

# Machine Characterization Based on an Abstract High Level Language Machine

*Rafael H. Saavedra-Barrera*<sup>†‡</sup>  
*Alan J. Smith*<sup>†</sup>  
*Eugene Miya*<sup>\*</sup>

## ABSTRACT

Runs of a benchmark or a suite of benchmarks are inadequate to either characterize a given machine or to predict the running time of some benchmark not included in the suite. Further, the observed results are quite sensitive to the nature of the benchmarks, and the relative performance of two machines can vary greatly depending on the benchmarks used. In this paper, we report on a new approach to benchmarking and machine characterization. The idea is to create and use a machine characterizer, which measures the performance of a given system in terms of a Fortran abstract machine. Fortran is used because of its relative simplicity and its wide use for scientific computation. The analyzer yields a set of parameters which characterize the system and spotlight its strong and weak points; each parameter provides the execution time for some primitive operation in Fortran.

We present measurements for a large number of machines ranging from small workstations to supercomputers. We then combine these measurements into groups of parameters which relate to specific aspects of the machine implementation, and use these groups to provide overall machine characterizations. We also define the concept of pershapes, which represent the level of performance of a machine for different types of computation. We introduce a metric based on pershapes that provides a quantitative way of measuring how similar two machines are in terms of their performance distributions. This metric is related to the extent to which pairs of machines have varying relative performance levels depending on which benchmark is used.

## 1. Introduction

One approach to comparing the CPU performance of different machines is to run a set of benchmarks on each. Benchmarking has the advantage that since real programs are being run on real machines, the results are valid, at least for that set of benchmarks; such results are much more believable than estimates produced from models of the system, no matter how detailed. To the extent that the benchmark set is representative of some target workload, the observed performance differences will reflect differences in practice.

Considerable effort has been expended to develop benchmark suites that are considered to reflect real workloads. Among them are the Livermore Loops [McM86], the NAS kernels [Bai85a, Bai85b]), and synthetic benchmarks (e.g. Dhrystone [Wei84, Wei88], Whetstone [Cur76]). Unfortunately, there are a number of shortcomings to benchmarking [Don87, Wor84]: (1) It is very

<sup>†</sup> Computer Science Division, EECS Department, University of California, Berkeley, California 94720.

<sup>‡</sup> Departamento de Ingenieria Electrica, Universidad Autonoma Metropolitana, México D.F., México.

<sup>\*</sup> NASA Ames Research Center, Moffett Field, Mountain View, California.

difficult to explain the benchmark results from the characteristics of the machines. (2) It is not clear how to combine individual measurements to obtain a meaningful evaluation of the various systems. (3) Given that there is almost never a good model of the machines being benchmarked, it is not possible to validate the results, nor to make predictions and/or extrapolations about expected performance for other programs. (4) Unless the benchmarks are tuned for each machine architecture, they may not take advantage of important architectural features. (5) The large variability in the performance of highly optimized computers is difficult to characterize with benchmarks. For example using benchmarks Harms et. al. found that the relative performance between the Fujitso VP-200 and the CRAY X-MP/22, varied from 0.41 to 5.39 on individual programs [Har88]; the ratio for the whole workload was only 1.12.

In this research, we present a new approach to characterizing machine performance. We do this via "narrow spectrum" benchmarking, by which we measure the performance of a machine on a large number of very specific operations, in our case, primitive operations in Fortran. This set of measurements characterizes each specific CPU. We separately analyze specific programs, ignoring at this stage of our research compiler optimizations and vector instructions. We can then combine the frequency of the primitive operations with their running times on various machines to predict the running time of any analyzed program on any analyzed machine. This approach also gives us considerable insight into both the machines and the programs, since the effects of individual parameters are immediately evident.

In this paper, we provide an overall presentation of this work, but concentrate on the specific issue of machine characterization: prediction of the execution time of benchmarks is done in [Saa88, Saa89]. Section 2 gives a somewhat more detailed overview of our research. In section 3, we describe the machine characterizer, and also the program analyzer. The parameters used to characterize a machine are explained in section 4. The methodology used to make measurements is presented in section 5, and the parameters derived from a number of machines are given in section 6. A comparison of machines is also provided in that section. The concepts of performance distributions (pershaps) and pershape distances between machines are given in section 7. Some unresolved issues are considered in section 8.

## 2. System Characterization and Performance Evaluation

The idea behind our approach is to distinguish between two different activities often ignored in machine evaluation; these are system characterization and performance evaluation. We define *system characterization* as an  $n$ -value vector where each component represents the performance of a particular operation ( $P_i$ ). This vector ( $\langle P_1, P_2, \dots, P_m \rangle$ ) fully describes the system at some level of abstraction. The parameters we use are a set of primitive operations, as found in the Fortran programming language, and are defined in section 4. We measure the values of the parameters using a *system characterizer*, which runs a set of 'software experiments', which detect, isolate and measure the performance of each basic operation. Using software experiments to measure each system allows us to validate our model and measurements by making predictions and checking the results with the execution of benchmarks and workloads. This approach is in contrast to studies which use a low-level machine architecture based model [Pou77].

The *performance evaluation* of a group of systems is the measurement of some number of properties during the execution of some workload. One property may be the total execution time to complete some job. It is important to note that the results depend, and are only valid, for the set of programs used in the evaluation. The evaluation includes not only the machine, but also the compiler, the operating system and the libraries. In this research we focus on the execution time of computationally intensive programs as our metric for evaluating different architectures.

### 2.1. A Linear Model for Program Execution

Our research is based on the assumption that the execution time of a program can be partitioned into independent time intervals, each corresponding to the execution of some operation of an abstract Fortran machine. The abstract Fortran machine (AFM) is used as a general model for all the machines, each executing the object code produced by their Fortran compilers. Thus,

each system represents a different implementation of the AFM. In this way, the AFM makes it possible to compare different architectures. As is shown in [Saav88], and to a lesser extent later in this paper, this assumption is reasonably accurate.

System designers use a similar approach, but at the hardware level, when they evaluate different implementations of the same architecture. In this case the model of the machine is defined by its instruction set, and they are interested in the mean instruction execution time [Mac84]. This quantity is equal to the sum of the mean nominal execution time, the mean pipeline delay caused by path conflicts and data dependencies, and the mean storage access delay caused by cache misses of instructions and operands. In our case, instead of having one single machine instruction to measure, we have a group of instructions corresponding to an abstract parameter. How each abstract parameter is implemented in each machine depends on its instruction set, compiler, and libraries. In fact, normally there will be several sequences of instructions implementing each parameter. Which particular sequence is generated by the compiler depends on the context in which the operation appears in the source program.

Our model of the total execution time is the following: Let  $\mathbf{P}_M = \langle P_1, P_2, \dots, P_n \rangle$  be the set of parameters that characterize the performance of machine  $M$ . Let  $\mathbf{C}_A = \langle C_1, C_2, \dots, C_n \rangle$  be the normalized dynamic distribution of operations in program  $A$ , and let  $C_{total}$  denote the total number of operations executed in program  $A$ . We obtain the expected execution time of program  $A$  on machine  $M$

$$T_{A,M} = C_{total} \sum_{i=1}^n C_i P_i = C_{total} \mathbf{C}_A \cdot \mathbf{P}_M \quad (1)$$

where

$$\sum_{i=1}^n C_i = 1$$

In general, given machines  $M_1, M_2, \dots, M_m$ , with characterizations  $\mathbf{P}_{M_1}, \mathbf{P}_{M_2}, \dots, \mathbf{P}_{M_m}$ , and a workload  $W$  formed by programs  $A_1, A_2, \dots, A_l$  with dynamic distributions  $\mathbf{C}_{A_1}, \mathbf{C}_{A_2}, \dots, \mathbf{C}_{A_l}$ , the expected execution time of machine  $M_k$  on workload  $W$  is

$$T_{W,M_k} = \sum_{j=1}^l C_{total,A_j} \mathbf{C}_{A_j} \cdot \mathbf{P}_{M_k} \quad (2)$$

where  $C_{total,A_j}$  is the total number of operations executed in program  $A_j$ .  $T_{W,M_i}$  provides a way to make a direct comparison between several machines with respect to workload  $W$ . The expected execution time of workload  $W$  on machine  $M_i$  is less than that of machine  $M_j$ , if  $T_{W,M_i} \leq T_{W,M_j}$ .

Using this model it is possible not only compare two different machine architectures using any workload, but also to explain their results in terms of the abstract parameters. Let  $\Phi_{M,A} = \langle \phi_1, \phi_2, \dots, \phi_n \rangle$  be the normalized distribution of the execution time for program  $A$  executed on machine  $M$ . Define:

$$\phi_i = \frac{C_{A,i} P_{M_k,i}}{\mathbf{C}_A \cdot \mathbf{P}_M}$$

Vector  $\Phi_{M,A}$  decomposes the total execution time in terms of each parameter and makes it possible to identify which operations are the most time consuming. We would expect that different machines will have different distributions, even for different implementations of the same architecture or/and different compilers. Once we have the machine characterizations, it is possible to study the effect of changes in the normalized dynamic distribution without writing real programs that correspond to these distributions, and in this way detect which parameters have a significant impact in the execution time for some machines.

An advantage of this scheme is that the  $l \cdot m$  machine-program combinations only require that each machine be measured once to obtain its characterization, and also that each program be analyzed once. Making an evaluation using normal benchmarking techniques requires the

execution of *l-m* programs. Moreover, once the machine has been measured, its characterization can be used at any time in the future for additional evaluations, in contrast to benchmarking in which access to the machine (same model, operating system, compiler, libraries) is needed for each new set of benchmarks.

## 2.2. Limits of the Linear Model

The only way in which the linear model can give acceptable results is if the following conditions hold: (1) The experimental measurements are representative of 'typical' occurrences of the parameters in real programs. (2) The errors caused by the low resolution and the intrusiveness of the measuring tools are small compared to the magnitude of the measurements. (3) Variability in the execution mean time caused by data dependencies, external concurrent activity, and nonreproducible conditions is small, and therefore does not significantly affect the results. In some cases, the above conditions cannot be satisfied, especially in highly pipelined machine where the execution time when there is a register dependency conflict is several times greater than the execution time without this delay. An example of this is the CYBER 205, where an add or multiply can take as little as 20 ns to execute, when the pipeline is full, or as much as 100 ns in the worst case [Ibb82]. If we consider the following two statements

$$\begin{aligned} X9 &= ((X1 + X2) * (X3 + X4)) + ((X5 + X6) * (X7 + X8)) \\ X6 &= ((X1 + X2) * X3 + X4) * X5 \end{aligned}$$

we find that the execution of the first statement takes approximately 360 ns, while the execution time of the second takes 400 ns. A simple linear model will estimate that the execution of the second statement will be less than that of the first statement, unless the model contains information on how the execution time is affected by data dependencies. Branching and interrupts also prevent the pipeline from working at peak speed. Although it is difficult to detect and measure how each machine will execute different statements, it is always possible to create new parameters that take into account data dependencies and measure the extra penalty in the execution time. In practice, the number of parameters cannot be expanded without limit.

## 2.3. Fortran and Other Programming Languages

The model presented above can also be applied to other general purpose languages. We chose Fortran instead of other programming language for the following reasons: 1) most large scale scientific computation, accounting for most of the CPU time on supercomputers, is done in Fortran; 2) the number of language constructs in Fortran is small; and 3) the execution time of most of the operations in Fortran does not depend on the value of the arguments. It is therefore natural to experiment first with a less complex programming language and test whether it is possible to make acceptable predictions. Most of the differences between Fortran and other general purpose languages do not prevent building an abstract machine model, although a model with a larger number of parameters and better experiments would be required.

## 3. Description of the System

In the last section we showed what we need in order to characterize machines using the linear model and how to use this information to make predictions about the execution time of programs. We have implemented (a) a system characterizer and assembled a library of machine characterizations ( $P_M$ ); (b) a program analyzer that generates the dynamic distribution ( $C_A$ ) and the total number of operations ( $C_{total}$ ) of Fortran programs; and (c) an execution predictor that takes  $P_M$ ,  $C_A$  and  $C_{total}$  and estimates the expected execution time of the applications. The complete process, characterization, analysis, and prediction is shown in figure 1. In the next two subsections we give an overview of the program analyzer and the execution predictor. A more in depth presentation of the system characterizer follows.

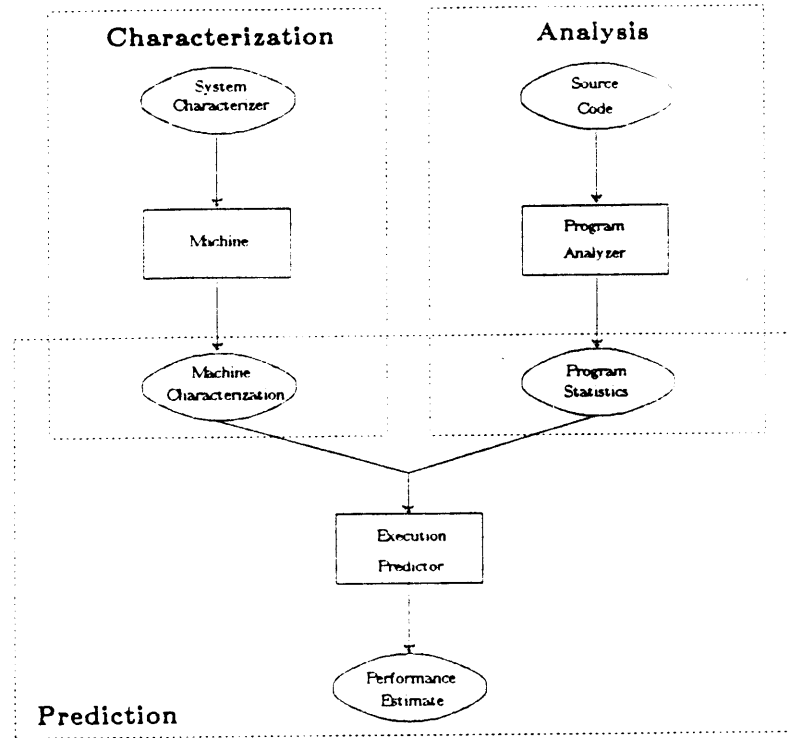


Figure 1: The process of characterization, analysis and prediction.

### 3.1. Program Analyzer

The program analyzer (PA) decomposes Fortran programs statically and dynamically in terms of the abstract parameters. This provides a uniform model for the execution of different applications. In addition both models, the performance model associated with machines, and the execution model associated with the applications are identical. Thus, it is possible using the dynamic distribution to compare different programs, putting the emphasis not on their syntactic or semantic properties, but in how they affect the performance of different systems.

The PA is basically the front end of a Fortran compiler. It takes as its input a Fortran program and after making a lexical and syntactical analysis, it outputs an instrumented version of the original program, from which we obtain the dynamic statistics. In addition, the PA also gives the static statistics for each parameters for each basic block. The operation of the program analyzer is shown in figure 2.

Let us number each of the basic blocks of the program  $j=1,2,\dots,m$ , and let  $s_{ij}$  ( $i=1,2,\dots,n$ ) designate the number of static occurrences of parameter  $P_i$  in block  $B_j$ . Matrix  $\mathbf{S}_A=[s_{ij}]$  of size  $n \times m$  represents the complete static statistics of the program. Let  $\mu_A=\langle \mu_1, \mu_2, \dots, \mu_j \rangle$  be the number of times each basic block is executed, then matrix  $\mathbf{D}_A=[d_{ij}]=[\mu_j \cdot s_{ij}]$  gives us the dynamic statistics. Matrix  $\mathbf{S}_A$  and vector  $\mu_A$  are obtained by parsing and instrumenting the source code. Vector  $\mathbf{C}_A$  (§ 2.1) and matrix  $\mathbf{D}_A$  are related by the following equations

$$C_i = \frac{\sum_{j=1}^m d_{i,j}}{C_{total}} \quad (3)$$

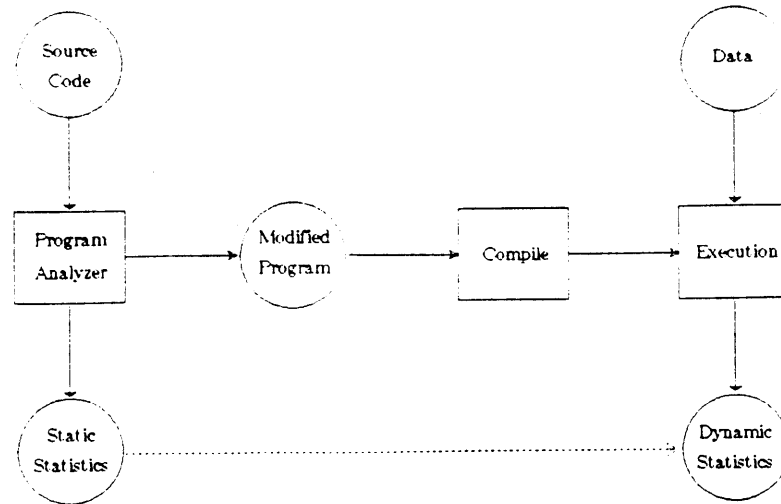


Figure 2: The static and dynamic statistics are obtained by parsing and instrumenting the source program.

---

and

$$C_{total} = \sum_{i=1}^n \sum_{j=1}^m d_{i,j}$$

The dynamic statistics are independent of the code generated by each compiler, and they only depend on the source code and the data used in the execution.

### 3.2. Execution Predictor

The execution predictor (EP) combines the machine characterization  $\mathbf{P}_M$  with the dynamic statistics  $\mathbf{C}_A$  to obtain estimates of the expected execution time of programs, using equations (1) and (3). The execution time is computed for each statement of the program, and this makes it possible to compute estimates for different parts of the program. In our system, this is done by inserting two special comments at the beginning and end of the particular region in which we are interested. There is no limit to the number of these regions as long as they are either disjoint or one is contained in the other. In addition to the expected execution time, the EP also reports the variance of the estimate, and the expected execution time per parameter along with its variance. This is done by measuring vector  $\sigma^2 \mathbf{P}_M = \langle \sigma^2 P_1, \sigma^2 P_2, \dots, \sigma^2 P_n \rangle$  with the system characterizer.

## 4. The Fortran Abstract Machine

By using a common parametric model, we are able to compare the performance of different architectures and make a fair comparison between them with respect to their execution of Fortran programs. How many parameters the system should have depends on how accurate we want our predictions to be, although this is limited by the resolution of our measuring tools; at some point increasing the number of parameters does not have any effect in improving our predictions. We are also limited in our accuracy by the fact that we do not analyze the code generated by the compiler, nor use information that is not contained in the source code of the program. In the next subsection we present the set of parameters in our model, and give a brief description of what they measure.

#### 4.1. Parameters in the System Characterizer

Each parameter of the model can be classified in one of the following broad categories: arithmetic and logical, procedure calls, array references, branching and iteration, and intrinsic functions. We decided which parameters to include in our model in an iterative manner. Initially we associated parameters with obvious basic operations, and after a first version of the system was running, new parameters were incorporated to distinguish between different uses and execution times of the 'same' abstract operation in the program. This was mainly the result of detecting a significant error between our predictions and real execution times. Although every basic operation in Fortran is characterized by some parameter, we have made some simplifications with some operations which were rarely executed in the benchmarks we used. It is straightforward to include new parameters in the model, and to write new experiments for the system characterizer. The parameters are classified in eighteen different groups according to the semantics of the operation; tables 1 and 2 present the 102 parameters.

#### 4.2. Global vs. Local Variables

Most operators are characterized by several parameters, depending on the operand types and sizes, and storage class (common / local). Global variables in Fortran (COMMON) are sometimes treated differently from local variables. In some compilers, variables stored in COMMONs are treated as components of a structure using a base-descriptor for each COMMON block which points to the first element of the COMMON. An operand is loaded by first adding an offset to the base-descriptor and then loading the operand. This way of treating simple variables makes the execution slower when they are allocated as global variables as opposed to local.

#### 4.3. Arithmetic and Logical Operations

Fortran is a language for scientific and numeric applications. For this reason the richness of the language lies in the arithmetic operators that it supports. In addition to the arithmetic operators, Fortran also provides six relational and six logical operators. Table 1 (groups 1-8) gives the fifty-six arithmetic parameters grouped by data type (real, complex, or integer), size (single or double precision) and storage class (local or global). Not all combinations are included; there is no double precision for integers or complex variables. In table 2 we find the logical and conditional parameters (groups 9-10). Each arithmetic mnemonic is formed by appending the first letter of the operation with the first letter of the data type (R, C, or I), plus a letter identifying the size of the operand (S or D), and at the end the first letter of the storage class (L or G).

The arithmetic operators defined in the system characterizer are: addition, multiplication, division (quotient), and exponentiation. The addition operator also includes subtraction. Addition (subtraction) between an array index and a constant is treated as a special parameter: most compilers add, at compile time, the constant to the base-descriptor of the array and eliminate the addition (see 4.6). In the case of exponentiation using a real base, we distinguish between two cases: one when the exponent is integer and the other when the exponent is real. When the exponent is integer, the result is computed by either executing the same number of multiplications as in the exponent (when this is small), or by binary decomposition. In the case of a real exponent, the result is computed using logarithms. If the base is integer, we have two cases, one with the exponent equal to two and another when the exponent is greater than two. Given that the number of exponentiations executed in most programs is small, these simplifications are sufficient.

Two groups of parameters require further explanation. One is the set of parameters that measure the overhead of the store operation (SRSL, SCRL, SISL, SRDL, SRSG, SCSG, SISG, and SRDG), and the other, what we called memory transfer parameters (TRSL, TCSL, TISL, TRDL, TRSG, TCSG, TISG, and TRDG). In Fortran loading an operand is not a separate operation, and so the time for a load is included in the execution time of the corresponding arithmetic or logical operation. A store, however, is a visible operation whose time can be measured. In some cases, the time for a store is negligible, (part of the store operation overlaps with the execution of the previous or/and following operation), while in other cases the time can be significant.

Table 1: Parameters in the system characterizer (part 1 of 2)

1 real operations (single, local)		4 real operations (single, global)	
01 SRSL	store	29 SRSG	store
02 ARSL	addition	30 ARSG	addition
03 MRSL	multiplication	31 MRSG	multiplication
04 DRSL	division	32 DRSG	division
05 ERSL	exponential ( $X^I$ )	33 ERSG	exponential ( $X^I$ )
06 XRSL	exponential ( $X^Y$ )	34 XRSG	exponential ( $X^Y$ )
07 TRSL	memory transfer	35 TRSG	memory transfer
2 complex operations, local operands		5 complex operations, global operands	
08 SCSL	store	36 SCSG	store
09 ACSL	addition	37 ACSG	addition
10 MCSL	multiplication	38 MCSG	multiplication
11 DCSL	division	39 DCSG	division
12 ECSL	exponential ( $X^I$ )	40 ECSG	exponential ( $X^I$ )
13 XCSL	exponential ( $X^Y$ )	41 XCSG	exponential ( $X^Y$ )
14 TCSL	memory transfer	42 TCSG	memory transfer
3 integer operations, local operands		6 integer operations, global operands	
15 SISL	store	43 SISG	store
16 AISL	addition	44 AISG	addition
17 MISL	multiplication	45 MISG	multiplication
18 DISL	division	46 DISG	division
19 EISL	exponential ( $I^2$ )	47 EISG	exponential ( $I^2$ )
20 XISL	exponential ( $I^J$ )	48 XISG	exponential ( $I^J$ )
21 TISL	memory transfer	49 TISG	memory transfer
7 real operations (double, local)		8 real operations (double, global)	
22 SRDL	store	50 SRDG	store
23 ARDL	addition	51 ARDG	addition
24 MRDL	multiplication	52 MRDG	multiplication
25 DRDL	division	53 DRDG	division
26 ERDL	exponential ( $X^I$ )	54 ERDG	exponential ( $X^I$ )
27 XRDL	exponential ( $X^Y$ )	55 XRDG	exponential ( $X^Y$ )
28 TRDL	memory transfer	56 TRDG	memory transfer

Table 1: The set of parameters measured by the system characterizer (part 1 of 2). Arithmetic operations are classified taking into account the type, width and storage class of their operands.

Loads are visible and must be accounted for in one case. In a memory transfer statement where there are no operators on the right hand side of the equal sign, the execution time of the statement cannot be explained just by the store operation. For these kind of statements we use the memory transfer parameters to distinguish them from normal statements that include expressions.



Table 2. Parameters in the system characterizer (part 2 of 2)

9 logical operations (local)		10 logical operations (global)	
57 ANDL	AND & OR	62 ANDG	AND & OR
58 CRSL	compare, real, single	63 CRSG	compare, real, single
59 CCSL	compare, complex	64 CCSG	compare, real, double
60 CISL	compare, integer, single	65 CISG	compare, integer, single
61 CRDL	compare, real, double	66 CRDG	compare, real, double
11 function call and arguments		13 branching parameters	
67 PROC	procedure call	69 GOTO	simple goto
68 AGRS	argument load	70 GCOM	computed goto
12 references to array elements		14 DO loop parameters	
71 ARR1	array 1 dimension	75 LOIN	loop initialization (step 1)
72 ARR2	array 2 dimensions	76 LOOV	loop overhead (step 1)
73 ARR3	array 3 dimensions	77 LOIX	loop initialization (step n)
74 IADD	array index addition	78 LOOX	loop overhead (step n)
15 intrinsic functions (real)		16 intrinsic functions (double)	
79 LOGS	logarithm	87 LOGD	logarithm
80 EXPS	exponential	88 EXPD	exponential
81 SINS	sine	89 SIND	sine
82 TANS	tangent	90 TAND	tangent
83 SQRS	square root	91 SQRD	square root
84 ABSS	absolute value	92 ABSD	absolute value
85 MODS	module	93 MODD	module
86 MAXS	max. and min.	94 MAXD	max. and min.
17 intrinsic functions (integer)		18 intrinsic functions (complex)	
95 ABSI	absolute value	98 LOGC	logarithm
96 MODI	module	99 EXPC	exponential
97 MAXI	max. and min.	100 SINC	sine
		101 SQRC	square root
		102 ABSC	absolute value

**Table 2:** The set of parameters measured by the system characterizer (part 2 of 2). Each standard intrinsic function in FORTRAN is represented by one parameter. For example, sine and cosine functions are characterized by the same parameter SINx (x can be S, D or C, depending on the data type of the argument).

#### 4.4. Procedure Calls and Arguments

Function and subroutine call overhead also significantly affect the execution time of programs. This overhead can be identified with three actions: passing the arguments between procedures, transferring control from the caller to the callee, and returning the result and control from the callee to the caller. In Fortran all the arguments are passed by reference including values computed by expressions, so the argument passing overhead is limited to setting up the reference.

We use two parameters which characterize function and procedure calls: the first measures the joint execution of the prologue and epilogue of the call (PROC), and the second the time it takes to load the address of an argument, either into registers, the static environment of the callee subprogram or in the execution stack (some systems use the execution stack or registers to pass arguments between subprogram units).

#### 4.5. Branch and Loop Control

Branches (table 2, group 13) affect the execution of a program in several ways. In pipelined machines, a penalty must be paid when a branch is taken and the target instruction has not been previously fetched<sup>1</sup>, since all partially executed instructions in the pipeline must be discarded and the new stream of instructions must be fetched [LEE84]. For a machine with a cache memory, a branch to an instruction that is not in the cache may increase the miss ratio by changing the locality of execution [SMI82].

Although Fortran has three different types of GO TO statements: unconditional, assigned, and computed, we characterize all three with only two parameters. We make a distinction between a direct jump (GOTO) and a computed jump (GCOM). In the first case the target of the jump is known by the compiler and it is normally implemented as an unconditional jump. In the second case the target depends on the value of some expression and it should be computed before the branch can be executed. In our model, parameter GOTO is used for unconditional GO TOs, and GCOM for computed and assigned GOTOs.

DO loops (group 14) are flow of control constructs that are widely used in scientific programs. The overheads incurred by this instruction are the time to set the initialization, limit and step of the control index, and the time it takes to update the index and check if the number of iterations has been completed. In addition, some machines implement loops with unit step differently from non-unit step loops. In some cases the loop is transformed to a loop with a unit step, which sometimes increases the execution time in non-unit loops. Four parameters (LOIN, LOOV, LOIX, LOOX) characterize loops with unit and non-unit increment.

Although Fortran has three different types of IF statements, neither of these has been found to need special parameters for their characterization. The block IF and the logical IF are decomposed in two parts: the evaluation of the predicate (arithmetic-logical expression) and a direct branch. The arithmetic IF is different only in the way it branches. We handle the branching part as a computed GOTO (GCOM).

#### 4.6. Array References

Array variables in expressions are treated as ordinary variables plus an additional overhead to compute the address of the element. We have three parameters (ARR1, ARR2, and ARR3) that characterize the dimension of an array reference (group 12). The overhead for variables in four and five dimensions is computed in the execution predictor using a linear combination of the three basic parameters. We found that most of the applications we examined had very few arrays with more than three dimensions and no examples of more than five.

As we mentioned in section 4.2, integer addition between a variable and a constant in an array index is considered a special operation. Initially we treated them as normal integer adds, but the predictions thus obtained were off significantly from the measured times. We fixed this problem by creating a new parameter (IADD) that measures the execution time of an add using an index and a constant as operands.

---

<sup>1</sup> Even in machines that have some kind of branch prediction circuitry, a penalty is incurred when the prediction is incorrect.

#### 4.7. Intrinsic Functions

Intrinsic functions form the last subset of parameters (groups 15-18). Although the number of times these instructions are executed in a program is small, and they only occur in some workloads, their execution times tend to be very large compared with that of a single arithmetic operation. We have twenty-four parameters in four groups which represent the intrinsic functions most used in scientific programs.

The execution time of an intrinsic function is not always constant and normally depends on the magnitude of the arguments. As an example, consider how the IBM 3090/200 computes the sine function [IBM87]. In the computation, the execution time of several steps depends not only in how large is the argument, but also in how small is its difference from the nearest multiple of  $\pi$ . Depending on the magnitude of this difference, a polynomial of degree one, three, five, or a table and additional arithmetic is needed to compute the result. The frequency of intrinsic functions is generally low, and the arguments unpredictable, and we have found that our assumption of constant execution time is a good enough approximation.

### 5. Machine Characterizer

The machine characterizer (MC) consists of 102 'software experiments' that measure the performance of each individual parameter needed to completely characterize a Fortran machine (see table 1). The MC is written as a Fortran program and runs from 200 seconds, on machine with good clock resolution, to 2000 seconds on machines with 1/60'th second clock resolution. We have run the MC on several machines ranging from low-end workstations to supercomputers. Each experiment tries to measure the execution time that each parameter takes to execute in 'typical' Fortran programs. This 'typical' execution time was obtained by looking at real programs and also by modifying those experiments that were identified as generating the biggest error in our predictions.

#### 5.1. Experiment Structure

Timing a benchmark is very different from making a detailed measurement of the parameters in the system characterizer. For some benchmarks the system clock is enough for timing purposes, and repetition of the measurements normally yields an insignificant variance in the averaged results. On the other hand, the measurement of the parameters in the system characterizer is more difficult due to a number of factors:

- The short execution time of most operations (20 ns - 10  $\mu$ s)
- The resolution of the measuring tools ( $\geq 1 \mu$ s)
- The difficulty of isolating the parameters using a program written in Fortran
- The intrusiveness of the measuring tools
- Variations in the hit ratio of the memory cache
- External events like interrupts, multiprogramming, and I/O activity
- The need to obtain repeatable results and accuracy

Most of our primitive operations have execution times of from ten to thousands of nanoseconds, and are implemented with a single or a small number of machine instructions. For this reason direct measurement is not possible, especially since our tests should work for many different architectures. In addition, the need to isolate an operation for measurement normally requires robust tests to avoid optimizations<sup>2</sup> from the compiler that would eliminate the operation

<sup>2</sup> Even when we compile without optimization, compilers try to apply some standard optimizing techniques, like constant folding, short-circuiting of logical expressions, and computing the address of an element in an array.

from the test and distort the results [CLA86]. Different techniques must be used, in particular avoiding the use of constants inside the test loops; using IF and GO TO instructions instead of the DO LOOP statements to control the execution of the test; and initializing variables in external procedures to avoid constant folding. Separate compilation of variable initialization procedures is used to make sure that the body of the test does not give enough information to the compiler to eliminate the operation being measured from inside the control test loop.

## 5.2. Test Structure and Measurement

The measurement tools we have are the system clock and the repeated execution of a sequence of statements. The resolution of the clock, the overhead of the timing routine and the overhead of the statements that control of measurements are the sources of error that we can control or work around. Variations in the hit ratio of the cache, interrupts, multiprogramming and I/O activity are more difficult to eliminate and measure (see § 8).

We use three different methods to measure the execution time of the parameters. The first is by *direct* measurement, i.e. executing some operation for some number of times and in different contexts. The second is with a *composite* measurement. In this case we execute a number of different operations and subtract the execution time of the known parameters to obtain the value of the one that is unknown. The third possibility is with an *indirect* measurement. Some parameters of the model are 'coupled'; it is not possible to execute one without executing the other. The way to measure one of the parameters is to run two or more tests with a different number of operations; the solution of a set of linear equations gives the correct result. Figure 3 shows the basic structure of our tests. This same structure is used in all the tests.

---

```

LIMIT = LIMIT0 * SPEEDUP * (TMAX - TMIN) / 2.
DO 4 K = 1, REPEAT
1  COUNTER = 1
   TIME0 = SECOND ()
2  IF (COUNTER .GT. LIMIT) GO TO 3
   ...
   body of the test
   ...
   COUNTER = COUNTER + 1
   GO TO 2
3  TIME1 = SECOND ()
   IF (TIME1 - TIME0 .GE. TMIN .AND. TIME1 - TIME0 .LE. TMAX) GO TO 4
   LIMIT = .5 * LIMIT * (TMAX - TMIN) / (TIME1 - TIME0)
   GO TO 1
4  SAMPLE(K) = TIME1 - TIME0
   CALL STAT (REPEAT, SAMPLE, AVE, VAR)

```

Figure 3: The basic structure of an experiment. The statement IF (TIME1 - TIME0 enforces the execution of each test for more than TMIN and less than TMAX seconds. If the execution is outside this interval a new value of LIMIT is computed and the test is repeated.

---

The sequence of statements to measure correspond to the 'body of the test'. These statements are executed for some number of times (LIMIT) and the execution time is measured (function SECOND). This time is called an observation. TMIN, and TMAX control the minimum and maximum time that each observation should run ( $t_{\min} \leq time_1 - time_0 \leq t_{\max}$ ). The two statements before the GO TO 1 enforce this condition. The DO loop is used to get several (REPEAT) observations to obtain a meaningful statistic. Because we don't know a priori how fast or slowly an operation executes in an arbitrary machine, we extrapolate by using the time it takes to run the test in the CRAY X-MP/48 and multiplying by their relative speeds. This is done using LIMIT0, which is the number of times the test runs in the CRAY X-MP/48, and SPEEDUP that

gives the relative speed of the machine. The relative speed is computed by running a small test at the beginning of the characterization.

### 5.3. Experimental Error and Confidence Intervals

There are many known sources for the variability of the CPU time [Cur75, Mer83] and consequently in our measurements. Some of these factors are: timer resolution of the clock, improper allocation of CPU cpu for I/O interrupt handling, cycle stealing, and changes in cache hit ratios due to interference with concurrent tasks. Small errors in the measurements have considerable impact in the predictions we make, and we must measure and compensate for them. We will proceed to derive expressions for the variance and the confidence intervals of the measurements. The following definitions are used in the analysis:

$T_{j_0}$	::=	CPU time before the observation (TIME0)
$T_{j_1}$	::=	CPU time after the observation (TIME1)
$C_{overhead}$	::=	overhead involved in the timing function
$IF_{overhead}$	::=	overhead involved in the if-loop control
$N_{limit}$	::=	number of times the body is executed (LIMIT)
$N_{repeat}$	::=	number of observations in the experiment (REPEAT)
$O_j$	::=	observation $j$ , equal to TIME1 - TIME0
$\hat{O}$	::=	sample mean of each observation (measurement)
$\hat{B}$	::=	sample mean time of one 'body of the test' execution
$\hat{P}_i$	::=	sample mean of parameter $i$
$\sigma^2$	::=	variance operator

We know that each observation  $O_j$  is equal to

$$O_j = T_{j_1} - T_{j_0}$$

then the mean value ( $\hat{O}$ ) and variance of these observations are

$$\hat{O} = \frac{1}{N_{repeat}} \sum_{j=1}^{N_{repeat}} O_j; \quad \sigma^2 O = \frac{1}{N_{repeat} - 1} \sum_{j=1}^{N_{repeat}} (O_j - \hat{O})^2 \quad (4)$$

Now the mean value of each observation is equal to the time it takes to execute the body of the test  $N_{limit}$  times, plus the overhead of the timing function ( $C_{overhead}$ ), and the extra instructions that control the test ( $IF_{overhead}$ ).

$$\hat{O} = N_{limit} (\hat{B} + IF_{overhead}) + C_{overhead}$$

where  $\hat{B}$  is the mean time it takes to execute once the body of the test. We can compute this value and the variance with the equations

$$\hat{B} = \frac{\hat{O} - C_{overhead}}{N_{limit}} - IF_{overhead}; \quad \sigma^2 B = \frac{\sigma^2 O + \sigma^2 C_{overhead}}{N_{limit}^2} + \sigma^2 IF_{overhead} \quad (5)$$

To obtain the mean value of parameter  $\hat{P}_i$  we need to know if the test is direct, composite or indirect. Let  $N$  be the number of times parameter  $\hat{P}_i$  is executed inside the body of the test, then the mean value and variance of parameter  $\hat{P}_i$  in a direct test are

$$\hat{P}_i = \frac{\hat{B}}{N}; \quad \sigma^2 P_i = \frac{\sigma^2 B}{N^2} \quad (6)$$

In a composite test we have

$$\hat{P}_i = \frac{\hat{B} - W_{extra}}{N}; \quad \sigma^2 P_i = \frac{\sigma^2 B + \sigma^2 W_{extra}}{N^2}$$

where  $W_{extra}$  is the additional work inside the body of the test or in the second test. In an indirect test  $\hat{P}_i$  is a function of several measurements.

$$\hat{P}_i = f(\hat{B}_1, \hat{B}_2, \dots, \hat{B}_n)$$

The normalized 90% confidence intervals are given by the expression

$$\left[ \hat{P}_i - t_{.95} \left( \frac{\sigma^2 P_i}{N_{repeat}} \right)^{1/2}, \hat{P}_i + t_{.95} \left( \frac{\sigma^2 P_i}{N_{repeat}} \right)^{1/2} \right] \quad (7)$$

where  $t_{.95}$  corresponds to the 95% percentile of the Student's  $t$  distribution. Looking at equations 4-7 we see that by increasing  $N$ ,  $N_{limit}$ , and  $N_{repeat}$ , we can reduce the variance in our measurements.

#### 5.4. Is the Minimum Better than the Average?

The measurements are obtained by computing the average of a number of observations. The sample represent the 'effective' execution time of the experiment plus an additional random variable that represents the 'noise' produced by concurrent activity and the resolution of our measuring tools. The error produced by the clock can be modeled as a random variable with mean zero and standard deviation of 1/6 times the resolution of the clock. The concurrent activity is a positive random variable that depends on the load of the system. The 'effective' execution time must be less than or equal to (ignoring clock resolution) any observation in the sample. Therefore instead of taking the average as our measurement, it might be better to take the minimum of the sample. However, what we are trying to characterize is the execution time of programs under real conditions, and in this case the average would correspond better to these 'normal' conditions.

Table 3: Estimates taking the average and the minimum

Machine	Real Time	Average	Error	Minimum	Error
Convex C-1	543 sec	551 sec	1.47 %	499 sec	8.10 %

We decided which approach was better by using both schemes (minimum and average) in the system characterizer and using the measurements for the prediction of the execution time of a set of ten programs. The results (table 3) for one of the machines tested showed that using the average produces a better estimate. The programs and both system characterizers were run under similar conditions.

#### 5.5. Reducing the Variance

Measuring small execution times from Fortran programs is difficult, especially for parameters that are measured using indirect tests. In a direct or composite test it is relatively easy to increase the number of operations executed inside the body of the test, and in this way reduce the variance (eq. 6). In the case of an indirect test, the parameters measured are coupled with other parameters and we cannot increase the execution of one without also increasing one or more of the others.

An example is the address computation for an array element. This parameter is measured by running some code using simple variables and then subtracting the running time from the execution time for the same code using array elements. Increasing the number of array elements in the test also increases the number of operations executed in the test. Moreover the variance of the difference of two random variables is the sum of the individual variances. Therefore for some parameters this produces an increase in variance, not a reduction. The way to reduce the variance in these cases is to increase the number of observations ( $N_{repeat}$ ) and/or the number of times the body is executed ( $N_{limit}$ ). How difficult it is to control the variance of some parameters can be seen in the following example. In this case parameters LOIN (loop initialization) and LOOV (loop overhead) are computed from the results of three tests using equations 8 and 9. Each  $\hat{B}_i$  represents the result obtained in one of the three tests used in the measurement of LOIN and LOOV. Even when the sample standard deviation is small ( $< 5\%$ ) for the  $\hat{B}_i$ s, in the case of the

Table 4: Mean and Standard Deviation			
Parameter	Mean ( $\mu$ )	Std. Dev. ( $\sigma$ )	$\sigma/\mu$ (%)
$\hat{B}_1$	69.9 $\mu$ s	1.49 $\mu$ s	2.13
$\hat{B}_2$	127.3 $\mu$ s	5.74 $\mu$ s	4.51
$\hat{B}_3$	107.4 $\mu$ s	4.53 $\mu$ s	4.22
LOIN	12.5 $\mu$ s	9.11 $\mu$ s	72.9
LOOV	1.99 $\mu$ s	0.73 $\mu$ s	36.8

Table 4: Relative magnitude of the standard deviation compared to the sample mean for parameters LOIN and LOOV. The  $\hat{B}_i$ s are experimental results used to compute the value of LOIN and LOOV. Each test consists of 5 observations executed for 1 second on a VAX-11/785.

DO loop parameters it is very large with respect to the mean. The three  $\hat{B}_i$  represent the performance in each of the three experiments used to obtain parameters LOIN and LOOV.

$$\text{LOIN} = 2\hat{B}_1 - \hat{B}_2; \quad \sigma^2\text{LOIN} = 4\sigma^2B_1 + \sigma^2B_2 \quad (8)$$

$$\text{LOOV} = \frac{\hat{B}_2 - \hat{B}_3}{N}; \quad \sigma^2\text{LOOV} = \frac{\sigma^2B_2 + \sigma^2B_3}{N^2} \quad (9)$$

### 5.8. The Effect of $N_{limit}$ and $N_{repeat}$ on the Variance

One question we haven't answered is what should be the magnitude of  $N_{limit}$  and  $N_{repeat}$  to obtain measurements which give a small  $\sigma/\mu$  ratio. These parameters are system dependent and are mainly affected by the resolution of the clock, the concurrent activity on the system, and the particular parameter being measured. We ran several experiments using different values for  $N_{repeat}$  and  $N_{limit}$  in several machines. Figure 4 shows the normalized confidence interval of ten parameters for values of  $N_{limit}$  such that the each test is run for at least 0.1, 0.2, 0.5, 1.0, 2.0 and 4.0 seconds on a Vax-11/780. We also obtained measurements for  $N_{repeat}$  equal to 5, 10 and 20 observations.

We see that for a fixed value of  $N_{repeat}$  the width of the confidence interval of our measurements decreases as the time of the test increases, but for small values of  $N_{repeat}$ , there is a limit to how much we can decrease the confidence interval by only increasing the time of the test ( $N_{limit}$ ). The reason for this is that by increasing the length of the test we reduce the variability due to short term variations in the concurrent activity of the system. However the probability of a change in the overall concurrent activity of the system increases with a larger test. This change may produce a greater variance if the size of the sample statistic is small. We see that the best results are obtained for 20 observations and 1 to 2 seconds for the duration of the test. In machines with good clock resolution acceptable results are obtained with 10 observations and .2 seconds for each test.

## 6. Measurements and Some Results

We have run the system characterizer on the machines shown in table 5. Of the fifteen systems, four are supercomputers, each implementing single precision floating point with 64 bits. On the other systems single precision variables are allocated using 32 bits. We gathered two sets of measurements for the SUN 3/260, one using the 68881 co-processor to execute floating point arithmetic, and another emulating the same functions in software. We also measured the effect of using different Fortran compilers, the VMS FORT compiler and the UNIX BSD F77 both running on the VAX-11/785, in both cases with Ultrix as the operating system. By using the characterizer we can quantify how much each parameter is affected by the addition of a new hardware feature or by changing the compiler.

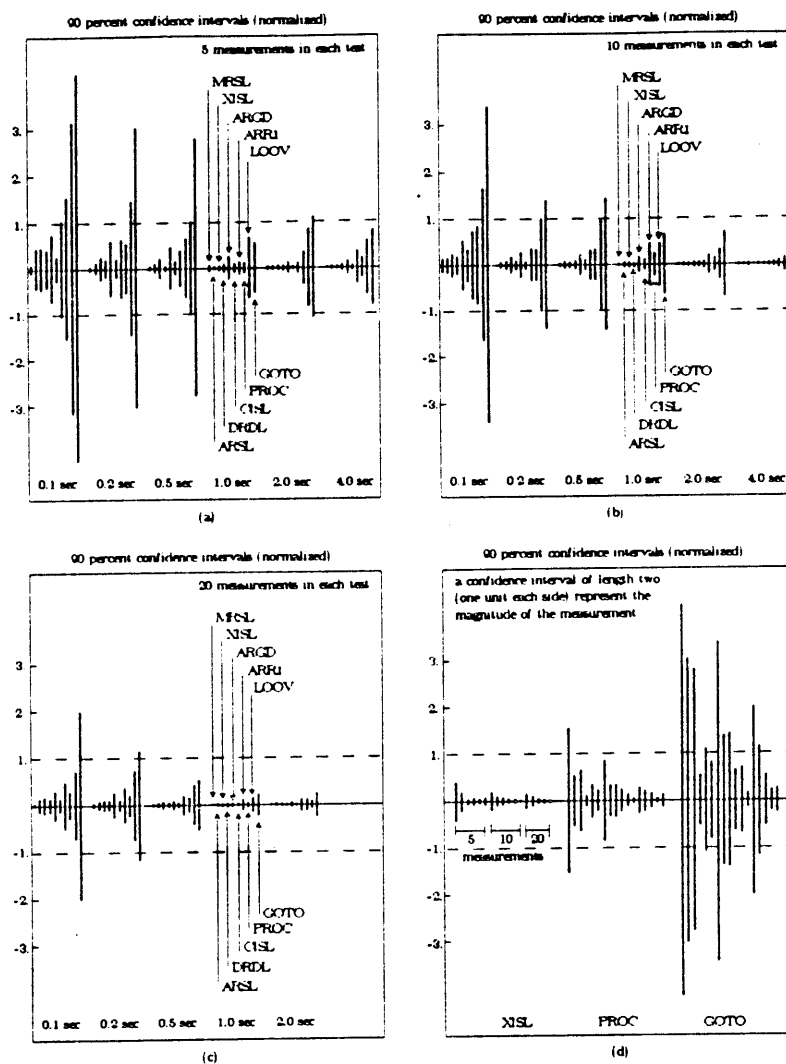


Figure 4: Normalized confidence intervals for ten different parameters. In (a), (b), and (c) we show how the length of the test ( $N_{limit}$ ) and the number of observations ( $N_{repeat}$ ) affect the confidence interval of the measurements, taken on a VAX 11/780. For a fixed number of observations, an increase in the execution time of the test tends to reduce the length of the confidence interval. Figure (d) shows for three parameters all their confidence intervals plotted together. All confidence intervals are normalized with respect to parameter  $P_i$ .

The measurements of all parameters are presented in the appendix in tables 10-14. The parameters are grouped according to tables 1 and 2, with all magnitudes in units of nanoseconds. Entries with magnitude '<1' represent parameters that were not detected by the characterizer. This happens when the execution time of the parameter is so small that most of its the execution overlapped with other operations; the total execution time of the program does not depend significantly on the occurrence of these parameters.

We can see some characteristics of the machines by looking at the results. For example, it is clear from the tables that the performance of the four supercomputers on the execution of double precision arithmetic is significantly lower than that for single precision. Single precision arithmetic operations and intrinsic functions take one order of magnitude less time to execute on these



Machine	Name/Location	Operating System	Compiler version	Memory	Integer single	Real	
						single	double
CRAY Y-MP/832	reynolds.arc.nasa.gov	UNICOS 4.0.8	CFT77 3.0	32 Mw	46	64	128
CRAY-2	navier.arc.nasa.gov	UNICOS 4.0.6	CFT77 3.0	128 Mw	46	64	128
CRAY X-MP/48	NASA Ames	COS 1.16	CFT 1.14	8 Mw	46	64	128
IBM 3090/200	cmsa.berkeley.edu	VM/CMS r.4	FORTRAN v2.3	32 MB	32	64	128
MIPS/1000	cassatt.berkeley.edu	UMIPS-BSD 2.1	F77 v1.21	16 MB	32	32	64
SUN 4/260	rosemary.berkeley.edu	SunOS r.4.0	F77	32 MB	32	32	64
VAX 8600	vangogh.berkeley.edu	UNIX 4.3 BSD	F77 v1.1	28 MB	32	32	64
VAX 3200	atlas.berkeley.edu	Ultrix 2.3	F77 v1.1	6 MB	32	32	64
VAX-11/785 (fort)	pioneer.arc.nasa.gov	Ultrix 3.0	Fort v4.7	16 MB	32	32	64
VAX-11/785 (f77)	pioneer.arc.nasa.gov	Ultrix 3.0	F77 v1.1	16 MB	32	32	64
VAX-11/780	wilbur.arc.nasa.gov	UNIX 4.3 BSD	F77 v2	4 MB	32	32	64
SUN 3/260 (f)	picasso.arc.nasa.gov	UNIX 4.2 r.3.2	F77 v1	16 MB	32	32	64
SUN 3/260	picasso.arc.nasa.gov	UNIX 4.2 r.3.2	F77 v1	16 MB	32	32	64
SUN 3/50	baal.berkeley.edu	UNIX 4.2 r.3.2	F77 v1	4 MB	32	32	64
IBM RT-PC/125	loki.berkeley.edu	ACIS 4.3	F77 v1	4 MB	32	32	64

Table 5: Characteristics of the machines. The size of the data type implementations are in number of bits. SUN 3/260 (f) uses the 68881 as a co-processor running at 20 MHz, while the CPU executes at 25 MHz. For the VAX-11/785 we used two FORTRAN compilers, the VAX FORT 4.7 and the Berkeley BSD 1.1.

machines than double precision. The greatest difference occurs on the IBM 3090/200 with double precision division. This operation takes almost 700ns using 64-bit operands, while the same operation with 128-bit operands takes around 75500ns. In contrast, the same operation takes less than 8000ns in any of the three CRAYs.

By looking at the results of the SUN 4/260 and SUN 3/260 (f), we can see the main differences between them. The greatest performance gap is found in floating point (real and complex) arithmetic, intrinsic functions, procedure calls and parameter passing. For integer arithmetic this difference is smaller.

It is also possible to compare our results with the numbers reported by manufacturers. However, this is no easy task given that our parameters may not map directly to a particular sequence of instructions and that there are many factors affecting the execution times of instructions. For example, on the 68020 the effective address calculation can take from zero to twenty-four cycles depending on the addressing mode and whether a prefetch instruction or/and an operand read is needed [Mot85, Mat87]. Nevertheless, table 6 shows timing estimates for four intrinsic functions (single precision) and also for the sequence of instructions implementing a procedure call. Included in the table are the measurements obtained with the system characterizer.

Table 6: Execution estimates vs. characterization results

	units	LOGS	EXPS	SINS	TANS	PROC
Timing Est.	cycles	672	598	482	574	113
	nsec	33600	29900	24100	28700	4420
Measurement	nsec	43799	28548	25790	31478	5034
Error		30.5%	4.5%	7.0%	9.7%	13.9%

For the intrinsic functions we assumed that the cycle time was 50 nsec (20MHz), and for procedure call 40 nsec (25MHz). We made some simplifying assumptions that are not necessarily valid for the SUN 3/260. We see that except for the logarithm function, our measurements are sufficiently closed to the timing estimates. This large difference is easily explained by looking at the code generated by the compiler; several additional instructions are included to determine, at

execution time, whether to compute  $\log(x)$ , or  $\log(x+1)$ .

The effect of different compilers can be seen in the results for the VAX-11/785. The FORT compiler produces code that is significantly faster for complex arithmetic and intrinsic functions, especially single precision intrinsics. There are some strange results in the case of the exponential operator. While the F77 code is between 2 and 5 times faster using a real base and an integer exponent, the FORT compiler is more than 4 times faster in the case of a real base and a real exponent. A similar situation occurs when the base is integer.

The fact that procedure calls are expensive operations on the VAX architecture can be corroborated when we compare the time it takes to execute this instruction on the VAX 8600 against either the MIPS/1000 or the SUN 4/260. A procedure call is approximately six times slower on the VAX 8600. This large gap is also found in the other VAX implementations, if we make the comparison against the SUN 3 or IBM RT-PC. This agrees with previous studies done on the VAX-11/780 that found that procedure calls take on the average 45.25 cycles to execute, while the average VAX instruction takes only 10.6 cycles [Eme82]. On the Whetstone benchmark 2.12% of the instructions executed are procedure calls and they represent 13% of the total execution time [Cla82].

### 6.1. A Reduced Representation of the Performance Measurements

The measurements obtained with the system characterizer makes it possible to compare different machine architectures either at the level of the parameters or by predicting the execution times of a set of programs using their parametric dynamic distributions. Predicting the execution time of a program is equivalent to reducing the set of basic measurements to a single number (the execution time) with the dynamic distribution acting as a weighting function. These two types of comparisons represent different extremes. On one side we have too much information with the raw measurements; it is difficult to identify those parameters that most affect performance without making reference to some particular workload. On the other extreme, a single number representing the execution time gives an illusion of precision by hiding the multidimensional aspects of program execution.

Therefore, it is convenient to represent the parameters in some 'reduced' form, in which overall performance is represented using a small number of dimensions, each associated with different aspects of the computation. In this way it is not only possible to compare the performance of a single operation or the overall performance with respect to a given workload, but also to focus on some particular mode of execution.

### 6.2. Combining Measurements and Selecting Weights

The two major issues when we reduce a large number of parameters into a smaller set are how to group the basic measurements, and how much weight to assign to each element.

For the first part we identified a small number of performance 'dimensions', each representing either a hardware or a software feature. These 'dimensions' should be as independent of each other as possible, and should reflect distinct components of the machine. A good selection of these new parameters will help us to better understand the behavior of the system. We use hardware, software and hybrid parameters. Integer addition is representative of the first group; trigonometric functions of the second; and floating point arithmetic, which in some machines is executed using special hardware and on others by software routines, belongs to the hybrid group.

The second issue, assigning weights to basic parameters, is a more difficult task, given that the impact of a parameter in the performance of a system is a function of the workload. However this workload dependency is not as serious a problem as in the case of reducing all parameters to a single number. The relative proportion of integer and floating point operations varies greatly from one program to another, but if we focus only on floating point, the relative distribution of these operations does not show the same degree of variability. We selected the weights based on extensive statistics of Fortran programs reported in the literature complemented with other statistics produced with our program analyzer [Knu71, Wei84, Saa88].

Table 7: Reduced Parameters

1 memory bandwidth (single)		2 memory bandwidth (double)	
TRSL	.125	TISL	.125
TRSG	.125	TISG	.125
3 integer addition		6 floating point addition	
AISL	.500	ARSL	.500
AISG	.500	ARSG	.500
4 integer multiplication		7 floating point multiplication	
MISL	.500	MRSL	.500
MISG	.500	MRSG	.500
5 integer arithmetic		8 floating point arithmetic	
DISL	.400	DRSL	.400
EISL	.090	ERSL	.090
XISL	.010	XRSL	.010
DISG	.400	DRSG	.400
EISG	.090	ERSG	.090
XISG	.010	XRSG	.010
9 complex precision arithmetic		10 double arithmetic	
ACSL	.325	ARDL	.325
MCSL	.125	MRDL	.125
DCSL	.040	DRDL	.040
ECSL	.008	ERDL	.008
XCSL	.002	XRDL	.002
ACSG	.325	ARDG	.325
MCSG	.125	MRDG	.125
DCSG	.040	DRDG	.040
ECSG	.008	ERDG	.008
XCSG	.002	XRDG	.002
11 intrinsic functions (single)		12 intrinsic functions (double)	
LOGS	.166	LOGC	.100
EXPS	.166	LOGD	.100
SINS	.166	EXPC	.100
TANS	.166	SINC	.100
SQRS	.166	SQRC	.100
MODS	.166	TAND	.100
MODD	.100		
13 logical operations		14 pipelining	
ANDL	.250	GOTO	.900
CRSL	.250	GCOM	.100
CCSL	.125	15 procedure calls	
CISL	.250	CALL	.750
CDRL	.125	ARGU	.250
16 address computation		17 iteration	
ARR1	.600	LOIN	.060
ARR2	.300	LOOV	.605
ARR3	.100	LOIX	.030
		LOOX	.305

Table 7: The seventeen reduced parameters, including basic measurements and their respective weights. With the exception of memory bandwidth, the sum of weights for each reduced parameter equals one. The sum of memory transfer weights equal .5, because the operation involves loading from memory and writing the results.

In table 7 we present the set of raw measurements and weights that formed each of the seventeen reduced parameters. Parameters characterizing hardware functional units are: integer addition and multiplication, logical operations, procedure calls, looping, and memory bandwidth (single and double precision). Software characteristics are represented by trigonometric functions

(single and double precision). Floating point, double precision and complex arithmetic, pipelining, and address computation belong to the hybrid class.

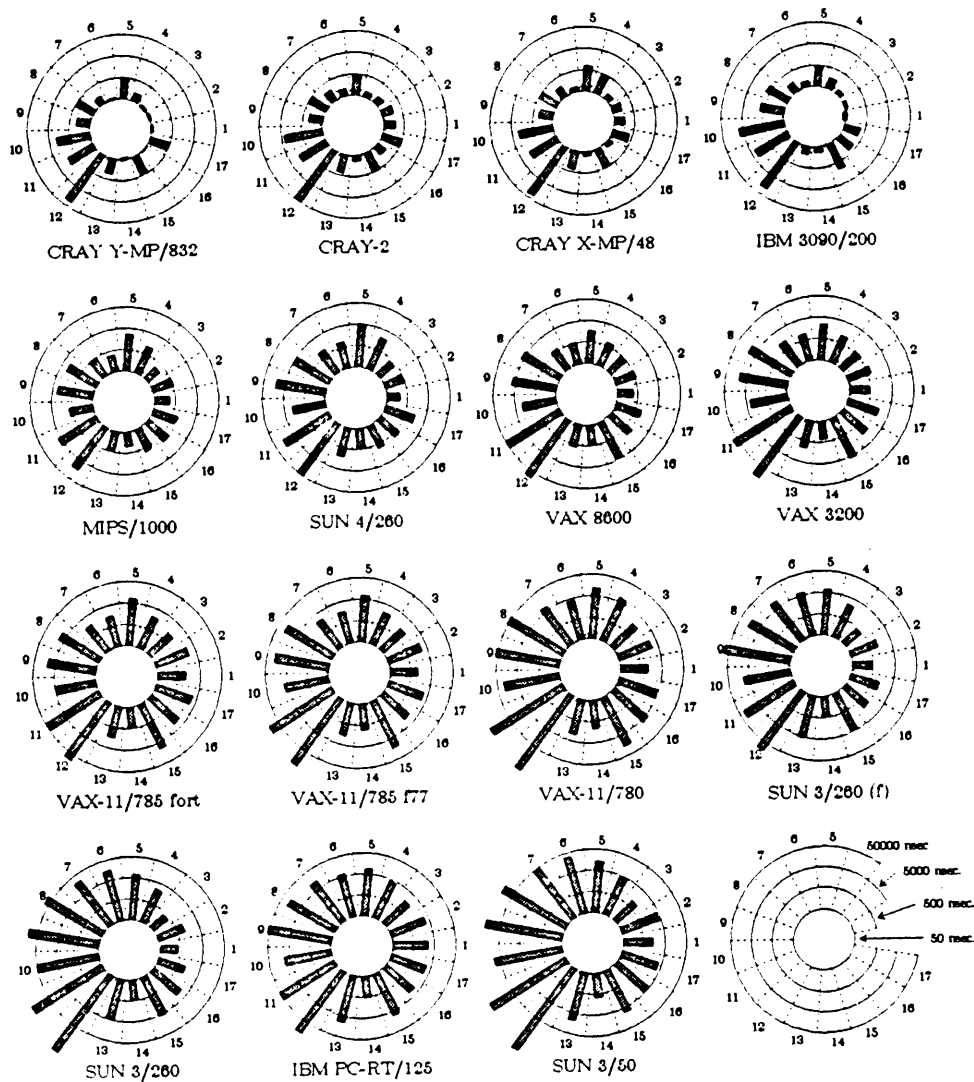


Figure 5: Performance of composite parameters. The three concentric circles represents 50, 500, 5000 and 50000 nanoseconds. Each Kiviat graph shows how different machines distribute their performance along different modes of execution.

### 6.3. Reduced Measurements and Kiviat Graphs

We present the same experimental measurements shown in tables 10-14 in the appendix in terms of the reduced parameters in figures 5 and 6. These results are also given in tables 15-17. In the tables, in addition to the magnitude we give the result normalized with respect to the shortest time, the CRAY Y-MP/832, and the VAX-11/780. For the VAX-11/780 we report the reciprocal. Both Kiviat graphs are logarithmic, with each circle representing a change of one order of magnitude with respect to its nearest neighbor. In figure 5 values are in units of nanoseconds with the circle closest to the center representing 50 nanoseconds. Quantities smaller than 50 nanoseconds are plotted in the direction of the center.

Sometimes it is convenient to express the performance distribution of the machine in terms of another machine which we defined as our standard unit of measure. The VAX-11/780 is usually arbitrarily rated as a 1 MIPS machine, and the performance of other machines is given in units of VAX-11/780 MIPS [Mip88]. We applied a similar transformation to the graphs in figure 5 to produce the Kiviat graphs of figure 6. Each dimension is normalized with respect to the VAX-11/780. In this case the smallest circle corresponds to a performance equal to one tenth of a VAX-11/780. As in figure 5, two adjacent circles have a separation of one order of magnitude.

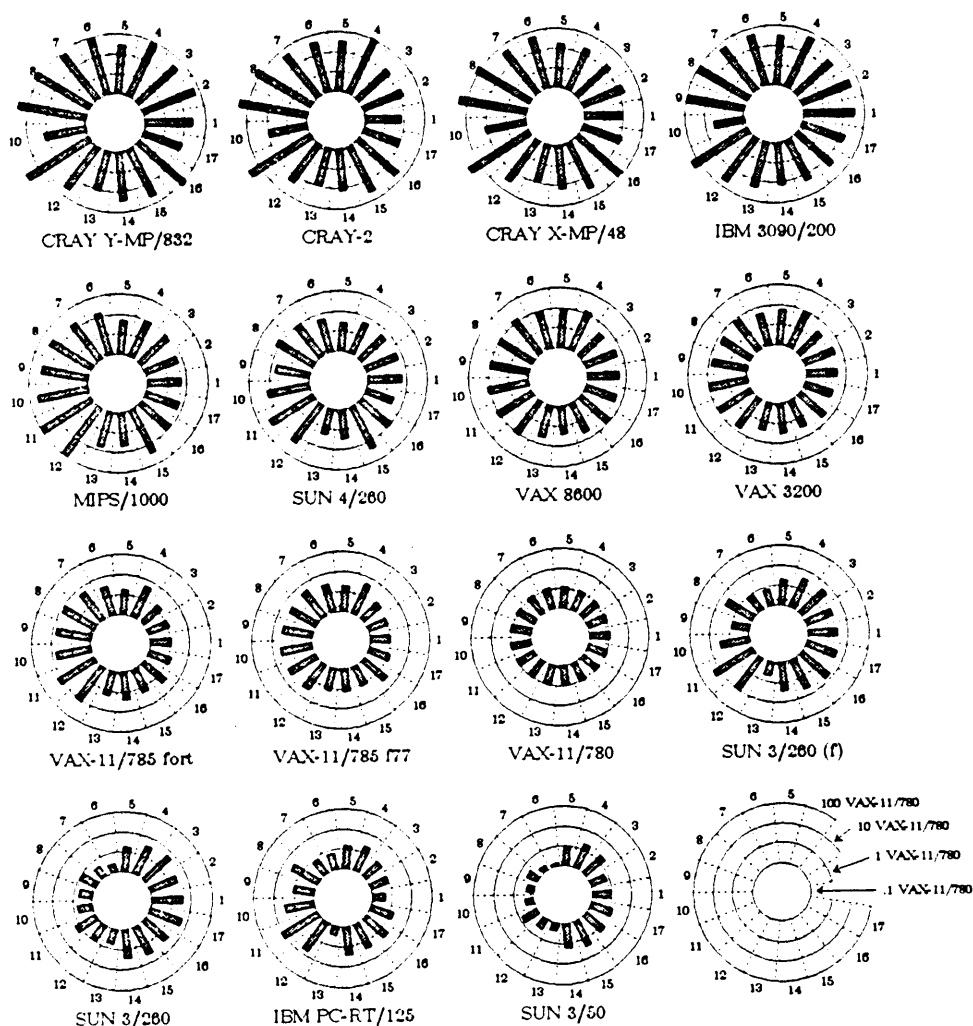


Figure 6: Performance of the reduced parameters with respect to the VAX-11/780. The concentric circles represents .1, 1, 10, and 100 times faster. The closest a performance shape (pershape) is to a circle, the closest the machine is to a VAX-11/780 in terms of how both machines distributed their performance along different computational modes.

Using the results from figures 6-7, and tables 15-17, we can identify several differences and similarities between the machines. The memory bandwidth results indicate that the only machines that show the same performance in single and double precision memory bandwidth are the CRAY Y-MP, the CRAY-2, and the CRAY X-MP. Although the single precision memory bandwidth in the IBM 3090 is faster than the CRAY Y-MP (34ns vs. 45ns), for double precision this situation is reversed (63ns vs. 40ns)<sup>3</sup>. The memory bandwidth reported here does not

<sup>3</sup> The difference between the memory bandwidth measured for the CRAY Y-MP/832 between single and double precision is a result of the measuring tools and the small execution times.

necessarily match the numbers given by the manufacturer. Our measurements characterize the execution time of a memory transfer assignment in a Fortran program, and for an arbitrary system this transfer is affected by the availability of registers, data cache, write buffer, and other circuitry that improves the data transfer between the CPU and memory.

We can see the effect of different compilers on the VAX-11/785. The difference in performance between the code produced by the two compilers is less than 10% in the cases of memory bandwidth with single precision, integer arithmetic, and DO loops. The FORT compiler code is 30% faster for real multiplication, 90% percent for complex arithmetic, 130% for real division and exponentials, and more than 3 times faster in intrinsic functions. On the other side, the F77 compiler code is less than 15% faster for integer division, and address computation. For intensive floating point programs, the code of the FORT compiler clearly outperforms the F77 compiler.

The effect of the floating point co-processor in the SUN 3/260 is also clear by looking at the results. Using the 68881 increases the performances of intrinsic functions by a factor of more than thirteen, and for floating point arithmetic by a factor from two to five.

The CRAY Y-MP/832 has the fastest times for floating point and complex arithmetic operations, function calls, array references, branching and single precision intrinsic functions (fig. 8 and tables 11-13). In particular, access to array elements is almost six times faster in the CRAY than in the IBM 3090/200 and 127 times faster than in the VAX-11/780. The CRAY machines are highly optimized for those parameters that are extensively used in scientific programs. The IBM 3090 is the fastest machine in double precision trigonometric functions, single precision memory bandwidth, and logical operations. The CRAY-2 shows similar performance to the CRAY X-MP and Y-MP in integer arithmetic (except integer addition), complex arithmetic, and procedure calls, but memory bandwidth shows a larger difference in both single and double precision.

A comparison between the MIPS/1000 and the SUN-4 shows better performance for the SUN-4 only in memory bandwidth and address computation, and the difference in all cases is less than 15%. The MIPS/1000 has an advantage of more than 75% in integer multiplication and arithmetic, floating point and complex arithmetic, intrinsic functions.

Floating point performance on the IBM RT-PC and SUN-3 (50 and 260) is slow compared with other machines and although a co-processor provides a significant improvement, their performance doesn't match the performance of comparable minicomputers. For example, the IBM 3090/200 is less than 12 times faster than SUN 3/50 (15 MHz) in integer addition, but 60 times faster than the SUN 3/260 (25 MHz) with 68881 (20 MHz) in floating point addition. The SUN 3/260 is between 4-6 times faster than the VAX-11/780 on procedure calls and array references, but the VAX-11/780 outperforms the SUN 3/260 on single precision floating point addition and multiplication.

## 7. Similar Performance Distributions (Performance Shapes)

Consider two machines  $M_X$  and  $M_Y$  that are identical except for the clock rates. These machines have the property that for any benchmark  $A$  their performance ratio (the execution time on one machine divided by the execution time on the other machine) is always a constant; thus, only one benchmark is sufficient to evaluate one against the other. For two arbitrary machines, however, this performance ratio can vary significantly for different benchmarks; it is possible to obtain a wide variety of performance ratios by running a sufficient set of benchmarks. Therefore it is important to quantify how different is the performance distribution of an arbitrary pair of machines and in this way determine how large we can expect the variability in the performance ratio to be when running a large sample of programs. This metric should group machines according to their performance 'shapes' and not by the magnitude of their performance parameters. A *performance shape* (pershape) is the Kiviat graph representing how performance is distributed along the different computational modes (reduced parameters). A pershape tells us not how large a parameter or set of parameters is with respect to other machines but how different machines distribute their performance. In the next subsection we present a metric that measures how similar are the absolute and normalized pershapes of two arbitrary machines.

### 7.1. A Metric for Performance Shapes

We would like a metric that captures the notion of similarity explained in the previous paragraph. By looking at figure 5 or 6, we clearly see that the pershape of the CRAY Y-MP/832 is very different than the pershapes of the VAX-11/780 or the SUN 3/50. But if we compare the CRAY Y-MP with the CRAY X-MP, or the VAX 8600 with the VAX 3200, we find that except for their relative sizes, the figures are very similar. It is this informal notion of similarity that we try to capture with the pershape distance.

First, there are several properties that we like our metric to satisfy in addition to the obvious properties required for distances. The pershape distance must be greater or equal to zero, and the distance from any machine to itself must be zero. It should satisfy the triangle inequality. It should be symmetric: the distance from  $A$  to  $B$  must be equal to the distance from  $B$  to  $A$ . One essential property for our metric is that if the performance of one machine is increased or decreased by the same quantity in all the dimensions the new distance does not change. By allowing two different pershapes vectors to have distance zero, we make the pershape distance a semi-metric<sup>4</sup> [Gil87]. Every parameter should have the same weight and any arbitrary permutation of the dimensions in both machines should not affect the distance. This means that our metric should be a function only of the relative performance of the machines and not of how we plot them. Making each dimension equally important intends to make the distance workload independent. The last property that we require is that if the performance in one dimension is changed in both machines by the same factor, their relative distance should not be affected.

The following discussion will give the rationale for allowing different pershapes vectors to have distance zero. It is important to understand that we are not trying to measure the difference in performance between two machines, but something completely different. We are interested in the variability of their expected performance. How fast one machine is compared to the other is always a function of the workload we use to evaluate them. What the pershape distance tries to measure is how large is the spectrum of possible comparative performance results when we use any possible workload composition. Therefore, given that two machines have a distance  $d$ , if in one machine we increase the performance of every dimension by the same factor ( $\lambda$ ), the distance should not be affected. Obviously, the machine will be faster or slower depending on whether  $\lambda$  is greater or less than 1, but the distribution of its performance remains the same. Therefore, its distance to any possible machine should not change. A similar situation happens when we add a constant to a random variable; the mean is affected, but the variance does not change.

Formally, let  $X = \langle x_1, x_2, \dots, x_n \rangle$  and  $Y = \langle y_1, y_2, \dots, y_n \rangle$  be two performance vectors in  $(0, \infty)^n$  representing the pershapes of machines  $M_X$  and  $M_Y$ . The metric

$$d(X, Y) = \left[ \frac{1}{n-1} \sum_{i=1}^n \left| \log\left(\frac{x_i}{y_i}\right) - \frac{1}{n} \sum_{j=1}^n \log\left(\frac{x_j}{y_j}\right) \right|^2 \right]^{1/2} \quad (10)$$

satisfies the following set of axioms:

- i)  $d(X, Y) \geq 0$
- ii)  $d(X, Y) = 0$  iff  $X = \lambda Y$  and  $\lambda > 0$
- iii)  $d(X, Y) = d(Y, X)$
- iv)  $d(X, Y) \leq d(X, Z) + d(Z, Y)$
- v)  $d(X_\sigma, Y_\sigma) = d(X, Y)$  for any arbitrary permutation  $\sigma$
- vi)  $d(\langle \lambda x_1, x_2, \dots, x_n \rangle, \langle \lambda y_1, y_2, \dots, y_n \rangle) = d(X, Y)$

Note that equation (10) is not the only possible distance satisfying the axioms: there are an infinity of different distance metrics with the same basic properties. The only metric property which

<sup>4</sup> In some textbooks, this is called a pseudo-metric [Kei85, Bou89]. We will not use the prefix 'semi-' or 'pseudo-' and simply refer to it as a metric.

provides any difficulty to verify is the triangle inequality. To verify axiom iv) we first rewrite equation (10) as follows:

$$d(X, Y) = \left[ \sum_{i=1}^n \left[ \frac{1}{(n-1)^{1/2}} \left[ \log(x_i) - \frac{1}{n} \sum_{j=1}^n \log(x_j) \right] - \frac{1}{(n-1)^{1/2}} \left[ \log(y_i) - \frac{1}{n} \sum_{j=1}^n \log(y_j) \right] \right]^2 \right]^{1/2} \quad (11)$$

then consider the mapping  $\phi : (0, \infty)^n \rightarrow \mathbb{R}^n$  defined by

$$\phi(x_i) = \frac{1}{(n-1)^{1/2}} \left[ \log(x_i) - \frac{1}{n} \sum_{j=1}^n \log(x_j) \right].$$

Now, if we replace  $\phi(X)$  and  $\phi(Y)$  in equation (11)

$$d(X, Y) = \left[ \sum_{i=1}^n (\phi(x_i) - \phi(y_i))^2 \right]^{1/2} \quad (12)$$

we obtain the Euclidean metric for  $\mathbb{R}^n$ , and the verification of the triangle inequality follows directly from the Cauchy-Schwarz inequality [Gil87].

In our presentation of the pershape distance, we did not specify whether vectors  $X$  and  $Y$  represent absolute or normalized pershapes. Computing a function on a normalized set of values does not always preserve some elementary properties. The output of the function may change when we normalize the inputs. It is important to see how our metric behaves when we normalize the set of reduced parameters.

Let  $X$  be an absolute pershape vector and  $X_Z$  a normalized vector obtained by dividing each component  $x_i$  of  $X$  with the same element in  $Z$

$$X_Z = \left\langle \frac{x_1}{z_1}, \frac{x_2}{z_2}, \dots, \frac{x_n}{z_n} \right\rangle$$

In linear algebra terms, normalizing vector  $X$  with respect to vector  $Z$  means applying a linear transformation  $T$  to vector  $X$ , such that the transformation matrix associated with  $T$  is diagonal. The matrix is zero everywhere except in the diagonal, with  $1/z_i$  as the diagonal element  $i$ . Now the normalized distance is given by

$$d(TX, TY) = d(X_Z, Y_Z) = d\left(\left\langle \frac{x_1}{z_1}, \dots, \frac{x_n}{z_n} \right\rangle, \left\langle \frac{y_1}{z_1}, \dots, \frac{y_n}{z_n} \right\rangle\right) = d(X, Y)$$

If we substitute the normalized parameters in equation (10), we see that the distance does not change. It is also easy to see that this property is enforced by axioms v) and vi). Thus, we say that distance  $d(X, Y)$  is isometric with respect to diagonal linear transformations.

In addition to measuring the distance between two performance vectors, the metric also gives information on which parameters will most affect the benchmark results between two machines. By ordering the terms inside of the first summation in equation (10), we find that the largest terms will be the ones that will contribute more to the summation, and therefore to the distance.

### 7.1.1. Similarity Results

Pershape distances were computed for all pairs of machines to detect which were the most and least similar machines. The most and least similar 25 are reported in table 8. The table shows that the most similar machines are the VAX 8600, VAX 3200, and the VAX-11/785, all using the F77 compiler. Other machines that are also close to each other are the SUN 3/50 and the SUN 3/260, both running without the 68881. The differences between these two machine are



the clock, the cache and the memory. The SUN 3/50 runs at 15 MHz, does not have a cache and uses standard memory chips. The SUN 3/260 runs at 25 MHz, has 64 Kbyte of virtual address write-back cache, and uses ECC for memory.

Most Similar Machines				Least Similar Machines			
	machine	machine	distance		machine	machine	distance
001	VAX 8600	VAX 3200	0.187	105	CRAY-2	SUN 3/200	1.753
002	VAX 8600	VAX-11/785 (f77)	0.214	104	CRAY X-MP/48	SUN 3/200	1.725
003	VAX 3200	VAX-11/785 (f77)	0.235	103	MIPS/1000	SUN 3/200	1.861
004	SUN 3/50	SUN 3/200	0.291	102	CRAY X-MP/48	SUN 3/50	1.848
005	VAX 3200	VAX-11/780	0.425	101	CRAY-2	SUN 3/50	1.847
006	VAX-11/785 (f77)	VAX-11/785 (fort)	0.432	100	MIPS/1000	SUN 3/50	1.591
007	CRAY Y-MP/832	CRAY X-MP/48	0.454	099	CRAY Y-MP/832	SUN 3/200	1.582
008	MIPS/1000	VAX-11/785 (fort)	0.478	098	VAX-11/785 (fort)	SUN 3/200	1.523
009	MIPS/1000	SUN 4/200	0.493	097	CRAY Y-MP/832	SUN 3/50	1.503
010	VAX 8600	VAX-11/780	0.498	096	VAX-11/785 (fort)	SUN 3/50	1.445
011	VAX 3200	VAX-11/785 (fort)	0.509	095	IBM 3090/200	SUN 3/200	1.434
012	VAX 8600	VAX-11/785 (fort)	0.516	094	IBM 3090/200	SUN 3/50	1.421
013	CRAY-2	CRAY X-MP/48	0.518	093	CRAY-2	SUN 3/200 (f)	1.420
014	VAX-11/785 (f77)	VAX-11/780	0.519	092	SUN 4/200	SUN 3/200	1.345
015	IBM RT-PC/125	SUN 3/200 (f)	0.522	091	CRAY X-MP/48	SUN 3/200 (f)	1.303
016	CRAY Y-MP/832	CRAY-2	0.532	090	VAX-11/785 (f77)	SUN 3/200	1.300
017	CRAY Y-MP/832	IBM 3090/200	0.861	089	VAX 8600	SUN 3/200	1.298
018	SUN 4/200	VAX-11/785 (fort)	0.863	088	SUN 4/200	SUN 3/50	1.286
019	VAX-11/785 (fort)	IBM RT-PC/125	0.872	087	CRAY-2	VAX-11/780	1.284
020	SUN 4/200	IBM RT-PC/125	0.884	086	VAX 8600	SUN 3/50	1.250
021	SUN 4/200	VAX 3200	0.712	085	CRAY Y-MP/832	SUN 3/200	1.242
022	SUN 4/200	VAX 8600	0.717	084	IBM RT-PC/125	SUN 3/200	1.233
023	SUN 4/200	VAX-11/785 (f77)	0.738	083	VAX-11/785 (f77)	SUN 3/50	1.231
024	MIPS/1000	VAX-11/785 (f77)	0.743	082	IBM 3090/200	SUN 3/200 (f)	1.228
025	MIPS/1000	VAX 8600	0.752	081	VAX 3200	SUN 3/200	1.203

Table 8: Pairs of machines with the smallest and largest pershape distance.

It is possible to use the results in table 8 to identify not only pairs of machines with similar pershapes, but also clusters of machines. Figure 7 illustrates one possible diagram showing for all the machines a bidirectional arrow joining the machines that have a distance less than 0.7. Different arrows are used to show how close the machines are. In the diagram we see three connected components, one formed by the supercomputers, another by the small workstations without floating point co-processors, and a large component mainly formed by two groups having a common neighbor. The closest of the two groups is formed by the machines implementing the VAX architecture and using the F77 compiler. The other group is formed by fast workstations. The VAX-11/785 using the FORT compiler acts a bridge between the two groups.

## 7.2. An Application of Pershape Distances

By using equation 10, it is possible not only to compute the distance between two machines but also to quantify which 'composite' parameters contribute most to unbalance the overall performance ratio between the two machines. In table 9 the execution times of nine programs are given for four of the machines. The table also includes the performance ratio between them, the maximum, minimum and geometric mean of their performance ratios, the maximum ratio of their relative performance and their pershape distance.

The programs used as benchmarks have different execution distributions and can be grouped in the following way: Shell, Erathostenes, and Baskett are integer programs; Alamos [Gri84, Sim87], Linpack [Don85, Don88], Livermore [McM86], and Mandelbrot are floating point intensive programs; Whetstone [Cur76] is a floating point and intrinsic function program; and the Smith benchmark [Smi89] mixes floating point, integer and logical operations. Baskett also executes a large proportion of function calls.

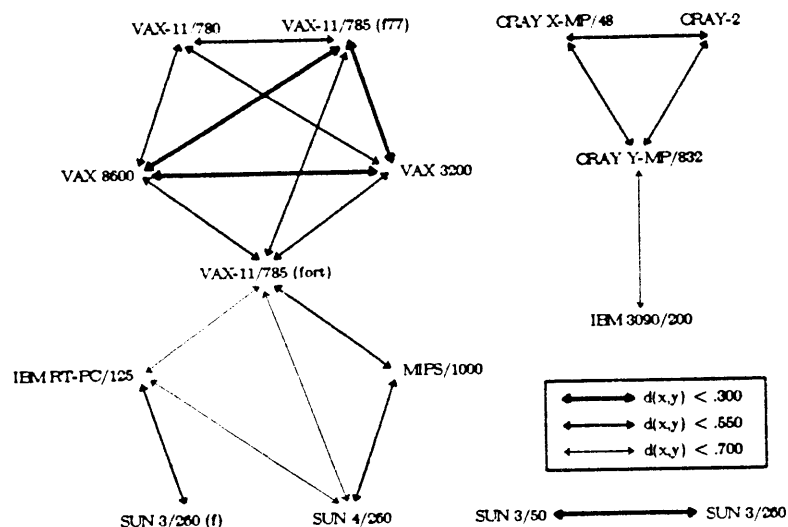


Figure 7: All machines with performance distance less than .700 are joined by a double arrow. The pershape distance identifies clusters of machines with similar performance distributions.

The results in the table show the relation between the pershape distance and the interval of possible benchmark results we can obtain when running a group of benchmarks. The pershape distance between the SUN 3/260 (without 68881) and the SUN 3/50 is only 0.29 and the interval of benchmark results is just 1.41. The difference between the smallest ratio (1.59) and the largest (2.25) is 41%. The same small distance is found between the IBM RT-PC and the SUN 3/260 (which uses a co-processor). Machines with large distance pershapes also give a large interval in the benchmark results, but the relation is not as clear as in the other cases. A possible explanation is that our program sample is not large enough, and certain types of operations that contribute to a large distance are not present in a large enough proportion to skew the benchmark results. The results do show that the SUN 3/50 can be 1.6 times slower than the SUN 3/260 (with 68881) in a predominantly integer benchmark, but 7.2 times slower in a benchmark with a high number of intrinsic functions. This is consistent with the performance ratios of the parameters representing integer operations and intrinsic functions. By looking at the distances between a group of machines, it is possible to identify which characteristics of the benchmarks will give a more complete evaluation of the systems. In contrast, programs that only exploit one of two characteristics will give skewed results.

## 8. Weak Points in the Characterizer

The current (new) version of the characterizer incorporates several additional parameters that were previously ignored. This has increased the number of parameters from 76 to 102. Complex variables and a better characterization of intrinsic functions form most of the new parameters. Even in this extended model there are several factors that have not been characterized.

- i) **Locality and Cache Memory:** code that exhibits different locality than our experiments affects the cache hit ratio and in consequence the access time for data and/or instructions. Measuring how the access time will be affected by different parts of the program will probably not be possible using a machine characterizer. By running some experiments with different degrees of locality we have found a variation of between four to ten percent.

program	SUN 3/260		IBM RT-PC III	SUN 3/50 IV	execution ratios					
	(f) I	II			II/I	III/I	IV/I	III/II	IV/II	IV/III
Alamos	1547.9 s	2838.9 s	3881.9 s	6273.2 s	1.83	2.51	4.05	1.37	2.21	1.62
Baskett	3.92 s	3.88 s	6.20 s	7.06 s	0.99	1.58	1.80	1.60	1.82	1.14
Erathostenes	0.64 s	0.64 s	1.10 s	0.90 s	1.00	1.72	1.59	1.72	1.59	0.93
Linpack	184.9 s	338.5 s	473.9 s	763.7 s	1.83	2.56	4.13	1.40	2.26	1.61
Livermore	507.1 s	1103.1 s	1610.1 s	2457.0 s	2.18	3.18	4.85	1.46	2.23	1.53
Mandelbrot	41.88 s	75.88 s	105.43 s	163.94 s	1.81	2.52	3.92	1.39	2.16	1.56
Shell	1.68 s	1.72 s	4.68 s	3.14 s	1.02	2.79	1.87	2.72	1.83	0.67
Smith	338.3 s	406.7 s	545.10 s	914.8 s	1.20	1.61	2.70	1.34	2.25	1.68
Whetstone	4.74 s	15.28 s	12.05 s	34.24 s	3.22	2.54	7.22	0.79	2.24	2.84
				minimum	0.99	1.58	1.59	0.79	1.59	0.67
				geom. mean	1.55	2.27	3.18	1.46	2.05	1.40
				maximum	3.22	3.18	7.22	2.72	2.26	2.84
				max/min	3.26	2.01	4.53	3.46	1.41	4.23
				d(x,y)	0.96	0.52	0.93	1.23	0.29	1.13

Table 9: Execution ratios between pair of machine and comparison against their performance distances. Machines with a small performance distance have less variability in their relative speed. The maximum and minimum entries correspond to the ratio of largest and smallest execution ratios. We give results for the SUN 3/200 with co-processor ((f) I), and without it (II). The roman numerals denote

- ii) Branching: The size of the branch affects the execution time by modifying the locus of execution. If the target of the branch is to a nonresident page this may involve a page fault and a context switch. A context switch normally involves flushing the cache and this forces a 'cold' start on cache references.
- iii) Hardware and/or Software Interlocks. In pipelined machines the time to produce a new result depends on the context in which the instruction is executed. This normally depends (in addition to the effective execution time) on the functional and data dependencies with respect to previously scheduled instructions. As in the previous two factors this is difficult to measure from a high-level program.
- iv) Machine Idioms. Special cases of some instructions are optimized to improve execution time. These idioms are used by the compiler whenever possible. Without knowledge of the architecture and the compiler, it is not possible to detect which are the idioms of a given machine. In machines with auto-increment and auto-decrement addressing modes, these modes may be used in statements like  $i = i + 1$ .
- v) 'Random' Noise Produced by Concurrent Activity. Although we address this problem in §§ 5.5 and 5.6, there is still a problem left when we run in a loaded system. A small increase in the load of the system tends to affect the measurements of some parameters, in particular array address computation, branches and loop overhead.
- vi) Optimization. In this study we only considered unoptimized code; the characterizer was compiled and run with optimization disabled. Even when it is not difficult to detect which optimizations are applied by the compiler, it is not clear how we can modify the execution time model to include optimized programs. Parsing a program and detecting which optimizations are possible and deciding for these which ones are going to be applied by a particular compiler, seems to require a 'super-optimizer'. It is outside of this research to write such program; we will try to develop other techniques to characterize optimization in the future.

## 9. Conclusions and Summary

In this paper we have presented a model for machine characterization based on a large number of high-level parameters representing operations for an abstract Fortran machine. This provides a uniform model in which machines with different architectures can be compared on equal terms. It is possible to detect differences and similarities between machines with respect to

individual parameters. In addition, we have presented a set of composite parameters that provide a more compact way of representing the effect of hardware or software features in the execution time of programs. Based on these composite parameters we presented the concept of performance shape to show how different machines distribute their possible performance in different ways. We defined a metric to measure the similarity between two pershapes and show how this distance can be used to classify machines and the metric's relation to the variation in benchmark results.

Using the characterization results or the reduced parameters, it is possible to make estimates for the execution time of programs and in this way study the sensitivity of the execution time with respect to variations in the workload. This last aspect will be presented in a forthcoming paper [Saa89]. We think that our approach will advance the state of the art of performance evaluation in several ways.

- (1) A uniform 'high level' model of the performance of computer systems allows us to make better comparisons between different architectures and identify their differences and similarities when the systems execute a common workload.
- (2) Using the characterization to predict performance provides us with a mechanism to validate our assumptions on how the execution time depends on individual components of the system.
- (3) With a uniform model that can be used for all machines sharing a common mode of computation, it is possible to define metrics that permit more extensive comparisons and in this way obtain a better understanding of the behavior of each system.
- (4) We can study the sensitivity of the system to changes in the workload, and in this way detect imbalances in the architectures.
- (5) The results obtained with the system characterizer give insight into the implementation of the CPU architecture, and the machine designers can use the results to improve future implementations.
- (6) Application programmers and users can identify the most time consuming parts of their programs and measure the impact of new 'improvements' on different systems.
- (7) For procurement purposes this is a less expensive and more flexible way of evaluating computer systems and new architectural features. Although the best way to evaluate a system is to run a real workload, a more extensive and intensive evaluation can be made using system characterizers to select a small number of computers for subsequent on-site evaluation.

In the last thirty years we have seen an explosion of new ideas in many field of computer science, but one problem that hasn't received much attention is how to make a fair comparison between two different architectures. Given the impact that computers have in all aspects of society we cannot afford to continue characterizing the performance of such complex systems using MIPS, MFLOPS or DHRYSTONES as our units of measure.

#### Acknowledgements

The material presented here is based on research supported in part by NASA under consortium agreement NCA2-128 and cooperative agreement NCC2-550, the Mexican Council for Science and Technology (CONACYT) under contract 49992, the National Science Foundation under grants CCR-8202591 and MIP-8713274, by the State of California under the MICRO program, and by the International Business Machines Corporation, Digital Equipment Corporation, Tandem, Hewlett-Packard, and Philips Research Laboratories / Signetics.

#### References

- [Bai85a] Bailey, D.H., Barton, J.T., "The NAS Kernel Benchmark Program", NASA Technical Memorandum 86711, August 1985.
- [Bai85b] Bailey, D.H., "NAS Kernel Benchmark Results", *Proc. First Int. Conf. on Supercomputing*, St. Petersburg, Florida, December 16-20, 1985, pp. 341-345.

- [Bou89] Bourbaki, N., *Elements of Mathematics: General Topology*, Springer Verlag, 1989.
- [Cla85] Clark, D.W., and Levy, "Measurement and Analysis of Instruction Set Usage in the VAX-11/780", *Proc. 9th Symposium on Computer Architecture*, April 1982. Translation Buffer: Simulation and Measurement", *Transactions on Computer Systems*, Vol. 3, No. 1, February 1985, pp. 31-62.
- [Cla85] Clark, D.W., and Emer, J.S. "Performance of the VAX-11/780 Translation Buffer: Simulation and Measurement", *Transactions on Computer Systems*, Vol. 3, No. 1, February 1985, pp. 31-62.
- [Cla86] Clapp, R.M., Duchesneau, L., Volz, R.A., Mudge, T.N., and Schultze T., " Toward Real-Time Performance Benchmarks for ADA", *Communications of the ACM*, Vol. 29, No. 8, August 1986, pp. 760-778.
- [Cur75] Currah B., "Some Causes of Variability in CPU Time", *Computer Measurement and Evaluation*, SHARE project, Vol. 3, 1975, pp. 389-392.
- [Cur76] Curnow, H.J., Wichmann, B.A., "A Synthetic Benchmark", *The Computer Journal*, Vol. 19, No.1, February 1976, pp. 43-49.
- [Don85] Dongarra, J.J., "Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment", *Computer Architecture News*, Vol. 13, No. 1, March 1985, pp. 3-11.
- [Don87] Dongarra, J.J., Martin, J., and Worlton J., "Computer Benchmarking: paths and pitfalls", *Computer*, Vol. 24, No. 7, July 1987, pp. 38-43.
- [Don88] Dongarra, J.J., "Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment", *Computer Architecture News*, Vol. 16, No. 1, March 1988, pp. 47-69.
- [Eme84] Emer, J.S. and Clark, D.W., "A Characterization of Processor Performance in the VAX-11/780", *Proceedings of the 11th Annual Symposium on Computer Architecture*, Ann Arbor, Michigan, June 1984.
- [Gil87] Giles, J.R., *Introduction to the Analysis of Metric Spaces*, (Australian Mathematical Society, lecture series 3; Cambridge University Press, 1987).
- [Gri84] Griffin, J.H., Simmons, M.L., "Los Alamos National Laboratory Computer Benchmarking 1983", Los Alamos Technical Report No. LA-10151-MS, June 1984.
- [Ibb82] Ibbett, R.N., *The Architecture of High Performance Computers* (Springer-Verlag, New York, 1982).
- [IBM87] *IBM 3090 VS Fortran v.2 Language and Library Reference*, SC26-4221-02, 1987.
- [Kel87] Kelley, J.L., *General Topology*. GTM: 27, Springer-Verlag, 1985.
- [Knu71] Knuth, D.E., "An Empirical Study of Fortran Programs", *Software-Practice and Experience*, Vol. 1, pp. 105-133 (1971).
- [Lee84] Lee, J.K.F., Smith, A.J., "Branch Prediction Strategies and Branch Target Buffer Design", *Computer*, Vol. 17, No. 1, January 1984, pp. 6-22.
- [Mac84] MacDougall, M.H., "Instruction-Level Program and Processor Modeling", *Computer*, Vol. 7 No. 14, July 1982, pp. 14-24.
- [McM86] McMahon, F.H., "The Livermore Fortran Kernels: A Computer Test of the Floating-Point Performance Range", Lawrence Livermore National Laboratory, UCRL-53745, December 1986.
- [Mer83] Merrill, H.W., "Repeatability and Variability of CPU timing in Large IBM Systems", *CMG Transactions*, Vol. 39, March 1983.
- [Mip88] MIPS Computer Systems, Inc. "Performance Brief CPU Benchmarks", Issue 3.5, October 1988.

- [Mot85] Motorola, Inc, *MC68020 32-Bit Microprocessor User's Manual*, Prentice-Hall, Inc, 1985.
- [Mot87] Motorola, Inc, *MC68881/MC68882 Floating-Point Coprocessor User's Manual*, Prentice-Hall, Inc, 1987.
- [Peu77] Peuto, B.L. and Shustek, L.J., "An Instruction Timing Model of CPU Performance", *The fourth Annual Symposium on Computer Architecture*, Vol. 5, No. 7, March 1977, pp. 165-178.
- [Saa88] Saavedra-Barrera, R.H., "Machine Characterization and Benchmark Performance Prediction", University of California, Berkeley, Technical Report No. UCB/CSD 88/437, June 1988.
- [Saa89] Saavedra-Barrera, R.H., and Smith A.J., "Scalar CPU Performance Evaluation via Benchmark Prediction", paper in preparation.
- [Sim87] Simmons, M.L. and Wasserman H.J., "Los Alamos National Laboratory Computer Benchmarking 1986", Los Alamos National Laboratory. LA-10898-MS, January 1987.
- [Smi82] Smith, A.J., "CPU Cache Memories", *ACM Computing Surveys*, Vol. 14, No. 3, September 1982, pp. 473-530.
- [Smi89] Smith, A.J., paper in preparation.
- [Wei84] Weicker, R.P., "Dhrystone: A Synthetic Systems Programming Benchmark", *Communications of the ACM*, Vol. 27, No. 10, October 1984.
- [Wor84] Worlton, J., "Understanding Supercomputer Benchmarks", *Datamation*, September 1, 1984, pp. 121-130.

## 10. Appendix

Group 1: Floating Point Arithmetic Operations (single, local)

machine	SRSL	ARSL	MRSL	DRSL	ERSL	XRSL	TRSL
CRAY Y-MP/832	13	46	111	210	660	4150	96
CRAY-2	39	70	101	250	78	4180	112
CRAY X-MP/432	82	76	154	357	91	5035	281
IBM 3090/200	1 <	82	140	684	129	4952	60
MIPS/1000	67	269	437	976	543	53018	499
SUN 4/260	104	755	788	2498	4724	60430	533
VAX 8600	72	425	575	1610	1097	217676	509
VAX 3200	262	805	999	2013	1847	361666	587
VAX-11/785 fort	263	1282	1524	3778	16305	82006	1799
VAX-11/785 f77	246	1371	1924	4034	3740	648082	2065
VAX-11/780	1086	3215	6739	9322	11041	2066420	1598
SUN 3/260 (f)	1978	5543	8709	11394	15998	58901	1293
SUN 3/260	1 <	13580	19118	23003	31612	2205175	1286
SUN 3/50	1 <	26420	40246	46476	60818	4743815	3076
IBM RT-PC/125	3639	5684	10715	12304	12437	231989	6235

Group 2: Floating Point Arithmetic Operations (complex, local)

machine	SCSL	ACSL	MCSL	DCSL	ECSL	XCSL	TCSL
CRAY Y-MP/832	30	85	267	497	818	10466	147
CRAY-2	32	110	221	386	48	17167	199
CRAY X-MP/432	63	124	271	511	1 <	13168	319
IBM 3090/200	26	215	679	3218	2940	13912	97
MIPS/1000	121	926	1727	12025	9004	72791	1097
SUN 4/260	1 <	8034	11808	29356	7561	130805	663
VAX 8600	275	1438	3523	39419	17876	326399	974
VAX 3200	792	2287	6925	47240	30134	510817	1072
VAX-11/785 fort	531	2653	7542	53236	26842	314924	3514
VAX-11/785 f77	1074	4717	10206	88085	83278	966246	4703
VAX-11/780	1319	9679	38202	328270	170337	3584596	3796
SUN 3/260 (f)	436	27270	83719	353726	133222	446378	1382
SUN 3/260	1 <	31812	109547	604151	183495	5417755	1095
SUN 3/50	265	63460	231185	1233098	453373	11405138	8310
IBM RT-PC/125	471	26969	47498	194262	183060	678778	5101

Group 3: Integer Arithmetic Operations (single, local)

machine	SISL	AISL	MISL	DISL	EISL	XISL	TISL
CRAY Y-MP/832	1 <	39	106	271	1113	1131	82
CRAY-2	1 <	61	62	324	126	131	114
CRAY X-MP/432	1 <	91	414	714	396	755	320
IBM 3090/200	1 <	76	143	439	163	358	73
MIPS/1000	1 <	227	945	2577	1111	2146	475
SUN 4/260	1 <	286	1634	3918	5882	7979	219
VAX 8600	1 <	357	628	1591	896	1883	462
VAX 3200	1 <	490	895	2206	1273	2562	750
VAX-11/785 fort	1 <	1002	1615	7292	1760	28928	2259
VAX-11/785 f77	1 <	1088	1789	7053	2309	5142	2182
VAX-11/780	1 <	1327	6024	10502	7779	15803	2186
SUN 3/260 (f)	1 <	298	2212	4011	13979	17174	393
SUN 3/260	1 <	237	2280	4119	14708	17398	251
SUN 3/50	1 <	813	3898	7039	29262	36348	856
IBM RT-PC/125	1 <	1497	3438	8837	4063	7581	2478

Table 10: Characterization results for Group 1-3. A value 1 < indicates that the parameter was not detected by the experiment.

Group 4: Floating Point Arithmetic Operations (double, local)

machine	SRDL	ARDL	MRDL	DRDL	ERDL	XRDL	TRDL
CRAY Y-MP/832	2	917	1628	5473	4804	108121	13
CRAY-2	1 <	1974	2752	7355	2072	194054	1 <
CRAY X-MP/432	89	1122	1812	6392	1073	138645	206
IBM 3090/200	1 <	424	964	75656	1493	48282	154
MIPS/1000	117	346	581	1558	838	49780	632
SUN 4/260	290	986	1228	4665	7046	133573	1058
VAX 8600	220	754	1725	5812	2841	208984	876
VAX 3200	276	1367	1896	4063	3256	353750	1178
VAX-11/785 fort	1047	2280	4243	7996	23368	177403	3920
VAX-11/785 f77	929	2893	5460	8921	9517	636736	5637
VAX-11/780	1142	10589	24687	48235	33181	2044644	5483
SUN 3/260 (f)	2159	5819	9272	11942	17793	112601	2610
SUN 3/260	1172	23804	49458	73051	46323	2482220	1364
SUN 3/50	2068	54245	100594	132284	110522	5504681	6682
IBM RT-PC/125	5765	6889	8125	12611	14110	200954	4823

Group 5: Floating Point Arithmetic Operations (single, global)

machine	SRSg	ARSg	MRSg	DRSg	ERSg	XRSg	TRSg
CRAY Y-MP/832	13	46	111	210	660	4150	96
CRAY-2	95	124	216	392	72	3811	506
CRAY X-MP/432	90	73	152	354	83	5052	287
IBM 3090/200	12	80	129	685	152	4901	60
MIPS/1000	70	268	435	973	536	50784	364
SUN 4/260	145	778	855	2573	4739	60387	573
VAX 8600	254	483	598	1600	1040	215039	408
VAX 3200	400	878	1076	2159	1770	361567	554
VAX-11/785 fort	998	1244	1501	3619	16318	81494	1430
VAX-11/785 f77	219	1378	1928	4049	3727	651381	2070
VAX-11/780	1517	3304	6646	9539	10649	2058123	1164
SUN 3/260 (f)	1948	5616	8945	11662	16018	58321	1157
SUN 3/260	1 <	13433	18920	22865	31453	2139622	716
SUN 3/50	1 <	26943	40586	47035	58663	4671805	5092
IBM RT-PC/125	3156	5629	9684	12287	13054	230580	6636

Group 6: Floating Point Arithmetic Operations (complex, global)

machine	SCSG	ACSG	MCSG	DCSG	ECSG	XCSG	TCSG
CRAY Y-MP/832	30	85	267	497	818	10466	147
CRAY-2	92	167	303	513	18	16738	512
CRAY X-MP/432	78	117	265	511	1 <	13177	335
IBM 3090/200	27	230	682	3162	2983	13965	102
MIPS/1000	121	927	1730	12049	8992	73007	1101
SUN 4/260	63	8027	12078	29703	7573	130146	664
VAX 8600	551	1544	3802	38787	17812	326228	1159
VAX 3200	519	2859	7334	46918	31358	511323	1599
VAX-11/785 fort	1055	3210	8827	52768	24348	298978	4393
VAX-11/785 f77	1144	4695	10039	87745	83649	962812	4680
VAX-11/780	1684	9778	35853	322237	168501	3679780	3297
SUN 3/260 (f)	1 <	27975	83103	352636	134477	453818	1519
SUN 3/260	3695	29210	107292	586288	191029	5307737	1 <
SUN 3/50	2383	63526	231688	1233524	448297	11359785	8016
IBM RT-PC/125	555	26948	47435	197036	182374	663827	5216

Table 11: Characterization results for Group 4-6. A value 1 < indicates that the parameter was not detected by the experiment.



Group 7: Integer Arithmetic Operations (single, global)

machine	SISG	AISG	MISG	DISG	EISG	XISG	TISG
CRAY Y-MP/832	1<	39	106	271	1113	1131	82
CRAY-2	1<	161	89	485	144	153	607
CRAY X-MP/432	1<	93	405	716	405	751	327
IBM 3090/200	1<	79	151	439	170	393	82
MIPS/1000	1<	227	942	2580	1110	2143	476
SUN 4/280	1<	421	1728	4022	6022	7972	252
VAX 8600	1<	522	606	1593	990	2010	622
VAX 3200	1<	594	1028	2202	1484	2852	826
VAX-11/785 fort	1<	1113	1888	7428	1857	29838	2269
VAX-11/785 f77	1<	1094	1788	7025	2279	5089	2167
VAX-11/780	1<	1616	7166	10731	8002	16036	2156
SUN 3/260 (f)	1<	438	2116	4015	14156	17771	427
SUN 3/260	1<	381	2121	4050	14212	16824	218
SUN 3/50	1<	937	3537	6887	29760	36609	738
IBM RT-PC/125	1<	1459	3422	8865	3956	7553	2438

Group 8: Floating Point Arithmetic Operations (double, global)

machine	SRDG	ARDG	MRDG	DRDG	ERDG	NRDG	TRDG
CRAY Y-MP/832	2	917	1626	5473	4804	108121	13
CRAY-2	1<	2051	2858	7360	2283	202494	302
CRAY X-MP/432	69	1122	1821	6342	1108	138935	222
IBM 3090/200	1<	421	963	73307	912	41150	154
MIPS/1000	108	349	587	1561	854	58483	679
SUN 4/280	278	961	1167	4586	7078	133796	1060
VAX 8600	252	796	1611	5905	2828	206974	817
VAX 3200	268	1515	2175	4238	3307	353997	1080
VAX-11/785 fort	1207	2202	4106	8044	23236	171685	3882
VAX-11/785 f77	968	2288	4498	7916	9192	636084	4461
VAX-11/780	1274	10848	24648	47719	34214	2024890	4551
SUN 3/260 (f)	2354	5700	9275	11817	18065	112638	2477
SUN 3/260	549	23648	49458	72612	45612	2562978	2664
SUN 3/50	2436	54749	100942	133431	110630	5495105	4046
IBM RT-PC/125	3799	6017	10228	13526	14088	203231	7612

Group 9,10: Conditional and Logical Parameters

machine	ANDL	CRSL	CCSL	CISL	CRDL	ANDG	CRSG	CCSG	CISG	CRDG
CRAY Y-MP/832	14	287	315	282	335	14	287	315	282	335
CRAY-2	38	95	237	96	1873	85	317	430	337	1963
CRAY X-MP/432	45	226	335	229	1243	46	228	322	231	1256
IBM 3090/200	104	94	106	73	214	111	138	171	161	259
MIPS/1000	185	471	375	337	603	183	474	404	335	602
SUN 4/280	310	1217	3767	236	1566	455	1333	4168	655	1585
VAX 8600	304	653	680	464	867	321	868	991	743	1022
VAX 3200	389	1127	1295	767	1603	412	1207	1371	844	1602
VAX-11/785 fort	954	1378	1727	1033	2649	1013	1747	2348	1116	2264
VAX-11/785 f77	769	1937	1966	1467	2578	768	1909	1937	1477	2629
VAX-11/780	1091	2823	2768	1987	3914	1100	3057	3674	2481	4573
SUN 3/260 (f)	414	6814	16257	329	7171	588	7093	15922	741	7057
SUN 3/260	394	5542	10938	332	10730	604	5728	11985	847	9765
SUN 3/50	803	13243	29047	559	22442	1399	14382	27969	1625	25800
IBM RT-PC/125	1005	16166	16033	2012	15919	1029	15430	15700	2130	15993

Table 12: Characterization results for Group 7-10. A value 1< indicates that the parameter was not detected by the experiment.

Group 11,12: Function Call, Arguments and References to Array Elements

machine	PROC	ARGU	ARR1	ARR2	ARR3	IADD
CRAY Y-MP/832	512	61	42	49	60	12
CRAY-2	574	40	122	159	200	1<
CRAY X-MP/432	583	73	59	104	148	2
IBM 3090/200	1162	70	128	410	746	17
MIPS/1000	797	139	523	1044	1592	1<
SUN 4/260	918	67	384	1004	1490	13
VAX 8600	4670	610	478	1223	2137	1<
VAX 3200	6991	957	668	1934	3316	1<
VAX-11/785 fort	11678	1515	1320	2897	5578	844
VAX-11/785 f77	16421	1526	995	2701	5057	32
VAX-11/780	19931	1783	2128	9592	18518	1<
SUN 3/260 (f)	5034	397	448	1661	2600	2
SUN 3/260	6548	594	990	3834	3484	1<
SUN 3/50	8836	1535	2042	6396	8759	100
IBM RT-PC/125	9395	991	2212	2406	4536	1<

Group 13,14: Branching and DO loop Parameters

machine	GOTO	GCOM	LOIN	LOOV	LOIN	LOOX
CRAY Y-MP/832	1<	406	1015	315	627	368
CRAY-2	15	692	1263	353	264	513
CRAY X-MP/432	25	483	966	180	1307	293
IBM 3090/200	38	460	660	130	952	353
MIPS/1000	137	1010	1938	417	1643	945
SUN 4/260	302	984	3378	1007	2320	1638
VAX 8600	262	1705	2540	396	6223	1070
VAX 3200	128	2117	3916	975	5336	1634
VAX-11/785 fort	277	1691	13042	972	11747	3124
VAX-11/785 f77	332	4262	8323	1621	7044	2788
VAX-11/780	588	4783	2525	2552	17363	4558
SUN 3/260 (f)	258	1742	2863	567	3256	1509
SUN 3/260	268	1694	1657	524	1957	1411
SUN 3/50	394	3001	6558	1976	5765	3776
IBM RT-PC/125	119	3395	11368	1236	5425	3396

Group 15: Intrinsic Functions (single precision)

machine	EXPS	LOGS	SINS	TANS	SQRS	ABSS	MODS	MAXS
CRAY Y-MP/832	1453	1314	1423	1514	1038	1<	265	177
CRAY-2	1980	1855	2067	2136	266	25	383	328
CRAY X-MP/432	1826	1627	1846	1985	1356	1<	318	200
IBM 3090/200	2893	2887	2805	4119	2534	37	1094	435
MIPS/1000	6612	5680	5751	5156	6745	61	7215	1470
SUN 4/260	13560	14197	12081	20338	14520	450	23141	4758
VAX 8600	67798	52587	42683	70577	23883	1285	26471	3275
VAX 3200	109786	77167	63001	99637	32436	2108	38300	4563
VAX-11/785 fort	27212	28438	39474	70494	22634	215	42421	4101
VAX-11/785 f77	201824	240223	109462	138871	56848	2996	88497	8302
VAX-11/780	690106	765909	468763	857151	177536	4230	186125	12234
SUN 3/260 (f)	43799	28548	25790	31478	12627	464	15571	15528
SUN 3/260	367032	443458	574151	686006	61509	1<	49869	18798
SUN 3/50	770610	950878	1272922	1512997	92447	4700	129932	47730
IBM RT-PC/125	27466	22327	23168	26511	7014	47189	179593	41101

Table 13: Characterization results for Group 11-15. A value 1< indicates that the parameter was not detected by the experiment.

Group 16: Intrinsic Functions (double precision)

machine	EXPD	LOGD	SIND	TAND	SQRD	ABSD	MODD	MAXD
CRAY Y-MP/832	51052	58111	32289	71186	8689	28	9581	2200
CRAY-2	88428	94288	67440	146037	12100	431	19508	1021
CRAY X-MP/432	70511	64914	37931	83390	9751	21	9459	727
IBM 3090/200	20471	21893	19390	28520	10193	70	78571	858
MIPS/1000	8565	7508	7997	7747	9330	39	6985	2385
SUN 4/260	22261	22220	21184	36096	27382	504	18995	6106
VAX 8600	67151	52267	41751	69792	23310	1755	24244	5083
VAX 3200	108029	79491	63237	101020	31103	3001	35793	7175
VAX-11/785 fort	51821	51081	97932	158473	30389	693	71349	7050
VAX-11/785 f77	203491	238536	107896	137069	55608	5906	84380	17867
VAX-11/780	701933	776686	467842	856357	179556	7504	176955	22079
SUN 3/260 (f)	46526	32093	28500	32619	13966	312	17058	19584
SUN 3/260	965293	1096555	1009146	1132512	94349	1<	60492	22428
SUN 3/50	2080362	2332882	2210543	2418819	175124	1<	150656	67713
IBM RT-PC/125	38928	34208	34753	37343	13901	11669	120334	44149

Groups 17,18: Intrinsic Functions (integer and complex)

machine	ABSI	MODI	MAXI	EXPC	LOGC	SINC	SQRC	ABSC
CRAY Y-MP/832	76	563	127	6093	4478	5027	4282	1784
CRAY-2	51	545	202	9290	7827	9244	3761	1618
CRAY X-MP/432	58	1644	192	7913	5755	6553	5020	2309
IBM 3090/200	93	541	399	6948	5384	7081	6188	2302
MIPS/1000	169	2607	1415	21639	20454	25382	17361	6955
SUN 4/260	1027	3261	2957	87132	46483	123282	73181	38489
VAX 8600	1381	2546	2983	168775	145238	282011	87857	47671
VAX 3200	1498	3640	3816	266255	233369	438531	118507	62425
VAX-11/785 fort	563	11022	3367	156792	81107	88997	102146	54562
VAX-11/785 f77	2897	8970	8096	458645	491650	760555	199835	113818
VAX-11/780	4262	16165	10719	1740767	1637191	2510877	626909	299780
SUN 3/260 (f)	1665	3721	3825	178628	196966	231844	232360	26185
SUN 3/260	3186	6529	3414	2722348	2339489	3656209	649596	168115
SUN 3/50	9953	15760	13886	5734016	5100953	7775319	1338978	364009
IBM RT-PC/125	2549	9232	8196	410201	233930	511218	380805	258205

Table 14: Characterization results for Group 16-18. A value 1< indicates that the parameter was not detected by the experiment.

Parameter 1: Memory bandwidth (single precision)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	44.5	1.30	1.00	22.27
CRAY-2	167.4	4.87	3.76	5.92
CRAY X-MP/48	151.8	4.41	3.41	6.53
IBM 3090/200	34.4	1.00	0.77	28.80
MIPS/1000	228.8	6.50	5.09	4.37
SUN 4/260	197.0	5.73	4.43	5.03
VAX 8600	250.1	7.27	5.62	3.96
VAX 3200	339.5	9.87	7.63	2.92
VAX-11/785 (fort)	969.8	28.19	21.79	1.02
VAX-11/785 (f77)	1004.5	29.20	22.57	0.99
VAX-11/780	990.8	28.80	22.27	1.00
SUN 3/260 (f)	408.8	11.89	9.19	2.42
SUN 3/260	308.9	8.98	6.94	3.21
IBM RT-PC/125	2223.3	64.63	49.96	0.45
Sun 3/50	1220.1	35.47	27.42	0.81

Parameter 2: Memory bandwidth (double precision)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	40.1	1.00	1.00	52.98
CRAY-2	168.9	4.21	4.21	12.58
CRAY X-MP/48	135.2	3.37	3.37	15.71
IBM 3090/200	63.4	1.58	1.58	33.51
MIPS/1000	438.6	10.94	10.94	4.84
SUN 4/260	430.7	10.74	10.74	4.93
VAX 8600	478.3	11.93	11.93	4.44
VAX 3200	616.0	15.36	15.36	3.45
VAX-11/785 (fort)	1963.7	48.97	48.97	1.08
VAX-11/785 (f77)	2282.2	56.91	56.91	0.93
VAX-11/780	2124.4	52.98	52.98	1.00
SUN 3/260 (f)	998.6	24.90	24.90	2.13
SUN 3/260	853.8	21.29	21.29	2.49
IBM RT-PC/125	2818.8	70.29	70.29	0.75
Sun 3/50	3381.7	84.33	84.33	0.63

Parameter 3: Integer addition

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	38.6	1.00	1.00	37.54
CRAY-2	110.8	2.87	2.87	13.08
CRAY X-MP/48	91.9	2.38	2.38	15.77
IBM 3090/200	77.3	2.00	2.00	18.74
MIPS/1000	227.4	5.89	5.89	6.37
SUN 4/260	353.4	9.16	9.16	4.10
VAX 8600	439.6	11.39	11.39	3.30
VAX 3200	541.9	14.04	14.04	2.67
VAX-11/785 (fort)	1057.3	27.39	27.39	1.37
VAX-11/785 (f77)	1108.2	28.71	28.71	1.31
VAX-11/780	1448.9	37.54	37.54	1.00
SUN 3/260 (f)	367.9	9.53	9.53	3.94
SUN 3/260	309.2	8.01	8.01	4.69
IBM RT-PC/125	1477.9	38.29	38.29	0.98
Sun 3/50	875.0	22.67	22.67	1.66

Parameter 6: Floating point addition (single)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	46.3	1.00	1.00	71.57
CRAY-2	97.0	2.10	2.10	34.13
CRAY X-MP/48	74.4	1.61	1.61	44.54
IBM 3090/200	81.2	1.75	1.75	40.81
MIPS/1000	268.4	5.80	5.80	12.35
SUN 4/260	786.3	16.55	16.55	4.32
VAX 8600	454.0	9.81	9.81	7.30
VAX 3200	841.3	18.17	18.17	3.94
VAX-11/785 (fort)	1263.3	27.29	27.29	2.62
VAX-11/785 (f77)	1391.5	30.06	30.06	2.38
VAX-11/780	3313.8	71.57	71.57	1.00
SUN 3/260 (f)	5579.7	120.51	120.51	0.59
SUN 3/260	13506.7	291.72	291.72	0.25
IBM RT-PC/125	5656.3	122.17	122.17	0.59
Sun 3/50	26681.8	576.28	576.28	0.13

Parameter 4: Integer multiplication

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	105.6	1.40	1.00	65.76
CRAY-2	75.7	1.00	0.72	91.74
CRAY X-MP/48	409.5	5.41	3.88	16.06
IBM 3090/200	147.2	1.95	1.40	47.18
MIPS/1000	943.9	12.47	8.94	7.36
SUN 4/260	1881.0	22.21	15.92	4.13
VAX 8600	617.1	8.15	5.84	11.26
VAX 3200	961.6	12.70	9.11	7.22
VAX-11/785 (fort)	1751.4	23.14	16.50	3.97
VAX-11/785 (f77)	1810.2	23.91	17.14	3.84
VAX-11/780	6944.7	91.74	65.76	1.00
SUN 3/260 (f)	2164.3	28.59	20.50	3.21
SUN 3/260	2200.8	29.07	20.84	3.16
IBM RT-PC/125	3430.1	45.31	32.48	2.03
Sun 3/50	3717.4	49.11	35.20	1.87

Parameter 7: Floating point multiplication (single)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	110.8	1.00	1.00	61.21
CRAY-2	158.2	1.43	1.43	42.87
CRAY X-MP/48	153.0	1.38	1.38	44.33
IBM 3090/200	134.5	1.22	1.22	50.43
MIPS/1000	436.0	3.94	3.94	15.56
SUN 4/260	821.4	7.41	7.41	8.26
VAX 8600	586.7	5.30	5.30	11.56
VAX 3200	1037.2	9.36	9.36	6.54
VAX-11/785 (fort)	1512.8	13.65	13.65	4.48
VAX-11/785 (f77)	1952.5	17.62	17.62	3.47
VAX-11/780	6782.1	61.21	61.21	1.00
SUN 3/260 (f)	8827.0	79.67	79.67	0.77
SUN 3/260	19018.9	171.85	171.85	0.36
IBM RT-PC/125	10199.2	92.05	92.05	0.67
Sun 3/50	40416.1	364.77	364.77	0.17

Table 15: Parameter values, and performance ratios with respect to fastest machine (m/f), CRAY Y-MP/832 (m/cray), and VAX-11/780 (v780/m) for reduced parameters 1-6. Numbers in column 'value' are given in nanoseconds.

Parameter 5: Integer arithmetic

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	430.3	1.25	1.00	23.01
CRAY-2	350.6	1.00	0.80	28.83
CRAY X-MP/48	650.1	1.88	1.50	15.34
IBM 3090/200	388.4	1.11	0.88	26.03
MIPS/1000	2305.5	6.58	5.25	4.39
SUN 4/260	4406.5	12.57	10.03	2.29
VAX 8600	1482.1	4.23	3.38	6.82
VAX 3200	2065.6	5.89	4.70	4.89
VAX-11/785 (fort)	6801.1	19.40	15.48	1.49
VAX-11/785 (f77)	6286.0	17.93	14.31	1.61
VAX-11/780	10108.7	28.83	23.01	1.00
SUN 3/260 (f)	6092.3	17.38	13.87	1.66
SUN 3/260	6212.8	17.72	14.14	1.63
IBM RT-PC/125	7953.7	22.69	18.11	1.27
Sun 3/50	11612.0	33.12	26.43	0.87

Parameter 8: Floating point arithmetic (single)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	389.4	1.08	1.00	133.02
CRAY-2	350.1	1.00	0.95	140.38
CRAY X-MP/48	400.7	1.15	1.09	122.63
IBM 3090/200	671.5	1.92	1.82	73.18
MIPS/1000	1914.7	5.47	5.18	25.86
SUN 4/260	4087.3	11.68	11.07	12.02
VAX 8600	5803.7	16.58	15.71	8.47
VAX 3200	9226.6	26.35	24.98	5.33
VAX-11/785 (fort)	7529.6	21.51	20.38	6.53
VAX-11/785 (f77)	17268.2	49.33	46.75	2.85
VAX-11/780	49138.4	140.38	133.02	1.00
SUN 3/260 (f)	13276.1	37.92	35.94	3.70
SUN 3/260	67471.1	192.72	182.65	0.73
IBM RT-PC/125	16756.0	47.86	45.36	2.93
Sun 3/50	142314.0	406.50	385.26	0.35

Parameter 9: Complex arithmetic

Machine	Speed	m/f	m/cray	v780/m
CRAY Y-MP/832	206.9	1.00	1.00	261.67
CRAY-2	242.8	1.18	1.18	222.98
CRAY X-MP/48	225.8	1.09	1.09	239.76
IBM 3090/200	661.9	3.20	3.20	81.79
MIPS/1000	2369.0	11.45	11.45	22.85
SUN 4/260	11087.1	53.59	53.59	4.88
VAX 8600	6295.2	30.43	30.43	8.60
VAX 3200	9041.8	43.70	43.70	5.99
VAX-11/785 (fort)	9547.4	46.15	46.15	5.67
VAX-11/785 (f77)	17514.6	84.65	84.65	3.09
VAX-11/780	54138.4	261.67	261.67	1.00
SUN 3/260 (f)	70687.1	341.65	341.65	0.77
SUN 3/260	113826.0	550.15	550.15	0.48
IBM RT-PC/125	50206.8	242.66	242.66	1.08
SUN 3/50	239606.0	1158.08	1158.08	0.23

Parameter 10: Double precision arithmetic

Machine	Speed	m/f	m/cray	v780/m
CRAY Y-MP/832	1846.2	1.00	1.00	12.63
CRAY-2	3229.7	1.75	1.75	7.22
CRAY X-MP/48	2127.7	1.15	1.15	10.96
IBM 3090/200	6626.8	3.59	3.59	3.52
MIPS/1000	673.1	0.37	0.37	34.65
SUN 4/260	1841.4	1.00	1.00	12.67
VAX 8600	2061.8	1.12	1.12	11.31
VAX 3200	2804.4	1.52	1.52	8.32
VAX-11/785 (fort)	4112.3	2.23	2.23	5.67
VAX-11/785 (f77)	5775.0	3.13	3.13	4.04
VAX-11/780	23323.1	12.63	12.63	1.00
SUN 3/260 (f)	7684.4	4.16	4.16	3.04
SUN 3/260	41992.5	22.73	22.73	0.56
IBM RT-PC/125	8380.8	4.54	4.54	2.78
SUN 3/50	80383.0	48.42	48.42	0.26

Parameter 11: Intrinsic functions (single)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	1167.9	1.00	1.00	449.60
CRAY-2	1447.8	1.24	1.24	362.68
CRAY X-MP/48	1492.8	1.28	1.28	351.74
IBM 3090/200	2722.1	2.33	2.33	192.90
MIPS/1000	6193.0	5.30	5.30	84.79
SUN 4/260	16306.2	13.96	13.96	32.20
VAX 8600	47333.0	40.53	40.53	11.09
VAX 3200	70054.6	59.98	59.98	7.50
VAX-11/785 (fort)	38445.3	32.92	32.92	13.66
VAX-11/785 (f77)	145176.0	124.31	124.31	3.62
VAX-11/780	525084.0	449.60	449.60	1.00
SUN 3/260 (f)	26302.2	22.52	22.52	19.97
SUN 3/260	363671.0	311.39	311.39	1.45
IBM RT-PC/125	47679.7	40.83	40.83	11.01
SUN 3/50	788297.0	674.97	674.97	0.67

Parameter 12: Intrinsic functions (double)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	25076.6	1.24	1.00	38.54
CRAY-2	45881.1	2.26	1.83	21.08
CRAY X-MP/48	30119.7	1.49	1.20	32.09
IBM 3090/200	20263.7	1.00	0.81	47.89
MIPS/1000	13296.7	0.66	0.53	72.68
SUN 4/260	47821.5	2.36	1.91	20.21
VAX 8600	94239.7	4.65	3.76	10.28
VAX 3200	147533.0	7.28	5.88	6.55
VAX-11/785 (fort)	88988.7	4.30	3.55	10.86
VAX-11/785 (f77)	285871.0	14.11	11.40	3.38
VAX-11/780	966457.0	47.69	38.54	1.00
SUN 3/260 (f)	101056.0	4.99	4.03	9.57
SUN 3/260	1372600.0	67.74	54.74	0.70
RT-PC.Aix	181562.0	8.96	7.24	5.32
SUN 3/50	2931770.0	144.68	116.91	0.33

Table 16: Parameter values, and performance ratios with respect to fastest machine (m/f), the CRAY Y-MP/832 (m/cray), and the VAX-11/780 (v780/m) for reduced parameters 7-12. Numbers in column 'value' are given in nanoseconds.

Parameter 5: Logical operations

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	227.1	2.10	1.00	9.82
CRAY-2	320.4	2.97	1.41	6.96
CRAY X-MP/48	322.2	2.98	1.42	6.92
IBM 3090/200	108.0	1.00	0.48	20.64
MIPS/1000	370.4	3.43	1.63	6.02
SUN 4/260	1107.2	10.25	4.88	2.01
VAX 8600	548.7	5.08	2.42	4.06
VAX 3200	933.1	8.64	4.11	2.39
VAX-11/785 (fort)	1388.3	12.86	6.11	1.61
VAX-11/785 (f77)	1650.1	15.28	7.27	1.35
Vax-11/780	2220.6	20.64	9.82	1.00
SUN 3/260 (f)	4817.7	44.61	21.22	0.46
SUN 3/260	4275.4	39.59	18.83	0.52
IBM RT-PC/125	8789.7	81.39	38.70	0.25
SUN 3/50	10087.6	93.40	44.42	0.22

Parameter 8: Pipelining

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	40.6	1.00	1.00	24.85
CRAY-2	83.1	2.05	2.05	12.14
CRAY X-MP/48	70.6	1.75	1.75	14.23
IBM 3090/200	79.8	1.97	1.97	12.64
MIPS/1000	224.2	5.52	5.52	4.50
SUN 4/260	370.8	9.13	9.13	2.72
VAX 8600	405.8	10.00	10.00	2.49
VAX 3200	327.0	8.06	8.06	3.09
VAX-11/785 (fort)	418.6	10.31	10.31	2.41
VAX-11/785 (f77)	800.7	19.72	19.72	1.26
VAX-11/780	1008.8	24.85	24.85	1.00
SUN 3/260 (f)	406.1	10.00	10.00	2.48
SUN 3/260	410.6	10.11	10.11	2.46
IBM RT-PC/125	446.7	11.00	11.00	2.26
SUN 3/50	654.7	16.13	16.13	1.54

Parameter 15: Procedure calls

Machine	Speed	m/f	m/cray	v780/m
CRAY Y-MP/832	309.3	1.00	1.00	38.47
CRAY-2	440.1	1.10	1.10	34.90
CRAY X-MP/48	455.7	1.14	1.14	33.71
IBM 3090/200	889.1	2.23	2.23	17.28
MIPS/1000	632.5	1.59	1.59	24.28
SUN 4/260	696.4	1.75	1.75	22.06
VAX 8600	3655.0	9.16	9.16	4.20
VAX 3200	5482.5	13.73	13.73	2.80
VAX-11/785 (fort)	8765.3	21.95	21.95	1.75
VAX-11/785 (f77)	12565.3	31.47	31.47	1.22
Vax.780.f77	15359.3	38.47	38.47	1.00
SUN 3/260 (f)	3875.0	9.71	9.71	3.96
SUN 3/260	5059.5	12.67	12.67	3.04
IBM RT-PC/125	7294.1	18.27	18.27	2.11
SUN 3/50	7012.4	17.56	17.56	2.19

Parameter 16: Address computation

Machine	Speed	m/f	m/cray	v780/m
CRAY Y-MP/832	45.6	1.00	1.00	127.28
CRAY-2	141.1	3.10	3.10	41.13
CRAY X-MP/48	81.6	1.79	1.79	71.13
IBM 3090/200	274.2	6.01	6.01	21.17
MIPS/1000	786.2	17.24	17.24	7.38
SUN 4/260	680.6	14.93	14.93	8.53
VAX 8600	867.6	19.03	19.03	6.69
VAX 3200	1312.4	28.78	28.78	4.42
VAX-11/785 (fort)	2219.0	48.66	48.66	2.62
VAX-11/785 (f77)	1941.8	42.58	42.58	2.99
VAX-11/780	5804.0	127.28	127.28	1.00
SUN 3/260 (f)	1027.1	22.52	22.52	5.65
SUN 3/260	2092.5	45.89	45.89	2.77
IBM RT-PC/125	2502.4	54.88	54.88	2.32
SUN 3/50	4020.0	88.16	88.16	1.45

Parameter 17: Iteration (DO loops)

Machine	value	m/f	m/cray	v780/m
CRAY Y-MP/832	382.3	1.50	1.00	9.78
CRAY-2	453.8	1.79	1.19	8.24
CRAY X-MP/48	295.3	1.16	0.77	12.66
IBM 3090/200	254.3	1.00	0.67	14.70
MIPS/1000	706.1	2.78	1.85	5.30
Sun 4/260	1380.9	5.43	3.61	2.71
VAX 8600	905.3	3.56	2.37	4.13
VAX 3200	1483.2	5.83	3.88	2.52
VAX-11/785 (fort)	2675.5	10.52	7.00	1.40
VAX-11/785 (f77)	2773.2	10.61	7.26	1.35
VAX-11/780	3739.0	14.70	9.78	1.00
SUN 3/260 (F)	1072.3	4.22	2.81	3.49
SUN 3/260	905.1	3.56	2.37	4.13
IBM RT-PC/125	2628.1	10.34	6.88	1.42
SUN 3/50	2913.6	11.46	7.62	1.28

Table 17: Parameter values, and performance ratios with respect to fastest machine (m/f), the CRAY Y-MP/832 (m/cray), and the VAX-11/780 (v780/m) for reduced parameters 13-17. Numbers in column 'value' are given in nanoseconds.