

**Speedy: An Integrated Performance
Extrapolation Tool for pC++ Programs**

**Bernd Mohr, Kesavan Shanmugam,
and Allen Malony**

**CIS-TR-95-06
February 1995**

**Department of Computer and Information Science
University of Oregon**

Speedy: An Integrated Performance Extrapolation Tool for pC++ Programs[†]

Bernd W. Mohr, Kesavan Shanmugam, Allen D. Malony

Department of Computer and Information Science,
University of Oregon, Eugene OR 97403, USA
{mohr, kesavans, malony}@cs.uoregon.edu

Abstract. Performance Extrapolation is the process of evaluating the performance of a parallel program in a target execution environment using performance information obtained for the same program in a different execution environment. Performance extrapolation techniques are suited for rapid performance tuning of parallel programs, particularly when the target environment is unavailable. This paper describes one such technique that was developed for data-parallel C++ programs written in the pC++ language. In pC++, the programmer can distribute a *collection* of objects to various processors and can have methods invoked on those objects execute in parallel. Using performance extrapolation in the development of pC++ applications allows tuning decisions to be made in advance of detailed execution performance measurements. The current pC++ language system includes \mathcal{T} , an integrated environment for analyzing and tuning the performance of pC++ programs. This paper presents *speedy*, a new addition to \mathcal{T} , that predicts the performance of pC++ programs on parallel machines using extrapolation techniques. *Speedy* applies the existing instrumentation support of \mathcal{T} to capture high-level event traces of a n -thread pC++ program run on a uniprocessor machine (made possible by pC++'s multithreaded runtime system) together with trace-driven simulation to predict the performance of the program run on a target n -processor machine. We describe how *speedy* works and how it is integrated into \mathcal{T} . We also show how *speedy* can be used to evaluate and tune a pC++ program for a given target environment.

Keywords: performance prediction, extrapolation, object-parallel programming, trace-driven simulation, performance debugging tools, and modeling.

1 Introduction

One of the foremost challenges for a parallel programmer is to achieve the best possible performance for an application on a parallel machine. For this purpose, the process of *performance debugging* (the iterative application of performance *diagnosis* [11] and *tuning*) is applied as an integral constituent part of a parallel program development methodology. Application of performance debugging in practice has invariably required the development of performance tools based on the measurement and analysis of actual parallel program execution. Parallel performance

[†] This research is supported by ARPA under Rome Labs contract AF 30602-92-C-0135 and Fort Huachuca contract ARMY DABT63-94-C-0029. The research is also supported by a NSF National Young Investigator (NYI) award.

environments [21] support performance debugging through program instrumentation, performance data analysis, and results presentation tools, but have often lacked in the level of their integration with parallel programming systems. However, recent efforts on developing portable high-level parallel language systems has motivated work in integrated program analysis environments where performance debugging concerns are more closely coupled with the language's use for program development [14,18]. Of particular interest is the incorporation of performance prediction support in the programming environment for giving feedback to the user on algorithm implementation or to the compiler on optimization strategies [6,7,15]. In most instances, however, there is a dependence on actual machine access for performance debugging, restricting the parallel programmer to consider optimization issues only for physically available machines. For parallel programs intended to be portable to a variety of parallel platforms, and scalable across different machine and problem size configurations, undertaking performance debugging for all potential cases, whether by empirical evaluation or prediction based on measurement, is usually not possible.

Ideally, an integrated program analysis environment for a parallel language system that provides support for performance debugging would include a means to predict performance where only limited access (if any) to the target system is given. The environment would measure only that performance data which is necessary from the execution environment and use high-level analysis to evaluate different program alternatives under different system configuration scenarios. In this manner, the environment would enable performance-driven parallel program design where algorithm choices could be considered early in the development process [25]. However, the user would demand a level of detail from predicted performance analysis comparable to that provided by measurements, which static prediction tools often cannot provide. Similarly, the user will be frustrated if the delay to generate predicted results is significantly greater than that with measurement-based experiments (hopefully, it is much less), a problem often faced by simulation systems that analyze program execution at too low a level. For example, the Proteus system [3,4] and the Wisconsin Wind Tunnel [22] have considerably advanced the efficiency and effectiveness of dynamic prediction techniques for architectural studies, but the overhead is prohibitively high to warrant their use for rapid and interactive performance debugging.

In this paper, we describe a performance prediction technique that combines high-level modeling with dynamic execution simulation to facilitate rapid performance debugging. The technique is one example of a general prediction methodology called *Performance Extrapolation* that estimates the performance of a parallel program in a target execution environment by using the performance data obtained from running the program in a different execution environment. In [24], we demonstrated that performance extrapolation is a viable process for parallel program performance debugging that can be applied effectively in situations where standard measurement techniques are restrictive or costly. From a practical standpoint, performance extrapolation methods must address the problem of how to achieve the comparative utility and accuracy of measurement-based analysis without incurring the expense of detailed dynamic simulation, but at the same time retaining the flexibility and robustness of model-based prediction techniques. However, with this even being the case, there remains the problem of how performance extrapolation can

be seamlessly integrated in a parallel language system, where it both leverages and complements the capabilities of the program analysis framework.

We have integrated our performance extrapolation techniques into the \mathcal{T} program analysis environment for pC++, a data-parallel C++ language system. In Section 2, we describe the pC++ language and the features of \mathcal{T} to show how the environment can easily be extended to provide support for predicting the performance of pC++ programs on parallel machines. The performance extrapolation approach to pC++ prediction is discussed in Section 3. In Section 4, we show how the performance extrapolation techniques have been integrated into \mathcal{T} , in the form of the *speedy* tool. *Speedy* has been used to evaluate and tune pC++ programs for different target environments and Section 5 presents some results from these experiments. The paper concludes with a discussion on future work.

2 pC++ and TAU

In this section, we give a brief overview of \mathcal{T} , (TAU, for **T**uning and **A**nalysis **U**tilities), a first prototype towards an integrated, portable program and performance analysis environment for pC++. pC++ is a language extension to C++ designed to allow programmers to compose distributed data structures with parallel execution semantics. The basic concept behind pC++ is the notion of a *distributed collection*, which is a type of concurrent aggregate “container class”.

More specifically, a *collection* is a structured set of objects which are distributed across the processing elements of the computer in a manner designed to be consistent with HPF [13]. To accomplish this, pC++ provides a simple mechanism to build *collections of objects* from a base *element* class. Member functions from this element class can be applied to the entire collection (or a subset) in parallel. This mechanism provides the user with a clean interface to *data-parallel* style operations by simply calling member functions of the base class. In addition, there is a mechanism for encapsulating SPMD style computation in a thread-based computing model that is both efficient and completely portable. To help the programmer build collections, the pC++ language includes a library of standard collection classes that may be used directly or subclassed. This includes classes such as *DistributedArray*, *DistributedMatrix*, *DistributedVector*, and *DistributedGrid*.

pC++ and its runtime system have been ported to several shared memory and distributed memory parallel systems, validating the system’s goal of portability. The ports include the Kendall Square Research KSR-1, Intel Paragon, TMC CM-5, IBM SP-1 / SP-2, Sequent Symmetry, SGI Challenge, Onyx, and PowerChallenge, BBN TC2000, Cray T3D, Meiko CS-2, and homogeneous clusters of UNIX workstations using PVM and MPI; work on porting pC++ to the Convex SPP is in progress. pC++ also has multi-threading support for running applications in a quasi-parallel mode on UNIX workstations; supported thread systems are Awesime [9], Pthreads, LWP, and the AT&T task library. This enables the testing and pre-evaluation of parallel pC++ applications in a familiar desktop environment. More details about the pC++ language and runtime system can be found in [1,16].

\mathcal{T} provides a collection of tools with user-friendly graphical interfaces to help a programmer analyze the performance of pC++ programs. Elements of the \mathcal{T} graphical interface represent objects of the pC++ programming model: collections, classes, methods, and functions. These language-level objects appear in all \mathcal{T} tools. By plan, \mathcal{T} was designed and developed in concert with the pC++ language system. It leverages pC++ language technology, especially in its use of the Sage++ toolkit [2] as an interface to the pC++ compiler for instrumentation and for accessing properties of program objects. \mathcal{T} is also integrated with the pC++ runtime system for profiling and tracing support. Because pC++ is intended to be portable, the tools are built to be portable as well. C++ and C are used to ensure portable and efficient implementation, and similar reasons led us to choose Tcl/Tk [20] for the graphical interface.

The \mathcal{T} tools are implemented as graphical *hypertools*. While the tools are distinct, providing unique capabilities, they can act in combination to provide enhanced functionality. If one tool needs a feature of another one, it sends a message to the other tool requesting it (e.g., display the source code for a specific function). With this design approach, the toolset can be easily extended. \mathcal{T} can also be retargeted to other programming environments, for instance those based on Fortran, which the Sage++ toolkit supports.

One important goal in \mathcal{T} 's development was to make the toolset as user-friendly as possible. For this purpose, many elements of the graphical user interface are analogous to *links* in *hypertext* systems: clicking on them brings up windows which describe the element in more detail. This allows the user to explore properties of the application by simply interacting with elements of most interest. The \mathcal{T} tools also support the concept of *global features*. If a global feature is invoked in any of the tools, it is automatically executed in all currently running tools. Examples of global features include locating information about a particular function or a class across all the tools. For good measure, \mathcal{T} comes with a full hypertext help system.

Figure 1 shows the pC++ programming environment and the associated \mathcal{T} tools architecture. The pC++ compiler frontend takes a user program and pC++ class library definitions (which provide predefined collection types) and parses them into an *abstract syntax tree* (AST). All access to the AST is done via the Sage++ library. Through command line switches, the user can choose to compile a program for profiling, tracing, and breakpoint debugging. In these cases, the instrumentor is invoked to do the necessary instrumentation in the AST. The pC++ backend transforms the AST into plain C++ with calls to the pC++ runtime system which supports both shared and distributed memory platforms. This C++ source code is then compiled and linked by the C++ compiler on the target system. The compilation and execution of pC++ programs can be controlled by *cosy* (**C**OMPile manager **S**tatus display); see 1 Figure 6, bottom. This tool provides a high-level graphical interface for setting compilation and execution parameters and selecting the parallel machine where a program will run.

The program and performance analysis environment is shown on the right side of Figure 1. They include the integrated \mathcal{T} tools, profiling and tracing support, and interfaces to stand-alone performance analysis tools developed

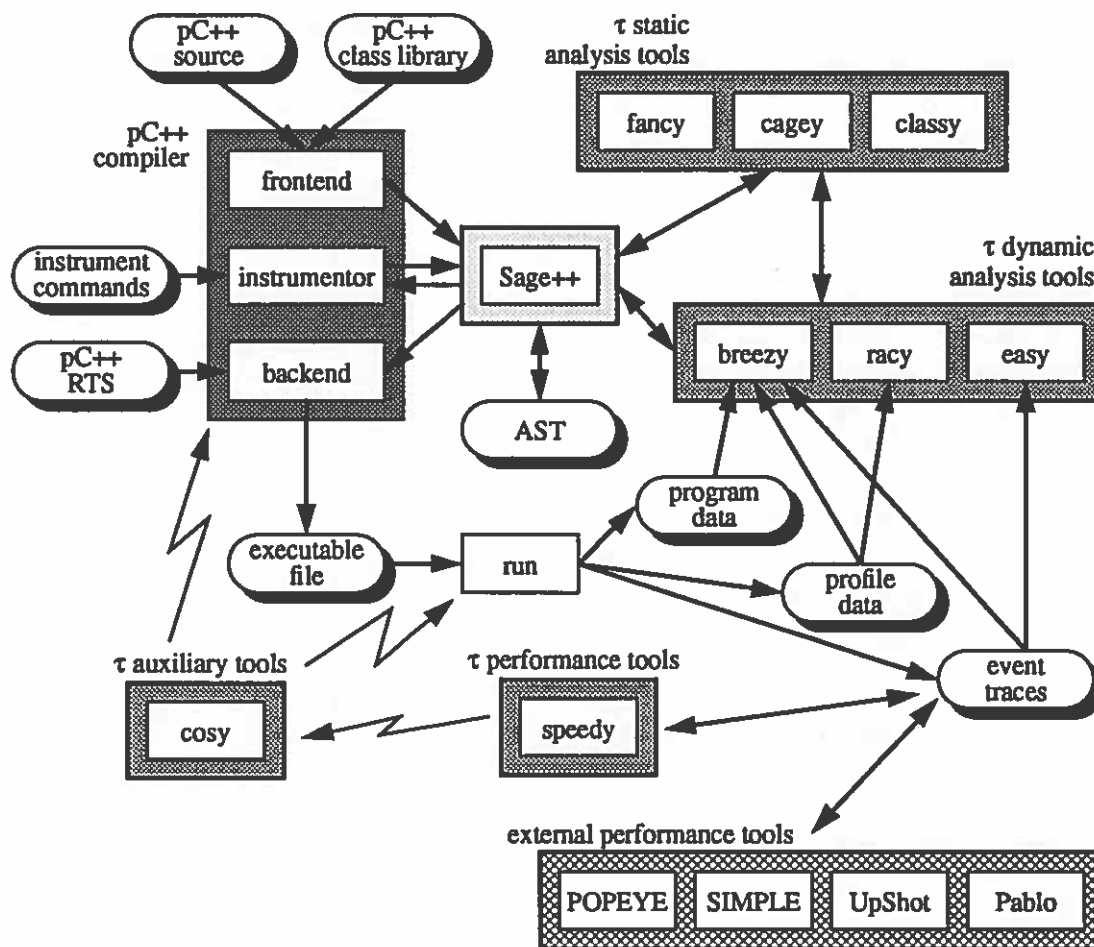


FIGURE 1. pC++ Programming Environment and τ Tools Architecture

partly by other groups [10,12,17,21]. The τ toolset provides support both for accessing static information about the program and for querying and analyzing dynamic data obtained from program execution.

2.1 Static Analysis Tools

One of the basic motivations behind using C++ as the base for a new parallel language is its proven support for developing and maintaining complex and large applications. However, to apply the C++ language capabilities effectively, users require support tools to manage and access source code at the level of programming abstractions. This is even more important for pC++ because of its data-parallel programming abstractions and SPMD execution paradigm.

Currently, τ provides three tools to enable the user to quickly get an overview of a large pC++ program and to navigate through it: the global function and method browser *fancy* (File ANd Class display), the static callgraph display *cagey* (CAll Graph Extended display), and the class hierarchy display *classy* (CLASS hierarchY browser). The tools are integrated with the dynamic analysis tools through the global features of τ , allowing the user to easily

find execution information about language objects. For instance, to locate the corresponding dynamic results (after a measurement has been made), the user only has to click on the object of interest (e.g., a function name in the callgraph display).

2.2 Dynamic Analysis Tools

Dynamic program analysis tools allow the user to explore and analyze program execution behavior. This can be done in three general ways. *Profiling* computes statistical information to summarize program behavior, allowing the user to find and focus quickly on the main bottlenecks of the parallel application. *Tracing* portrays the execution as a sequence of abstract *events* that can be used to determine various properties of time-based behavior. *Breakpoint debugging* allows a user to stop the program at selected points and query the contents of program state. For all analysis modes, the most critical factor for the user is how the high-level program semantics are related to the measurement results. \mathcal{T} helps in presenting the results in terms of pC++ language objects and in supporting global features that allow the user to locate the corresponding routine in the callgraph or source text by simply clicking on the related measurement result or state objects.

\mathcal{T} 's dynamic tools currently include an execution profile data browser called *racy* (**R**outine and **d**ata **A**ccess profile **d**isplay), an event trace browser called *easy* (**E**vent **A**nd **S**tate **d**isplay), and a breakpoint debugger called *breezy* (**B**reakpoint **E**xecutive **E**nvironment for **v**isualization and **d**ata **d**isplay). To generate the dynamic execution data for these tools, profiling and tracing instrumentation and measurement support has been implemented in pC++.

A more detailed discussion of the \mathcal{T} tools can be found in [5,16,18].

3 ExtraP - A Performance Extrapolation Tool for pC++

ExtraP is a performance extrapolation system for pC++ that has been integrated into \mathcal{T} in the form of *speedy*. This section explains the modeling approach of extrapolation and the techniques used by *ExtraP*.

3.1 Performance Extrapolation

Performance extrapolation (Figure 2) is the process of obtaining the performance information PI_1 of a parallel program for an execution environment E_1 and using PI_1 to predict the performance information PI_2^p (the superscript p indicates a predicted quantity) of the same program in a different execution environment E_2 . The performance information PI_2^p is then used to compute the predicted performance metrics of the program in E_2 , PM_2^p . This process can be considered as a translation or extrapolation of PI_1 to PI_2^p using the knowledge about E_1 and its similarities to and differences from E_2 .

As used above, an execution environment embodies the collection of compiler, runtime system, and architectural features that interact to influence the performance of a parallel program. Notice also how performance information is

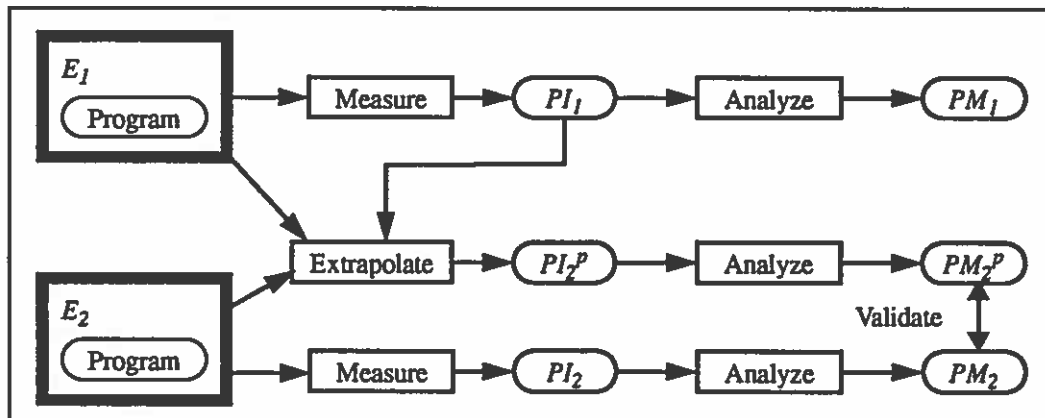


FIGURE 2. Performance Extrapolation

differentiated from performance metric. A performance *metric* is typically obtained by analyzing the static and dynamic performance *information* about a parallel program's execution. For example, the trace data obtained by running a parallel program can be condensed to calculate the execution time of the program and other metrics.

3.2 A Performance Extrapolation Technique for pC++

We have developed a technique that extrapolates the performance of a n -thread pC++ program from a 1-processor execution to a n -processor execution. In this technique, a n -thread pC++ program is executed on a single processor using a non-preemptive threads package. Important high-level events including remote accesses and barriers are recorded along with timestamps during the program run in a trace file. The instrumented runtime system is configured such that these remote accesses are treated as taking place instantaneously and the threads are seen to be released from a barrier as soon as the last thread enters it. Such a trace file captures the order of events in a pC++ program along with the computation time between the events, but leaves the actual timing of the events for later extrapolation analysis.

The events are then sorted on a per thread basis, adjusting their timestamps to reflect concurrent execution. This is possible because the non-preemptive threads package switches the threads only at synchronization points and global barriers are the only synchronization used by pC++ programs. This imposes a regular structure to the trace file where each thread records events between the exit from a barrier and an entry into another without being affected by another thread. The sorted trace files look as if they were obtained from a n -thread, n -processor run, except that they lack certain features of a real parallel execution. For example, the timings for remote accesses and barriers are absent in these trace files. A trace-driven simulation using these trace files attempts to model such features and predict the events as they would have occurred in a real n -processor execution environment. The extrapolated trace files are then used to obtain various performance metrics related to the pC++ program. The technique is depicted in Figure 3. For more details refer to [23,24]. The next section explains the various models used for trace-driven simulation.

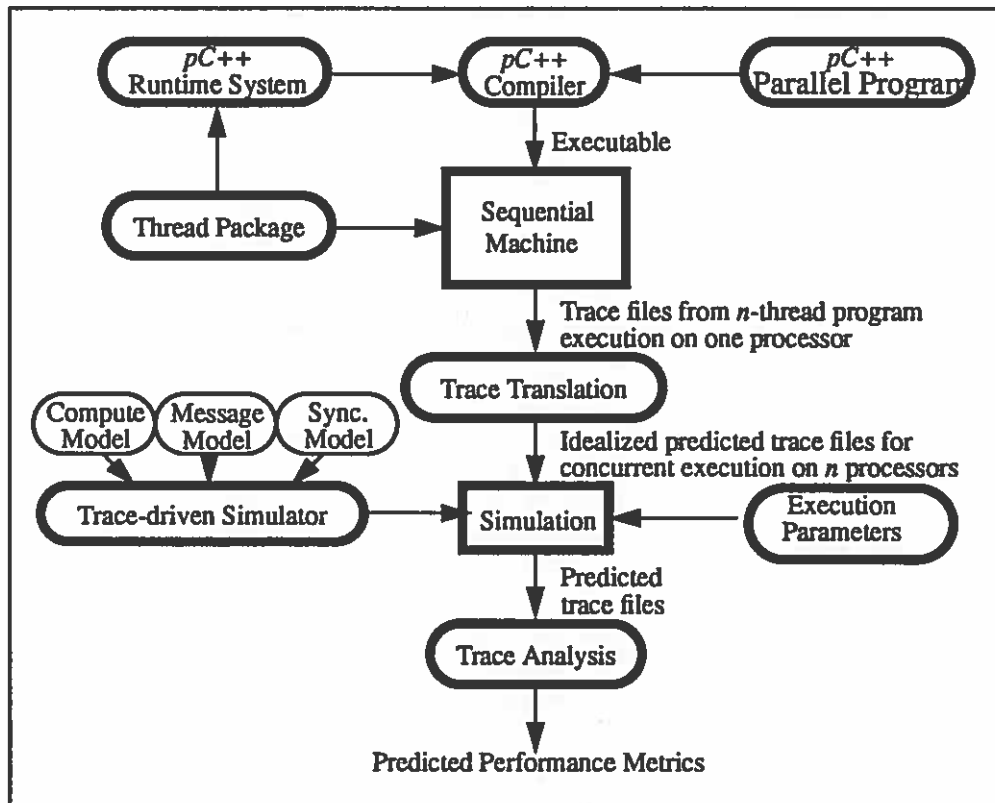


FIGURE 3. A Performance Extrapolation Technique for pC++

3.3 Simulation Architecture and Models

The trace-driven simulation is the heart of the pC++ performance extrapolation. The simulation system consists of three main components:

- Processor model
- Remote data access model
- Barrier model

The processor model uses a simple ratio of processor speeds to scale the computation time between events from the measurement environment to the target environment. It is also responsible for choosing a policy for servicing remote data references; when a request for data is submitted by a remote thread, a processor can service it in three different ways:

No interrupt: In this case, no messages are handled during the time between events. Messages are processed only when a thread waits for a barrier release or a remote data access reply.

Interrupt: In this case, an arrival of a message for a particular thread interrupts its computation. After the message is processed, the thread resumes its computation.

Poll: The scaled computation time between events is split into smaller chunks, and at the end of each chunk, the thread processes messages that have been received during that time.

The remote data access model determines how a remote data access is translated into messages and how it is handled by the various components in the system. During the simulation each remote access in the program is modelled as a remote request for data from one thread to the thread that “owns” the data (pC++ follows the “owner computes” model [8,13]). The owner thread services the request and returns the data to the requesting thread. This is equivalent to how the pC++ system operates in distributed memory environments. Hence, messages are the natural representation for the remote access protocol in the simulation. Figure 4 graphically depicts the how the remote data

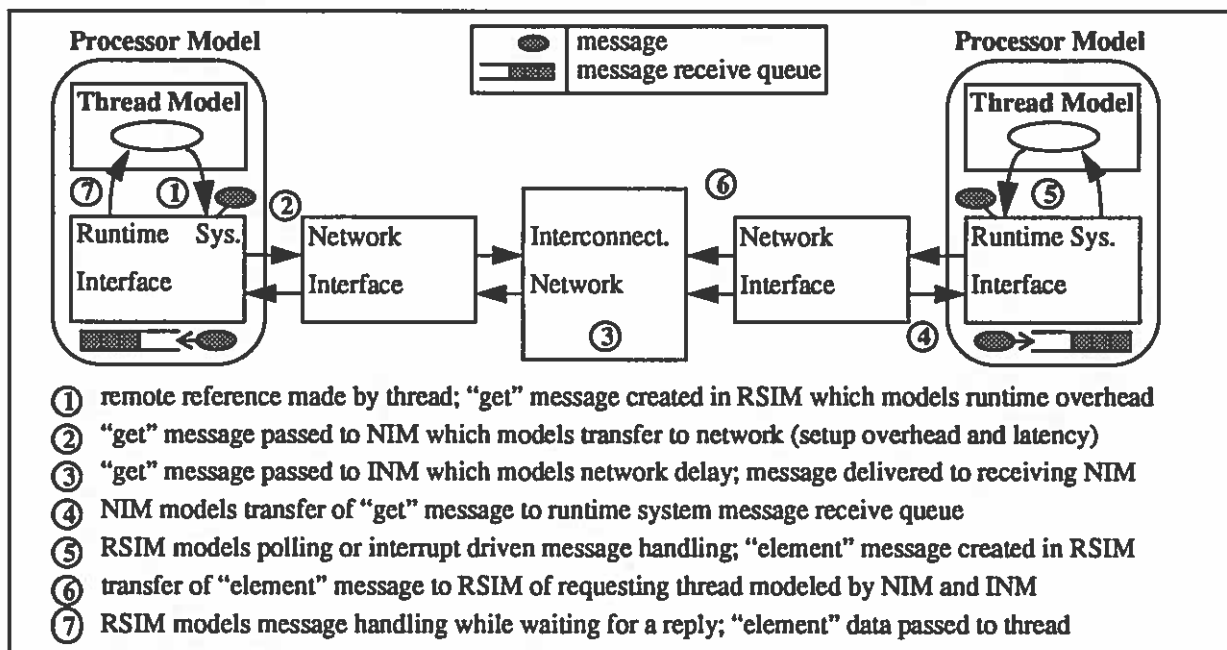


FIGURE 4. Remote Data Access Model

accesses are processed in the simulation using messages. The remote data access model itself is composed of the following subcomponents:

- runtime system model,
- network interface model and
- interconnect network model

Each of these models have various parameters that represent the characteristics of the remote data access model. For example, the interconnect network model includes the latency, bandwidth of the target platform as its parameters. For a complete list of parameters, refer to [23,24].

ExtraP uses a linear, master-slave barrier model to handle the barrier events. Thread 0 acts as the master thread while all the other threads are slaves. Every slave thread entering a barrier sends a message to the master thread and waits for a release message from the master thread to continue to the next data-parallel phase. The master thread waits for messages from all the slaves and then sends release messages to all of them. For distributed memory systems, the

pC++ runtime system must continue to service remote data access messages that arrive at a processor even when the threads that run on that processor have reached the barrier. This is also true in the simulation. The parameters in the barrier model can be controlled so that hardware barriers or barriers implemented through shared memory can be easily represented. The linear barrier model delivers an upper bound on barrier synchronization times. We can easily substitute other barrier algorithms (e.g. logarithmic) if a more accurate simulation of barrier operation is required.

All the three models described above and the three submodels of Remote Data Access model have a variety of parameters that can be tuned to match a specific target environment. For example, Table 1 lists the parameters used in the barrier model and their sample values. The next section explains how these parameters can be set and the extrapolation experiment carried out using \mathcal{T} . A new addition to \mathcal{T} called *speedy* interacts with *ExtraP* to perform the necessary extrapolation experiments. This integration of *ExtraP* with \mathcal{T} is of paramount importance because *ExtraP* is intended to be used as part of a program analysis environment to provide a performance debugging methodology. In [24] we showed how *ExtraP* can give good feedback about the performance of a pC++ program as the user attempts to tune the code to different target environments.

TABLE 1. Parameters for the barrier model

Parameter	Description	Example
<i>EntryTime</i>	Time for each thread to enter a barrier.	5.0 μ sec
<i>ExitTime</i>	Time for each thread to exit the barrier after it has been lowered.	5.0 μ sec
<i>CheckTime</i>	Delay incurred by the master thread every time it checks if all the threads have reached the barrier.	2.0 μ sec
<i>ExitCheckTime</i>	Delay incurred by a slave thread every time it checks to see if the master has released the barrier.	2.0 μ sec
<i>ModelTime</i>	Time taken by the master thread to start lowering the barrier after all the slaves have reached the barrier.	10.0 μ sec
<i>BarrierByMsgs</i>	1 - use actual messages for barrier synchronization. The message transfer time will contribute to the barrier time. 0 - do not use actual messages for barrier synchronization.	1
<i>BarrierMsgSize</i>	Size of a message used for barrier synchronization.	16

4 Integrating TAU and ExtraP

ExtraP is integrated with in two ways:

First, for generating the traces needed for the simulation, *ExtraP* uses the extensive built-in event tracing system of the pC++. The pC++ compiler has an integrated instrumentor module which allows the user to selectively insert trace recording instructions in the application program. In addition, there are instrumented versions of the predefined pC++ class/collection library and the pC++ runtime system. The inserted event markers in the different modules of the

system are assigned to *event classes* (e.g., user, runtime, collections,...) which can be activated or deactivated separately at runtime. Event tracing is fully operational on all parallel computer systems supported by pC++. This has several advantages for an *ExtraP* user. The traces used for simulation can be analyzed with all the event trace browsers supported by \mathcal{T} (currently *easy*, *Pablo* [21], *SIMPLE* [17], and *upshot* [12]). As the *ExtraP* model is based on the operational characteristics of pC++ event classes, the user can also generate semantically equivalent traces on real parallel computer systems for comparing to or validating the extrapolation results. Finally, the users can use \mathcal{T} 's integrated performance analysis tool, *racy*, to analyze their program execution and compare these to the simulated results. For example, in Figure 5, *racy* displays performance results for the pC++ Poisson benchmark executed on a 8 processor SGI PowerChallenge. *Racy* measures and displays function execution time profiles (shown on the left) and local/remote data access ratios (on the right).

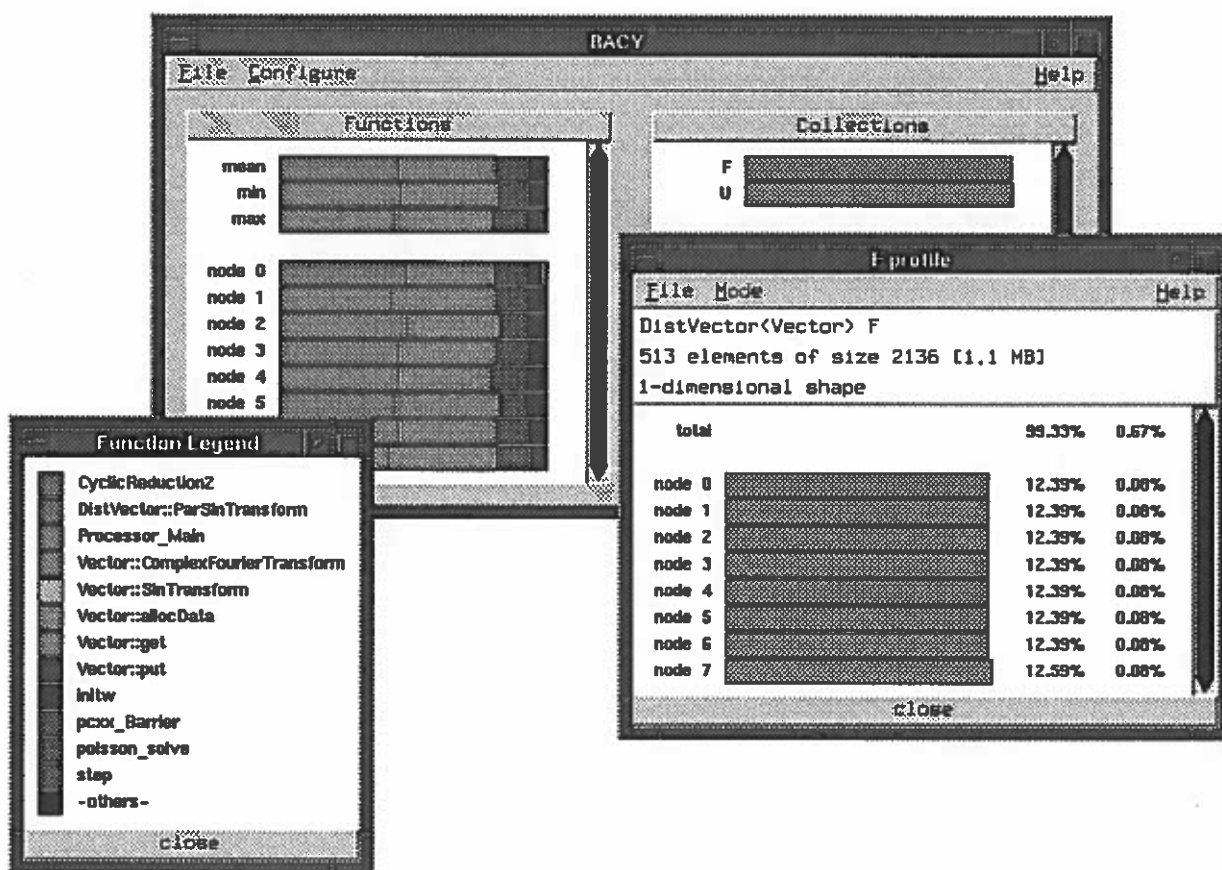


FIGURE 5. RACY Performance Analysis Display for Poisson Benchmark

Second, the actual extrapolation experiments can be controlled through a new \mathcal{T} tool *speedy* (Speedup and Parallel Execution Extrapolation Display). *Speedy* is a graphical tool which has been incorporated into the \mathcal{T} environment. Pressing the *speedy* button in the TAU main control window (see Figure 6, top), brings up its main control panel (see Figure 7). Here, the user can control the compilation of the specified pC++ object program, specify the parameters for the extrapolation model and the experiment, execute the experiment, and finally view the experiment results. *Speedy*

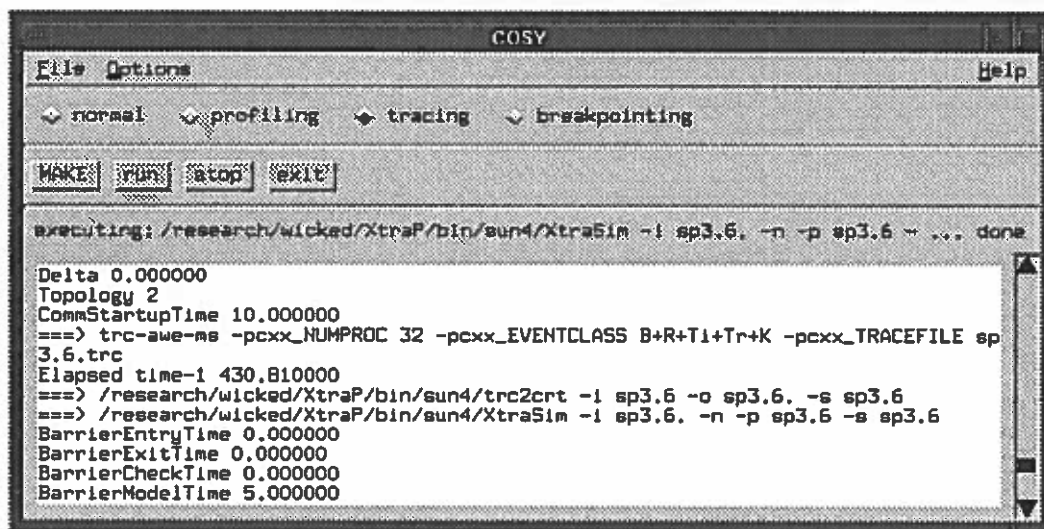
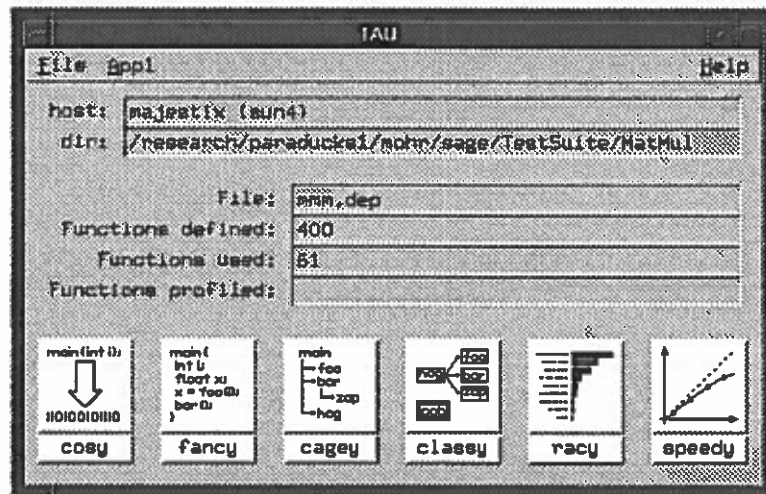


FIGURE 6. TAU Main Control Window and COSY

uses *cosy* (see Figure 6, bottom) for automatically performing the necessary compilation, execution, trace processing, and extrapolation commands. *Speedy* also automatically keeps track of all parameters by storing them in *experiment description files* and managing all necessary trace and experiment control files. By loading a former experiment description file into *speedy*, the user can re-execute the experiment or just reuse some of the parameter specifications.

In Figure 7, the user specified an experiment where the value of the parameter “Number of Processors” is stepping through powers of two starting from one to thirty-two. After each iteration of the extrapolation, the execution time graph is updated; a speedup graph can also be selected. The experiment and extrapolation model parameters can be entered and viewed through the *ExtraP* parameter file viewer (see Figure 8). Numerical parameters can either be entered directly into the input entry or manipulated through a slider bar and increment / decrement buttons. Parameters with discrete values can be specified through a pull-down menu. In Figure 8, the viewer displays the parameters associated with the modeling of the target processor of the target machine. Other modeling parameter




FIGURE 7. SPEEDY Main Control Panel

groups can be displayed by hitting one of the buttons in the top of the viewer window. Besides the five parameter groups described in Section 3.3, the group “General” allows to set parameters controlling the generation and post-processing of the execution traces. Discrete parameter values (like `ProcessMsgType` in the picture) can be changed through a pop-up menu, numerical values for parameters can either directly typed in or manipulated with the slider or the increment buttons in the lower half of the screen.

5 Experiments

In this section, we will see how *speedy* can be used during the development cycle of pC++ programs. *Speedy* allows the programmer to explore various design choices during the development process itself. In particular, we will show how the effects of various data distributions can be compared using *speedy*. *Speedy* also provides a framework in which the user can study different parts of the program and their contributions to the performance of the program. In this sense, *speedy* can be used a profile prediction tool for parallel programs. *Speedy* allows all of this to be done from a workstation environment without ever having to run the programs on target machines.

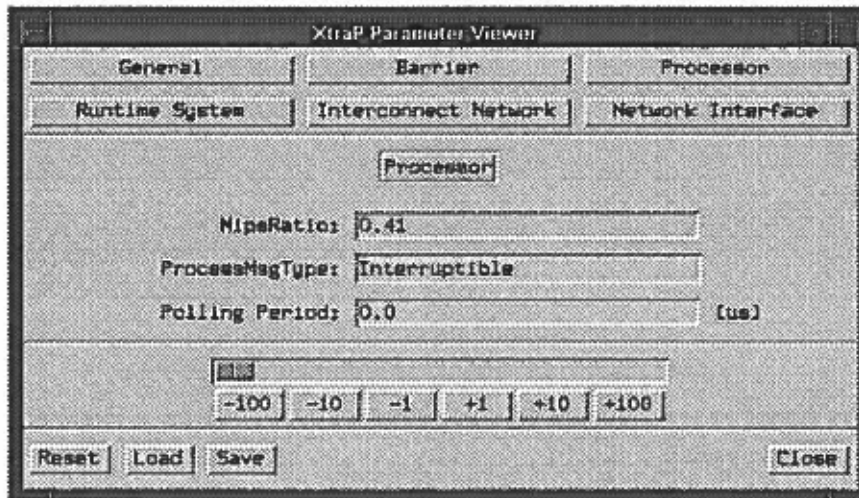


FIGURE 8. ExtraP Parameter File Viewer

Our experiments were done on a Sun workstation with a non-preemptive threads package called Awesime. The simulation parameters were tuned to match a CM-5 machine [14,19,24].

Our first experiment is designed to show how various design choices can be made using *speedy*. We chose *MatMul*, a simple matrix multiplication program written in pC++, where the two-dimensional data distributions can be changed to improve the performance. We ran the experiments with two distributions: (BLOCK, BLOCK), where the matrix is divided into blocks of equal size and each processor is assigned a block, and (BLOCK, WHOLE), where the rows of the matrix are equally divided between the processors. The predicted results are shown in Figure 9. The stepwise

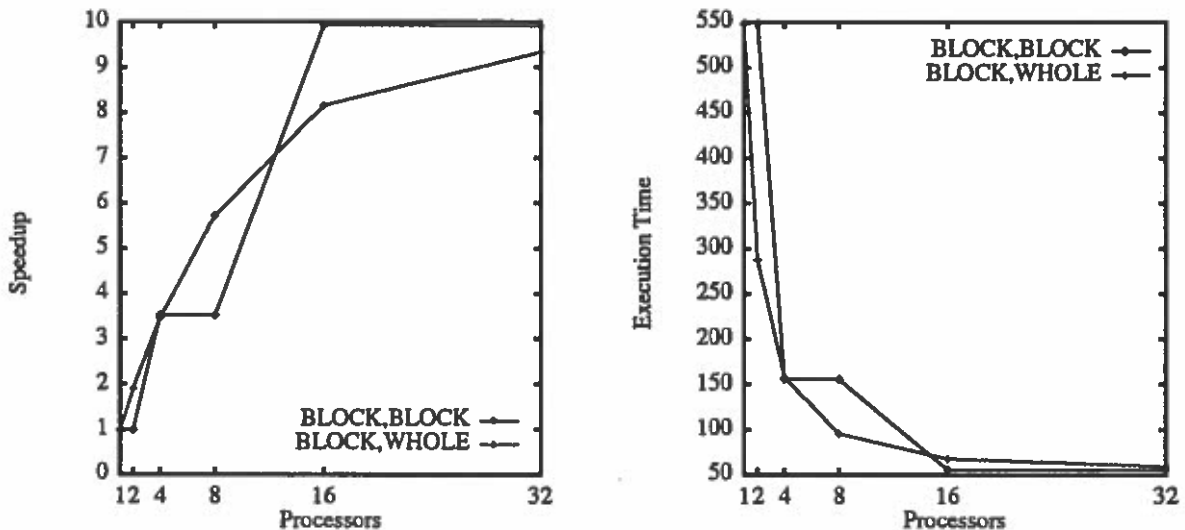


FIGURE 9. Predicted results for MatMul

behavior of (BLOCK, BLOCK) distribution arises from the fact that the distribution changes only when the number of processors is a perfect square. So after 4 processors, the next improvement in performance will only be for 9 processors, and then for 16 and so on (however, the 9 processor case is not a sample point in our experiment). The results also show that even with the stepwise behavior (BLOCK, BLOCK) starts winning after 16 processors. Such crossover point information is very useful for the programmer during the development process. Validation against the actual CM-5 performance was accurate [24].

The goal of our next experiment is to show how *speedy* can be used to selectively study various portions of the program. Such profile information is useful when the programmer wants to tune parts of the program for a particular machine. We used *Poisson* as the test case. It is a fast poisson solver written in pC++ which uses FFT based sine transforms together with a cyclic reduction algorithm to solve PDEs. We used *speedy* to selectively instrument the code for the transforms and cyclic reduction. After extrapolating the performance to a CM-5 architecture, *speedy* predicted the results shown in Figure 10. While the code for sine transforms scales up very well with a speedup of

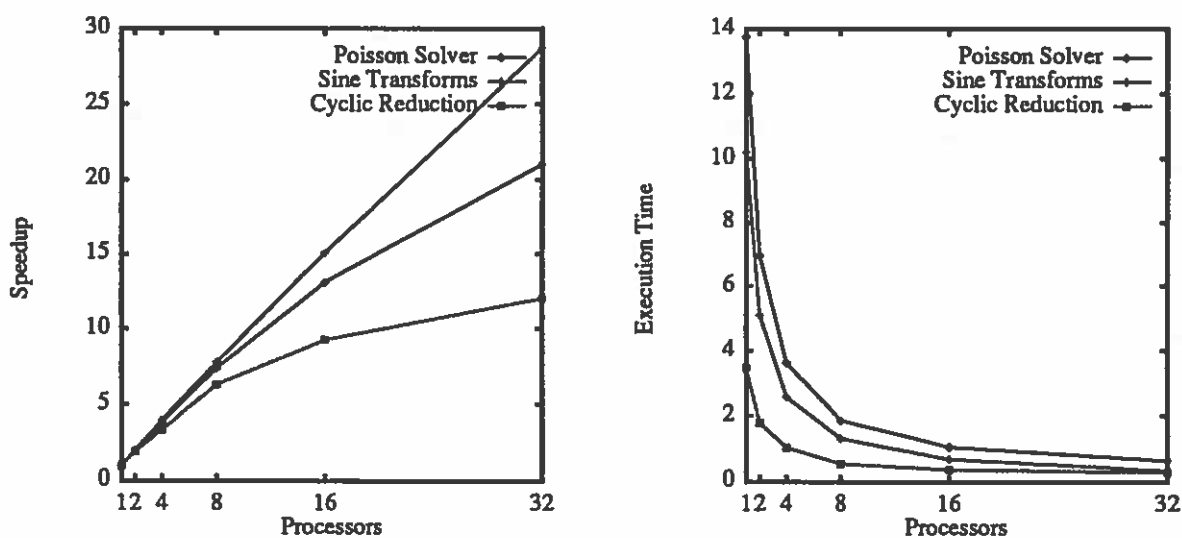


FIGURE 10. Predicted results for Poisson

28.69 for 32 processors, the speedup curve for cyclic reduction starts to flatten after 16 processors. A further study of the trace files revealed that there are no remote accesses in the sine transform part of the poisson solver which accounts for its near-linear speedup. In contrast, the number of remote accesses in cyclic reduction increases as the number of processors thus affecting the performance badly. The overall speedup for poisson solver is predicted to be in between that of sine transforms and cyclic reduction. This experiment tells us that to improve the performance of Poisson, we must tune Cyclic reduction first because it is the bottleneck. *Speedy* can be used in this way to locate the bottlenecks in a program. The performance behavior observed using *speedy* is consistent with actual results [16].

6 Conclusion

The *speedy* and *ExtraP* tools are representative of the level of parallel performance evaluation support that is expected to be available for high-level parallel language systems and to be integrated in program analysis environments where a performance engineered code development process is desired. The requirement for performance prediction as part of this process, at least in pC++, is driven by the need to evaluate parallel codes that are intended to be ported to different execution platforms. Furthermore, the integration aspects (e.g., of merging *ExtraP* into \mathcal{T}) are of key importance as the extrapolation techniques must leverage the sophisticated compiler, runtime system, and tool infrastructure to make the application of performance prediction in parallel code development feasible. Our future work will concentrate on making the *ExtraP* technology more robust with additional models so that the different target system environments can be better represented. We also intend to extend the capabilities of the *speedy* tool to provide more support for automated performance experimentation and to better link the \mathcal{T} analysis and visualization tools to the performance data that *ExtraP* produces.

7 For more information...

Documentation, technical papers, and source code for pC++, Sage++, and \mathcal{T} are available via anonymous FTP from `ftp.extreme.indiana.edu` in the directory `~ftp/pub/sage`, or via the World Wide Web at the URLs `http://www.extreme.indiana.edu/sage` and `http://www.cs.uoregon.edu/paracomp/tau`.

8 References

- [1] F. Bodin, P. Beckman, D. Gannon, S. Yang, S. Kesavan, A. Malony, B. Mohr, Implementing a Parallel C++ Runtime System for Scalable Parallel Systems, Proceedings Supercomputing 93, IEEE Computer Society and ACM SIGARCH, pages 588-597, November 1993.
- [2] F. Bodin, P. Beckman, D. Gannon, J. Gotwals, S. Narayana, S. Srinivas, B. Winnicka, Sage++: An Object Oriented Toolkit and Class Library for Building Fortran and C++ Restructuring Tools, Proceedings Oonski '94, Oregon, 1994.
- [3] E. A. Brewer, C. N. Dellarocas, A. Colbrook, W. E. Wehl, Proteus: A High Performance Parallel Architecture Simulator, Technical Report MIT/LCS/TR-516, MIT Laboratory of Computer Science, September 1991.
- [4] E. A. Brewer, W. E. Wehl, Developing Parallel Applications Using High-Performance Simulation, Proc. ACM/ONR Workshop on Parallel and Distributed Debugging, pages 158-168, May 1993.
- [5] D. Brown, S. Hackstadt, A. Malony, B. Mohr, Program Analysis Environments for Parallel Language Systems: The TAU Environment, Proc. of the Workshop on Environments and Tools For Parallel Scientific Computing, Townsend, Tennessee, pages. 162-171, May 1994.
- [6] M. E. Crovella and T. J. LeBlanc, Parallel Performance Prediction Using Lost Cycles Analysis, Proc. Supercomputing 94, IEEE Computer Society and ACM, pages 600-609, November 1994..

- [7] T. Fahringer, Automatic Performance Prediction for Parallel Programs on Massively Parallel Computers, Ph.D. Thesis, University of Vienna, September 1993.
- [8] D. Gannon and J. K. Lee, Object Oriented Parallelism: *pC++* Ideas and Experiments, Proc. Japan Society of Parallel Processing, pages 13-23, 1991.
- [9] D. C. Grunwald, A Users Guide to AWESIME: An Object Oriented Parallel Programming and Simulation System, Technical Report 552-91, Department of Computer Science, University of Colorado at Boulder, November 1991.
- [10] S. Hackstadt, A. Malony, Next-Generation Parallel Performance Visualization: A Prototyping Environment for Visualization Development, Proc. Parallel Architectures and Languages Europe, (PARLE), Athens, Greece, 1994.
- [11] R. Helm, A. D. Malony and S. F. Fickas, Capturing and Automating Performance Diagnosis: The Poirot Approach, Proc. International Parallel Processing Symposium
- [12] V. Herrarte, E. Lusk, Studying Parallel Program Behavior with Upshot, Technical Report ANL-91/15, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.
- [13] High Performance Fortran Forum, High Performance Fortran Language Specification version 1.0, TR92225, Center for Research on Parallel Computation, Rice University, January 1993.
- [14] S. Hiranandani, K. Kennedy, C.-W. Tseng, S. Warren, The D Editor: A New Interactive Parallel Programming Tool, Proc. Supercomputing '94, IEEE Computer Society Press, pages 733-742, November 1994.
- [15] J. Kohn and W. Williams, ATExpert, Journal of Parallel and Distributed Computing, Vol. 18, 1993, pages 205-222.
- [16] A. Malony, B. Mohr, P. Beckman, D. Gannon, S. Yang, F. Bodin, Performance Analysis of *pC++*: A Portable Data-Parallel Programming System for Scalable Parallel Computers, Proc. 8th Int. Parallel Processing Symb. (IPPS), Mexico, IEEE Computer Society Press, pages 75-85, April 1994.
- [17] B. Mohr, Standardization of Event Traces Considered Harmful or Is an Implementation of Object-Independent Event Trace Monitoring and Analysis Systems Possible?, Proc. CNRS-NSF Workshop on Environments and Tools For Parallel Scientific Computing, St. Hilaire du Touvet, France, Elsevier, Advances in Parallel Computing, Vol. 6, pages 103-124, 1993.
- [18] B. Mohr, D. Brown, A. Malony, TAU: A Portable Parallel Program Analysis Environment for *pC++*, Proc. of CONPAR 94 - VAPP VI, Linz, Austria, Springer Verlag, LNCS 854, pages 29-40, September 1994.
- [19] On-line information on CM-5, CM-5 Technical Summary, <http://www.think.com/Prod-Serv/Products/cmmd.html>, November 1992.
- [20] J. Ousterhout, Tcl and the Tk Toolkit, Addison-Wesley, 1994.
- [21] D. A. Reed, R. D. Olson, R. A. Aydt, T. M. Madhyasta, T. Birkett, D. W. Jensen, B. A.A. Nazief, B. K. Totty, Scalable Performance Environments for Parallel Systems. Proc. 6th Distributed Memory Computing Conference, IEEE Computer Society Press, pages 562-569, 1991.
- [22] S. K. Reinhardt, M. D. Hill, J. R. Larus, A. R. Lebeck, J. C. Lewis and D. A. Wood, The Wisconsin Wind Tunnel: Virtual Prototyping of Parallel Computers, Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, pages 48-60, May 1993.
- [23] K. Shanmugam, Performance Extrapolation of Parallel Programs, Master's Thesis, Department of Computer and Information Science, University of Oregon, June 1994.

- [24] K. Shanmugam, A. D. Malony, Performance Extrapolation of Parallel Programs, Submitted to ICPP 95.
- [25] H. Wabnig and G. Haring, PAPS - The Parallel Program Performance Prediction Toolset, Computer Performance Evaluation - Modelling Techniques and Tools, Lecture Notes in Computer Science 794, Springer-Verlag, pages 284-304, 1994.