

MEASURING HIGH PERFORMANCE COMPUTING PRODUCTIVITY

Stuart Faulk¹
John Gustafson²
Philip Johnson³
Adam Porter⁴
Walter Tichy⁵
Lawrence Votta⁶

Abstract

One key to improving high performance computing (HPC) productivity is to find better ways to measure it. We define productivity in terms of mission goals, i.e. greater productivity means that more science is accomplished with less cost and effort. Traditional software productivity metrics and computing benchmarks have proven inadequate for assessing or predicting such end-to-end productivity. In this paper we introduce a new approach to measuring productivity in HPC applications that addresses both development time and execution time. Our goal is to develop a public repository of effective productivity benchmarks that anyone in the HPC community can apply to assess or predict productivity.

Key words:

1 Introduction

While computer performance has improved dramatically, real productivity in terms of the science accomplished with these ever-faster machines has not kept pace. Indeed, scientists are finding it increasingly costly and time-consuming to write, port, or rewrite their software to take advantage of the new hardware. While machine performance remains a critical productivity driver for high performance computing (HPC) applications, software development time increasingly dominates hardware speed as the primary productivity bottleneck. Traditional benchmarks do not provide a way to measure productivity.

Removing or ameliorating productivity bottlenecks in next-generation HPC systems is a key objective of the Defense Advanced Research Projects Agency (DARPA) High Productivity Computing Systems (HPCS) program. This objective has been characterized as the long-term goal of improving software productivity “at a rate commensurate with improvements in the underlying hardware” (Kepner, 2004).

Stating the goal this way reflects the scientific community’s experience that improved hardware performance often does not yield similar improvements in important measures such as total cost, effort, or time to solution. Addressing these larger productivity issues will require addressing productivity bottlenecks in software development, project management, and system administration in addition to hardware performance.

Our work focuses on developing an effective approach to characterizing and measuring HPCS software productivity. Before we can hope to address software productivity problems, we must agree on what those problems are and how to determine whether we have solved them. In short, we must agree on what we mean by “productivity” in HPCS applications, and how such productivity can be measured, before we can rationally address the problem of how best to improve it. Goals of the work include the following.

¹COMPUTER AND INFORMATION SCIENCE, UNIVERSITY OF OREGON, EUGENE, OR 97405, USA (FAULK@CS.UOREGON.EDU)

²ECS INC., PLEASANTON, CA 94588, USA

³COLLABORATIVE SOFTWARE DEVELOPMENT LABORATORY, UNIVERSITY OF HAWAII, HONOLULU, HI 96822, USA

⁴COMPUTER SCIENCE DEPARTMENT, UNIVERSITY OF MARYLAND, MD 20742, USA

⁵INFORMATICS, UNIVERSITY OF KARLSRUHE, 76128 KARLSRUHE, GERMANY

⁶SUN MICROSYSTEMS INC., MENLO PARK, CA 94025, USA

1. *Develop a common definition of HPCS productivity* that the HPCS developers, suppliers, and buyers (e.g. government agencies) can agree on. For our purposes, this means a definition that is consistent with the mission-level view that greater productivity means that more science is accomplished with less cost and effort.
2. *Develop effective measures of HPCS productivity* that encompass the overall development process – design time as well as execution time. In particular, we seek to develop measures that apply to a wide range of development environments and broadly across HPC application domains (e.g. weather prediction, fluid dynamics, nuclear applications, mechanical design, etc.) to assess, compare, or predict productivity.
3. *Provide productivity measurement capabilities* to guide productivity improvement for both hardware and software developers. Effective, objective measurement provides the basis for systematic productivity improvement. We seek to provide common, public benchmarks, and metrics to use in assessing and improving productivity.

Both HPCS developers and buyers have traditionally used standardized benchmarks (e.g. LINPACK) to guide development choices. System developers use the benchmark results to guide platform development and subsequently demonstrate the speed of their machines. Buyers traditionally use such benchmarks to predict computation times and choose among competing platforms. However, the benchmarks and corresponding metrics employed to date have proven to be decreasingly effective predictors of end-to-end productivity. Traditional benchmarks focus almost entirely on hardware speed. Thus, they typically attempt to predict only execution-time productivity, ignoring development time. Further, they do not measure other properties of an application that matter to users: reliability, repeatability, portability, reusability, maintainability, etc.

In this paper we describe our approach to creating a new type of standardized benchmark that encompasses (1) the breadth of design time and execution time activities as well as (2) the productivity contributions of both functional and non-functional requirements. In addition to defining a canonical computation problem, these “productivity benchmarks” seek to characterize an end-to-end productivity problem by capturing the representative context of the computation. This includes the overall process in terms of canonical workflows (Kepner, 2004) as well as the developmental attributes that contribute to the overall value of the software to its stakeholders. We will call such a multidimensional productivity benchmark a “productivity benchmark suite” (PBS).

Briefly, a PBS comprises a canonical problem in the context of behavioral and developmental requirements representative of a particular HPC domain. In addition to the functional and non-functional requirements, the PBS will provide targeted metrics and tools for measuring productivity in terms of overall costs and benefits across the development cycle. The goal is to create a set of benchmarking capabilities that, when applied, will exercise and measure not only the execution efficiency of a platform on a particular class of HPC problems, but all the dimensions of development that contribute to the value of a solution.

Our long-range goal is to develop a public repository of well-validated PBSs that are representative of the productivity challenges in each distinct HPC domain. Platform developers or buyers can apply these PBSs to assess and predict productivity of particular HPC platforms on their domains of interest, in particular the following.

1. **HPCS buyers:** buyers of HPCS platforms are faced with the problem of predicting which vendor’s system will provide the best “bang for the buck” in their application domain. This question embraces the total cost of ownership including software development, execution, operation, maintenance, and so on. Buyers will be able to apply benchmarks from the repository to answer specific questions about how different platforms or development strategies can be expected to affect their productivity.
2. **HPCS developers:** HPCS platform vendors seek to develop systems that will improve productivity for their customer base. To do this, they must understand what the productivity problems are in their customer’s application domains, what properties of the platform affect those productivity issues, and how to measure the results of platform design changes. Vendors will be able to apply benchmarks to guide architectural design.

A goal of our work is to provide the HPCS community with the capabilities needed to determine whether new technologies will effectively address critical productivity problems before vendors deploy those technologies in next-generation platforms. This will reduce the risk that next-generation high performance systems will fail to meet DARPA’s HPCS productivity goals.

Our overall approach is empirical in the sense that we will develop PBSs based on observations of real developers and their productivity problems with actual high performance applications. We will ensure validity of the benchmarks and metrics through carefully controlled experiments.

The following sections describe the planned approach and expected results in detail. First, we describe a frame-

work for reasoning about HPCS productivity and illustrate its application with a real application. We then describe ongoing work in developing a framework for PBSs and our approach to productivity measurement. We illustrate a new approach to measuring a system's developmental qualities (e.g. maintainability). Finally, we discuss ongoing work in creating and validating productivity benchmarks.

2 Reasoning about HPCS Productivity

The economic definition of productivity is the “output per unit-of-work”. For example, “worker productivity” denotes the value of goods and services produced in a period divided by the hours of labor used to produce them. While the units vary depending on the realm of discourse, all definitions of productivity share the underlying concept that increasing productivity means producing more of value with less work.

Conceptually, the economic definition of productivity is both simple and intuitively appealing. It is also consistent with the objectives of the HPCS program, i.e. providing platforms that allow scientists to accomplish more science in less time at lower cost. Unfortunately, the economic definition has proven difficult to apply to software development. This is true not just for HPC applications, but for the broad classes of software applications more typically addressed in software engineering productivity studies (e.g. Boehm and Wolverton, 1980; Boehm, 1981).

The mutable, intangible nature of both the processes and the products of software development make the outputs or units-of-work difficult to define or measure. The default has been to choose metrics that are relatively easy to measure, but that bear only a loose relationship to the value of what is produced.

Our goal is to address these issues by providing a framework for characterizing and measuring the perceived value of the output to system stakeholders. We define the output to include any properties of the system that consume work and have stakeholder value, including those that have no direct physical analog in the code (e.g. usability).

In the remainder of this section we describe the conceptual framework for our approach. We first discuss some of the historical difficulties and pitfalls associated with software productivity metrics. We then describe a general framework for modeling a software development's output value based on the stakeholder values of developmental properties and requirements. We argue that the proposed framework is sufficiently flexible to address productivity issues across development time and run time. We then walk through an example taken from requirements for a high performance application that illustrates consistency of the model with a real HPC application.

2.1 PITFALLS OF TRADITIONAL PRODUCTIVITY MEASURES

Software productivity has been and remains a core mission of software engineering. None the less, problems in adequately measuring software productivity, much less predicting or improving it, have proven among the most intractable in the field. While there is general agreement that software productivity should reflect the economic definition, output value per unit-of-work, there is little agreement on how to define the outputs or the units of work. In general, industry experience has been that those software development properties that we can most easily and precisely measure (e.g. the number of lines of code produced) have little direct relationship to the system's stakeholder value while many properties with established value (e.g. maintainability) have no established, effective metric.

The long-running controversy over the most common software productivity metric, the number of source lines of code (SLOC) developed, illustrates the types of problems that arise. SLOC became one of the first commonly accepted productivity metrics (Boehm, 1981) and remains in regular use (Boehm et al., 1995). However, if applied incautiously, SLOC introduces a number of uncertainties and potential paradoxes. The number of lines of code necessary to implement a particular functionality will vary greatly from one programming language to another. Worse, they can vary inversely with the quality of the programmer and efficiency of the code (i.e. better programmers do more with less code). If used without careful controls, SLOC productivity measures may indicate that productivity has decreased when it has actually increased (Jones, 1986). In short, creating lines of code is a necessary but not sufficient condition to creating software of value. There is no predictable relationship between SLOC (or similar metrics such as function points; Albrecht and Gaffney, 1983) to critical code properties that enhance value such as maintainability, portability, reusability, security, and so on.

Paradoxes and uncertainty likewise attend the use of common units-of-work such as labor hours or, more typically in software, man-months. Brooks (1995) first articulated some of the problems in using conventional units of work to measure or predict software productivity in his book, *The Mythical Man-Month*; most famously, he observed that adding manpower to a late software project typically made it later. Subsequent empirical software engineering research at the National Aeronautics and Space Administration (NASA) Goddard Software Engineering Laboratory (e.g. Porter et al., 1995) underscored the diversity of factors that affect the cost and quality of software production. Likewise, detailed time-motion studies of software developers (Perry et al., 1996) suggest that even the people doing the work do not accurately predict

or even recall exactly what they spend their time on. Such work indicates that making reasonably accurate correlations between units-of-work and produced value requires careful, detailed empirical study of the development process in a controlled context.

Similar issues arise with productivity measures in HPC. Traditionally, measures of machine performance such as the peak number of floating-point operations per second (FLOPS) are used as a predictive measure of HPC output. Historically, such measures had validity where the total time to solution was overwhelmingly constrained by program execution time, and where the program execution time was overwhelmingly dominated by the time spent performing floating-point operations. As the effort needed for developing, porting, rewriting, and tuning the software has increased, the relevance of such measures has correspondingly decreased. (Floating-point arithmetic also no longer dominates execution time even in the most compute-intensive programs, yet the focus on FLOPS persists.) Likewise, as program properties other than execution performance (e.g. portability, maintainability, etc.) have become more important production values, the need for metrics that account for these properties has increased as well.

Early productivity studies observed the phenomenon that Weinberg and Schulman (1974) characterized as “what-you-measure-is-what-you-get” (WYMIWYG). Repeated studies have consistently confirmed this observation’s predictive power. If one measures productivity in terms of the number of lines of code, programmers will produce reams of code. If one measures HPCS productivity in terms of FLOPS, then we will have processors that show very high arithmetic rates on carefully chosen benchmarks. In neither case will we necessarily obtain any improved productivity in the sense of more functions implemented or more science done per dollar spent.

An important corollary is that the reverse also holds true. In short, if we want to get something, we should measure it. This principle dictates, in particular, that if we want developmental properties such as portability, interoperability, maintainability, etc., in addition to execution time properties such as performance and accuracy, then we need to measure all of those properties as directly as possible. Even if we ignore past pitfalls, this principle suggests that we cannot use a single metric that obscures the productivity contribution of each critical property. Rather, we will need measures that explicitly address development output as comprising multiple, diverse properties of value.

2.2 AN HPCS PRODUCTIVITY FRAMEWORK

Intuitively, each application development has associated with it some value to its stakeholders. The overall value

will typically be due to a number of different properties of the code, its execution behavior, and its development process. Exactly which properties have value and how much value is attributed to them will vary from one type of application to the next and from one class of stakeholder to the next. In particular, we observe the following.

1. The overall value is a function of a number of distinct properties of the static work products, the code, and the development process itself. In addition to accuracy and timeliness of results, these may include run-time properties such as security, availability, and locality. It may also include desirable developmental properties such as portability, maintainability, or reusability or even organizational concerns such as consistency with organizational standards or the amount of legacy code reused.
2. Exactly which attributes of the system are important will vary from one type of application to the next, from one class of stakeholders to another, and possibly even from one run of the application to another.
3. The relative importance or priority of each attribute may likewise vary by application type, stakeholder class, and so on.
4. The relative values of individual properties as well as the total value of a given solution may change over (calendar) time (Snir and Bade, 2004).

In short, the value associated with any particular development is typically a function of a number of different properties that we can expect to vary from one development effort to the next or even one run of the program to another. We can capture this by representing the total relative value as a vector over the values of the properties of interest using the following framework. We begin by associating with each property of interest i :

1. a metric of completion C_i ;
2. a relative value weight v_i .

Briefly, the metric of completion C_i is a measure of the degree to which the realization of the property i meets stakeholder goals for that property. The relative value weight v_i represents the importance of the property i relative to the other properties of interest. The value of some set of properties $i = 1$ to n is given by the vector:

$$V_A = (v_1 C_1, v_2 C_2, \dots, v_n C_n). \quad (1)$$

Assuming independence and that we can normalize each of the $v_i C_i$ to a common metric (e.g. labor or cost), we can express the total value as the sum.

$$V_A = v_1 C_1 + v_2 C_2 + \dots + v_n C_n. \quad (2)$$

For example, we could calibrate each C_i such that $C_i = 1$ whenever property i meets its design goals, and v_i gives the relative importance of property i expressed as a percentage such that $\sum_{i=1}^n (v_i) = 100$ and $V_A = 100$ exactly when all the C_i are satisfied. Relative productivity is then given by the value produced divided by the work consumed to produce it:

$$P = V_A / W. \quad (3)$$

Equivalently, we can say that the greater the value of V_A for a given amount of work, the higher the productivity. This corresponds to our intuitive view that greater productivity implies greater value per unit of work.

Both equations (1) and (2) are useful. Equation (1) has the advantage that the individual contributions of each property of interest to the total value are clear and that we need not attempt to normalize over different kinds of properties. However, equation (2) has the advantage of summarizing the overall value in a way that might be compared across projects if normalized over a common metric like cost.

We anticipate further refinement to our equations to reflect significant dependences. In the general case, the degree of completeness and value of one system property or even a set of properties may depend on others. For example, it does not make sense to talk about the “value” of properties such as maintainability or portability if the code does not do what it is supposed to. Further, the values of individual properties, as well as the program as a whole, are likely to change over time. A solution today is typically worth more than the same solution tomorrow. Thus, the value weight v_i associated with each property i may be a function of calendar time t as can V_A itself. Understanding the nature of these relationships, their significance, and how best to represent them will require careful study of different classes of HPCS applications.

2.3 APPLICABILITY OF THE VALUE MODEL

By design, our value function must be used in the context of a computing application that establishes the value space of interest. For productivity benchmarking, this context will be given by the PBS. The definition of the PBS will include the definitions of the properties of interest, corresponding metrics of completion, and representative value weights. Appropriate properties and values will be obtained from empirical studies of representative development efforts in the application area of the benchmark; i.e. for a benchmark simulating behavior of a weather code, representative properties, values, and completion criteria would be gleaned from the weather simulation community. This framework will then be tailored to the intended

use of the benchmark. If particular properties are irrelevant, they can be omitted (equivalently, given zero values).

The basis for constructing such productivity models as well as the applicability of our framework can be illustrated by considering an example from a real high performance application. Tables 1 and 2 are taken from the Software Development Plan (SDP) for Virtual Prototyping and Accelerated Testing of DoD Composite Material Combat Systems (VPATC, 2003) and are part of the definition of system requirements and constraints.

Table 1 gives a subset of the critical operational issues (COI) and measures of effectiveness and suitability (MOE&S) defined for the VPATC system. Table 2 defines what are called the system’s critical technical parameters (CTPs). It maps the technical parameter values that should be measured back to MOE&S in Table 1 and gives both the minimal and optimal criteria for the CTPs for the scheduled sequences system testing milestones. This sequence begins with the system acceptance test (SAT), follows with the Alpha and Beta tests, and finishes with the initial operational test and evaluation (IOT&E).

Clearly, the value of the VPATC system is not a function of its execution performance alone or even of its computational behavior as a whole. A given implementation of the system requirements will be acceptable only if it meets the minimum criteria for the COI as detailed in Table 2. For example, a system will be acceptable only if its execution performance meets the criteria that “fixed speedup exceeds 60% of optimum on 32 or more processors” and the code meets the portability requirement of “running on three HPC platforms with the same valid results.” Thus, we observe the following.

1. To satisfy requirements, the system must meet development goals for a number of distinct properties concurrently with functional requirements.
2. The properties of interest span development and execution time. They include (static) developmental properties (e.g. COI 4: maintainability and adaptability), execution time properties (e.g. COI 1: performance), and organizational properties (e.g. COI 5: training and technology transfer).
3. Associated with each property is some testable metric of completion. For example, portability is measured based on the number of HPC platforms that will yield valid results when running the same code.
4. Notions of value, while implicit, are present in rudimentary form. For example, it is clear that an implementation of the system that meets the “optimum objectives” listed in Table 2 has greater value than an implementation that meets only “minimum objectives.” A rough value scale could be constructed based on the number of objectives satisfying optimum criteria.

Table 1
Critical operational issues and measures of effectiveness and suitability.

COI	COI title	COI description	MOE&S
1	Performance	Does the project provide computational results that are accurate, stable, and reliable in a portable scalable environment?	Scalable performance improvement over current systems Software robustness
2	Interoperability	Does the project application code's derived data integrate with reusable software components and scientific visualization techniques?	Sharing of project developed software resources among local and remote users
3	SOS portfolio interoperability	Does the underlying framework enable the exchange of results with other simulations?	Sharing of project developed software simulation results in real-time among system-of-systems simulation users
4	Maintainability and adaptability	Does the project software adhere to standards and accepted practices, utilize standard languages and libraries, utilize common visualization tools, and provide adequate documentation?	Portable project developed software products across existing and future HPC platforms Reliability of application software products Maintainability of application software products
5	Training and technology transfer	Does the project software reside in a catalogued repository, utilize standard languages, and provide adequate documentation?	Software dissemination and technology transfer
6	Usability	Does the project software's human-computer interface support ease of learning, ease of use, effectiveness and efficiency, and user satisfaction?	Maintainability of application software products
7	Security	Are appropriate access controls in place to safeguard the intellectual property rights and security concerns associated with the project software?	Security

Clearly, we do not expect that every system specification will express properties and values so clearly or can be as easily mapped to our productivity model. The VPATC specification is, in our experience, unusual in its clarity and specificity, representative of current best practices. None the less, the format and placeholders for COI and MOE&S reflect government standards and are required for all similar DoD development. Clearly most codes will have similar kinds of requirements, although it may take more detective work (e.g. interviews and observation) to characterize them as precisely.

3 Developing a Productivity Benchmark Suite

We observe that reasoning about or measuring productivity in the VPATC domain requires considering a wide range of different kinds of system properties. Clearly, a development effort that meets the optimum criteria for the same effort, time, and cost as one that meets the minimum criteria would be considered more productive.

Thus, these properties should be considered part of the "output" that one must measure to assess productivity.

The same reasoning would apply to developing an effective benchmark for assessing productivity in VPATC's application domain. If, for example, we wanted to use a productivity benchmark to answer a question of the form "will we achieve greater productivity on a VPATC application using HPC platform A or platform B?" then the benchmark must incorporate the kinds of properties, values, and metrics we observe in the VPATC specification. In addition to defining a computational problem that exercises the hardware in the same manner the VPATC simulation does, the benchmark would need to define analogous requirements for the critical system properties (COI) including interoperability, usability, security, maintainability, adaptability, and so on. Fully executing the benchmark would require solving the computational problem in a manner that satisfies all of these functional and non-functional requirements. Measuring productivity against the benchmark would require measuring the extent to which each of the requirements had been satisfied against the effort expended.

Table 2
Critical technical parameters.

CTP	Test event	Evaluation optimum objectives	Evaluation minimum criteria
Scalable software suites: Demonstrate reduction in clock time as a function of increased central processing units (CPUs)	SAT	Determine the effective software code architecture for construction of scalable composite material predictors and dynamic models	Same
	Alpha	Fixed speedup exceeds 50% of optimum on 16 processors	Fixed speedup exceeds 40% of optimum on 8 processors
	Beta	Fixed speedup exceeds 60% of optimum on 32 processors	Fixed speedup exceeds 50% of optimum on 16 processors
	IOT&E	Fixed speedup exceeds 70% of optimum on 64 processors	Fixed speedup exceeds 60% of optimum on 32+ processors
Portable, reusable application software:	SAT	Architecture of code determined, and approaches for parallel execution analyzed	Same
Software applications behave the same and produce similar results, within an acceptable margin of error, on a variety of scalable HPC platforms	Alpha	Codes run on two HPC platforms with valid results	Same
	Beta	Codes run on three HPC platforms with same valid results	Codes run on two HPC platforms with same valid results
	IOT&E	Codes run on four or more HPC platforms with same valid results	Codes run on three HPC platforms with same valid results
Portable, reusable application software:	SAT	Reusable software components identified.	Determine the architecture for pre- and post-processing
Code provides for data output	Alpha	Data output in Tecplot format and XDMF format for custom interface.	Data output in Tecplot format.
	Beta	Data output in Tecplot format and XDMF format for custom interface.	Data output in Tecplot format.
	IOT&E	Data output in Tecplot format and XDMF format for custom interface.	Same.
Stable, accurate and robust software:	SAT	Determine the architecture for pre- and post-processing	Same
Interface with all required external software products and codes	Alpha	Software stores data in XML/HDF format and supplies interfaces to PETSc and the Scalable Parallel Direct Solver Library for Sparse Symmetric Positive Definite Systems (PSPASES)	Software stores data in XML/HDF format and supplies interfaces to the Portable Extensible Toolkit for Scientific Computation (PETSC)
	Beta	Software stores data in XML/HDF format and supplies interfaces to PETSc and PSPASES	Same
	IOT&E	“ “	Same

Notionally, this characterizes the content and use of a productivity benchmark suite. We view a PBS as a publicly available package that effectively represents the development challenges characteristic of a particular high performance application domain. Each such package defines a computational problem in a context that simulates the characteristics and constraints of a typical application in a particular HPC domain. The challenge is to define the context sufficiently so that:

1. the context adequately characterizes value space of the end-to-end requirements and goals of a real application domain;
2. it is possible to make meaningful comparisons in productivity measures between distinct applications of the benchmark.

Our approach is based on the empirical derivation of canonical workflows (Kepner, 2004) and purpose-based benchmarks (Gustafson, 2004). Together with the associated non-functional requirements, value function, and metrics, these sufficiently constrain the problem that different developers should be able to apply the benchmark and generate productivity measures that can be meaningfully compared. This will provide a basis for a public repository of commonly applicable HPCS benchmarks. The key components are as follows.

Canonical workflows. Briefly, canonical workflows are used to characterize and constrain the process context of a productivity benchmark. A canonical workflow characterizes both the development process and the execution workflow associated with creating and using a HPC application to meet an overall set of mission goals. It characterizes the process steps and work products associated with characteristic development paradigms in the HPC community.

Purpose-based benchmarks. Purpose-based benchmarks (PBBs) are described in detail in a separate paper in this issue (Gustafson, 2004). Briefly, PBBs are computational problems that accurately embody the design and execution time challenges of real applications in a domain. Unlike traditional benchmarks, PBBs are designed to exercise both the development process and the development platform in essentially the same manner (with reduced size) that real development problems do in a particular application domain.

Non-functional requirements. The benchmark will include representative execution time and developmental requirements with their associated metrics of completion and effectiveness. These include any requirements on the development process, administration, static-design, and runtime behavior characteristic of the application domain.

Characteristic value function. Associated with the requirements is a representative value function (e.g. in the

form of the value function (1) or (2) above). The value function characterizes a value proposition (i.e. relative values of the different requirements) associated with applications in the domain interest.

Productivity metrics and tools. A set of standardized metrics, algorithms, and tools for measuring productivity associated with both development time and execution time activities and goals. These are discussed further in subsequent sections of this paper.

Our approach to building PBSs is empirical in the sense of being derived either from observation of real developers or from carefully controlled experiments. For example, we will obtain the application properties of interest and their relative values directly from developers in particular HPC domains like by direct inquiry or by observation. The empirical methods we plan to apply to particular parts of the problem are described in context.

3.1 MEASURING DEVELOPMENT

There are two major goals of empirical measurement in the context of our benchmarks.

1. **Characterization.** Initially, the primary goal of measuring the processes and products of HPCS applications is to better understand what actually happens during such development. Measurement supports identification of potential problems and bottlenecks in HPCS development, clarification of the similarities and differences between the various workflows for development, and the potential creation of predictive models for required resources and product quality.
2. **Control.** Once a baseline set of measures has been obtained and used to characterize HPCS development, the use of measurement can begin to support project management activities. In this application, measures taken from the project requirements or from in-process development can be compared to measures obtained from prior development efforts or used as input to the predictive models generated from these measures. These comparisons and model outputs can be used to help guide the new development. Possible forms of guidance include the need for new or different kinds of resources, the appropriateness of the given workflow chosen for the goals of the project, and approaches to improving the quality of the system. Of course, the measures taken during every development can feed back into the characterization process to provide better understanding and modeling of HPCS development.

We will measure HPCS development in both qualitative and quantitative ways. Our measurement techniques will

include structured interviews, time and motion studies, and automated measurement. Each of these techniques has different strengths and weaknesses. By employing all of these techniques, we can ameliorate the weaknesses present in each form and improve the overall validity of the results.

3.1.1 Structured Interviews. In structured interviews, a researcher talks directly with members of the development team, recording data using notebooks, audio tape, or video tape to learn more about the developer's view of the development process and its strengths and weaknesses. Structured interviews are useful for general characterization of a workflow, gaining insight into the kinds of quantitative measures that would be useful to collect, and collecting examples of process problems and solutions.

Some advantages of the structured interview process are that it is relatively inexpensive to carry out and does not require extensive access of the researcher into the development process. However, it suffers from the fact that the way a developer recalls a development situation could vary in significant ways from the reality of that development situation when it actually occurred. In addition, social or political pressures can influence the way a developer represents development obstacles or problems. Over time, a developer might simply forget or not perceive significant influences on development.

In many cases, however, we will be interested in data, observations, and lessons learned from projects that have already completed. For such projects, structured interviews are the only way to recover much of the information of interest (e.g. perspectives of the time spent on different activities or the relative value of different system properties). Here the use of carefully structured interviews allows us to generalize over the data collected. This is the approach currently being used in a set of retrospective studies of HPC applications being led by Doug Post of Los Alamos National Laboratory (Post and Kendall, 2004). Results from these project retrospectives will be used to help characterize application domains in terms of characteristic requirements and canonical work flows.

3.1.2 Time and Motion Studies. In time and motion studies, also known as "naturalistic observation", the observer spends time "shadowing" one or more members of the development team, recording the tasks to perform and the time intervals during which the tasks are performed. Developer logs can augment direct observations.

Time and motion studies have the advantages of supporting the development of fine-grained models of how developers spend their time, and surfacing issues in development process that might not otherwise be perceived by developers. The data that are collected are thus of generally higher quality and fidelity than that collected by struc-

ured interviews. The disadvantages of time and motion studies are access and cost. The organization and developers must be willing to allow a researcher to monitor their behavior over extended periods and measures must be taken to prevent that monitoring from interfering with the developer's behavior. In addition, the technique is very expensive, requiring the researcher to essentially work full-time at the institution carrying out data collection. Because of this, time and motion studies are usually restricted to just a few days of data collection. This means that only a small period of development during any given project can be monitored.

Several examples of software development time and motion studies appear in Perry et al. (1996). In that work, the authors summarized the results of three studies of developer time usage at Lucent Technologies.

The first experiment was an after-the-fact analysis of one software developer's personal diary. This diary recorded that developer's work-related activities over a three-year period. This analysis helped the authors create initial hypotheses about time usage they tested in a later experiment. For example, analysis of the diary suggested that, much more than the authors expected, developers were blocked because important information was not immediately available to them.

The authors then conducted a second experiment to examine their newly generated hypotheses. To do this they created a data collection instrument accurate to about half an hour of time. Next, they recruited a number of Lucent developers who agreed to fill out the data collection forms each day. From these data, the authors made several observations. First, the subjects typically worked on two features simultaneously, so they could context switch when they became blocked. Secondly, even though context switching allows developers to keep working when blocked on one project, the cost is that individual projects take longer to finish. This was particularly important in the telephony industry where time-to-market pressures were severe.

Finally, the authors conducted a third study to compare self-reported time data with observed time data. The goal was to understand how accurate self-reported data are. This is important because it is much cheaper to collect data this way than to hire an external reporter. This study suggested that, contrary to some people's beliefs, developers did not consciously misrepresent data. However, there were some systematic sources of bias. For example, some developers tended to lump short work interruptions (e.g. a colleague drops in for a 10-min technical discussion) in with the activity they were conducting when interrupted.

Lessons learned from these studies are being applied to development and validation of our productivity benchmarks to ensure accurate representation of workflows and accurate data on units of work.

3.1.3 Automated Measurement. A third form of measurement involves collection of data using the artifacts of development itself. For example, if development uses a source code control system, then the logs from this system can be analyzed to understand the patterns of developer interaction with the source files over time. The system itself can be generated at various points during its development to recover measures of its size or complexity.

Automated measurement has the advantage of collecting more objective measures of the process and products of development that are not filtered through the perceptions of developers. Automated measurement also has the advantage of being relatively low cost: researchers are not required to be on site, and developers do not have to deal with the potential short-term or long-term intrusion of a researcher in their daily activities. The disadvantage of automated measurement is its incompleteness relative to the other methods: many development activities cannot be reconstructed by analysis of the evidence left by tool usage.

An example of automated measurement is the work by Johnson et al. (2003) on Hackstat. Hackstat is a system for automated metrics collection and analysis which provides custom “sensors” that are attached to developer tools, such as their editor, build tool, configuration management system, and so forth. The sensors unobtrusively monitor development and collect data about the process and products of development.

Some aspects of HPCS productivity measurement are similar to software productivity measurement in general. For example, automated metrics collection can provide an indication of how much time is spent actively modifying source code files. Automated tools can also monitor activities such as the invocation of test cases, and the sequence and nature of developer interactions with a configuration management repository. The tools can automatically gather data on the overall size of the system (expressed as numbers of non-comment lines of code, and/or number of methods, and/or number of compilation units), or complexity measures (such as measures of coupling between or cohesion within modules, and the characteristics of the inheritance hierarchy).

However, other aspects of HPCS productivity measurement will be required due to the specialized nature of this domain. For example, many workflows for HPCS application software development involve the initial development of a serial version of the system and measurement of its performance characteristics. Following the establishment of this baseline performance level, the next development stage involves implementing a parallelized version, typically using packages such as MPI or OpenMP. Following this, experimental runs are performed to compute measures such as Speedup, which indicates how much faster (or slower) the system executes as the number

of processors allocated to the problem increases. Based upon the Speedup curve or other measures of parallel performance, the developer may decide to reimplement the parallelized version. By automatically monitoring the invocations of tools for compilation, execution, and performance profiling, it is possible to measure the time spent in each of these phases of HPC development.

3.2 MEASURING DEVELOPMENTAL QUALITIES (“ILITIES”)

Obtaining any sort of fine-grained measures of productivity using our approach requires that we be able to provide relatively precise measures of progress, completion or other figures of merit for developmental properties such as maintainability and portability. While such properties frequently appear in software requirements, the conventional wisdom is that they cannot be effectively measured. Where these properties are part of a system’s acceptance criteria, industry standard practice is to judge them up or down by some form of peer review (e.g. Fagan inspections; Fagan, 1976).

While our value equation certainly permits all-or-none completion criteria (i.e. value zero or one for one or more C_i in equation (2)), their use coarsens any measure of relative productivity. In many cases, users will be interested in questions of degree, e.g. how much more portability do we get using development platform A than platform B? Further, if such properties account for a significant portion of a developer’s productive effort, we may lose the precision needed to make meaningful productivity comparisons. Thus, it is important to develop effective metrics for developmental properties.

In the following, we illustrate one approach to providing just such a metric for the property commonly referred to as “maintainability”. While we do not expect the specific approach described to generalize to other developmental properties, it does illustrate that (1) conventional wisdom is wrong, at least in this case, and (2) there are some general principles that may be useful in developing metrics for similar kinds of properties.

3.2.1 Example: A Maintainability Metric. Our goal is to develop an architectural design for our software system that is “maintainable”. By “maintainable”, we mean that the system is relatively easy to change for expected types of changes.¹ In the general case, a single architectural design cannot be equally easy to change for all types of changes so the designer must choose which kinds of changes will be accommodated easily and which will not (e.g. cause dovetailing changes or “break” the architecture).

In short, before it makes sense to talk about (much less measure) maintainability, we must be precise about what we mean by the word. Any given architecture may be

more maintainable than another relative to one set of changes but not to another. Thus, we must first answer the question: “maintainable relative to what?” We answer this question by specifying a list of the following.

1. Anticipated types of changes. We list changes that potentially impact the architectural design. These should not be highly detailed; rather they should identify classes of change such as: “it is expected that the Doppler radar will be replaced with a new model over the lifetime of the aircraft” or “it is likely that the pattern-matching algorithm will be replaced with a more efficient version”.
2. Relative priority of each type of change. Specifying the priorities of different changes gives a measure of their relative value and allows the architect to make appropriate trade-offs if necessary.

We now develop a metric that characterizes a design’s maintainability relative to our list of expected changes. We characterize priorities as low, med, or high, and assign to them the respective weights 1, 3, and 9. Our architectural design strategy applies information hiding (Parnas, 1972). That is, the designer seeks to (1) encapsulate each piece of information that is likely to change in exactly one module and (2) decompose the modules such that if two items are likely to change independently, then they will be encapsulated in different modules.

To the finished design, we apply a simple pass/fail completion criterion. If at most one module must be changed in response to each an anticipated change, then the design passes (receiving value 1); otherwise, the design fails (receiving value 0). The value of each change is then given by its priority times its completion value and the maintainability is given by the sum over the list of anticipated changes.

We can now calculate the maintainability metric associated with a particular architecture by “playing” the set of anticipated changes against the design. That is, each change is treated as a scenario that a reviewer applies to the system. If the change can be made by making changes to only one module (i.e. the design is consistent with the information-hiding principle) then it passes for that change.

The result is a metric with granularity proportional to the number of anticipated changes that characterize “maintainability”. Further, it is clear that, by its construction, the metric measures what the stakeholders mean by modifiability. Modifiability is defined to mean that anticipated changes are confined to one module. We then generate the change metric from the definitions of anticipated changes and produce the metric by simulating those changes.

There is, of course, considerable latitude in how we assign weights or values, as for example in the following.

1. We could construct a simple metric based on giving a value of “1” for each change confined to a single module and “0” otherwise, i.e. simply the number of changes that can be accommodated easily from a list of changes.
2. In practice, we have constructed metrics that include not only the priority but also the likelihood associated with each change. Lower weights can be assigned to changes of lower likelihood so the metric gives higher value to architectures that address the likely changes first.
3. We can express the maintainability metric as a value in terms of the amount of rework or cost associated with making changes. Rather than pass/fail, we evaluate each change against an architecture by estimating the amount of rework (and cost) to implement the change; a standard estimation exercise for change requests. We then use the total rework (weighted as desired) as the relative metric of design maintainability. This actually gives the measure in a form directly translatable to productivity.

An objective of our empirical studies will be to develop comparable metrics and methods of measurement for other properties interest. These will be empirically validated and included in the purpose-based benchmarks.

4 Benchmark Development and Execution

We are in the process of developing an initial set of benchmarks, metrics, and tools to validate our conceptual approach to productivity measurement for HPCS. Our development approach is iterative:

1. identify community of developers who will execute benchmarks;
2. develop productivity measurement infrastructure appropriate for that community, e.g. define benchmarks, define workflows and corresponding functional and non-functional requirements, create and install measurement instruments and analysis techniques;
3. observe and measure developers as they execute the benchmark using the previously defined infrastructure components;
4. analyze benchmark performance and evaluate and improve infrastructure.

Identifying the developer community. We have identified two initial developer communities who will participate in our studies. One is made up of computer science graduate students at the University of Maryland. The other

is a group of professional software developers working remotely from Russia. We are working with these two communities because each presents different experimental design characteristics and cost/benefit trade-offs for our research. In particular, these two communities represent different trade-offs between internal validity, external validity, cost, and data quality. That is, the student community provides us with an *in vitro* experimental setting. Here we have substantial control over many aspects of observable behavior: the programming tasks, programming environment, outside influences, observation methods used, etc. In addition, the costs associated with observing students are relatively low, making prototyping more feasible.

As with most *in vitro* situations, however, control comes at the price of representativeness. The students are not usually professional developers (although some are), so the tasks must be restricted in their time length and complexity, etc. To gather data that are more representative of the complex software development workplace, we need access to an *in vivo* experimental setting. The professional developer community provides this. This setting allows for less control because we must be careful not to overly interfere with the developers. On the other hand, the data we do collect are likely to be more relevant to our overall goal of understanding HPCS productivity (stronger external validity).

Develop productivity measurement infrastructure. Here we have begun by identifying a general PBB. This problem involves writing software to compute “optimal” designs for a weight-bearing truss with certain material characteristics (Gustafson, 2004). The original problem was specified with considerable detail. We are using this definition as given for the professional community, but are working to scale down the benchmark for the student community.

Each developer community will be following a different canonical workflow. At a high level, the professional community is following the Enterprise Workflow, while the students will be following the Lone Researcher Workflow (see Kepner overview). As part of this work, we are also defining the functional and non-functional requirements for the benchmark.

As discussed in Section Measuring Development, we are developing low-level data collection mechanisms (sensors) to capture developer activities unobtrusively. One novel way in which we will use these data is to develop low-level event traces (opening/closing files, running the compiler, etc.) and match them to higher-level actions defined by the canonical workflow. This will help us to better understand iteration, backtracking, and time usage within the overall process flows to help us better understand bottlenecks in the workflows.

Execute the purpose-based benchmark. Benchmark execution begins with an inoculation step. Participants are informed of the experimental procedures, explained about all data we will collect and why, informed of their right to withdraw from the study at any time, and are asked to provide informed consent. We then explain all the functional and non-functional requirements required of the implemented benchmark. For the student community we will also conduct a tutorial on the engineering and mathematical concepts underlying mechanical trusses. Our goal here is to avoid any miscommunication, or apprehensions about the experiment. Initially, participants will also be asked to maintain some manual record to help us calibrate our automated techniques. Finally, participants will implement the benchmark in the context of the appropriate canonical workflow. Along with measuring overall productivity we also expect to identify and measure current bottlenecks in development process, to highlight differences between activity- and performance-based benchmarks, and to begin analyzing how different machine architectures address existing bottlenecks.

We will be executing our first pre-pilot study using CS graduate students in the spring semester 2004. We will simultaneously begin studying the professional developers.

Analyze benchmark performance, and evaluate and improve infrastructure. As we begin to obtain data from these studies, we will continuously monitor and improve our infrastructure.

5 Open Issues

Our approach to characterizing and measuring productivity is consistent with DARPA’s programmatic goals of:

- addressing both development time and execution time productivity issues;
- providing benchmarks suitable for all the major classes of HPCS applications;
- measuring productivity in terms of the actual value created;
- addressing strategic development concerns outside of a single life cycle such as the value of transferable, portable, and reusable software products;
- finding measures of productivity that users and vendors will agree on.

The (necessarily) open-ended framework allows us to define the output of development to include any properties necessary to characterize “value” for that development or class of developments. Indeed, any result to which we attach value and for which we can devise a measure can be included. Further, value expressions can be expanded

or contracted as needed to meet new situations by adding new terms or removing existing ones.

However, the approach also will require substantial work to make it operational and effective.

1. The properties of interest and relative values for different application areas are not currently known. These will have to be established empirically through stakeholder interviews and review of existing codes or specifications.
2. For any property of interest, we will have to establish appropriate and effective metrics of completion (or comparable metrics of scale). For many properties such as portability, maintainability, etc., there is currently no agreed upon measure or even process for determining relative figures of merit. For each such property, we will either have to develop new approaches or fall back on imprecise gradation determined by review.
3. Our value equations (1) and (2) are only a first approximation. It is unclear yet to what extent we will need to address dependences or how best to incorporate time into the model. In-depth examination of real development situations will be needed to make this clear.
4. Currently benchmarks do not address the range of properties we have discussed. Effective procedures will need to be developed for incorporating the properties of interest (e.g. as requirements), metrics, and measurement processes or tools into the benchmark specification. Likewise, we will need to develop directions on how to tailor the benchmarks and interpret the results of their execution.

6 Summary

Ensuring that next-generation HPC platforms significantly improve real productivity in terms of the science accomplished will require new approaches to characterizing, measuring, and predicting productivity. Current productivity metrics and benchmarks fall short in several ways. Traditional software metrics focus on measurable outputs but often bear little relationship to the actual value of what is produced. Common benchmarks tend to focus on machine performance, ignoring the growing bottlenecks associated with software development.

Our goal is to establish, apply, and validate an effective approach to assessing and predicting productivity that spans both development and execution time. We seek to provide these capabilities in a form that supports platform buyers in choosing the best system and platform developers in providing technology that addresses real productivity problems.

In this paper, we have described an empirical approach to understanding and addressing HPCS productivity. While it is clear that software development is increasingly a productivity issue in many HPC systems, few specifics are known. Before such problems can be addressed we need a better understanding of precisely what kinds of problems are occurring, where these problems occur in the process, how these differ from one process to the next, and how they vary from one type of HPC application domain to the next. We plan to collect such data through careful observation of (1) real projects in common development domains and (2) experimental development efforts on canonical benchmark problems. Data from real development efforts will allow us to understand and catalog problems, requirements, and constraints characterizing different types of HPC applications. We can then define representative requirements and value propositions for different domains.

Carefully controlled experiments will help us better understand precisely where developers spend their time and how different platform features might increase the efficiency those activities. From this we expect to develop detailed canonical workflows representative of different development environments.

Detailed knowledge of problem characteristics, requirements, values, and workflows will be combined to develop tailored productivity benchmarks for key HPCS domains. These benchmarks will provide not only a representative computation problem, but representative non-functional requirements as well. This includes domain-characteristic requirements for properties such as portability, reliability, maintainability, etc., along with appropriate measurement techniques. As a whole, each benchmark will exercise the entire development process across a value space appropriate to the domain and provide metrics and tools for measuring productivity throughout the development cycle.

7 ACKNOWLEDGMENTS

This material is based upon work supported by DARPA under Contract No. NBCH3039002. We thank Mr Dale Shires of the US Army Research Laboratory High Performance Computing Division for permission to use portions of the VPATC (2003) specification. We also thank Dr Doug Post of the Los Alamos National Laboratory and Mr Andy Mark of the High Performance Computing Modernization Program Office for their help in obtaining the specification. Thanks also to Bill Walster, Dolores Shaffer, and our anonymous reviewers for their careful reviews and thoughtful comments.

AUTHOR BIOGRAPHIES

Stuart Faulk received his BA from Cornell University and his PhD in computer science from the University of North Carolina at Chapel Hill. He is currently on faculty at the University of Oregon with the Department of Computer and Information Science. He serves as teaching faculty in the Oregon Master of Software Engineering program. Previously, he served as principle investigator for the Consortium Requirements Engineering project at the Software Productivity Consortium and as project head for the Software Cost Reduction Project at the US Naval Research Laboratory. Current research interests include empirical methods and software engineering for HPC systems, value-based software engineering, and software product lines.

John Gustafson received his BS degree from Caltech, and his MS and PhD degrees from Iowa State University, all in applied mathematics. He has been a software engineer at the Jet Propulsion Laboratory in Pasadena, a senior staff scientist and product development manager at Floating Point Systems, and a member of the technical staff at Sandia National Laboratories in Albuquerque, NM. He founded the Scalable Computing Laboratory within Ames Laboratory, USDOE, and led computational science research efforts there for 10 years. He joined Sun Microsystems in 2000, and led Sun's HPCS program. He is currently Chief Technical Officer for Cluster Enterprise Solutions. His research areas are HPC, performance analysis, parallel architectures, numerical analysis, algorithms, and computer graphics.

Philip M. Johnson is a Professor of Information and Computer Sciences at the University of Hawaii and Director of the Collaborative Software Development Laboratory (<http://csdl.ics.hawaii.edu/>). He received BS degrees in both biology and computer science from the University of Michigan, and MS and PhD degrees in computer science from the University of Massachusetts. His research interests include empirical software engineering, computer supported cooperative work, and high tech entrepreneurship.

Adam A. Porter earned his BS degree summa cum laude in computer science from the California State University at Dominguez Hills, Carson, California in 1986. In 1988 and 1991 he earned his MS and PhD degrees from the University of California at Irvine. Currently an associate professor, he has been with the department of Computer Science and the Institute for Advanced Computer Studies at the University of Maryland since 1991. He is a winner of the National Science Foundation Fac-

ulty Early Career Development Award and the Dean's Award for Teaching Excellence in the College of Computer, Mathematics, and Physical Sciences. His current research interests include empirical methods for identifying and eliminating bottlenecks in industrial development processes, experimental evaluation of fundamental software engineering hypotheses, and development of tools that demonstrably improve the software development process. Dr Porter is a member of the ACM and IEEE Computer Society.

Walter F. Tichy has been professor of Computer Science at the University Karlsruhe, Germany, since 1986. Previously, he was senior scientist at Carnegie Group, Inc., in Pittsburgh, PA and served six years on the faculty of Computer Science at Purdue University in West Lafayette, IN. His primary research interests are software engineering and parallelism. He is currently directing research on a variety of topics, including empirical software engineering, autonomic computing, software configuration management, cluster computing, optimizing compilers for parallel computers, and optoelectronic interconnects. He has consulted widely for industry. He earned an MS and a PhD in computer science from Carnegie Mellon University in 1976 and 1980, respectively. He is director of the Forschungszentrum Informatik, a technology transfer institute. He is co-founder of ParTec AG, a company specializing in software for computer clusters. He was program co-chair for the 25th International Conference on Software Engineering (2003). Dr Tichy is a member of ACM, GI, and the IEEE Computer Society.

Lawrence Votta received his BS degree in physics from the University of Maryland, College Park, MD in 1973, and his PhD degree in physics from the Massachusetts Institute of Technology, Cambridge, MA in 1979. He is a Distinguished Engineer at Sun Microsystems Inc., working to improve the software and system reliability and availability of Sun's products while pursuing his research interest in high availability computing and empirical software engineering. He works in the Sun Customer Advocates for Reliability (SunCARE) group providing technical leadership of the Reliability, Availability and Serviceability (RAS) and productivity analysis teams for Sun's HPCS program.

NOTE

- 1 "Maintainability" has no precise definition and is used in the industry to cover virtually any kind of change to software following deployment. Our definition focuses on a one common aspect of what is usually meant by "maintainability".

References

- Albrecht, A. and Gaffney, J. 1983. Software function, source lines of code, and development effort prediction: a software science validation. *IEEE Transactions on Software Engineering* 9(6):639–648.
- Boehm, B. 1981. *Software Engineering Economics*, Prentice-Hall, Englewood Cliffs, NJ.
- Boehm, B. and Wolverton, R. W. 1980. Software cost modeling: some lessons learned. *Journal of Systems and Software* 1(3):195–201.
- Boehm, B., Clark, B., Horowitz, E., Westland, C., Madachy, R., and Selby, R. 1995. Cost models for future software life cycle processes: COCOMO 2.0. *Annals of Software Engineering* 1:57–94.
- Brooks, F. P. 1995. *The Mythical Man-Month: Essays on Software Engineering*, anniversary edition, Addison-Wesley, New York.
- Fagan, M. 1976. Design and code inspections to reduce errors in program development. *IBM Systems Journal* 15(3):182–211.
- Gause, D. and Weinberg, G. 1989. *Exploring Requirements: Quality Before Design*, Dorset House, London.
- Gustafson, J. 2004. Purpose-based benchmarks. *International Journal of High Performance Computing and Applications* 18(4).
- Johnson, P., Kou, H., Agustin, J., Chan, C., Moore, C. A., Miglani, J., Zhen, S., and Doane, W. E. 2003. Beyond the personal software process: metrics collection and analysis for the differently disciplined. *Proceedings of the 2003 International Conference on Software Engineering*, Portland, OR, May.
- Jones, C. 1986. *Programming Productivity*, McGraw-Hill, New York.
- Kepner, J. 2004. HPC productivity: an overarching view. *International Journal of High Performance Computing and Applications* 18(4).
- Kitchenham, B. A. 1992. Empirical studies of assumptions that underlie software cost estimation. *Information and Software Technology* 34(4):211–218.
- Parnas, D. L. 1972. On the criteria to be used in decomposing systems into modules. *Communications of the ACM* 15(12):1053–1058.
- Perry, D., Staudenmayer, N., and Votta, L. 1996. Understanding and improving time usage in software development. *Trends in Software: Software Process*, A. Fugetta and A. Wolf, editors, Wiley, New York.
- Porter, A., Votta, L., and Basil, V. 1995. Comparing detection methods for software requirement inspections: a replicated experiment. *IEEE Transactions on Software Engineering* 21(6):563–575.
- Post, D. and Kendall, R. 2004. Software project management and quality engineering practices for complex, coupled, multiphysics, massively parallel computation simulations: lessons learned from ASCI. *International Journal of High Performance Computing and Applications* 18(4).
- Snir, M. and Bade, D. 2004. A framework for measuring super-computer productivity. *International Journal of High Performance Computing and Applications* 18(4).
- VPATC. 2003. Software Development Plan for Virtual Prototyping and Accelerated Testing of DoD Composite Material Combat Systems (VPATC), SOS-3, April 2003.
- Weinberg, G. and Schulman, E. 1974. Goals and performance in computer programming. *Human Factors* 16(1):70–77.