.00	3.56 %
			--------	--------
			448.00	24.95 %

Baseline Set 2:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	96	192.00	10.69 %
2	4.00	48	192.00	10.69 %
3	8.00	24	192.00	10.69 %
4	16.00	12	192.00	10.69 %
5	32.00	2	64.00	3.56 %
			--------	--------
			832.00	46.34 %

The only differences between these 2 task sets are the workloads for Tasks 1-4. Otherwise, they were run on the same compiler and target processor. Notice that the requested workload utilization is much higher for the second baseline task set than for the first. Now, compare the results of Experiment 1 for the two baseline task sets. The first set of results is a duplicate of those presented for Systems Designers XD Ada in Appendix F. The test number in this and other sample outputs indicates the number of times the workload rate has been incremented to arrive at the given workload in a series of tests.

Set 1, Last test with no missed/skipped deadlines:

`Test 33 characteristics:`

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	544.00	2	1088.00	60.60 %
			1472.00	81.99 %

Set 2, Last test with no missed/skipped deadlines:

`Test 24 characteristics:`

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	96	192.00	10.69 %
2	4.00	48	192.00	10.69 %
3	8.00	24	192.00	10.69 %
4	16.00	12	192.00	10.69 %
5	400.00	2	800.00	44.56 %
			1568.00	87.34 %

As expected, increasing the workloads of Tasks 1-4, and thus starting from a higher baseline workload utilization, allowed Set 2 to achieve a higher overall workload utilization at the expense of the performance of Task 5. The overhead for Tasks 1-4 is the same for both task sets while in Set 2 Tasks 1-4 contribute more to the workload utilization because they have more Kilo-Whets per period. Task 5 is the primary contributor of overhead for the task set.

In comparing these results, there are two major points to look for. The first is the frequency of Task 5. In Set 1 it is 544 Hertz while in Set 2 it is only 400 Hertz. This means that Set 2 is performing significantly less overhead than Set 1 (all of the difference can be attributed to Task 5). Second, Set 2 is able to perform more Kilo-Whets per second than Set 1. Because of this, it has an overall workload utilization of 87.34% which is higher than Set 1 at 81.99%.

This phenomenon explains some of the "surprising" or "anomalous" results that were experienced when comparing Hartstone results for different compilers or processor boards. When starting with the *same* baseline task set, i was sometimes the case that a compiler or board that had a low utilization for the baseline (prc :essed the workload faster) had a lower breakdown utilization even though we considered it a better compiler or faster board. This was due to the fact that the frequency of Task 5 was driven much higher for the better compiler or faster board, and as a consequence, the overhead became proportionally higher as well. Section 3.8 gives an example of how a processor board can influence utilization. The moral here is that breakdown utilization is not a sufficient metric to compare two compilers or processor boards. Guidelines for setting the

starting baseline to achieve the most meaningful results are given in the users guide [Donohoe 90].

3.4. Experiment 1

The above examples are results from Experiment 1. This experiment starts with a baseline task set, and increases the frequency of the highest frequency task (Task 5) for each new test until a task misses a deadline. The frequencies of the other tasks and the per-period workloads of all tasks do not change. The amount by which the frequency increases must preserve the harmonic nature of the task set frequencies: this means a minimum increase by an amount equal to the frequency of Task 4. Primarily, Experiment 1 tests the system's ability to handle a fine granularity of time and to switch rapidly between processes. As a result of this, Experiment 1 is affected most by the clock and delay statement resolution (see Appendix B) as well as task switching overhead.

It is important to start with a baseline task set in Experiment 1 that will not attempt to run a test where the period of Task 5 is close to either the delay or clock resolution. Those compilers with a clock resolution of 100 or even 10 milliseconds cannot support the baseline task sets shown above. In the previous example, Task 5 is running at a rate of 400 Hertz. Therefore, its period is 2.5 ms which is not measurable except with a fine-grained clock. In a number of compilers, the clock granularity effect overwhelmed the context switching time.

For those compilers which have a delay or clock resolution of more than 1 ms, it will probably be necessary to modify the Hartstone task set so that the limits of the clock granularity are not hit early. *Scaling* an experiment involves modifying the baseline task set by changing the frequencies and/or workloads to increase or decrease the amount of work. In running the tests documented in the appendices, the rule of thumb used was that an experiment was not properly scaled if Task 5 missed its deadlines before Task 1. Since Task 1 is the lowest priority task in the set, and because the task set is harmonic, it seems reasonable that it should miss its deadlines first. This will not be the case, however, if Task 5's period approaches the clock granularity before the CPU is fully loaded. Adding delay and clock resolutions into our previous analysis of Task 5 raises the following issue: In order to meet its deadline, Task 5 must have sufficient time to perform the requested workload, incur scheduling overhead (including any on behalf of lower priority tasks), start late due to the delay statement resolution, and have the completion time incorrect by as much as the clock resolution. Those compilers with both an imprecise delay and clock resolution will be more affected than those where only one is imprecise.

Even when the delay and clock resolution are not an issue, it is still possible to have Task 5 missing some of its deadlines before Task 1. The overhead experienced by Task 5 may vary greatly and it may only miss during the worst cases at first. These tend to be at the beginning of the major cycle when Task 5 is apt to be interrupted by multiple delays expiring on behalf of lower priority tasks. Task 5 may miss its first deadline of the major cycle, have to skip one, and then succeed with time to spare until the next major cycle or half-cycle. Load shedding, where a task skips at least its next deadline upon discovery that it has missed the current one, provides additional time for lower priority tasks to run.

To illustrate this point, consider these two baseline task sets. These sets were run on Systems Designers XD Ada using the configuration presented in the appendices. In both cases the raw speed is 1795 KWIPS and the step size is 1.78%.

Baseline Set 3:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
			-------	--------
			320.00	17.82 %

Baseline Set 4:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	1.78 %
2	4.00	8	32.00	1.78 %
3	8.00	4	32.00	1.78 %
4	16.00	2	32.00	1.78 %
5	32.00	2	64.00	3.56 %
			-------	--------
			192.00	10.69 %

The only difference between the baseline task sets shown above is the amount of work being performed by Tasks 1-4. In the following results, Task 5 began missing its deadlines first for both of the task sets.

Set 3, Last test with no missed/skipped deadlines:

Test 34 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	560.00	2	1120.00	62.38 %
			-------	--------
			1376.00	76.64 %

Set 4, Last test with no missed/skipped deadlines:

`Test 34 characteristics:`

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	1.78 %
2	4.00	8	32.00	1.78 %
3	8.00	4	32.00	1.78 %
4	16.00	2	32.00	1.78 %
5	560.00	2	1120.00	62.38 %
			-------	--------
			1248.00	69.51 %

In both cases, Task 5 began missing deadlines on the next test while Tasks 1-4 ran without missing any. This is attributed to the expiring delays and context switching time at the beginning of the major cycle. Notice that the determining factor in both sets of results is the maximum frequency achievable by Task 5 (560 Hertz), and not the requested workload utilization. Breakdown is attributed to runtime system overhead rather than clock or delay resolution because the clock resolution for this compiler is more than ten times as precise as the period of Task 5. It is important to remember that workload utilization figures may be arbitrarily low in the event that Task 5 misses its deadlines before Task 1. There is very little difference between Set 3 and Set 1 shown earlier. By increasing the workloads of Tasks 1-4 by 50% over those in Set 3, Set 1 demonstrates a Task 1 breakdown while still reflecting a Task 5 frequency of 544 Hertz.

3.5. Experiment 2

Experiment 2 starts with the baseline task set after which all the frequencies are increased by a factor of 1.1, then 1.2, then 1.3, and so on for each new test until a deadline is missed. The per-period workloads of all tasks do not change. This scaling preserves the harmonic frequencies. As with Experiment 1, this sequence is affected by the clock and delay statement resolution. By contrast with Experiment 1, the increasing rates of work affect all tasks, not just one.

Since Experiment 2 increases the frequency of Task 5 at a much slower rate than Experiment 1, it will be possible to return the baseline task set to its original size, if it had to be scaled up for Experiment 1 to prevent Task 5 from missing its deadlines before Task 1. See, for example, the results for System Designers XD Ada in Appendix F.

3.6. Experiment 3

In Experiment 3, the workload of each task is increased by 1 Kilo-Whetstone per period for each new test, continuing until a deadline is missed. The task frequencies do not change. Therefore, if the baseline set is within the range which can be supported with the delay and clock resolution applicable to the compiler, the remainder of the tests will be as well. If an imprecise delay or clock had forced a scaling back of the baseline set in order to run Experiments 1 and 2, it may be possible to scale the frequencies of the task set back up to the baseline sets shown earlier.

Experiment 3 increases the total workload rate without increasing the system overhead as in the other experiments. It is likely that the best total utilization figures will be achieved with Experiment 3. Unless the baseline set was scaled back dramatically for Experiments 1 and 2, the total number of deadlines per second which are being met in Experiment 3 will be considerably lower than in the first two experiments.

Depending on the raw speed of the benchmark, Experiment 3 may have a very large step size. This is in part due to the indivisible nature of a Kilo-Whestone in the benchmark. Also, since one Kilo-Whetstone is added to the workload of each task, the number of Kilo-Whetstones added per second is relative to the frequencies of each task. If the raw speed is low, this will be a significant increase in the experiment step size. In the results presented in the appendices, an attempt was made to keep step sizes below 4% and, where possible, below 2%. With a large step size, it is difficult to ascertain exactly where the breakdown point is. In one case, a task set succeeded at 94% and failed at 98% with a step size of 4%. With further testing, it was possible to determine that the actual breakdown point was about 96%. A large step size cannot be remedied by reducing the workloads of the task set as was possible in Experiments 1 and 2, but only by reducing the frequencies. This will reduce the total number of Kilo-Whetstones per second added in each new test.

3.7. Experiment 4

Starting with the baseline task set, new tasks with the same frequency, workload, and priority as the "middle" task, Task 3, of the baseline set are added until a deadline is missed. The frequencies and workloads of the baseline task set do not change. This sequence tests the system's ability to handle a large number of tasks (and a large number of delay expirations at each major cycle).

Experiment 4 adds both overhead and Kilo-Whetstones per second to the executing task set. It also increases the amount of blocking (preemption by activities of lower priority tasks) incurred by the task set. This can be seen in 2 ways. First, Tasks 4 and 5 will be blocked waiting for the servicing of delay expirations on behalf of those tasks executing at the priority of Task 3. Once several tasks have been added, this becomes a significant amount of time. Second, since priorities are no longer unique to tasks, the performance of the ever-increasing number of tasks running at the priority of Task 3 is of interest. Similarly, the breakdown point is of interest because, again, Task 1 might not miss its deadlines before higher priority tasks. This is because many tasks have concurrent delay expirations and this may cause breakdowns at the beginning of a major cycle because of delay handling and queue management.

3.8. Changing the CPU

Earlier in this section, and in Appendix F, results are given for Experiment 1 running on a Motorola MVME133A single-board computer (20 MHz MC68020 CPU) under Systems Designers XD Ada. Consider the following baseline task set, tested on a Motorola MVME133, (12.5 MHz MC68020 CPU) using the same cross-compilation system.

Baseline Set on MVME133A:

```
Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.39
```

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	32.00	2	64.00	3.56 %
			448.00	24.95 %

```
Experiment step size:    1.78 %
```

Baseline Set on MVME133:

```
Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1121.84
```

```
Test 1 characteristics:
```

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	8.56 %
2	4.00	24	96.00	8.56 %
3	8.00	12	96.00	8.56 %
4	16.00	6	96.00	8.56 %
5	32.00	2	64.00	5.70 %
			448.00	39.93 %

```
Experiment step size:    2.85 %
```

Notice that the lower KWIPS rate for the MVME133 increases the requested workload utilization to almost 40%. As shown earlier, this implies a higher maximum achievable utilization. Here are the results of these tests:

MVME133A, Last test with no missed/skipped deadlines:

Test 33 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	544.00	2	1088.00	60.60 %
			1472.00	81.99 %

MVME133, Last test with no missed/skipped deadlines:

Test 16 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	8.56 %
2	4.00	24	96.00	8.56 %
3	8.00	12	96.00	8.56 %
4	16.00	6	96.00	8.56 %
5	272.00	2	544.00	48.49 %
			928.00	82.72 %

As expected, the faster MVME133A board, achieved a considerably higher frequency for Task 5. Also, the MVME133A board was able to perform more Kilo-Whets per second than the MVME133 board. The total utilization, however is about the same.

3.9. General Observations

The compilers with the best Calendar.Clock and delay statement resolutions tended to yield the better Hartstone results, in terms of utilization, deadlines met per second, and frequency of Task 5. The effects of context switching overhead and processing delay expirations were second order effects. These factors were an important determinant of Hartstone performance. To a lesser extent the results were also influenced by the efficiency of the mathematical library; this factor showed up in the reported raw speed for the experiments. The PIWG benchmark results reported in Appendix B were good indicators of Hartstone performance. While it is possible to have poor Hartstone results even with good clock and delay statement resolution (because of interrupts, context switching overheads, and queue management), it is not possible have good Hartstone results with poor clock and delay statement resolution. A good deal of variation was observed that could not be explained solely by PIWG results.

3.10. Hartstone Results

To obtain the results documented in this report, the guidelines recommended in the Hartstone User's Guide [Donohoe 90] were followed (i.e., the benchmark was first run as-is on all the systems tested and the step size and task set utilization of the baseline test was examined). If these values were too high or too low, the frequencies or workloads (or both) of the baseline task set were changed to bring them in line with the guidelines.

All of the cross-compilers tested required modification of the baseline task set for at least one experiment. The guideline that the workload utilization of the baseline set should be in the range of 10 percent to 30 percent was broken only in the following cases: Verdix and VMS Ada required baseline workload utilizations for Experiment 1 of 56.88 and 69.64 percent respectively. This was done to prevent Task 5 (as the lowest priority task) from missing its deadlines first. In each of these cases the workloads of Tasks 1-4 were raised until Task 5 no longer missed its deadlines before Task 1. In Experiment 2, VMS Ada required a baseline workload resolution of 51.46 percent for the same reason as Experiment 1. Similar scaling back was done for TeleSoft Ada with less success. Despite modifications to the baseline set, the higher priority tasks would still miss deadlines before the lower priority ones. See Appendix H for details.

The principal reasons for scaling the frequencies and workloads up or down were:

- an imprecise Calendar.Clock or delay statement
- a low or high raw speed
- a test where Task 5 missed deadlines first
- a large step size

As discussed earlier, the main effect of an imprecise clock or delay statement is the imposition of an upper bound on a task's frequency at a value lower, possibly much lower, than that achievable in the presence of high precision. It should be noted that the timing accuracy is not necessarily a function of the hardware, but rather a function of the software that uses the hardware clocks. This is demonstrated by the quite different results obtained from the three compilers running on the same Motorola 68020 board.

4. Summary and Conclusions

4.1. Hartstone Conclusions

Having run the Hartstone benchmark against seven different compilers, we continue to believe that it is a highly useful tool for evaluating the real-time characteristics of Ada compilers and runtime systems. Like any benchmark program, the Hartstone benchmark has limitations. Unlike most benchmarks, it rarely can be run "as is" and will not produce meaningful results without some interpretation. The purpose of this report has been to provide enough variety of results so that users can develop their own testing regimes using the benchmark.

One of the primary lessons that we have learned in applying the program to several compilers (as opposed to just writing the program in the first place) is that the utilization figure is not the only meaningful metric. The utilization was meant to give an overall picture of the total overhead of a compilation system. Since the utilization is 100% in the theoretical model of a harmonic task set with rate monotonic assignment of priorities, it was expected that the overhead would be reflected as one minus the utilization. What we failed to adequately anticipate is that the utilization would be largely dependent on the precision of the timing services as well.

Utilization may be misleading for the following reason. The starting point of the experimentation must be carefully chosen based on the precision of the timing (so that the highest priority task does not start to miss its deadlines first). This means that for a compilation system with imprecise timing services and high Whetstone throughput, the Hartstone benchmark must be modified to start out at a much higher utilization than a compilation system with precise timing and low Whetstone throughput. But both systems may have similar breakdown utilization. What is more important than utilization in many cases is the number of deadlines that can be processed per second by the highest frequency task and the number of deadlines that can be processed per second by the total set of tasks.

Thus we have found that the frequency of the highest-frequency task and the number of deadlines processed per second are two useful metrics. All three of these measures give an indication of the overhead of the runtime system. They all provide measures of the amount of useful work that can be accomplished in the face of significant context switching. The user would like to see high workloads as well as high utilization and high context switching. The figure of merit described in [Weiderman 89] that is a function of the utilizations of the four experiments turns out to be too simplistic. More careful analysis is required.

One particularly nasty implication of using utilization as the metric and a uniform workload as a starting point is that the better compilers may have breakdowns at lower utilizations. This is because compilers generating more efficient code will have higher throughput for the baseline task set. In Experiment 1 this allows the highest-frequency task to be pushed to the limit of the clock resolution. This in turn causes Task 5 breakdown at a low utilization whereas a compiler with less throughput will start at a higher utilization for the baseline task set and will a have Task 1 breakdown with higher utilization. Consequently, we have downplayed the importance of utili-

zation alone and have adjusted the baseline workloads to better balance the increasing workload so that the frequency of the high frequency task does not dominate the results by always causing breakdown.

Another minor adjustment from the original specification of the benchmark is that instead of running the program for a minimum of ten seconds we have realized that it is more important to make the length of the run proportional to the length of the major cycle. If the major cycle is one second or less, then a 10 second run appears to be adequate. If the major cycle is more than one second, then the length of a Hartstone run should be extended proportionately. (It should be noted that this guideline was not met in some of the results presented in this report.)

4.2. Observations on Ada Compilation Systems

The most striking observation that can be made about Ada compilation systems is the wide variability with respect to their timing behavior. Timing resolution of the system-provided clock varied by three orders of magnitude even for different compilers running on the same target. Timing resolution and behavior of the delay statement (granularity, overall accuracy, and accuracy near zero) were highly correlated with the timing behavior of the system clock. Both had a significant impact on the Hartstone results. The conclusion here is that applications with high precision timing requirements need to be careful about the compilers selected.

A second major conclusion is that subtle bugs can and do lurk in the Ada runtime systems with regard to timing behavior. It is extremely difficult to force the runtime code into every possible timing scenario. The Hartstone benchmark found runtime bugs in several of the compilation systems tested. These were usually manifested in missed deadlines by high priority, high frequency tasks before they reached the limits imposed by the clock resolution and overhead. The harmonic nature of the Hartstone tasks causes many deadlines to expire at the same time. This seemed to be the feature most likely to trigger a bug, and did so even in one of the better compilation systems tested. There remain behaviors of implementations that have not been adequately explained. This has been found to be a fruitful area for further conformance testing by the Ada Compiler Validation Capability (ACVC).

A third major conclusion is that the support services, particularly for embedded computing, vary greatly from one compilation system to the next and should be carefully evaluated. Areas in which we noticed variability are the availability of a vendor supplied mathematics library and, if provided, the naming of functions in that library, quality of error messages and documentation, ease of configurability of the runtime system, size of the runtime image, speed of output to the host system, speed of downloading, and the number and names of numeric data representations. In a number of these areas the 9X language revision process, is providing some help in standardization.

Finally, and to end on a positive note, there are reasons to be sanguine about the prospects for using Ada in hard real-time applications. We found that many of the Ada compilation systems could reliably meet hard deadlines with low overhead. Periodic tasks of over 500 Hz proved feasible even with significant baseline processing for some compilation systems. We believe that

these results effectively put to rest the often heard canard that there are no Ada implementations that can effectively handle deadline-driven computing with tasking. We have shown that the overhead of tasking in the face of significant task switching and delay expirations can be well under 10% of workload across a variety of targets.

4.3. Use of Hartstone

We believe that the Hartstone benchmark can be a useful tool for reducing the risk of developing real-time programs in Ada. We would like to promote its use by program managers who need to select a compiler, by application developers who need to use a compiler, and by compiler developers who need to develop and sell compilers.

We would like this report to be a template for reporting results. We would like to provide support for Hartstone and encourage others to develop variants which suit their purposes. We would like to see variants written in other languages, variants that use different scheduling strategies (for example cyclic scheduling), and variants that use workloads other than Whetstones. In order to support this activity, the SEI is distributing the Hartstone benchmark electronically and providing an electronic bulletin board for interaction among its users. Further details of this activity can be found in Appendix J.

With respect to the flexibility of Hartstone as a prototyping tool it should be noted that performing the same type of experimentation with a cyclic executive model would be a daunting task. Instead of changing only the workloads, frequencies, and priorities, the whole fabric of the application executive would have to be changed for each new experiment. In that sense, it should be possible to demonstrate one of the primary advantages of a rate monotonic scheduling philosophy by comparing a series of Hartstone experiments using both rate monotonic as well as cyclic scheduling.

4.4. Future work

The results, analysis, and conclusions presented in this report only address the problem of independent, periodic, harmonic tasks. The original Hartstone concept paper [Weiderman 89] called for relaxing these restrictions in subsequent series of tests. While the current Hartstone benchmark answers a number of questions about the abilities of Ada compilations systems, it is still too simple a test to inspire a high degree of confidence in the abilities of Ada compilation systems to handle more complex applications in which there are interrupts being processed, communication between tasks, and variability in processing times. Investigation into these areas is expected to be a fruitful endeavor and will undoubtedly yield as many insights and surprises as the current studies.

References

[Baker 90]   Baker, T.
                  Fixing some Time-Related Problems in Ada.
                  Ada Letters - Special Edition on the Third International Workshop on Real-
                      Time Ada Issues (1989) X(4):136-143, 1990.

[Borger 89]   Borger, M., Klein, M., Veltre, R.
                  Real-Time Software Engineering in Ada: Observations and Guidelines.
                  Technical Report CMU/SEI-89-TR-22, DTIC: ADA219020, Software Engineer-
                      ing Institute, Carnegie Mellon University, Pittsburgh, PA 15213, Septem-
                      ber, 1989.

[Clapp 86]   Clapp, Russell M. et al.
                  Toward Real-Time Performance Benchmarks for Ada.
                  Communications of the ACM 29(8):760-778, August, 1986.

[Curnow 76]   Curnow, H.J. and Wichmann, B.A.
                  A Synthetic Benchmark.
                  Computer Journal 19(1):43-49, January, 1976.

[Donohoe 90]   Donohoe, P., Shapiro, R., Weiderman, N.
                  Hartstone Benchmark User's Guide, Version 1.0.
                  March, 1990

[Harbaugh 84]   Harbaugh, S. and Forakis, J.
                  Timing Studies Using a Synthetic Whetstone Benchmark.
                  Ada Letters 4(2):23-34, 1984.

[Liu 73]   Liu, C.L. and Layland, J.W.
                  Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environ-
                      ment.
                  Journal of the Association of Computing Machinery 20(1):46-61, January,
                      1973.

[LRM 83]   United States Department of Defense.
                  Reference Manual for the Ada Programming Language
                  American National Standards Institute, New York, 1983.

[Pollack 90]   Pollack, R. H., and Campbell, D. J.
                  Clock Resolution and the PIWG Benchmark Suite.
                  Ada Letters - Special Edition on Ada Performance Issues X(3):91-97, 1990.

[Roy 90]   Roy, D., and Gupta, L.
                  PIWG Analysis Methodology.
                  Ada Letters - Special Edition on Ada Performance Issues X(3):217-229, 1990.

[Sha 89]   Sha, L. and Goodenough, J.B.
                  Real-Time Scheduling Theory and Ada.
                  Technical Report CMU/SEI-89-TR-14, DTIC: ADA211397, Software Engineer-
                      ing Institute, Carnegie Mellon University, Pittsburgh, PA 15213, April,
                      1989.

[Weiderman 89]   Weiderman, Nelson.
Hartstone: Synthetic Benchmark Requirements for Hard Real-Time Applications.
Technical Report CMU/SEI-89-TR-23, DTIC: ADA219326, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213, June, 1989.

[Wichmann 88]   Wichmann, B.A.
Validation Code for the Whetstone Benchmark.
Technical Report DITC 107/88, National Physical Laboratory, Teddington, Middlesex, UK, March, 1988.

Appendix A: Rationale for the Hartstone Periodicity Paradigm

The implementation of task periodicity in the Hartstone benchmark is based on the paradigm exhibited in Section 9.6 of the Ada Language Reference Manual [LRM 83], a version of which is shown below.

```
declare
  use Calendar;
  -- Period is a global constant of type Duration
  Next_Start : Time := Clock + Period;
begin
  loop
    Next_Delay := Next_Start - Clock;
    delay Next_Delay;
    -- do some work
    Next_Start := Next_Start + Period;
  end loop;
end;
```

In this implementation, a task computes its next scheduled activation time, suspends itself until that time, then "wakes up" and begins executing its assigned workload. On completion of the workload, a new activation time is computed. In the case of a Hartstone periodic task, the deadline for completion of the workload is the new activation time; failure to complete the assigned work before the next scheduled activation time is a missed deadline for the task. This is a highly portable implementation of an Ada periodic task, but a well known flaw is the fact that the task could be preempted between the reading of Calendar.Clock and the issuance of the **delay** statement (see, for example, [Borger 89, Baker 90]). If that happens, the execution of the **delay** statement (and the subsequent expiration of the delay) is postponed for the duration of the preemption. Long preemption will cause a long increase in the actual time of restart for the periodic task. While this is a problem in general for the periodic Ada paradigm, it can be demonstrated that it is *not* a problem for the harmonic Hartstone task set.

The Hartstone task set (in the Version 1.0 Ada implementation) has the following properties:

- The tasks have harmonic frequencies (i.e., the frequency of any task in the set is an integral multiple of the frequency of all lower-frequency tasks).

- The tasks have been assigned unique rate monotonic scheduling priorities [Liu 73], [Sha 89].

- A higher-priority task will preempt the execution of any lower-priority task.

- All tasks are independent, in the sense that they do not explicitly communicate or synchronize with one another.

- Task execution times vary from task to task but the execution time of any one task remains constant.

- All tasks are initially activated simultaneously and remain in phase.

- The number of machine instructions between the reading of the clock and the issuance of the delay is small.

Any preemption suffered by a task, τ_i, is caused by the execution of higher-priority tasks. (We assume that any blocking caused by expiration of delays by lower priority tasks is small, but it is covered by the same argument as preemption in any event.) Since all tasks of higher priority than τ_i will always be scheduled for activation simultaneously with τ_i, task τ_i is delayed in its starting time by at least the sum of the computation time of all the higher priority tasks. (Some of the high frequency tasks could execute more than once during this period of preemption.) Depending on its execution time, τ_i may be preempted at other times during its cycle as well. Thus, even though task τ_i is ready to execute at the beginning of its period, the earliest it can actually start is the scheduled time plus the preemption time, as shown in the first timeline diagram of Figure A-1. As before, the task does not begin executing until time t_a; it then gets preempted during execution at time t_b and doesn't regain control of the CPU until time t_c.
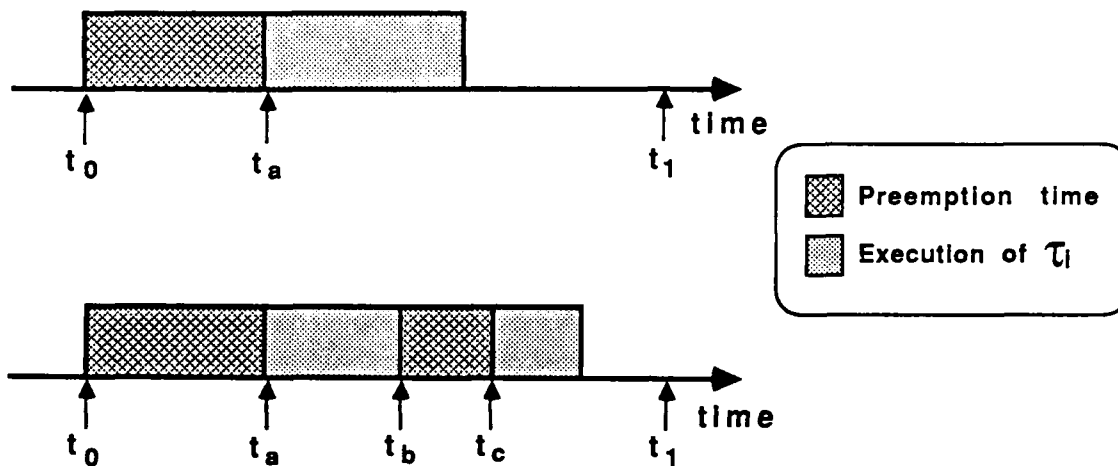
Figure A-1: Preemption at Start of Period

The task is scheduled to execute at time t_0; because of preemption and the harmonic nature of this task set, it actually starts at time t_a. This situation will repeat at time t_1, the start of the task's next period, and so on for all of its periods. Depending on the execution time of τ_i, and the frequency of the higher-priority tasks, task τ_i may also be preempted during its execution, as shown in the second timeline of Figure A-1.

Task τ_i may experience multiple intervals of preemption during the course of its period. If one or more intervals of preemption occur between the time the clock is read and the time the delay is started (refer to this as the window of vulnerability), then this postponed start-of-delay will result in a postponed expiry-of-delay. The question is, what is the effect of this delay postponement on the ability of this task to meet its deadline? We will examine two cases: (1) one interval of preemption occurs in the window of vulnerability or (2) more than one interval of preemption occurs in the window of vulnerability. Case 2 is extremely unlikely due to the fact that the window of vulnerability is only a few machine instructions, but it must be considered nevertheless.

In both cases the amount of time that the expiry-of-delay is postponed is equal to the total amount of preemption time experienced in the window of vulnerability. If the task would have been preempted by higher priority tasks for the entire duration of the delay postponement, then the postponement would not have an effect because the task could not have started even if there were no postponement (see first timeline of Figure A-2). If the task would not have experienced preemption for the entire duration of the postponement, then there would be intervals of time that the task is wasting, i.e., it would have been doing useful work if it had not been suspended due to the delay postponement (see second timeline of Figure A-2).
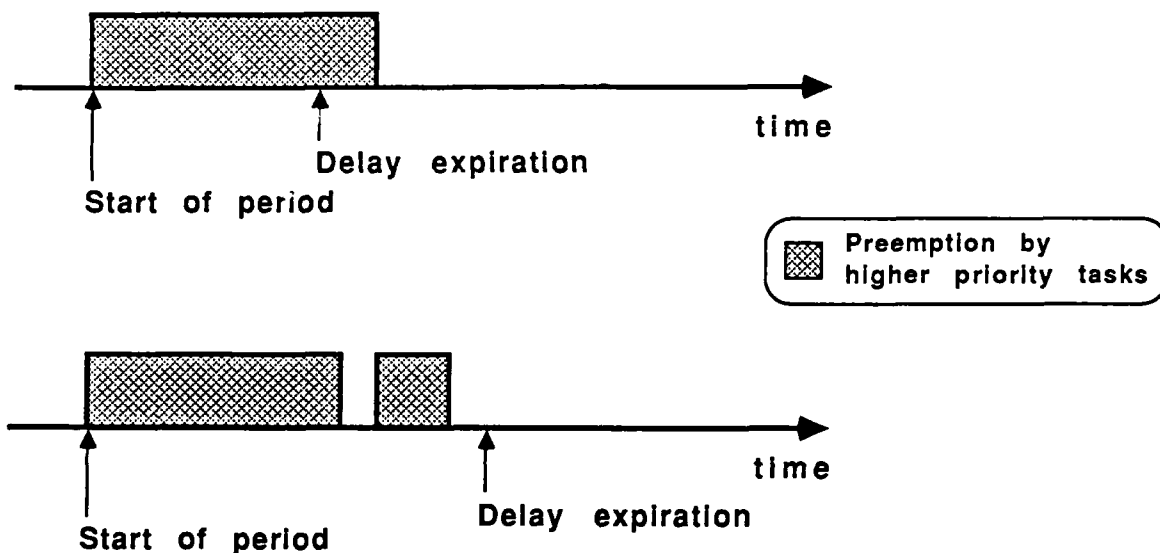
Figure A-2: Preemption at Start of Period

Clearly, wasting time will affect the task's ability to meet its deadline. The issue then boils down to determining the circumstances under which delay postponement causes wasted time and the duration of wasted time. The following argument under Case-1 will show that if there is only a single interval of preemption during the window of vulnerability, then there is no wasted time. The argument under Case-2 will show that if there are multiple intervals of preemption during the window of vulnerability, then the wasted time can not be longer than the amount of execution time it takes to perform all instructions from the reading of the clock to the issuance of the delay (i.e., the amount of time associated with the window of vulnerability, if the window executed in a non-interruptable section. We will refer to this amount of execution time as the window execution time).

Case-1: A Single Interval of Preemption

Since tasks are in phase and each has a constant execution time, the longest interval of preemp-

tion that a task can experience during the course of its period is the interval that starts at the beginning of its period. Since, by assumption, a single interval of preemption causes postponement, postponement is no longer than the *longest* single interval. Since we know an instance of the *longest* single interval of preemption occurs at the *beginning* of the period, the postponed expiry-of-delay must occur within the first interval of preemption. Hence there is no wasted time as illustrated in the first timeline of Figure A-2.

Case-2: Multiple Intervals of Preemption

If there are multiple intervals of preemption, then the total separation time between these multiple intervals must be less than the window execution time. Consider an example.

If there were two intervals of preemption, the separation time between these intervals must be less than the window execution time. If the separation time was greater than the window execution time, then the delay would have been issued within the separation time and consequently there would only be one interval of preemption. The same argument can be made for three intervals of preemption, which have two intervals of separation. If the separation time is greater than the window execution time, the the delay would be issued in one of the separation intervals. If the delay is issued during the second separation interval, then there would only be two intervals of preemption (not three). If the delay is issued during the first separation interval, then there would only be one interval of preemption (not three).

Now consider two subcases. In the first subcase the total amount of preemption due to multiple intervals is shorter than the interval of preemption at the beginning of the task's period. In this case there is no wasted time for the same reason as in Case-1. In the second subcase we show that the delay caused by the multiple intervals of preemption is the same as multiple intervals of interruption at the beginning of the task's period during which there is very little opportunity to execute.

In the second subcase, the multiple intervals of preemption are greater than the initial interval of preemption at the beginning of the task's period. Now let's examine this multi-interval preemption. First, observe that this multi-interval preemption pattern repeats. Since we have a totally harmonic task set with constant execution times, all preemption patterns repeat including this particular one. Also, the first occurrence of this pattern must be relatively close to the beginning of the period. For example, if the first interval in the multi-interval preemption is due to task τ_k with higher priority and frequency than τ_i, then the first occurrence of this multi-interval preemption pattern starts with the first execution of τ_k. The first execution of τ_k is part of the first interval of preemption that starts at the beginning of τ_i's period. Therefore this multi-period preemption pattern is at least partially overlapping the first interval of preemption that starts at the beginning of τ_i's period. As such, the multiple interval of preemption extends beyond the initial interval of preemption at the beginning of the task's period, albeit with some extremely short gaps. Now if we look for potential wasted time we find that the wasted time can be at most the separation time in the multi-window preemption pattern. But we know that this is less than the window execution time, so the amount of wasted time is only a few instructions and unlikely to cause any missed deadlines.

It should be noted that this argument is heavily dependent on the repetitive nature of the harmonic task set and will not suffice for a non-harmonic task set. It should also be noted that a "delay until" construct (in place of a simple delay) will essentially combine the clock reading and the scheduling operation and will cure the flaw. This is an Ada 9X issue which will be addressed in the next version of the Ada standard.

Appendix B: Supporting Benchmark Results

B.1. Resolution of Calendar.Clock

The results presented here were obtained by running the PIWG A000090 test. In all cases the default Calendar.Clock resolution is shown. It is possible to change the default value for some cross-compilers, but, in general, this requires a user to modify an assembly language module that controls the operation of a programmable timer, and re-integrate it into the runtime system. As can be seen from the table, and Figure B-1, there is substantial variation in the clock resolutions of the compilers tested, ranging from 31 microseconds to a tenth of a second.

Compiler	Target Processor	Calendar.Clock Resolution
DDC-I DACS-80386PM 4.3(3.1a)	Intel 80386	1 millisecond
DEC VAX Ada 2.1	MicroVAX 3200	10 milliseconds
Rational Environment	Rational R1000	31 microseconds
Systems Designers XD Ada 1.0	Motorola 68020	122 microseconds
Tartan Ada VMS/1750A 2.14	Fairchild 9450	61 microseconds
TeleSoft TeleGen2 3.22a	Motorola 68020	100 milliseconds
Verdix VADS 5.7	Motorola 68020	10 milliseconds

B.2. Resolution of the Delay Statement

The results illustrated in this section were obtained from a benchmark program that was derived from the PIWG Y000001 test. In the modified version, the program requests discrete integral delays of 0, 1, 2, through 40 milliseconds (enough to fit on the graphs shown). The graphs of the test results, shown in Figures B-2 through B-8 plot the actual delay versus the delay requested by the test. The forty-five degree line represents the ideal case: the duration of the delay is exactly equal to that requested. The delay values plotted are accurate to the nearest millisecond. This graphical resolution is sufficient to illustrate the diversity of implementations of the delay statement. In the "stair-step" graphs, filled and unfilled circles represent the points of discontinuity, with the filled circle being the actual delay experienced for that particular requested value.

In all cases, the default delay statement resolution was measured. It is possible to change the default value for some cross-compilers, but, in general, this requires a user to modify an assembly language module that controls the operation of a programmable timer, and re-integrate it into the runtime system. For a discussion of the patterns exhibited in the various graphs, see the "Delay and Scheduling Measurements" section of [Clapp 86].
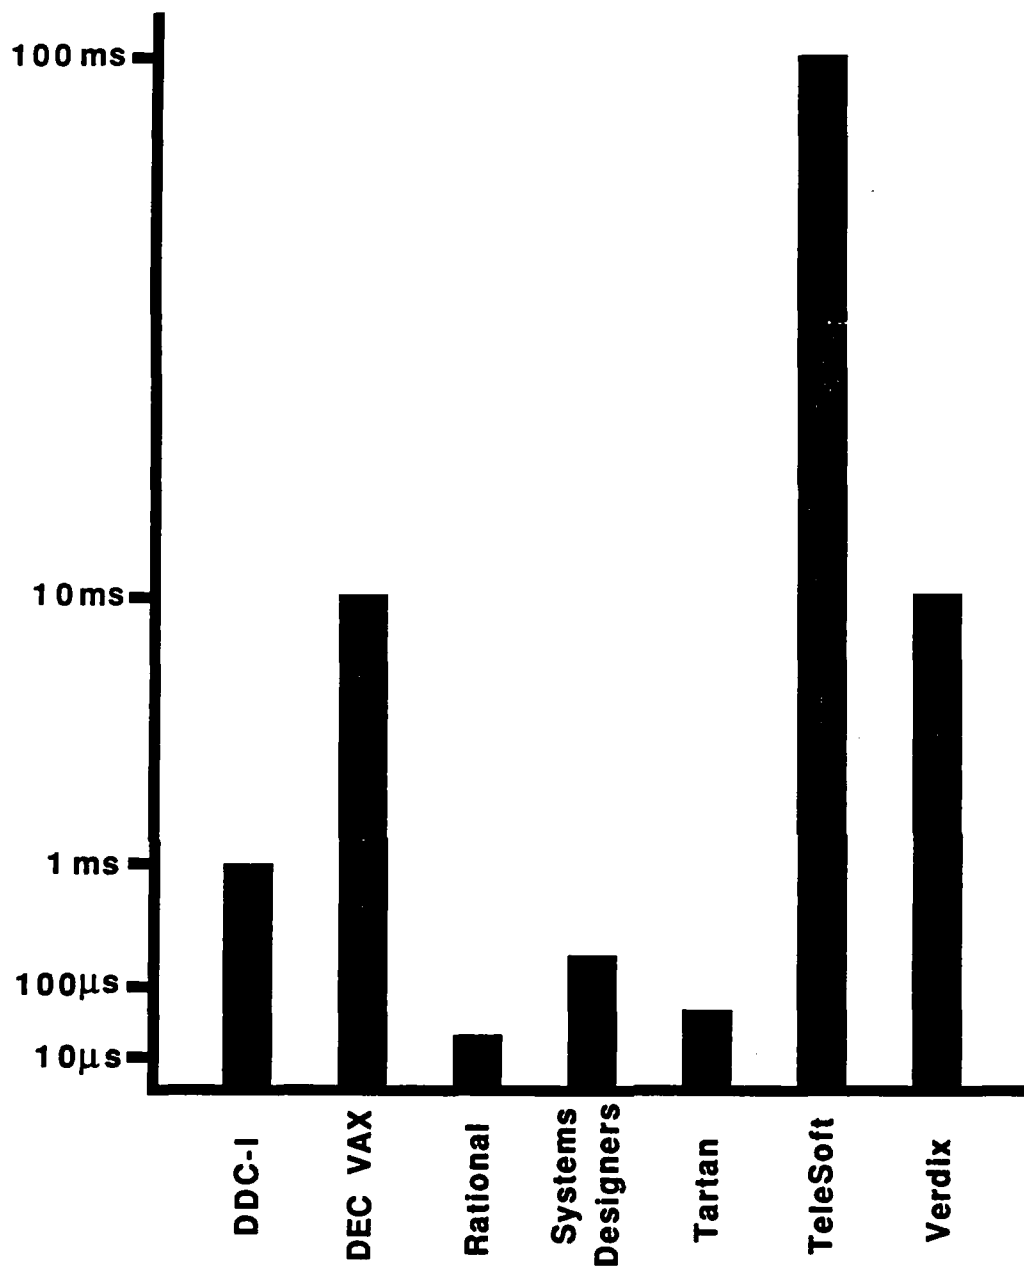
Figure B-1: Default Calendar.Clock Resolution

Host: DEC MicroVAX II, running VAX/VMS 5.3-1
Target: Intel iSBC 386/116-M02   (16 MHz i80386)
Compiler: DDC-I DACS-80386PM 4.3(3.1a)
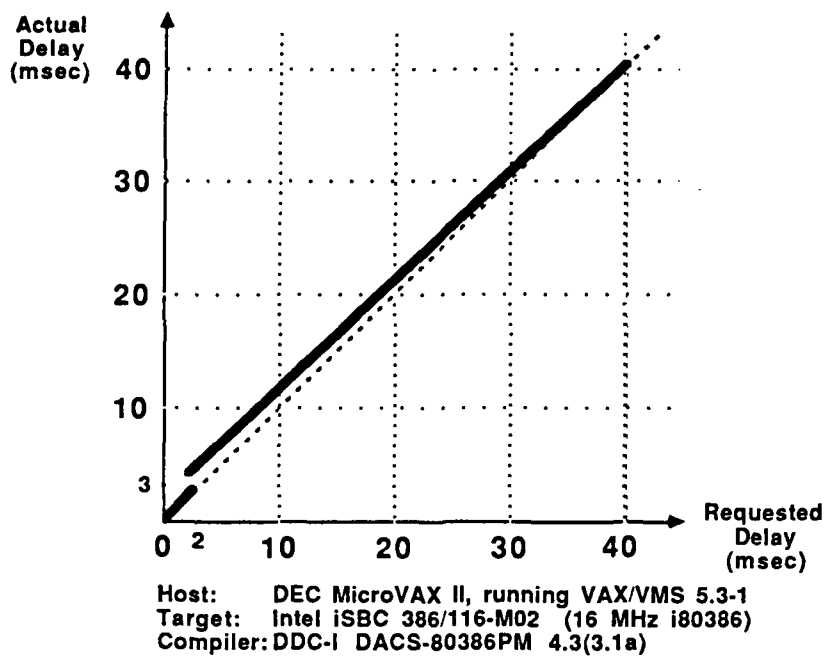
Figure B-2: Default Delay Statement Resolution - DDC-I Ada

Host: DEC MicroVAX 3200, running VAX/VMS 5.3-1
Target: Same as host
Compiler: DEC VAX Ada 2.1
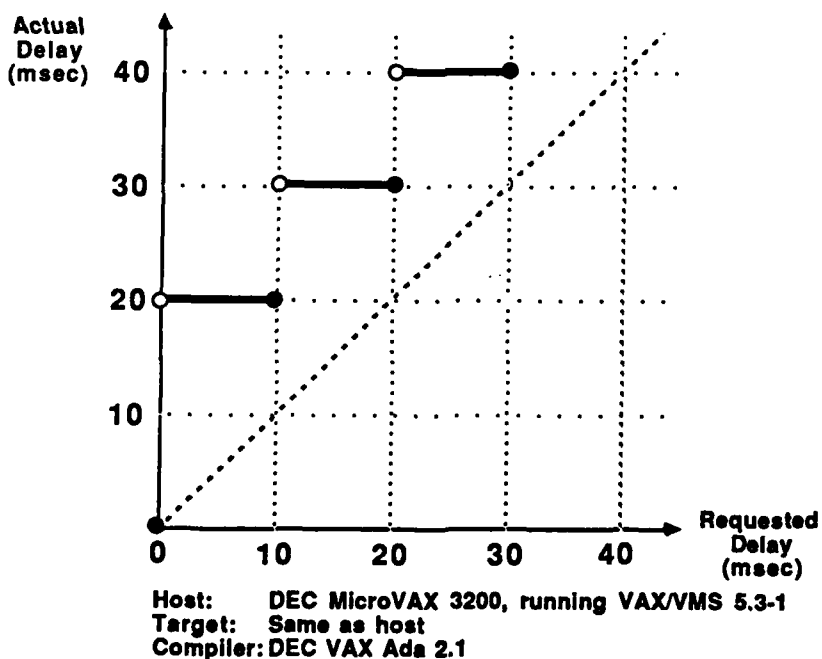
Figure B-3: Default Delay Statement Resolution - DEC VAX Ada

Host: Rational R1000, Series 200, Model 20
Target: Same as host
Compiler: Rational Environment, Delta 1 Release
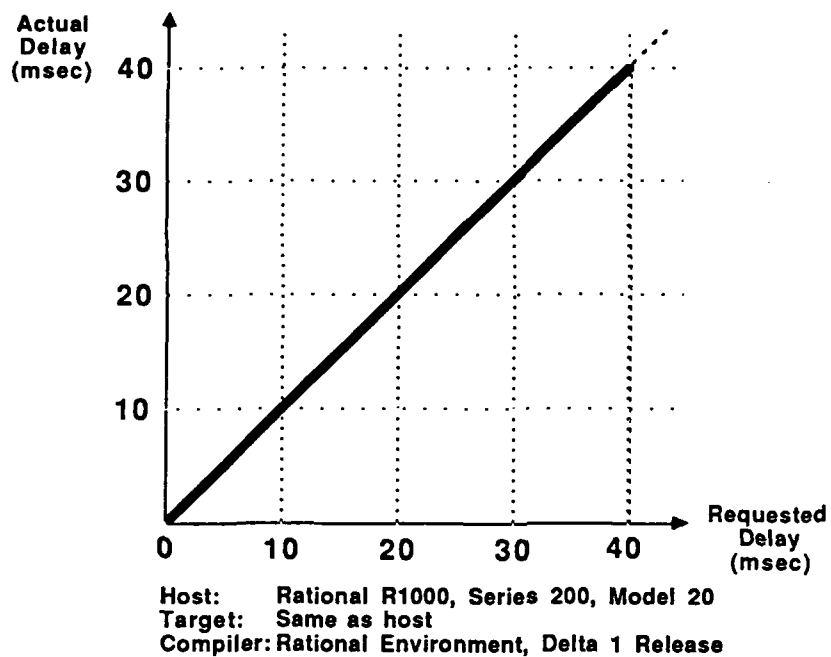
Figure B-4: Default Delay Statement Resolution - Rational Environment

Host: DEC MicroVAX 3200, running VAX/VMS 5.3-1
Target: Motorola MVME133A (20 MHz MC68020)
Compiler: Systems Designers XD Ada 1.0
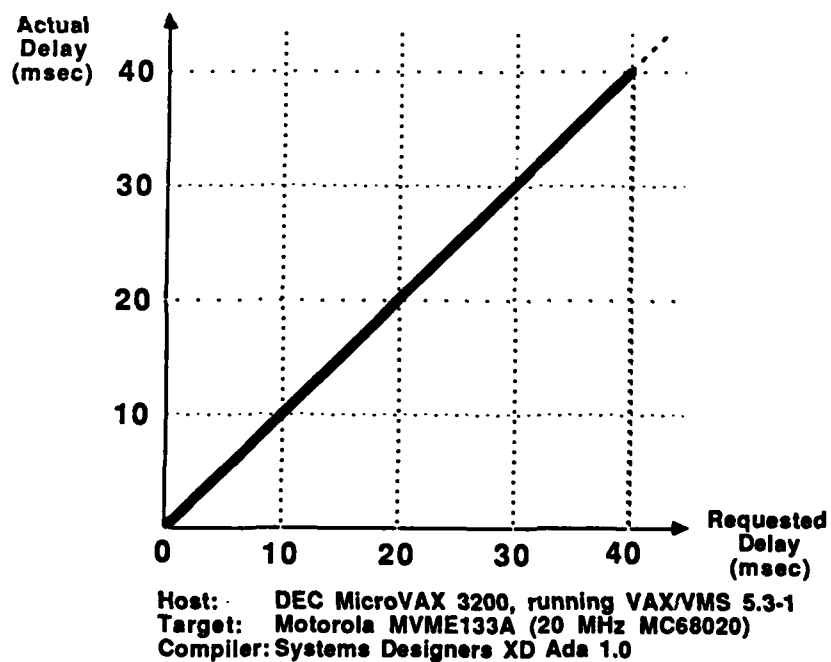
Figure B-5: Default Delay Statement Resolution - XD Ada

Host: DEC MicroVAX 3200, running VAX/VMS 5.3-1
Target: Fairchild SBC-50   (15 MHz F9450)
Compiler: Tartan Ada VMS/1750A 2.14
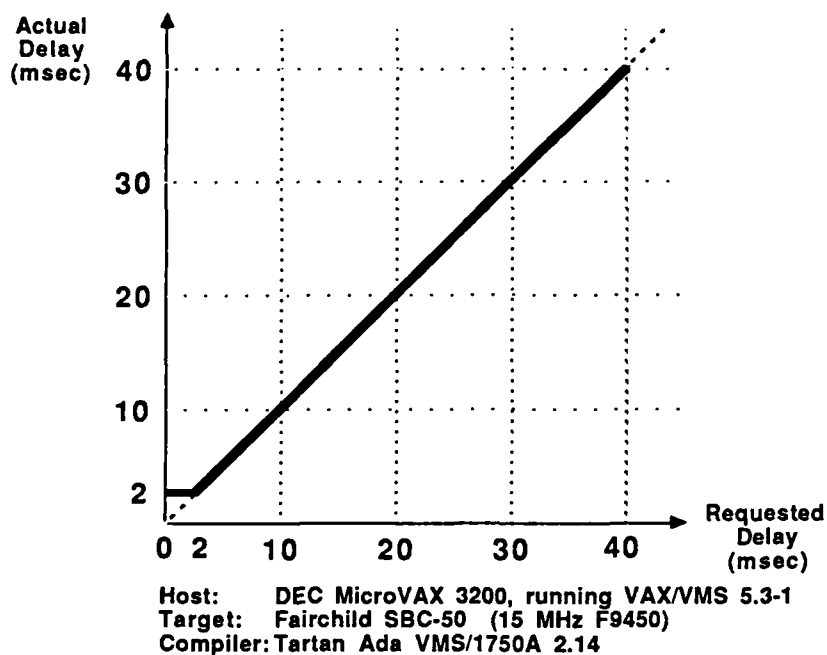
Figure B-6:   Default Delay Statement Resolution - Tartan Ada

Host: DEC MicroVAX 3200, running VAX/VMS 5.3-1
Target: Motorola MVME133A   (20 MHz MC68020)
Compiler: TeleSoft TeleGen2 3.22a
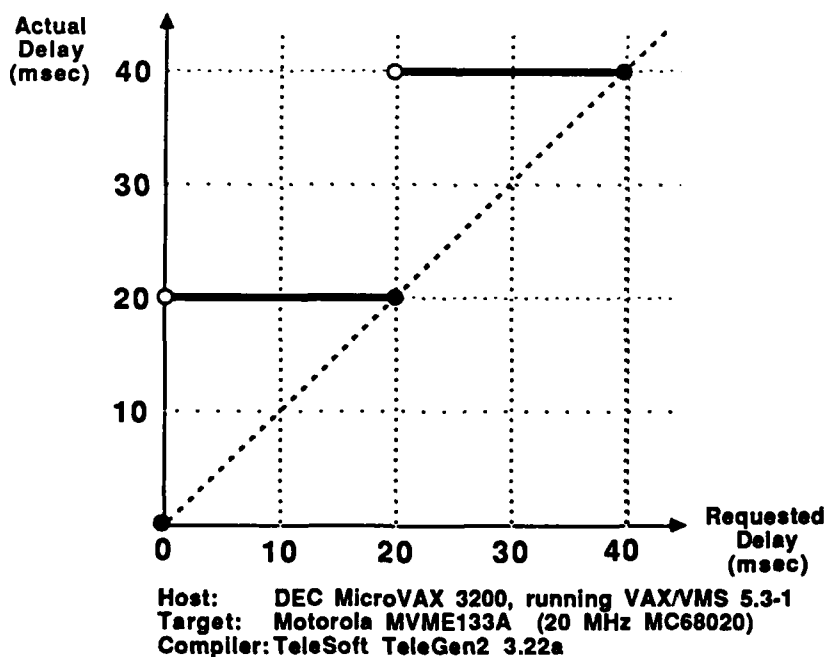
Figure B-7:   Default Delay Statement Resolution - TeleGen2 Ada

Actual
Delay
(msec) 40 ·······································○──────●·······
                                                  ·
        30 ·····················○──────●·······
                               ·
        20 ·········○──────●·······
                   ·
        10 ○──────●·······
        ●
         0   10   20   30   40   Requested
                                      Delay
                                     (msec)

Host: DEC MicroVAX 3200, running VAX/VMS 5.3-1
Target: Motorola MVME133A (20 MHz MC68020)
Compiler: Verdix VADS 5.7
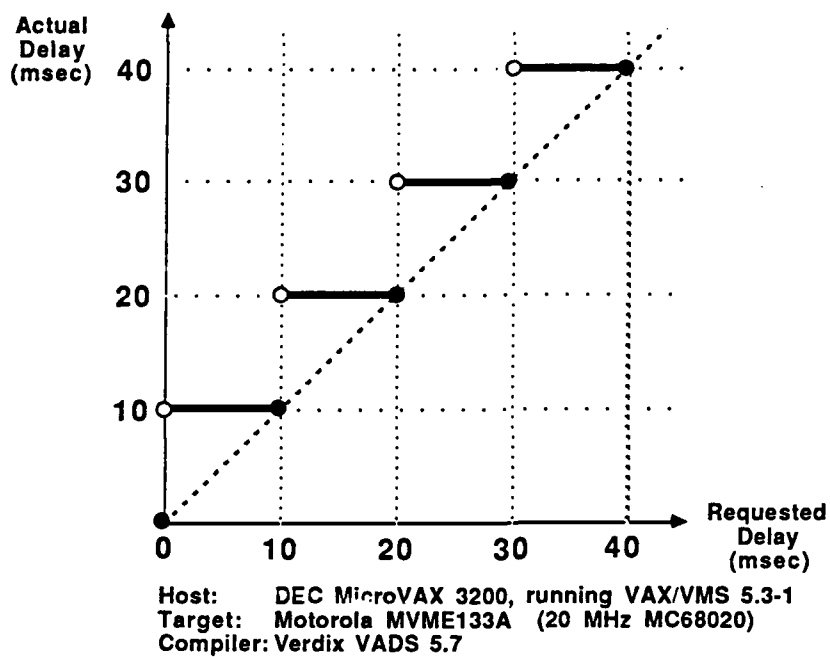
Figure B-8: Default Delay Statement Resolution - Verdix Ada

B.3. Whetstone Benchmark Results

These results are included to provide an independent measure of processing speed which is comparable to the measurement provided by the Small_Whetstone benchmark used in Hartstone.[3] They are from PIWG test A000092, for compilation systems that provide a mathematical library, and from PIWG test A000093 for those that do not. For compilers that do provide a mathematical library, this benchmark (and a listing of the generated assembly code) provides some indication of the efficiency of the implementation of the library functions. The tests were compiled with full optimization in effect and no runtime checks suppressed.

Compiler	Target Processor	Kilo-Whetstones Per Second
DDC-I DACS-80386PM 4.3(3.1a)	19 MHz Intel 80386	839
DEC VAX Ada 2.1	MicroVAX 3200	3890
Rational Environment	Rational R1000	764
Systems Designers XD Ada 1.0	20 MHz Motorola 68020	2181
Tartan Ada VMS/1750A 2.14	15 MHz Fairchild 9450	434
TeleSoft TeleGen2 3.22a	20 MHz Motorola 68020	794
Verdix VADS 5.7	20 MHz Motorola 68020	395

Table B-1: Whetsone Benchmark Results

B.4. Calendar.Clock Calling Overhead

The benchmark to produce these results comes from the University of Michigan benchmark suite.

Compiler	Target Processor	Calling Overhead (microseconds)
DDC-I DACS-80386PM 4.3(3.1a)	16 MHz Intel 80386	19
DEC VAX Ada 2.1	MicroVAX 3200	80
Rational Environment	Rational 1000	74
Systems Designers XD Ada 1.0	20 MHz Motorola 68020	19
Tartan Ada VMS/1750A 2.14	15 MHz Fairchild 9450	128
TeleSoft TeleGen2 3.22a	20 MHz Motorola 68020	121
Verdix VADS 5.7	20 MHz Motorola 68020	114

Table B-2: Calendar.Clock Calling Overhead

---

[3]Work is in progress to make the Ada language version of Small_Whetstone an official, controlled, benchmark similar to Whetstone.

B.5. Procedure Calling Overhead

The overhead of making the call to the Small_Whetstone is included in the Hartstone raw speed measurement. The following table gives some indication of what that overhead is for the various compilers tested in this study. Results were obtained by the running the PIWG P000005 test, which measures the call and return time of a procedure that is in a separately-compiled package and which is passed a single integer value of mode In. The tests were compiled with full optimization in effect and no runtime checks suppressed.

Compiler	Target Processor	P000005 (microseconds)
DDC-I DACS-80386PM 4.3(3.1a)	16 MHz Intel 80386	14
DEC VAX Ada 2.1	MicroVAX 3200	10
Rational Environment	Rational 1000	4
Systems Designers XD Ada 1.0	20 MHz Motorola 68020	1
Tartan Ada VMS/1750A 2.14	15 MHz Fairchild 9450	18
TeleSoft TeleGen2 3.22a	20 MHz Motorola 68020	4
Verdix VADS 5.7	20 MHz Motorola 68020	5

Table B-3: Procedure Calling Overhead

Hartstone Results

The following appendices give the results obtained from running the Hartstone benchmark against each of seven compilers. For each appendix, the first page gives the configuration tested and describes any changes that had to be made to the benchmark to make it run. The first page also describes the nature of any error conditions or anomalous conditions that occurred during the runs. The second page of each appendix is a graphical summary of the results of the four experiments.

The first bar of the graph for each experiment represents the workload processed at the beginning of the experiment (the baseline task set), the workload processed just before breakdown (before any deadlines are missed), and the raw speed (the workload processed without any tasking or deadlines). While these bars are drawn in terms of KWIPS, it is possible to get a rough feel for the utilization as the non-white proportion of the bar. If the white portion of the bar is a large proportion bar, the breakdown utilization was low. If there is no white portion of the bar, the breakdown utilization reached almost 100%. Because of scaling, the proportions are not linear.

The second bar for each experiment represents the number of deadlines processed per second by the task set just before breakdown. The third bar represents frequency of the highest frequency task which is also the number of deadlines processed per second by the highest frequency task. The ratio of bar 3 to bar 2 gives the proportion of the deadlines that can be attributed to the highest-frequency task. If bar 3 is close to bar 2 it means that Task 5 processed a large percentage of the deadlines. If bar 3 is small compared to bar 2, it means that the high frequency task played a lesser role in the results. For Experiment 4, there is a fourth bar representing the number of tasks in total before breakdown occurred.

The scales used on the graphs are unconventional. They are neither linear nor exponential. Linear scales would have prevented representing all the data on one graph because of the difference in magnitude of the numbers, while an exponential scale would have compressed the data too much. Instead we have chosen to represent the data using five separate linear scales. From 0 to 200 is one linear scale, from 200 to 400 is another linear scale with half the units per inch, from 400 to 800 is another linear scale with half the previous units per inch, and so forth up to 3200. This is probably not a statistically acceptable technique, but it does allow the presentation of a large amount of information on a single page in a way that the computer output does not.

From the third page and following are the computer outputs for each of the four experiments. For each experiment, there is output for the starting workload and result (test 1), the output for the test just before any deadlines are missed, the output for the test when deadlines are first missed, and output for the last test. There is also a summary page for each experiment giving the important metrics for the experiment.

Appendix C: Summary Results: DDC-I

The following is the host-target configuration used for generating the results reported here:

HOST: MicroVAX II running VAX/VMS, release 5.1-1

CROSS-COMPILER: DDC-I DACS-80386PM, release 4.3(3.1a), ACVC 1.10

TARGET: Intel iSBC 386/116-M02: 16 MHz i80386 with 16 MHz i80387 math co-processor; 2Mb RAM; 64Kb RAM cache. The board User's Guide says:

> The processor supports slower memory and I/O devices by inserting wait states. All on-board I/O operations (except for operations involving the math coprocessor or the MPC) require at least eight wait states for execution. On-board DRAM accesses require from two to three wait states, while EPROM accesses require eight wait states. DRAM accesses while in the pipelined mode require two wait states. Cache hits are zero wait state accesses.

Full optimization was specified for all compilations. No checks were suppressed. For the DDC-I compiler, the default resolution of Calendar.Clock is 1 millisecond (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-2. The mathematical library is named Math_Pack and the logarithmic function used by the Small_Whetstone procedure is named "Log". The renames clause was removed from workload body since it will not work with a Log function which requires multiple parameters. **Delay** statements had to be inserted into package Experiment to slow down the summary output at the end of an experiment, particularly in Experiment 4, where the r_sults of many tasks have to be displayed. Without the **delay** statements, the output became scrambled and meaningless. (Increasing the size of the ALTYPEAHD buffer on the VMS host did not solve the problem.) An attempt was made to change the Calendar.Clock resolution from its default value of 1 millisecond to 100 microseconds (this can be done by modifying a qualifier of the link command), but the change caused the PIWG and SEI **delay** statement resolution tests, and the PIWG clock resolution test, to crash with a protection exception. A test of Hartstone, with the 100-microsecond timer value in effect, showed no improvement in the results of Experiment 1, so all Hartstone experiments were run with the default 1 millisecond resolution in effect. Linking and downloading programs took considerably longer with DDC-I than with other cross-compilers.

In all 4 experiments the results were as predicted. Task 1 began missing its deadlines before the other tasks. At no time did a test fail when a later test succeeded. The workloads were reduced to keep the baseline utilization within the guidelines. In Experiment 3, the frequencies were scaled back to obtain an adequate step size.

Compiler: DDC-I DACS-80386 PM, R4.3(3.1a)
Target: Intel iSBC 386/116-M02
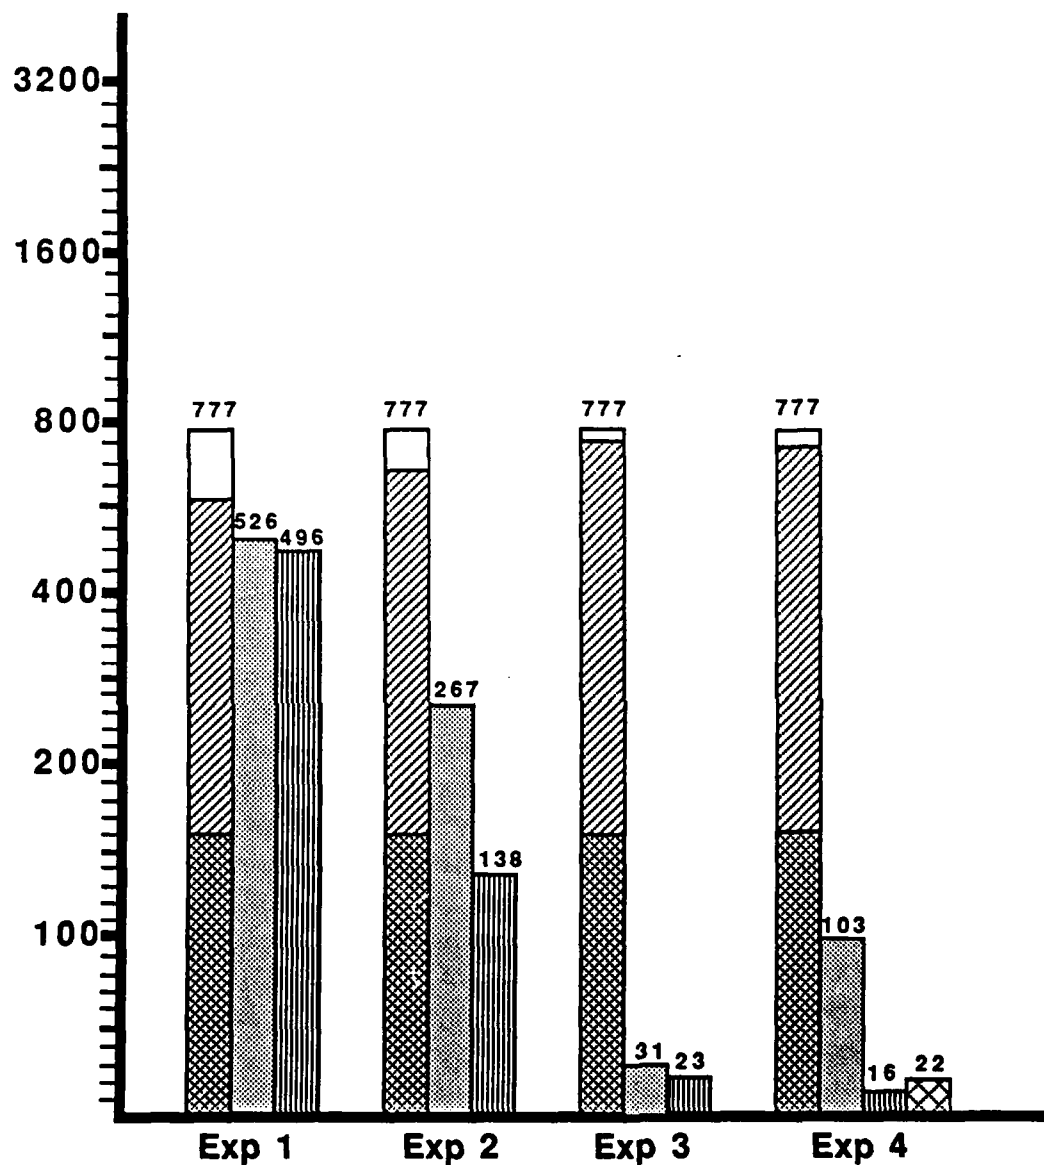
First Bar
- ☐ Raw throughput (KWIPS)
- ▨ Workload before breakdown (KWIPS)
- ▧ Baseline workload (KWIPS)

Second Bar ▦ Deadlines per second before breakdown

Third Bar ▥ Frequency of highest frequency task

Fourth Bar ⧄ Number of tasks before breakdown

C.1. Experiment 1: DDC-I

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   776.99

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	24	48.00	6.18 %
2	4.00	12	48.00	6.18 %
3	8.00	6	48.00	6.18 %
4	16.00	3	48.00	6.18 %
5	32.00	1	32.00	4.12 %
			-------	--------
			224.00	28.83 %

Experiment step size: 2.06 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

===

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   776.99

Test 26 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	24	48.00	6.18 %
2	4.00	12	48.00	6.18 %
3	8.00	6	48.00	6.18 %
4	16.00	3	48.00	6.18 %
5	432.00	1	432.00	55.60 %
			-------	--------
			624.00	80.31 %

Experiment step size: 2.06 %

Test 26 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	2.315	4318	0	0	0.000

==

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   776.99

Test 27 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	24	48.00	6.18 %
2	4.00	12	48.00	6.18 %
3	8.00	6	48.00	6.18 %
4	16.00	3	48.00	6.18 %
5	448.00	1	448.00	57.66 %
			-------	--------
			640.00	82.37 %

Experiment step size: 2.06 %

Test 27 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	3	9	8	151.259
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	2.232	4489	0	0	0.000

===

Final test performed:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 776.99

Test 30 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	24	48.00	6.18 %
2	4.00	12	48.00	6.18 %
3	8.00	6	48.00	6.18 %
4	16.00	3	48.00	6.18 %
5	496.00	1	496.00	63.84 %
			--------	--------
			688.00	88.55 %

Experiment step size: 2.06 %

Test 30 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	4	16	1895.752
2	250.000	0	20	20	113.623
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	2.016	4965	0	0	0.000

===

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : DDC-I DACS-80386PM 4.3(3.1a) VAX/VMS -> Intel i80386
Target    : Intel iSBC 386/116-M02 (16 MHz i80386 & 16 MHz i80387)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 26 of Experiment 1

    Raw (non-tasking) benchmark speed in KWIPS:  776.99

    Full task set:

        Total       Deadlines       Task Set        Total
        Tasks       Per Second      Utilization     KWIPS
          5           462.00          80.31 %        624.00

    Highest-frequency task:

        Period      Deadlines         Task          Task
        (msec)      Per Second      Utilization     KWIPS
         2.315        432.00          55.60 %        432.00

    Experiment step size:    2.06 %


================================================================


            END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

C.2. Experiment 2: DDC-I

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 777.23

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	4.12 %
2	4.00	8	32.00	4.12 %
3	8.00	4	32.00	4.12 %
4	16.00	2	32.00	4.12 %
5	32.00	1	32.00	4.12 %
			-------	--------
			160.00	20.59 %

Experiment step size: 2.06 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

===

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.23

Test 34 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	8.60	16	137.60	17.70 %
2	17.20	8	137.60	17.70 %
3	34.40	4	137.60	17.70 %
4	68.80	2	137.60	17.70 %
5	137.60	1	137.60	17.70 %
			--------	--------
			688.00	88.52 %

Experiment step size: 2.06 %

Test 34 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	116.279	87	0	0	0.000
2	58.140	173	0	0	0.000
3	29.070	345	0	0	0.000
4	14.535	689	0	0	0.000
5	7.267	1377	0	0	0.000

==

Test when deadlines first missed/skipped:

═══

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.23

Test 35 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	8.80	16	140.80	18.12 %
2	17.60	8	140.80	18.12 %
3	35.20	4	140.80	18.12 %
4	70.40	2	140.80	18.12 %
5	140.80	1	140.80	18.12 %
			-------	--------
			704.00	90.58 %

Experiment step size: 2.06 %

Test 35 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	113.636	1	44	44	39.035
2	56.818	177	0	0	0.000
3	28.409	353	0	0	0.000
4	14.205	704	0	0	0.000
5	7.102	1413	0	0	0.000

═══

Final test performed:
    See preceding summary of test 35

==

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : DDC-I DACS-80386PM 4.3(3.1a) VAX/VMS -> Intel i80386
Target   : Intel iSBC 386/116-M02 (16 MHz i80386 & 16 MHz i80387)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 34 of Experiment 2

    Raw (non-tasking) benchmark speed in KWIPS: 777.23

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   266.60   88.52 %   688.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
         7.267   137.60   17.70 %   137.60

    Experiment step size: 2.06 %

==

                END OF HARTSTONE BENCHMARK SUMMARY RESULTS

C.3. Experiment 3: DDC-I

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 777.17

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	4.12 %
2	2.00	16	32.00	4.12 %
3	4.00	8	32.00	4.12 %
4	8.00	4	32.00	4.12 %
5	16.00	2	32.00	4.12 %
			-------	--------
			160.00	20.59 %

Experiment step size: 3.99 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.17

Test 20 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	51	51.00	6.56 %
2	2.00	35	70.00	9.01 %
3	4.00	27	108.00	13.90 %
4	8.00	23	184.00	23.68 %
5	16.00	21	336.00	43.23 %
			-------	--------
			749.00	96.38 %

Experiment step size: 3.99 %

Test 20 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Test when deadlines first missed/skipped:

══

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.17

Test 21 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	52	52.00	6.69 %
2	2.00	36	72.00	9.26 %
3	4.00	28	112.00	14.41 %
4	8.00	24	192.00	24.70 %
5	16.00	22	352.00	45.29 %
			-------	--------
			780.00	100.36 %

Experiment step size: 3.99 %

Test 21 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	5	5	987.500
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

══

Final test performed:
    See preceding summary of test 21

==

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : DDC-I DACS-80386PM 4.3(3.1a) VAX/VMS -> Intel i80386
Target   : Intel iSBC 386/116-M02 (16 MHz i80386 & 16 MHz i80387)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 20 of Experiment 3

    Raw (non-tasking) benchmark speed in KWIPS: 777.17

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   31.00   96.38 %   749.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
        62.500   16.00   43.23 %   336.00

    Experiment step size: 3.99 %

==

                END OF HARTSTONE BENCHMARK SUMMARY RESULTS

C.4. Experiment 4: DDC-I

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.05

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	4.12 %
2	2.00	16	32.00	4.12 %
3	4.00	8	32.00	4.12 %
4	8.00	4	32.00	4.12 %
5	16.00	2	32.00	4.12 %
			160.00	20.59 %

Experiment step size: 4.12 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.05

Test 19 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	4.12 %
2	2.00	16	32.00	4.12 %
3	4.00	8	32.00	4.12 %
4	8.00	4	32.00	4.12 %
5	16.00	2	32.00	4.12 %
6	4.00	8	32.00	4.12 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
23	4.00	8	32.00	4.12 %
			-------	--------
			736.00	94.72 %

Experiment step size: 4.12 %

Test 19 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
23	250.000	40	0	0	0.000

===

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   777.05

Test 20 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	4.12 %
2	2.00	16	32.00	4.12 %
3	4.00	8	32.00	4.12 %
4	8.00	4	32.00	4.12 %
5	16.00	2	32.00	4.12 %
6	4.00	8	32.00	4.12 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
24	4.00	8	32.00	4.12 %
			768.00	98.84 %

Experiment step size: 4.12 %

Test 20 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9014.648
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
24	250.000	40	0	0	0.000

===

Final test performed:

```
=========================================================================
```

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 777.05

Test 21 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	4.12 %
2	2.00	16	32.00	4.12 %
3	4.00	8	32.00	4.12 %
4	8.00	4	32.00	4.12 %
5	16.00	2	32.00	4.12 %
6	4.00	8	32.00	4.12 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
25	4.00	8	32.00	4.12 %
			-------	--------
			800.00	102.95 %

Experiment step size: 4.12 %

```
-------------------------------------------------------------------------
```

Test 21 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9043.945
2	500.000	0	1	19	9500.977
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
25	250.000	40	0	0	0.000

```
=========================================================================
```

```
===============================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : DDC-I DACS-80386PM 4.3(3.1a) VAX/VMS -> Intel i80386
Target    : Intel iSBC 386/116-M02 (16 MHz i80386 & 16 MHz i80387)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 19 of Experiment 4

   Raw (non-tasking) benchmark speed in KWIPS:   777.05

   Full task set:

      Total        Deadlines       Task Set        Total
      Tasks        Per Second      Utilization     KWIPS
       23          103.00           94.72 %        736.00

   Highest-frequency task:

      Period       Deadlines        Task           Task
      (msec)       Per Second      Utilization     KWIPS
      62.500        16.00            4.12 %         32.00

   Experiment step size:    4.12 %

===============================================================================


                END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

Appendix D: Summary Results: DEC - VAX/VMS

The following is the host-target configuration used for generating the results reported here:

HOST:	DEC MicroVAX 3200 running VAX/VMS, release 5.3-1
COMPILER:	DEC VAX Ada, Version 2.1, ACVC 1.10
TARGET:	Same as host

Full optimization was specified for all compilations. No checks were suppressed. For the DEC VAX Ada compiler, the resolution of Calendar.Clock is 10 milliseconds (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-3. The mathematical library is named Float_Math_Lib and the logarithmic function used by the Small_Whetstone procedure is named "Log". The Hartstone process was run with VMS priority 31 (the highest real-time priority) on a quiet system. There was no VMS interference that appeared to cause premature missing of deadlines.

Initial plans called for rerunning the tests using DEC VAXELN 4.0 and DEC VAXELN Ada 2.0. This was not done due to time constraints, but we expect the results would have been similar to the ones presented here because the clock and delay precision are the same for VAXELN Ada as they are for VAX Ada.

In all 4 experiments the results were as predicted. Task 1 began missing its deadlines before the other tasks. At no time did a test fail when a later test succeeded. The frequencies were scaled back due to the imprecise clock and delay statement. The workloads were increased to keep the baseline utilization within the guidelines. In Experiment 1, the workloads were increased substantially to prevent premature completion of the experiment with Task 5 missing its deadlines first. This resulted in a baseline utilization in excess of the guidelines.

Compiler: DEC VAX Ada, V2.1
Target: DEC MicroVAX 3100
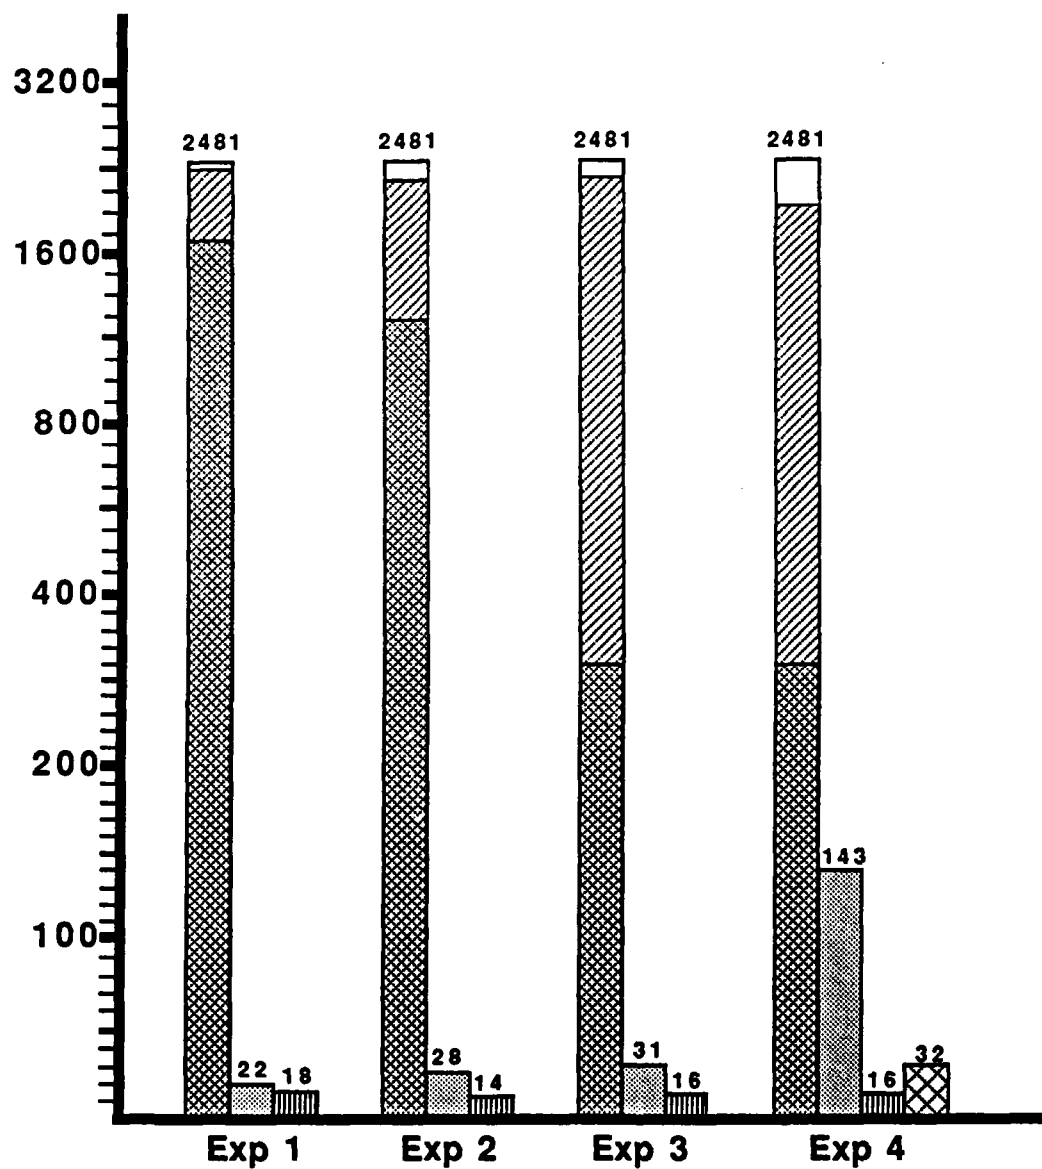
First Bar
- ☐ Raw throughput (KWIPS)
- ▨ Workload before breakdown (KWIPS)
- ▩ Baseline workload (KWIPS)

Second Bar ▦ Deadlines per second before breakdown

Third Bar ▥ Frequency of highest frequency task

Fourth Bar ⊠ Number of tasks before breakdown

D.1. Experiment 1: DEC - VAX/VMS

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.25	1536	384.00	15.48 %
2	0.50	768	384.00	15.48 %
3	1.00	384	384.00	15.48 %
4	2.00	192	384.00	15.48 %
5	4.00	48	192.00	7.74 %
			-------	--------
			1728.00	69.64 %

Experiment step size: 3.87 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	4000.000	3	0	0	0.000
2	2000.000	5	0	0	0.000
3	1000.000	10	0	0	0.000
4	500.000	20	0	0	0.000
5	250.000	40	0	0	0.000

==

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 8 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.25	1536	384.00	15.48 %
2	0.50	768	384.00	15.48 %
3	1.00	384	384.00	15.48 %
4	2.00	192	384.00	15.48 %
5	18.00	48	864.00	34.82 %
			-------	--------
			2400.00	96.72 %

Experiment step size: 3.87 %

Test 8 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	4000.000	3	0	0	0.000
2	2000.000	5	0	0	0.000
3	1000.000	10	0	0	0.000
4	500.000	20	0	0	0.000
5	55.556	183	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 9 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.25	1536	384.00	15.48 %
2	0.50	768	384.00	15.48 %
3	1.00	384	384.00	15.48 %
4	2.00	192	384.00	15.48 %
5	20.00	48	960.00	38.69 %
			-------	--------
			2496.00	100.59 %

Experiment step size: 3.87 %

Test 9 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	4000.000	1	1	1	3742.188
2	2000.000	5	0	0	0.000
3	1000.000	10	0	0	0.000
4	500.000	20	0	0	0.000
5	50.000	214	0	0	0.000

==

Final test performed:
    See preceding summary of test 9

===

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : DEC VAX Ada, release 2.1
Target   : DEC MicroVAX 3200, running VMS 5.3-1

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 8 of Experiment 1

    Raw (non-tasking) benchmark speed in KWIPS: 2481.41

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   21.75   96.72 %   2400.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
        55.556   18.00   34.82 %   864.00

    Experiment step size: 3.87 %

===

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

D.2. Experiment 2: DEC - VAX/VMS

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

```
=======================================================================

Experiment:    EXPERIMENT_2
Completion on: Miss/skip  25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2487.59

Test 1 characteristics:

    Task   Frequency   Kilo-Whets   Kilo-Whets   Requested Workload
    No.     (Hertz)    per period   per second        Utilization
     1       0.50         512         256.00           10.29 %
     2       1.00         256         256.00           10.29 %
     3       2.00         128         256.00           10.29 %
     4       4.00          64         256.00           10.29 %
     5       8.00          32         256.00           10.29 %
                                      -------          --------
                                      1280.00           51.46 %

Experiment step size:   5.15 %


-----------------------------------------------------------------------

Test 1 results:

Test duration (seconds):  10.0

    Task     Period       Met        Missed       Skipped       Average
    No.     in msecs   Deadlines   Deadlines    Deadlines     Late (msec)
     1     2000.000        5           0            0            0.000
     2     1000.000       10           0            0            0.000
     3      500.000       20           0            0            0.000
     4      250.000       40           0            0            0.000
     5      125.000       80           0            0            0.000


=======================================================================
```

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2487.59

Test 9 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.90	512	460.80	18.52 %
2	1.80	256	460.80	18.52 %
3	3.60	128	460.80	18.52 %
4	7.20	64	460.80	18.52 %
5	14.40	32	460.80	18.52 %
			2304.00	92.62 %

Experiment step size: 5.15 %

Test 9 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1111.111	9	0	0	0.000
2	555.556	19	0	0	0.000
3	277.778	37	0	0	0.000
4	138.889	73	0	0	0.000
5	69.444	143	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2487.59

Test 10 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.95	512	486.40	19.55 %
2	1.90	256	486.40	19.55 %
3	3.80	128	486.40	19.55 %
4	7.60	64	486.40	19.55 %
5	15.20	32	486.40	19.55 %
			-------	--------
			2432.00	97.77 %

Experiment step size: 5.15 %

Test 10 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1052.631	2	4	4	340.820
2	526.316	20	0	0	0.000
3	263.158	38	0	0	0.000
4	131.579	77	0	0	0.000
5	65.789	151	0	0	0.000

==

Final test performed:

══

Experiment: EXPERIMENT_2
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2487.59

Test 11 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	512	512.00	20.58 %
2	2.00	256	512.00	20.58 %
3	4.00	128	512.00	20.58 %
4	8.00	64	512.00	20.58 %
5	16.00	32	512.00	20.58 %
			-------	--------
			2560.00	102.91 %

Experiment step size: 5.15 %

Test 11 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	5	5	536.328
2	500.000	18	1	1	121.094
3	250.000	40	0	0	0.000
4	125.000	78	1	1	50.781
5	62.500	157	1	2	83.984

══

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : DEC VAX Ada, release 2.1
Target    : DEC MicroVAX 3200, running VMS 5.3-1

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 9 of Experiment 2

    Raw (non-tasking) benchmark speed in KWIPS: 2487.59

    Full task set:

        Total       Deadlines      Task Set      Total
        Tasks       Per Second     Utilization   KWIPS
          5           27.90          92.62 %      2304.00

    Highest-frequency task:

        Period      Deadlines       Task         Task
        (msec)      Per Second      Utilization   KWIPS
        69.444        14.40          18.52 %      460.80

    Experiment step size:    5.15 %

================================================================


          END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

D.3. Experiment 3: DEC - VAX/VMS

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	2.58 %
2	2.00	32	64.00	2.58 %
3	4.00	16	64.00	2.58 %
4	8.00	8	64.00	2.58 %
5	16.00	4	64.00	2.58 %
			--------	--------
			320.00	12.90 %

Experiment step size: 1.25 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 66 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	129	129.00	5.20 %
2	2.00	97	194.00	7.82 %
3	4.00	81	324.00	13.06 %
4	8.00	73	584.00	23.54 %
5	16.00	69	1104.00	44.49 %
			-------	--------
			2335.00	94.10 %

Experiment step size: 1.25 %

Test 66 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 67 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	130	130.00	5.24 %
2	2.00	98	196.00	7.90 %
3	4.00	82	328.00	13.22 %
4	8.00	74	592.00	23.86 %
5	16.00	70	1120.00	45.14 %
			--------	--------
			2366.00	95.35 %

Experiment step size: 1.25 %

Test 67 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	1	5	4	393.750
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Final test performed:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 71 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	134	134.00	5.40 %
2	2.00	102	204.00	8.22 %
3	4.00	86	344.00	13.86 %
4	8.00	78	624.00	25.15 %
5	16.00	74	1184.00	47.71 %
			-------	--------
			2490.00	100.35 %

Experiment step size: 1.25 %

Test 71 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	3	7	2341.146
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

```
===============================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : DEC VAX Ada, release 2.1
Target    : DEC MicroVAX 3200, running VMS 5.3-1

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 66 of Experiment 3

   Raw (non-tasking) benchmark speed in KWIPS: 2481.41

   Full task set:

      Total       Deadlines       Task Set        Total
      Tasks       Per Second      Utilization     KWIPS
        5           31.00           94.10 %        2335.00

   Highest-frequency task:

      Period      Deadlines        Task           Task
      (msec)      Per Second       Utilization    KWIPS
      62.500        16.00           44.49 %        1104.00

   Experiment step size:   1.25 %

===============================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

D.4. Experiment 4: DEC - VAX/VMS

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	2.58 %
2	2.00	32	64.00	2.58 %
3	4.00	16	64.00	2.58 %
4	8.00	8	64.00	2.58 %
5	16.00	4	64.00	2.58 %
			-------	--------
			320.00	12.90 %

Experiment step size: 2.58 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 29 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	2.58 %
2	2.00	32	64.00	2.58 %
3	4.00	16	64.00	2.58 %
4	8.00	8	64.00	2.58 %
5	16.00	4	64.00	2.58 %
6	4.00	16	64.00	2.58 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
33	4.00	16	64.00	2.58 %
			2112.00	85.11 %

Experiment step size: 2.58 %

Test 29 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
33	250.000	40	0	0	0.000

==

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 30 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	2.58 %
2	2.00	32	64.00	2.58 %
3	4.00	16	64.00	2.58 %
4	8.00	8	64.00	2.58 %
5	16.00	4	64.00	2.58 %
6	4.00	16	64.00	2.58 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
34	4.00	16	64.00	2.58 %
			-------	-------
			2176.00	87.69 %

Experiment step size: 2.58 %

Test 30 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	4	6	812.500
2	500.000	10	5	5	242.188
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
34	250.000	40	0	0	0.000

===

Final test performed:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 25.0 percent of deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 2481.41

Test 35 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	2.58 %
2	2.00	32	64.00	2.58 %
3	4.00	16	64.00	2.58 %
4	8.00	8	64.00	2.58 %
5	16.00	4	64.00	2.58 %
6	4.00	16	64.00	2.58 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
39	4.00	16	64.00	2.58 %
			2496.00	100.59 %

Experiment step size: 2.58 %

Test 35 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9054.688
2	500.000	0	1	19	9523.438
3	250.000	30	5	5	42.188
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	36	2	2	35.156
7	250.000	32	4	4	39.063
8	250.000	32	4	4	25.391
9	250.000	24	8	8	29.297
10	250.000	28	6	6	23.438
11	250.000	24	8	8	26.367
12	250.000	34	3	3	31.250
13	250.000	22	9	9	26.910
14	250.000	30	5	5	35.938
15	250.000	30	5	5	42.188
16	250.000	30	5	5	37.500
17	250.000	26	7	7	35.714
18	250.000	30	5	5	28.125
19	250.000	36	2	2	46.875

20	250.000	28	6	6	35.156
21	250.000	34	3	3	46.875
22	250.000	25	8	8	19.531
23	250.000	30	5	5	29.688
24	250.000	30	5	5	35.938
25	250.000	28	6	6	40.365
26	250.000	28	6	6	42.969
27	250.000	36	2	2	42.969
28	250.000	36	2	2	39.063
29	250.000	29	6	6	32.552
30	250.000	30	5	5	31.250
31	250.000	30	5	5	20.313
32	250.000	34	3	3	39.063
33	250.000	38	1	1	7.813
34	250.000	34	3	3	31.250
35	250.000	34	3	3	46.875
36	250.000	40	0	0	0.000
37	250.000	38	1	1	23.438
38	250.000	32	4	4	25.391
39	250.000	28	6	6	23.438

===

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : DEC VAX Ada, release 2.1
Target    : DEC MicroVAX 3200, running VMS 5.3-1

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 29 of Experiment 4

   Raw (non-tasking) benchmark speed in KWIPS: 2481.41

   Full task set:

      Total        Deadlines        Task Set        Total
      Tasks        Per Second      Utilization      KWIPS
       33           143.00           85.11 %       2112.00

   Highest-frequency task:

      Period       Deadlines         Task           Task
      (msec)       Per Second      Utilization      KWIPS
      62.500         16.00            2.58 %        64.00

   Experiment step size:   2.58 %

================================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

Appendix E: Summary Results: Rational

The following is the host-target configuration used for generating the results reported here:

HOST: Rational R1000, Series 200, Model 20

COMPILER: Rational Environment, Delta 1 release

TARGET: Same as host

Full optimization was specified for all compilations. No checks were suppressed. For the Rational Environment, the resolution of Calendar.Clock is 31 microseconds (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-4. A mathematical library was not provided with this release of the Rational Environment so one was adapted from the mathematical functions provided in PIWG test A000093. This library was named PIWG_Math_Lib and the logarithmic function used by the Small_Whetstone procedure was named "Log".

In order to create a quiet system in which Hartstone would run, the following 4 jobs were removed from the Rational system:

- Archive Server
- Print Spooler
- FTP Server
- Print Queue Server

The final action taken before the benchmark tests were run was to delay the SNAPSHOT daemon by 24hrs. Once the tests were complete, the SNAPSHOT daemon was restored to normal operation.

In none of the 4 experiments were the results as predicted. In Experiments 1 and 3, a test failed when a later test succeeded. Although the frequencies were scaled back for Experiments 1 and 2, Task 5 missed its deadlines first. This was not expected given a precise clock and delay statement. In Experiment 4, the added tasks (which have a priority equal to Task 3) missed their deadlines before Task 2. Since this is a self-targetted compiler, the missed deadlines could be attributed to activities of the operating system that we were unable to shut off. This theory has not been confirmed.

Compiler: Rational, Delta 1
Target: Rational R1000, Series 200, Model 20
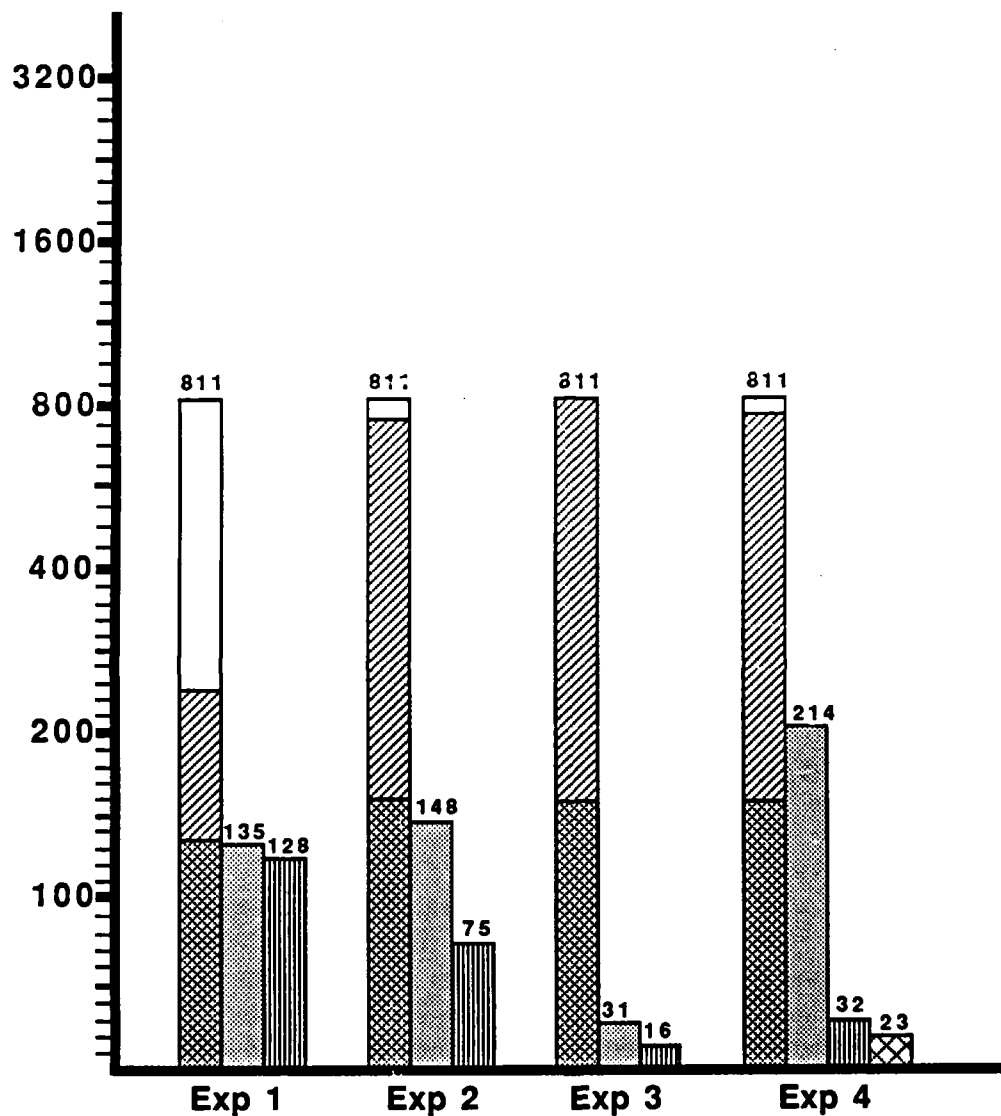
First Bar

	Raw throughput (KWIPS)
	Workload before breakdown (KWIPS)
	Baseline workload (KWIPS)

Second Bar   Deadlines per second before breakdown

Third Bar   Frequency of highest frequency task

Fourth Bar   Number of tasks before breakdown

CMU/SEI-90-TR-7

E.1. Experiment 1: Rational

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   810.54

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	64	32.00	3.95 %
2	1.00	32	32.00	3.95 %
3	2.00	16	32.00	3.95 %
4	4.00	8	32.00	3.95 %
5	8.00	1	8.00	0.99 %
			-------	--------
			136.00	16.78 %

Experiment step size: 0.49 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

==

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   810.54

Test 31 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	64	32.00	3.95 %
2	1.00	32	32.00	3.95 %
3	2.00	16	32.00	3.95 %
4	4.00	8	32.00	3.95 %
5	128.00	1	128.00	15.79 %
			-------	-------
			256.00	31.58 %

Experiment step size: 0.49 %

Test 31 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	7.812	1280	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   810.54

Test 16 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	64	32.00	3.95 %
2	1.00	32	32.00	3.95 %
3	2.00	16	32.00	3.95 %
4	4.00	8	32.00	3.95 %
5	68.00	1	68.00	8.39 %
			--------	--------
			196.00	24.18 %

Experiment step size: 0.49 %

Test 16 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	38	1	1	192.841
5	14.706	640	1	40	578.461

==

Final test performed:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 810.54

Test 125 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	64	32.00	3.95 %
2	1.00	32	32.00	3.95 %
3	2.00	16	32.00	3.95 %
4	4.00	8	32.00	3.95 %
5	504.00	1	504.00	62.18 %
			-------	--------
			632.00	77.97 %

Experiment step size: 0.49 %

Test 125 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	1.984	4888	72	81	0.387

===

```
===============================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Rational Environment, Delta 1 release
Target    : Rational R1000, Series 200, Model 20

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 31 of Experiment 1

   Raw (non-tasking) benchmark speed in KWIPS:  810.54

   Full task set:

      Total      Deadlines     Task Set      Total
      Tasks      Per Second    Utilization   KWIPS
        5          135.50        31.58 %      256.00

   Highest-frequency task:

      Period     Deadlines       Task        Task
      (msec)     Per Second    Utilization   KWIPS
       7.812      128.00         15.79 %      128.00

   Experiment step size:   0.49 %

===============================================================


          END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

E.2. Experiment 2: Rational

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.41

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	32	32.00	3.95 %
2	2.00	16	32.00	3.95 %
3	4.00	8	32.00	3.95 %
4	8.00	4	32.00	3.95 %
5	16.00	2	32.00	3.95 %
			-------	--------
			160.00	19.77 %

Experiment step size: 1.98 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

═══

Experiment: EXPERIMENT_2
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.41

Test 38 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	4.70	32	150.40	18.58 %
2	9.40	16	150.40	18.58 %
3	18.80	8	150.40	18.58 %
4	37.60	4	150.40	18.58 %
5	75.20	2	150.40	18.58 %
			-------	--------
			752.00	92.91 %

Experiment step size: 1.98 %

Test 38 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	212.766	48	0	0	0.000
2	106.383	95	0	0	0.000
3	53.191	189	0	0	0.000
4	26.596	377	0	0	0.000
5	13.298	753	0	0	0.000

═══

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_2
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 809.41

Test 39 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	4.80	32	153.60	18.98 %
2	9.60	16	153.60	18.98 %
3	19.20	8	153.60	18.98 %
4	38.40	4	153.60	18.98 %
5	76.80	2	153.60	18.98 %
			-------	--------
			768.00	94.88 %

Experiment step size: 1.98 %

Test 39 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	208.333	49	0	0	0.000
2	104.167	97	0	0	0.000
3	52.083	193	0	0	0.000
4	26.042	385	0	0	0.000
5	13.021	767	1	1	0.092

===

Final test performed:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.41

Test 42 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	5.10	32	163.20	20.16 %
2	10.20	16	163.20	20.16 %
3	20.40	8	163.20	20.16 %
4	40.80	4	163.20	20.16 %
5	81.60	2	163.20	20.16 %
			-------	--------
			816.00	100.81 %

Experiment step size: 1.98 %

Test 42 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	196.078	0	26	26	69.985
2	98.039	103	0	0	0.000
3	49.020	205	0	0	0.000
4	24.510	409	0	0	0.000
5	12.255	817	0	0	0.000

==

```
========================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Rational Environment, Delta 1 release
Target    : Rational R1000, Series 200, Model 20

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 38 of Experiment 2

   Raw (non-tasking) benchmark speed in KWIPS:  809.41

   Full task set:

       Total       Deadlines       Task Set       Total
       Tasks       Per Second      Utilization    KWIPS
         5          145.70           92.91 %       752.00

   Highest-frequency task:

       Period      Deadlines         Task         Task
       (msec)      Per Second      Utilization    KWIPS
       13.298        75.20           18.58 %       150.40

   Experiment step size:   1.98 %

========================================================================


              END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

E.3. Experiment 3: Rational

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.69

Test 1 characteristics:

```
    Task    Frequency   Kilo-Whets   Kilo-Whets   Requested Workload
    No.      (Hertz)    per period   per second       Utilization
     1        1.00          32          32.00           3.95 %
     2        2.00          16          32.00           3.95 %
     3        4.00           8          32.00           3.95 %
     4        8.00           4          32.00           3.95 %
     5       16.00           2          32.00           3.95 %
                                      --------        --------
                                       160.00          19.76 %
```

Experiment step size: 3.83 %

Test 1 results:

Test duration (seconds): 10.0

```
    Task     Period        Met        Missed      Skipped      Average
    No.     in msecs    Deadlines   Deadlines   Deadlines    Late (msec)
     1     1000.000        10          0           0          0.000
     2      500.000        20          0           0          0.000
     3      250.000        40          0           0          0.000
     4      125.000        80          0           0          0.000
     5       62.500       160          0           0          0.000
```

===

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.69

Test 22 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	53	53.00	6.55 %
2	2.00	37	74.00	9.14 %
3	4.00	29	116.00	14.33 %
4	8.00	25	200.00	24.70 %
5	16.00	23	368.00	45.45 %
			-------	--------
			811.00	100.16 %

Experiment step size: 3.83 %

Test 22 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.69

Test 17 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	48	48.00	5.93 %
2	2.00	32	64.00	7.90 %
3	4.00	24	96.00	11.86 %
4	8.00	20	160.00	19.76 %
5	16.00	18	288.00	35.57 %
			656.00	81.02 %

Experiment step size: 3.83 %

Test 17 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	78	1	1	22.003
5	62.500	158	1	1	39.093

==

Final test performed:
    See preceding summary of test 22

==

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : Rational Environment, Delta 1 release
Target   : Rational R1000, Series 200, Model 20

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 22 of Experiment 3

    Raw (non-tasking) benchmark speed in KWIPS: 809.69

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   31.00   100.16 %   811.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
        62.500   16.00   45.45 %   368.00

    Experiment step size: 3.83 %

==

                END OF HARTSTONE BENCHMARK SUMMARY RESULTS

E.4. Experiment 4: Rational

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 809.92

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	3.95 %
2	4.00	8	32.00	3.95 %
3	8.00	4	32.00	3.95 %
4	16.00	2	32.00	3.95 %
5	32.00	1	32.00	3.95 %
			-------	--------
			160.00	19.76 %

Experiment step size: 3.95 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

===

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.92

Test 20 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	3.95 %
2	4.00	8	32.00	3.95 %
3	8.00	4	32.00	3.95 %
4	16.00	2	32.00	3.95 %
5	32.00	1	32.00	3.95 %
6	8.00	4	32.00	3.95 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
24	8.00	4	32.00	3.95 %
			-------	--------
			768.00	94.82 %

Experiment step size: 3.95 %

Test 20 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000
6	125.000	80	0	0	0.000
.		.	.	.	.
.		.	.	.	.
.		.	.	.	.
24	125.000	80	0	0	0.000

==

Test when deadlines first missed/skipped:

═══

Experiment: EXPERIMENT_4
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 809.92

Test 21 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	3.95 %
2	4.00	8	32.00	3.95 %
3	8.00	4	32.00	3.95 %
4	16.00	2	32.00	3.95 %
5	32.00	1	32.00	3.95 %
6	8.00	4	32.00	3.95 %
.	.	.	.	.
.	.	.	.	.
25	8.00	4	32.00	3.95 %
			-------	--------
			800.00	98.78 %

Experiment step size: 3.95 %

Test 21 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	10	10	320.999
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000
6	125.000	80	0	0	0.000
7	125.000	80	0	0	0.000
8	125.000	80	0	0	0.000
9	125.000	80	0	0	0.000
10	125.000	78	1	1	5.219
11	125.000	80	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
15	125.000	80	0	0	0.000
16	125.000	78	1	1	4.944
17	125.000	80	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
25	125.000	80	0	0	0.000

═══

Final test performed:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 150 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   809.92

Test 22 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	16	32.00	3.95 %
2	4.00	8	32.00	3.95 %
3	8.00	4	32.00	3.95 %
4	16.00	2	32.00	3.95 %
5	32.00	1	32.00	3.95 %
6	8.00	4	32.00	3.95 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
26	8.00	4	32.00	3.95 %
			--------	--------
			832.00	102.73 %

Experiment step size: 3.95 %

Test 22 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	1	5	14	1344.714
2	250.000	0	20	20	125.893
3	125.000	78	1	1	5.188
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000
6	125.000	78	1	1	10.010
7	125.000	80	0	0	0.000
8	125.000	78	1	1	10.803
9	125.000	80	0	0	0.000
10	125.000	76	2	2	6.744
11	125.000	76	2	2	8.026
12	125.000	80	0	0	0.000
13	125.000	80	0	0	0.000
14	125.000	78	1	1	8.881
15	125.000	76	2	2	8.270
16	125.000	78	1	1	8.698
17	125.000	76	2	2	6.775
18	125.000	80	0	0	0.000
19	125.000	78	1	1	9.521
20	125.000	78	1	1	4.303

21	125.000	80	0	0	0.000
22	125.000	76	2	2	8.621
23	125.000	78	1	1	11.536
24	125.000	78	1	1	5.341
25	125.000	78	1	1	4.456
26	125.000	78	1	1	8.209

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Rational Environment, Delta 1 release
Target    : Rational R1000, Series 200, Model 20

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 20 of Experiment 4

   Raw (non-tasking) benchmark speed in KWIPS:  809.92

   Full task set:

        Total      Deadlines      Task Set       Total
        Tasks      Per Second    Utilization     KWIPS
         24         214.00         94.82 %       768.00

   Highest-frequency task:

        Period     Deadlines        Task          Task
        (msec)     Per Second    Utilization      KWIPS
        31.250       32.00          3.95 %        32.00

   Experiment step size:   3.95 %

================================================================


            END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

Appendix F: Summary Results: Systems Designers

The following is the host-target configuration used for generating the results reported here:

HOST:	DEC MicroVAX 3200 running VAX/VMS, release 5.3-1
CROSS-COMPILER:	Systems Designers XD Ada, Version 1.0, ACVC 1.10
TARGET:	Motorola MVME133A: 20 MHz MC68020 CPU with 20 MHz MC68881 Floating-Point Co-processor; one wait state; 1Mb RAM; 256-byte on-chip instruction cache

Full optimization was specified for all compilations. No checks were suppressed. For the XD Ada compiler, the resolution of Calendar.Clock is 122 microseconds (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-5. The mathematical library is named Float_Math_Lib and the logarithmic function used by the Small_Whetstone procedure is named "Log".

In all 4 experiments the results were as predicted. Task 1 began missing its deadlines before the other tasks. At no time did a test fail when a later test succeeded. The workloads were increased slightly in Experiment 1 to prevent a premature completion of the experiment with Task 5 missing its deadlines first.

Compiler: Systems Designers XD Ada, V1.0
Target: Motorola MVME133A (20MHz MC68020)
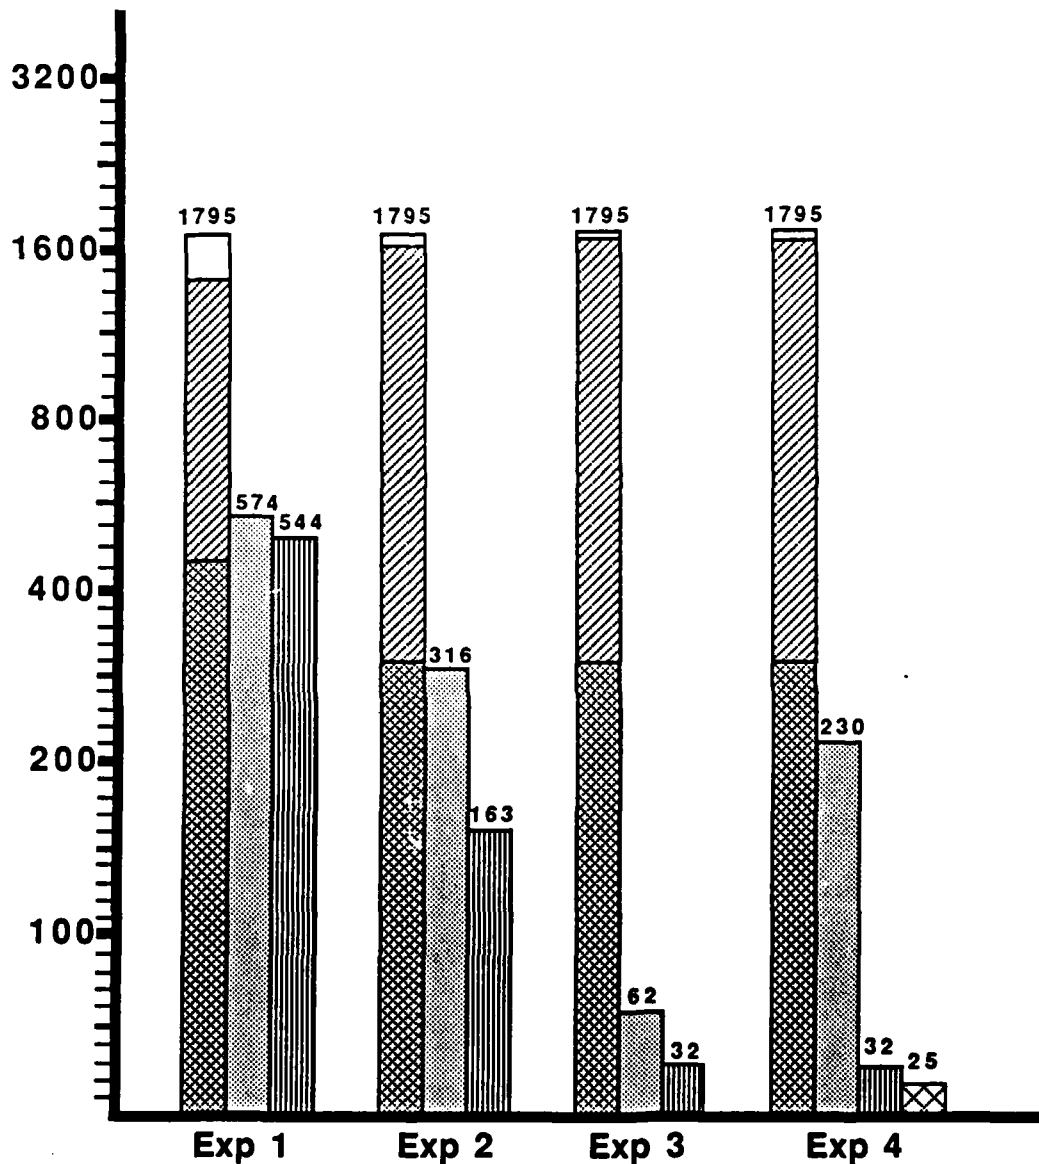
First Bar	☐	**Raw throughput (KWIPS)**
	▨	**Workload before breakdown (KWIPS)**
	▨	**Baseline workload (KWIPS)**
Second Bar	▨	**Deadlines per second before breakdown**
Third Bar	▥	**Frequency of highest frequency task**
Fourth Bar	▧	**Number of tasks before breakdown**

F.1. Experiment 1: Systems Designers

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.39

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	32.00	2	64.00	3.56 %
			-------	--------
			448.00	24.95 %

Experiment step size: 1.78 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

==

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.39

Test 33 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	544.00	2	1088.00	60.60 %
			-------	--------
			1472.00	81.99 %

Experiment step size: 1.78 %

Test 33 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	1.838	5440	0	0	0.000

===

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.39

Test 34 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	560.00	2	1120.00	62.38 %
			1504.00	83.77 %

Experiment step size: 1.78 %

Test 34 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	10	10	224.768
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	1.786	5600	0	0	0.000

==

Final test performed:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.39

Test 36 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	48	96.00	5.35 %
2	4.00	24	96.00	5.35 %
3	8.00	12	96.00	5.35 %
4	16.00	6	96.00	5.35 %
5	592.00	2	1184.00	65.95 %
			-------	--------
			1568.00	87.33 %

Experiment step size: 1.78 %

Test 36 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	6	14	835.907
2	250.000	8	16	16	84.835
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	1.689	5908	6	6	0.020

==

```
=========================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Systems Designers XD Ada 1.0 VAX/VMS -> MC68020
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 33 of Experiment 1

   Raw (non-tasking) benchmark speed in KWIPS: 1795.39

   Full task set:

        Total        Deadlines       Task Set        Total
        Tasks        Per Second      Utilization      KWIPS
          5            574.00          81.99 %        1472.00

   Highest-frequency task:

        Period       Deadlines         Task           Task
        (msec)       Per Second      Utilization      KWIPS
         1.838         544.00          60.60 %        1088.00

   Experiment step size:   1.78 %

=========================================================================


            END  OF  HARTSTONE  BENCHMARK  SUMMARY  RESULTS
```

F.2. Experiment 2: Systems Designers

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.45

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
			320.00	17.82 %

Experiment step size: 1.78 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

==

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.45

Test 42 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	10.20	32	326.40	18.18 %
2	20.40	16	326.40	18.18 %
3	40.80	8	326.40	18.18 %
4	81.60	4	326.40	18.18 %
5	163.20	2	326.40	18.18 %
			-------	--------
			1632.00	90.90 %

Experiment step size: 1.78 %

Test 42 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	98.039	102	0	0	0.000
2	49.020	204	0	0	0.000
3	24.510	408	0	0	0.000
4	12.255	816	0	0	0.000
5	6.127	1632	0	0	0.000

===

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.45

Test 43 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	10.40	32	332.80	18.54 %
2	20.80	16	332.80	18.54 %
3	41.60	8	332.80	18.54 %
4	83.20	4	332.80	18.54 %
5	166.40	2	332.80	18.54 %
			-------	--------
			1664.00	92.68 %

Experiment step size: 1.78 %

Test 43 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	96.154	53	26	26	20.747
2	48.077	209	0	0	0.000
3	24.038	417	0	0	0.000
4	12.019	833	0	0	0.000
5	6.010	1665	0	0	0.000

==

Final test performed:
    See preceding summary of test 43

===

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : Systems Designers XD Ada 1.0 VAX/VMS -> MC68020
Target   : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 42 of Experiment 2

    Raw (non-tasking) benchmark speed in KWIPS: 1795.45

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   316.20   90.90 %   1632.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
         6.127   163.20   18.18 %   326.40

    Experiment step size: 1.78 %

===

              END OF HARTSTONE BENCHMARK SUMMARY RESULTS

F.3. Experiment 3: Systems Designers

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
			-------	--------
			320.00	17.82 %

Experiment step size: 3.45 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

==

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 24 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	55	110.00	6.13 %
2	4.00	39	156.00	8.69 %
3	8.00	31	248.00	13.81 %
4	16.00	27	432.00	24.06 %
5	32.00	25	800.00	44.56 %
			1746.00	97.25 %

Experiment step size: 3.45 %

Test 24 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

==

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 25 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	56	112.00	6.24 %
2	4.00	40	160.00	8.91 %
3	8.00	32	256.00	14.26 %
4	16.00	28	448.00	24.95 %
5	32.00	26	832.00	46.34 %
			-------	--------
			1808.00	100.70 %

Experiment step size: 3.45 %

Test 25 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	10	10	465.985
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

===

Final test performed:
    See preceding summary of test 25

==

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : Systems Designers XD Ada 1.0 VAX/VMS -> MC68020
Target   : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 24 of Experiment 3

    Raw (non-tasking) benchmark speed in KWIPS: 1795.43

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   62.00   97.25 %   1746.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
        31.250   32.00   44.56 %   800.00

    Experiment step size: 3.45 %

==

                END OF HARTSTONE BENCHMARK SUMMARY RESULTS

F.4. Experiment 4: Systems Designers

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
			-------	--------
			320.00	17.82 %

Experiment step size: 3.56 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000

===

Last test with no missed/skipped deadlines:

```
=================================================================
```

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 22 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
6	8.00	8	64.00	3.56 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
26	8.00	8	64.00	3.56 %
			1664.00	92.68 %

Experiment step size: 3.56 %

```
-----------------------------------------------------------------
```

Test 22 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	20	0	0	0.000
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000
6	125.000	80	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
26	125.000	80	0	0	0.000

```
=================================================================
```

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 23 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.00	32	64.00	3.56 %
2	4.00	16	64.00	3.56 %
3	8.00	8	64.00	3.56 %
4	16.00	4	64.00	3.56 %
5	32.00	2	64.00	3.56 %
6	8.00	8	64.00	3.56 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
27	8.00	8	64.00	3.56 %
			--------	--------
			1728.00	96.24 %

Experiment step size: 3.56 %

Test 23 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	500.000	0	10	10	246.747
2	250.000	40	0	0	0.000
3	125.000	80	0	0	0.000
4	62.500	160	0	0	0.000
5	31.250	320	0	0	0.000
6	125.000	80	0	0	0.000
.		.	.	.	.
.		.	.	.	.
.		.	.	.	.
27	125.000	80	0	0	0.000

===

Final test performed:

```
=======================================================================

Experiment:    EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 1795.43

Test 24 characteristics:

     Task   Frequency   Kilo-Whets   Kilo-Whets   Requested Workload
     No.     (Hertz)    per period   per second      Utilization
      1        2.00         32          64.00           3.56 %
      2        4.00         16          64.00           3.56 %
      3        8.00          8          64.00           3.56 %
      4       16.00          4          64.00           3.56 %
      5       32.00          2          64.00           3.56 %
      6        8.00          8          64.00           3.56 %
      .         .            .            .               .
      .         .            .            .               .
      .         .            .            .               .

     28        8.00          8          64.00           3.56 %
                                       -------         --------
                                       1792.00          99.81 %

Experiment step size:   3.56 %

-----------------------------------------------------------------------

Test 24 results:

Test duration (seconds):  10.0

     Task     Period      Met        Missed       Skipped      Average
     No.    in msecs   Deadlines   Deadlines    Deadlines    Late (msec)
      1     500.000        0           4            16        1997.498
      2     250.000        0          20            20         124.005
      3     125.000       80           0             0           0.000
      4      62.500      160           0             0           0.000
      5      31.250      320           0             0           0.000
      6     125.000       80           0             0           0.000
      .                    .           .             .             .
      .                    .           .             .             .
      .                    .           .             .             .

     28     125.000       80           0             0           0.000

=======================================================================
```

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Systems Designers XD Ada 1.0 VAX/VMS -> MC68020
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 22 of Experiment 4

    Raw (non-tasking) benchmark speed in KWIPS: 1795.43

    Full task set:

        Total       Deadlines     Task Set       Total
        Tasks       Per Second    Utilization    KWIPS
         26          230.00         92.68 %      1664.00

    Highest-frequency task:

        Period      Deadlines       Task          Task
        (msec)      Per Second    Utilization     KWIPS
        31.250        32.00         3.56 %        64.00

    Experiment step size:   3.56 %


================================================================


             END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

Appendix G: Summary Results: Tartan

The following is the host-target configuration used for generating the results reported here:

HOST:	DEC MicroVAX 3200 running VAX/VMS, release 5.3-1
CROSS-COMPILER:	Tartan VMS/1750A, Version 2.14, ACVC 1.10
TARGET:	Fairchild SBC-50: 15 MHz F9450, no wait states; 128Kb RAM

Full optimization was specified for all compilations. No checks were suppressed. For the Tartan compiler, the resolution of Calendar.Clock is 61 microseconds (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-6. A mathematical library was not provided with this release of the compiler so one was adapted from the mathematical functions provided in PIWG test A000093. This library was named PIWG_Math_Lib and the logarithmic function used by the Small_Whetstone procedure was named "Log".

Hartstone would not compile successfully under Tartan Ada. The following error message resulted from the compilation of Periodic_Tasks_Body.

```
TARTAN Ada VMS/1750A,V2.14 Copyright 1989, Tartan Laboratories

*** First diagnostic is on line 1

   1|
     ^1
*** 1 Fatal Error 9999: An internal error has occurred in the Tartan Ada
***    compiler.  Please notify Tartan at (412) 856-3600.  The internal
***    error code is 2YFROWCOZGV%C.
```

This error was traced back to the Pragma Inline of the Small_Whetstone procedure in Workload_Spec. This pragma was deleted and Hartstone then compiled successfully.

In Experiments 3 and 4 the results were as predicted. Task 1 began missing its deadlines before the other tasks. At no time did a test fail when a later test succeeded. The frequencies and workloads were reduced to keep the baseline utilization within the guidelines. In Experiment 3 the frequencies were scaled back to obtain an adequate step size. In Experiments 1 and 2, the breakdown pattern was quite unexpected. All of the tasks missed their deadlines, and by a substantial amount (several seconds). Such a dramatic breakdown after only a small percentage increase over the preceding test is unusual. These results have not yet been explained. Further scaling back of the frequencies did not produce better results.

Compiler: Tartan VMS/1750A, V2.14
Target: Fairchild SBC-50 (15MHz F9450)
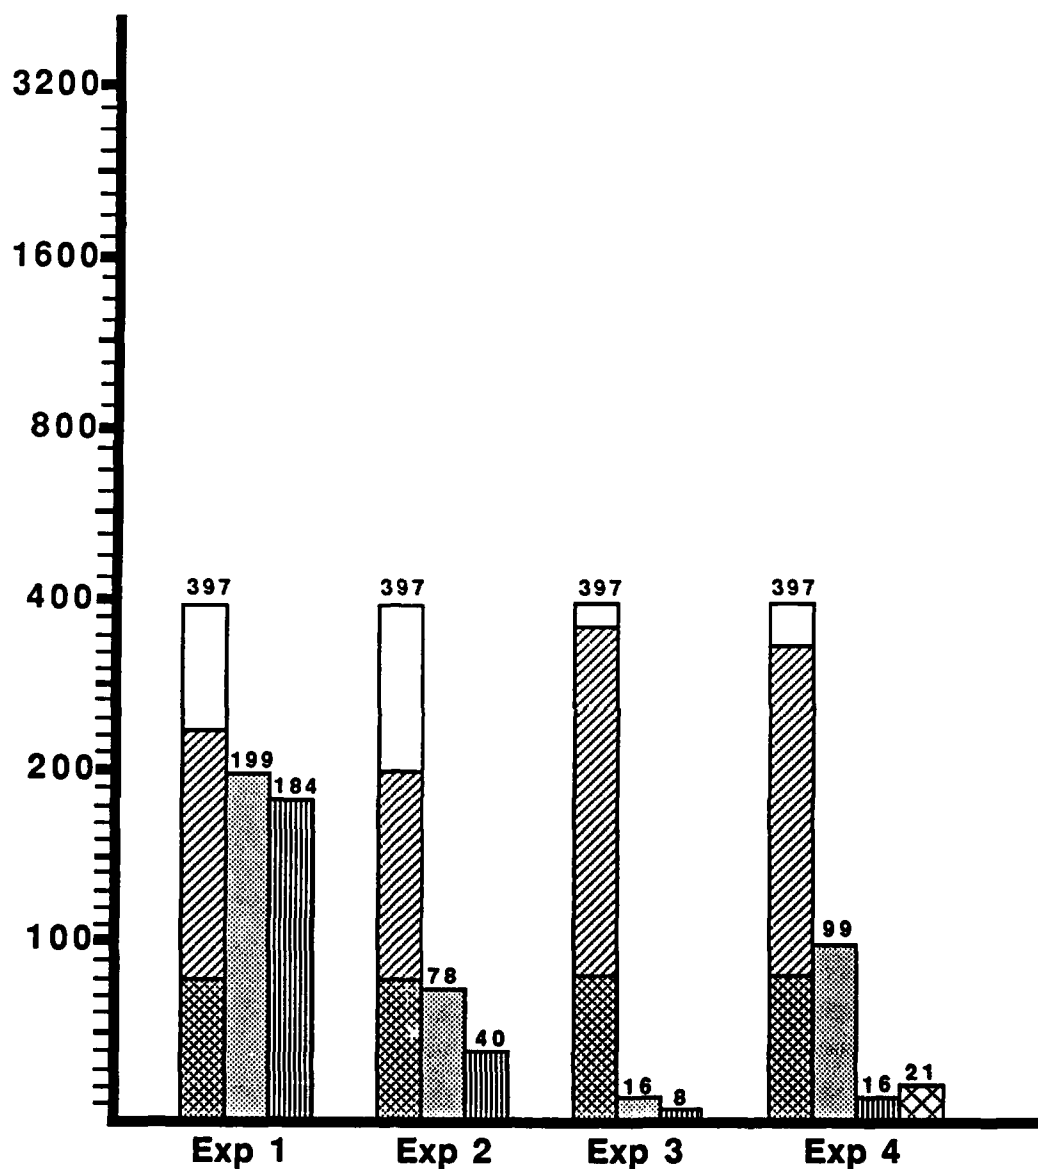
First Bar
☐ Raw throughput (KWIPS)
▨ Workload before breakdown (KWIPS)
▦ Baseline workload (KWIPS)

Second Bar ▨ Deadlines per second before breakdown

Third Bar ▥ Frequency of highest frequency task

Fourth Bar ▨ Number of tasks before breakdown

G.1. Experiment 1: Tartan

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.63

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	16.00	1	16.00	4.03 %
			-------	--------
			80.00	20.17 %

Experiment step size: 2.02 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.63

Test 22 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	184.00	1	184.00	46.39 %
			-------	--------
			248.00	62.53 %

Experiment step size: 2.02 %

Test 22 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	5.435	1841	0	0	0.000

===

Test when deadlines first missed/skipped:

═══

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.63

Test 23 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	192.00	1	192.00	48.41 %
			-------	--------
			256.00	64.54 %

Experiment step size: 2.02 %

Test 23 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	12369.506
2	500.000	1	1	18	12327.818
3	250.000	1	1	38	12799.133
4	125.000	1	2	77	6447.021
5	5.208	21	3	1904	4368.591

═══

Final test performed:
    See preceding summary of test 23

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Tartan VAX/VMS -> 1750A, release 2.14
Target    : Fairchild SBC-50 (15 MHz F9450)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 22 of Experiment 1

    Raw (non-tasking) benchmark speed in KWIPS:  396.63

    Full task set:

        Total      Deadlines      Task Set      Total
        Tasks      Per Second     Utilization   KWIPS
          5         199.00          62.53 %      248.00

    Highest-frequency task:

        Period     Deadlines        Task        Task
        (msec)     Per Second     Utilization   KWIPS
         5.435      184.00          46.39 %      184.00

    Experiment step size:   2.02 %

================================================================


             END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

G.2. Experiment 2: Tartan

```
                    HARTSTONE BENCHMARK SUMMARY RESULTS


Baseline test:

================================================================

Experiment:     EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.62

Test 1 characteristics:

        Task    Frequency   Kilo-Whets    Kilo-Whets    Requested Workload
        No.     (Hertz)     per period    per second        Utilization
         1       1.00           16           16.00           4.03 %
         2       2.00            8           16.00           4.03 %
         3       4.00            4           16.00           4.03 %
         4       8.00            2           16.00           4.03 %
         5      16.00            1           16.00           4.03 %
                                             -------         --------
                                              80.00           20.17 %

Experiment step size:   2.02 %

------------------------------------------------------------------------

Test 1 results:

Test duration (seconds):  10.0

        Task    Period        Met         Missed      Skipped      Average
        No.     in msecs   Deadlines .  Deadlines    Deadlines    Late (msec)
         1     1000.000       10           0            0           0.000
         2      500.000       20           0            0           0.000
         3      250.000       40           0            0           0.000
         4      125.000       80           0            0           0.000
         5       62.500      160           0            0           0.000


================================================================
```

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.62

Test 16 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.50	16	40.00	10.09 %
2	5.00	8	40.00	10.09 %
3	10.00	4	40.00	10.09 %
4	20.00	2	40.00	10.09 %
5	40.00	1	40.00	10.09 %
			-------	--------
			200.00	50.43 %

Experiment step size: 2.02 %

Test 16 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	400.000	26	0	0	0.000
2	200.000	51	0	0	0.000
3	100.000	101	0	0	0.000
4	50.000	201	0	0	0.000
5	25.000	401	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.62

Test 17 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	2.60	16	41.60	10.49 %
2	5.20	8	41.60	10.49 %
3	10.40	4	41.60	10.49 %
4	20.80	2	41.60	10.49 %
5	41.60	1	41.60	10.49 %
			208.00	52.44 %

Experiment step size: 2.02 %

Test 17 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	384.615	18	1	8	6066.284
2	192.308	35	1	17	6410.278
3	96.154	70	1	34	6484.497
4	48.077	141	1	67	6485.352
5	24.038	281	1	135	6532.775

==

Final test performed:
    See preceding summary of test 17

==

Benchmark : Hartstone Benchmark, Version 1.0
Compiler : Tartan VAX/VMS -> 1750A, release 2.14
Target   : Fairchild SBC-50 (15 MHz F9450)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 16 of Experiment 2

    Raw (non-tasking) benchmark speed in KWIPS: 396.62

    Full task set:

        Total   Deadlines   Task Set   Total
        Tasks   Per Second   Utilization   KWIPS
          5   77.50   50.43 %   200.00

    Highest-frequency task:

        Period   Deadlines   Task   Task
        (msec)   Per Second   Utilization   KWIPS
        25.000   40.00   10.09 %   40.00

    Experiment step size: 2.02 %

==

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

G.3. Experiment 3: Tartan

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.62

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	32	16.00	4.03 %
2	1.00	16	16.00	4.03 %
3	2.00	8	16.00	4.03 %
4	4.00	4	16.00	4.03 %
5	8.00	2	16.00	4.03 %
			-------	--------
			80.00	20.17 %

Experiment step size: 3.91 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

===

Last test with no missed/skipped deadlines:

================= ==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whatstone Instructions Per Second (KWIPS): 396.62

Test 20 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	51	25.50	6.43 %
2	1.00	35	35.00	8.82 %
3	2.00	27	54.00	13.61 %
4	4.00	23	92.00	23.20 %
5	8.00	21	168.00	42.36 %
			-------	--------
			374.50	94.42 %

Experiment step size: 3.91 %

Test 20 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

==

Test when deadlines first missed/skipped:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.62

Test 21 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	52	26.00	6.56 %
2	1.00	36	36.00	9.08 %
3	2.00	28	56.00	14.12 %
4	4.00	24	96.00	24.20 %
5	8.00	22	176.00	44.37 %
			-------	--------
			390.00	98.33 %

Experiment step size: 3.91 %

Test 21 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	0	3	2	624.919
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

===

Final test performed:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.62

Test 22 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	53	26.50	6.68 %
2	1.00	37	37.00	9.33 %
3	2.00	29	58.00	14.62 %
4	4.00	25	100.00	25.21 %
5	8.00	23	184.00	46.39 %
			-------	--------
			405.50	102.24 %

Experiment step size: 3.91 %

Test 22 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	0	2	3	2511.780
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

===

```
================================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Tartan VAX/VMS -> 1750A, release 2.14
Target    : Fairchild SBC-50 (15 MHz F9450)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 20 of Experiment 3

    Raw (non-tasking) benchmark speed in KWIPS:  396.62

    Full task set:

        Total       Deadlines       Task Set        Total
        Tasks       Per Second      Utilization     KWIPS
          5           15.50           94.42 %        374.50

    Highest-frequency task:

        Period      Deadlines         Task          Task
        (msec)      Per Second      Utilization      KWIPS
        125.000        8.00           42.36 %        168.00

    Experiment step size:    3.91 %

================================================================================


            END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

G.4. Experiment 4: Tartan

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.63

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	16.00	1	16.00	4.03 %
			80.00	20.17 %

Experiment step size: 4.03 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

```
================================================================
```

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.63

Test 18 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	16.00	1	16.00	4.03 %
6	4.00	4	16.00	4.03 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
22	4.00	4	16.00	4.03 %
			352.00	88.75 %

Experiment step size: 4.03 %

```
----------------------------------------------------------------
```

Test 18 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
22	250.000	40	0	0	0.000

```
================================================================
```

Test when deadlines first missed/skipped:

```
================================================================
```

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   396.63

Test 19 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	16.00	1	16.00	4.03 %
6	4.00	4	16.00	4.03 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
23	4.00	4	16.00	4.03 %
			-------	--------
			368.00	92.78 %

Experiment step size: 4.03 %

```
------------------------------------------------------------------
```

Test 19 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	4	6	1245.178
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
23	250.000	40	0	0	0.000

```
================================================================
```

Final test performed:

===

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 396.63

Test 21 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	4.03 %
2	2.00	8	16.00	4.03 %
3	4.00	4	16.00	4.03 %
4	8.00	2	16.00	4.03 %
5	16.00	1	16.00	4.03 %
6	4.00	4	16.00	4.03 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
25	4.00	4	16.00	4.03 %
			-------	--------
			400.00	100.85 %

Experiment step size: 4.03 %

Test 21 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9053.467
2	500.000	0	2	18	4131.622
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	36	2	2	14.709
7	250.000	32	4	4	9.338
8	250.000	38	1	1	16.907
9	250.000	40	0	0	0.000
10	250.000	38	1	1	16.602
11	250.000	40	0	0	0.000
12	250.000	40	0	0	0.000
13	250.000	34	3	3	5.188
14	250.000	40	0	0	0.000
15	250.000	36	2	2	1.831
16	250.000	38	1	1	17.090
17	250.000	38	1	1	12.268
18	250.000	38	1	1	16.968
19	250.000	40	0	0	0.000
20	250.000	34	3	3	11.882

21	250.000	38	1	1	12.878
22	250.000	38	1	1	12.451
23	250.000	40	0	0	0.000
24	250.000	34	3	3	11.719
25	250.000	36	2	2	16.632

```
===============================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Tartan VAX/VMS -> 1750A, release 2.14
Target    : Fairchild SBC-50 (15 MHz F9450)

Characteristics of best test for this experiment:   (no missed/skipped
deadlines)

   Test 18 of Experiment 4

   Raw (non-tasking) benchmark speed in KWIPS:   396.63

   Full task set:

       Total         Deadlines      Task Set        Total
       Tasks         Per Second     Utilization     KWIPS
        22            99.00          88.75 %        352.00

   Highest-frequency task:

       Period        Deadlines       Task           Task
       (msec)        Per Second     Utilization     KWIPS
       62.500         16.00          4.03 %          16.00

   Experiment step size:   4.03 %


===============================================================================


           END OF  HARTSTONE  BENCHMARK  SUMMARY  RESULTS
```

Appendix H: Summary Results: TeleSoft

The following is the host-target configuration used for generating the results reported here:

HOST: DEC MicroVAX 3200 running VAX/VMS, release 5.3-1
CROSS-COMPILER: TeleSoft TeleGen2, release 3.22a, ACVC 1.9
TARGET: Motorola MVME133A: 20 MHz MC68020 CPU with 20 MHz
                      MC68881 Floating-Point Co-processor; one wait
                      state; 1Mb RAM; 256-byte on-chip instruction cache

Full optimization was specified for all compilations. No checks were suppressed. For the TeleSoft compiler, the resolution of Calendar.Clock is 100 milliseconds (see Appendix B) and Float'Digits is 6. The characteristics of the **delay** statement are shown in Figure B-7. The mathematical library is named Math_Library and the logarithmic function used by the Small_Whetstone procedure is named "Ln"; this was renamed to "Log" in package Workload.

The original intent was to use the 3.23 release of TeleGen2 for Hartstone testing; however Hartstone could not be made to work under this release. Compiling and linking were no problem, but when downloaded, Hartstone would always fail with the following message:

```
>>> Unhandled exception:  CONSTRAINT_ERROR  (Access Check)
    Raised in   Experiment.??? at address 3075A
```

According to the link map, this address is somewhere in the "ofm/cgsker" library component. Scaling back the baseline frequencies and workloads had no effect; the error was consistently reproducible. With the same compilation command and the same linker options file (except for the HEAP2 area, a 3.23 feature), the same Hartstone code compiled and linked without problems under the older 3.22a release of the compiler, and also produced actual results when downloaded. Thus release 3.22a is the version of TeleGen2 for which results are reported here.

Because of the 100 ms Calendar.Clock resolution and the 20 ms **delay** statement resolution of the TeleGen2 3.22a compiler, the baseline test missed deadlines when run "as-is". When the baseline frequencies and workloads were scaled back by cutting them in half, the baseline test still missed deadlines, as shown in the first listing of Experiment 1 results. Further scaling back allowed the test to proceed beyond the baseline test, but the highest-frequency task began missing deadlines before any others, resulting in the kind of inverted task breakdown pattern described in the user's guide. This result is shown in the second listing for Experiment 1. Several more attempts were made to get a "normal" task breakdown pattern, but without success. Even when the highest-frequency task was given a minimal amount to do (e.g., 16 Kilo-Whetstones in a one-second period) and the other tasks were given much larger workloads in an attempt to get them to miss deadlines first, the test still resulted in higher-priority tasks missing deadlines before the lower-priority tasks. An example of this is shown in the third listing. Because of these problems, and the necessity to run Hartstone on other cross-compilation systems, no further testing of the TeleSoft compiler was done. The anomalous results have been reported to TeleSoft.

H.1. Experiment 1: TeleSoft

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   684.95

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	2.34 %
2	2.00	8	16.00	2.34 %
3	4.00	4	16.00	2.34 %
4	8.00	2	16.00	2.34 %
5	16.00	1	16.00	2.34 %
			-------	--------
			80.00	11.68 %

Experiment step size: 1.17 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	72	4	4	12.939
5	62.500	112	23	25	26.006

===

Last test with no missed/skipped deadlines:
    Not applicable

Test when deadlines first missed/skipped:
    See preceding summary of test 1

Final test performed:
    See preceding summary of test 1

```
====================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : TeleSoft TeleGen2 Ada 3.22a VAX/VMS -> MC68020
Target    : Motorola MVME133A (20 MHz MC68020 + 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Not applicable

====================================================================


            END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

H.2. Modified Experiment 1: TeleSoft

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   684.90

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.06250	1024	64.00	9.34 %
2	0.12500	512	64.00	9.34 %
3	0.25000	256	64.00	9.34 %
4	0.50000	128	64.00	9.34 %
5	1.00000	32	32.00	4.67 %
			-------	--------
			288.00	42.05 %

Experiment step size: 2 34 %

Test 1 results:

Test duration (seconds): 64.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	16000.000	4	0	0	0.000
2	8000.000	8	0	0	0.000
3	4000.000	16	0	0	0.000
4	2000.000	30	1	1	100.098
5	1000.000	60	2	2	5n0.000

==

Last test with no missed/skipped deadlines:
   Not applicable

Test when deadlines first missud/skipped:
   See preceding summary of test 1

Final test performed:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   684.90

Test 5 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.06250	1024	64.00	9.34 %
2	0.12500	512	64.00	9.34 %
3	0.25000	256	64.00	9.34 %
4	0.50000	128	64.00	9.34 %
5	3.00000	32	96.00	14.02 %
			-------	--------
			352.00	51.39 %

Experiment step size: 2.34 %

Test 5 results:

Test duration (seconds):   64.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	16000.000	4	0	0	0.000
2	8000.000	8	0	0	0.000
3	4000.000	16	0	0	0.000
4	2000.000	30	1	1	300.293
5	333.333	136	22	35	323.453

==

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : TeleSoft TeleGen2 Ada 3.22a VAX/VMS -> MC68020
Target    : Motorola MVME133A (20 MHz MC68020 + 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Not applicable

================================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

H.3. Modified Experiment 1: TeleSoft

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   684.90

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.06250	2048	128.00	18.69 %
2	0.12500	1024	128.00	18.69 %
3	0.25000	512	128.00	18.69 %
4	0.50000	256	128.00	18.69 %
5	1.00000	16	16.00	2.34 %
			--------	--------
			528.00	77.09 %

Experiment step size: 1.17 %

Test 1 results:

Test duration (seconds):   64.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	16000.000	4	0	0	0.000
2	8000.000	8	0	0	0.000
3	4000.000	8	4	4	1249.756
4	2000.000	19	6	7	1150.146
5	1000.000	36	11	17	1199.751

==

Last test with no missed/skipped deadlines:
    Not applicable

Test when deadlines first missed/skipped:
    See preceding summary of test 1

Final test performed:

```
================================================================
```

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   684.90

Test 2 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.06250	2048	128.00	18.69 %
2	0.12500	1024	128.00	18.69 %
3	0.25000	512	128.00	18.69 %
4	0.50000	256	128.00	18.69 %
5	1.50000	16	24.00	3.50 %
			-------	--------
			536.00	78.26 %

Experiment step size: 1.17 %

```
----------------------------------------------------------------
```

Test 2 results:

Test duration (seconds): 64.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	16000.000	4	0	0	0.000
2	8000.000	8	0	0	0.000
3	4000.000	12	2	2	299.805
4	2000.000	22	4	6	1974.976
5	666.667	46	17	34	966.625

```
================================================================
```

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : TeleSoft TeleGen2 Ada 3.22a VAX/VMS -> MC68020
Target    : Motorola MVME133A (20 MHz MC68020 + 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Not applicable

================================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

Appendix I: Summary Results: Verdix

The following is the host-target configuration used for generating the results reported here:

HOST:	DEC MicroVAX 3200 running VAX/VMS, release 5.3-1
CROSS-COMPILER:	Verdix VADS, Version 5.7, ACVC 1.10
TARGET:	Motorola MVME133A: 20 MHz MC68020 CPU with 20 MHz MC68881 Floating-Point Co-processor; one wait state; 1Mb RAM; 256-byte on-chip instruction cache

Full optimization was specified for all compilations. No checks were suppressed. For the VADS compiler, the resolution of Calendar.Clock is 61 microseconds (see Appendix B) and Float'Digits is 15. The characteristics of the **delay** statement are shown in Figure B-8. Because of the performance penalty of doing double-precision computations, the floating-point computations in the Workload package were done in type Short_Float, which has a Short_Float'Digits value of 6. The mathematical library used was the (unsupported, as yet) Generic_Elementary_Functions package from the Verdix PUBLICLIB directory. The logarithmic function (used by the Small_Whetstone procedure) of this library is named Log. The renames clause must be removed from workload body since it will not work with a Log function which requires multiple parameters.

In all 4 experiments the results were as predicted. Task 1 began missing its deadlines before the other tasks. At no time did a test fail when a later test succeeded. The frequencies were scaled back due to the imprecise delay statement. The workloads were decreased to keep the baseline utilization within the guidelines. In Experiment 1, the workloads were increased substantially to prevent premature completion of the experiment with Task 5 missing its deadlines first. This resulted in a baseline utilization in excess of the guidelines. In Experiment 3, the frequencies were decreased further to obtain an adequate step size.

Compiler: Verdix VADS, V5.7
Target: Motorola MVME133A (20MHz MC68020)

First Bar
- ☐ Raw throughput (KWIPS)
- ▨ Workload before breakdown (KWIPS)
- ▧ Baseline workload (KWIPS)

Second Bar ▦ Deadlines per second before breakdown

Third Bar ▥ Frequency of highest frequency task

Fourth Bar ▨ Number of tasks before breakdown

I.1. Experiment 1: Verdix

```
                    HARTSTONE BENCHMARK SUMMARY RESULTS


Baseline test:


=======================================================================

Experiment:    EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 1 characteristics:

     Task    Frequency   Kilo-Whets    Kilo-Whets    Requested Workload
     No.     (Hertz)     per period    per second        Utilization
      1       1.00          64           64.00           12.64 %
      2       2.00          32           64.00           12.64 %
      3       4.00          16           64.00           12.64 %
      4       8.00           8           64.00           12.64 %
      5      16.00           2           32.00            6.32 %
                                        -------          --------
                                        288.00           56.88 %


Experiment step size:    3.16 %


-----------------------------------------------------------------------


Test 1 results:

Test duration (seconds):  10.0

     Task     Period        Met        Missed      Skipped       Average
     No.     in msecs    Deadlines    Deadlines   Deadlines    Late (msec)
      1     1000.000        10           0           0           0.000
      2      500.000        20           0           0           0.000
      3      250.000        40           0           0           0.000
      4      125.000        80           0           0           0.000
      5       62.500       160           0           0           0.000


=======================================================================
```

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 10 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	12.64 %
2	2.00	32	64.00	12.64 %
3	4.00	16	64.00	12.64 %
4	8.00	8	64.00	12.64 %
5	88.00	2	176.00	34.76 %
			-------	--------
			432.00	85.32 %

Experiment step size: 3.16 %

Test 10 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	11.364	881	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 11 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	12.64 %
2	2.00	32	64.00	12.64 %
3	4.00	16	64.00	12.64 %
4	8.00	8	64.00	12.64 %
5	96.00	2	192.00	37.92 %
			-------	--------
			448.00	88.48 %

Experiment step size: 3.16 %

Test 11 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	5	5	370.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	10.417	961	0	0	0.000

==

Final test performed:

==

Experiment: EXPERIMENT_1
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 13 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	64	64.00	12.64 %
2	2.00	32	64.00	12.64 %
3	4.00	16	64.00	12.64 %
4	8.00	8	64.00	12.64 %
5	112.00	2	224.00	44.24 %
			-------	--------
			480.00	94.80 %

Experiment step size: 3.16 %

Test 13 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	4	6	1122.500
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	8.929	1040	40	40	0.350

==

```
================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Verdix VAX/VMS -> MC68020, release 5.7
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 10 of Experiment 1

    Raw (non-tasking) benchmark speed in KWIPS:  506.33

    Full task set:

        Total       Deadlines       Task Set        Total
        Tasks       Per Second      Utilization     KWIPS
          5           103.00          85.32 %        432.00

    Highest-frequency task:

        Period      Deadlines         Task          Task
        (msec)      Per Second      Utilization     KWIPS
        11.364        88.00           34.76 %        176.00

    Experiment step size:   3.16 %

================================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

I.2. Experiment 2: Verdix

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	3.16 %
2	2.00	8	16.00	3.16 %
3	4.00	4	16.00	3.16 %
4	8.00	2	16.00	3.16 %
5	16.00	1	16.00	3.16 %
			-------	--------
			80.00	15.80 %

Experiment step size: 1.58 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

===

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 41 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	5.00	16	80.00	15.80 %
2	10.00	8	80.00	15.80 %
3	20.00	4	80.00	15.80 %
4	40.00	2	80.00	15.80 %
5	80.00	1	80.00	15.80 %
			-------	--------
			400.00	79.00 %

Experiment step size: 1.58 %

Test 41 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	200.000	50	0	0	0.000
2	100.000	100	0	0	0.000
3	50.000	201	0	0	0.000
4	25.000	401	0	0	0.000
5	12.500	800	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 42 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	5.10	16	81.60	16.12 %
2	10.20	8	81.60	16.12 %
3	20.40	4	81.60	16.12 %
4	40.80	2	81.60	16.12 %
5	81.60	1	81.60	16.12 %
			408.00	80.58 %

Experiment step size: 1.58 %

Test 42 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	196.078	13	19	19	41.526
2	98.039	102	0	0	0.000
3	49.020	205	0	0	0.000
4	24.510	409	0	0	0.000
5	12.255	817	0	0	0.000

==

Final test performed:

==

Experiment: EXPERIMENT_2
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 44 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	5.30	16	84.80	16.75 %
2	10.60	8	84.80	16.75 %
3	21.20	4	84.80	16.75 %
4	42.40	2	84.80	16.75 %
5	84.80	1	84.80	16.75 %
			424.00	83.74 %

Experiment step size: 1.58 %

Test 44 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	188.679	0	27	27	48.630
2	94.340	106	0	0	0.000
3	47.170	213	0	0	0.000
4	23.585	424	0	0	0.000
5	11.792	849	0	0	0.000

==

```
==================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Verdix VAX/VMS -> MC68020, release 5.7
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

   Test 41 of Experiment 2

   Raw (non-tasking) benchmark speed in KWIPS:  506.33

   Full task set:

      Total      Deadlines       Task Set       Total
      Tasks      Per Second      Utilization    KWIPS
        5          155.00          79.00 %       400.00

   Highest-frequency task:

      Period     Deadlines         Task          Task
      (msec)     Per Second      Utilization     KWIPS
      12.500       80.00           15.80 %        80.00

   Experiment step size:    1.58 %

==================================================================


               END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

I.3. Experiment 3: Verdix

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

===

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	32	16.00	3.16 %
2	1.00	16	16.00	3.16 %
3	2.00	8	16.00	3.16 %
4	4.00	4	16.00	3.16 %
5	8.00	2	16.00	3.16 %
			-------	--------
			80.00	15.80 %

Experiment step size: 3.06 %

Test 1 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

===

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 506.33

Test 27 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	58	29.00	5.73 %
2	1.00	42	42.00	8.30 %
3	2.00	34	68.00	13.43 %
4	4.00	30	120.00	23.70 %
5	8.00	28	224.00	44.24 %
			483.00	95.39 %

Experiment step size: 3.06 %

Test 27 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	5	0	0	0.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 28 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	59	29.50	5.83 %
2	1.00	43	43.00	8.49 %
3	2.00	35	70.00	13.83 %
4	4.00	31	124.00	24.49 %
5	8.00	29	232.00	45.82 %
			498.50	98.45 %

Experiment step size: 3.06 %

Test 28 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	0	3	2	626.667
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

==

Final test performed:

==

Experiment: EXPERIMENT_3
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 29 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	0.50	60	30.00	5.93 %
2	1.00	44	44.00	8.69 %
3	2.00	36	72.00	14.22 %
4	4.00	32	128.00	25.28 %
5	8.00	30	240.00	47.40 %
			514.00	101.52 %

Experiment step size: 3.06 %

Test 29 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	2000.000	0	2	3	2495.000
2	1000.000	10	0	0	0.000
3	500.000	20	0	0	0.000
4	250.000	40	0	0	0.000
5	125.000	80	0	0	0.000

==

```
===================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Verdix VAX/VMS -> MC68020, release 5.7
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 27 of Experiment 3

    Raw (non-tasking) benchmark speed in KWIPS:  506.33

    Full task set:

        Total       Deadlines       Task Set        Total
        Tasks       Per Second      Utilization     KWIPS
          5            15.50           95.39 %       483.00

    Highest-frequency task:

        Period      Deadlines         Task          Task
        (msec)      Per Second      Utilization     KWIPS
        125.000        8.00           44.24 %       224.00

    Experiment step size:   3.06 %

===================================================================


            END OF HARTSTONE BENCHMARK SUMMARY RESULTS
```

I.4. Experiment 4: Verdix

HARTSTONE BENCHMARK SUMMARY RESULTS

Baseline test:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 1 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	3.16 %
2	2.00	8	16.00	3.16 %
3	4.00	4	16.00	3.16 %
4	8.00	2	16.00	3.16 %
5	16.00	1	16.00	3.16 %
			-------	--------
			80.00	15.80 %

Experiment step size: 3.16 %

Test 1 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000

==

Last test with no missed/skipped deadlines:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 23 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	3.16 %
2	2.00	8	16.00	3.16 %
3	4.00	4	16.00	3.16 %
4	8.00	2	16.00	3.16 %
5	16.00	1	16.00	3.16 %
6	4.00	4	16.00	3.16 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
27	4.00	4	16.00	3.16 %
			--------	--------
			432.00	85.32 %

Experiment step size: 3.16 %

Test 23 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	10	0	0	0.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
27	250.000	40	0	0	0.000

==

Test when deadlines first missed/skipped:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS): 506.33

Test 24 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	3.16 %
2	2.00	8	16.00	3.16 %
3	4.00	4	16.00	3.16 %
4	8.00	2	16.00	3.16 %
5	16.00	1	16.00	3.16 %
6	4.00	4	16.00	3.16 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
28	4.00	4	16.00	3.16 %
			448.00	88.48 %

Experiment step size: 3.16 %

Test 24 results:

Test duration (seconds): 10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9010.000
2	500.000	20	0	0	0.000
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	40	0	0	0.000
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
28	250.000	40	0	0	0.000

==

Final test performed:

==

Experiment: EXPERIMENT_4
Completion on: Miss/skip 50 deadlines

Raw speed in Kilo-Whetstone Instructions Per Second (KWIPS):   506.33

Test 25 characteristics:

Task No.	Frequency (Hertz)	Kilo-Whets per period	Kilo-Whets per second	Requested Workload Utilization
1	1.00	16	16.00	3.16 %
2	2.00	8	16.00	3.16 %
3	4.00	4	16.00	3.16 %
4	8.00	2	16.00	3.16 %
5	16.00	1	16.00	3.16 %
6	4.00	4	16.00	3.16 %
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
29	4.00	4	16.00	3.16 %
			-------	--------
			464.00	91.64 %

Experiment step size: 3.16 %

Test 25 results:

Test duration (seconds):   10.0

Task No.	Period in msecs	Met Deadlines	Missed Deadlines	Skipped Deadlines	Average Late (msec)
1	1000.000	0	1	9	9000.000
2	500.000	0	4	16	1987.500
3	250.000	40	0	0	0.000
4	125.000	80	0	0	0.000
5	62.500	160	0	0	0.000
6	250.000	34	3	3	216.667
7	250.000	40	0	0	0.000
8	250.000	36	2	2	210.000
9	250.000	40	0	0	0.000
10	250.000	40	0	0	0.000
11	250.000	40	0	0	0.000
12	250.000	40	0	0	0.000
13	250.000	40	0	0	0.000
14	250.000	40	0	0	0.000
15	250.000	40	0	0	0.000
16	250.000	40	0	0	0.000
17	250.000	40	0	0	0.000
18	250.000	40	0	0	0.000
19	250.000	38	1	1	210.000

20	250.000	40	0	0	0.000
21	250.000	40	0	0	0.000
22	250.000	40	0	0	0.000
23	250.000	32	4	4	220.000
24	250.000	40	0	0	0.000
25	250.000	40	0	0	0.000
26	250.000	40	0	0	0.000
27	250.000	36	2	2	215.000
28	250.000	40	0	0	0.000
29	250.000	40	0	0	0.000

```
========================================================================

Benchmark : Hartstone Benchmark, Version 1.0
Compiler  : Verdix VAX/VMS -> MC68020, release 5.7
Target    : Motorola MVME133A (20 MHz MC68020 & 20 MHz MC68881)

Characteristics of best test for this experiment:
(no missed/skipped deadlines)

    Test 23 of Experiment 4

    Raw (non-tasking) benchmark speed in KWIPS:  506.33

    Full task set:

        Total       Deadlines       Task Set        Total
        Tasks       Per Second      Utilization     KWIPS
         27          119.00          85.32 %         432.00

    Highest-frequency task:

        Period      Deadlines          Task         Task
        (msec)      Per Second      Utilization     KWIPS
        62.500        16.00           3.16 %         16.00

    Experiment step size:    3.16 %

========================================================================
```

END OF HARTSTONE BENCHMARK SUMMARY RESULTS

Appendix J: Obtaining Hartstone Source Code and Information

Hartstone source code and supporting documentation can be obtained from the Real-Time Embedded Systems Testbed (REST) Project at the Software Engineering Institute in a number of different ways. Full details can be obtained by sending a request for information to the electronic mail or postal address listed below.

Electronic mail requests should be sent to the following Internet address:

<div align="center">

HARTSTONE-INFO@SEI.CMU.EDU

</div>

Electronic mail received at this address will automatically return to the sender instructions on all available distribution mechanisms.

For people who do not have Internet access, the address to send information requests to is:

REST Transition Services
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
Phone: (412) 268-7700

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION		1b. RESTRICTIVE MARKINGS		
UNCLASSIFIED		NONE		

2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT
N/A	APPROVED FOR PUBLIC RELEASE
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	DISTRIBUTION UNLIMITED
N/A	

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
CMU/SEI-90-TR-7	ESD-90-TR-208

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
SOFTWARE ENGINEERING INSTITUTE	SEI	SEI JOINT PROGRAM OFFICE

6c. ADDRESS (City, State and ZIP Code)	7b. ADDRESS (City, State and ZIP Code)
CARNEGIE MELLON UNIVERSITY PITTSBURGH, PA 15213	ESD/XRS1 HANSCOM AIR FORCE BASE, MA 01731

8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
SEI JOINT PROGRAM OFFICE	SEI JPO	F1962885C0003

8c. ADDRESS (City, State and ZIP Code)	10. SOURCE OF FUNDING NOS.			
CARNEGIE MELLON UNIVERSITY SOFTWARE ENGINEERING INSTITUTE JPO PITTSBURGH, PA 15213	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.
		N/A	N/A	N/A

11. TITLE (Include Security Classification)
Hartstone Benchmark Results and Analysis

12. PERSONAL AUTHOR(S)
Patrick Donohoe, Ruth Shapiro, Nelson Weiderman

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Yr., Mo., Day)	15. PAGE COUNT
FINAL	FROM _____ TO _____	June 1990	188

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR.	

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Hartstone is a series of timing requirements for testing a system's ability to handle hard real-time applications. It is specified as a set of processes with well-defined workloads and timing constraints. The name HARTSTONE derives from HArd Real Time and the fact that the workloads are based on the well-known Whetstone benchmark. This report describes the results obtained by running Version 1.0 of the Hartstone benchmark, an Ada implementation of one of the requirements, on a number of compiler/target processor combinations. The characteristics and expected behavior of the benchmark are described, actual results are presented and analyzed, and the lessons learned about the compilers and processors, and the benchmark itself, are discussed. Nothing in this report should be taken as an endorsement of, or an indictment of, a particular product. Users of Ada technology are encouraged to experiment with the Hartstone benchmark relative to their own particular application requirements.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS ☒	UNCLASSIFIED, UNLIMITED

22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE NUMBER (Include Area Code)	22c. OFFICE SYMBOL
KARL SHINGLER	(412) 268-7630	SEI JPO

DD FORM 1473, 83 APR EDITION OF 1 JAN 73 IS OBSOLETE.