

An Environment of Conducting Families of Software Engineering Experiments

Lorin Hochstein³
Taiga Nakamura¹
Forrest Shull²
Nico Zazworka¹
Martin Voelp¹
Marvin V. Zelkowitz^{1 2}
Victor R. Basili^{1 2}

University of Maryland¹
Computer Science Department
College Park, Maryland 20742

Fraunhofer Center for Experimental Software Engineering²
College Park, 20742 MD, USA

University of Nebraska – Lincoln³
Computer Science and Engineering
Lincoln, Nebraska 68588-0115

May 2007

Technical Report CS-TR-4877
UMIACS-TR-2007-28

An Environment of Conducting Families of

Software Engineering Experiments

Abstract

The classroom is a valuable resource for conducting software engineering experiments. However, coordinating a family of experiments in classroom environments presents a number of challenges to researchers. This paper describes an environment that simplifies the process of collecting, managing and sanitizing data from classroom experiments, while minimizing disruption to natural subject behavior. We have successfully used this environment to study the impact of parallel programming languages on programmer productivity at multiple universities across the United States.

Table of Contents

Abstract	3
Table of Contents	3
1. Introduction	4
2. Classroom as software engineering lab	5
3. The Experiment Manager framework	7
4. Current status	13
5. Supported Analyses	15
6. Evaluation	18
7. Related work	19
8. Conclusions	19
Acknowledgement	19
References	19

1. Introduction

Scientific research advances by the creation of new theories and methods, followed by an experimental paradigm to either validate those theories and methods or to offer alternative models, which may be more appropriate and accurate. Computer science is no different from other sciences, and the field has been moving to adopt such experimental approaches [ZE98]. However, in much of computer science, and in software engineering in particular, the experimental model poses difficulties possibly unique among the sciences. Software engineering is concerned about the appropriate models applicable to the development of large software systems. As such it involves the study of numerous programmers and other professionals over long periods of time.

The major experimental approach accepted in scientific research is the replicated study. However, by being expensive to produce, software is not amenable to such studies. Typically case studies of a single development are monitored, and after many such studies, general trends can be observed. This was the major approach of the NASA Software Engineering Laboratory in the period from 1976 to 2001 [BA02]. Although this approach is commonly used in various fields of research involving humans, such as clinical study in medical science, conducting many studies is still difficult and expensive, which is often a major obstacle for good software engineering research.

Instead of large replications, running studies in classrooms using students as subjects has become a favorite approach for trying out new development techniques [SJ05]. Even though a conclusion drawn from student subjects cannot always be generalized to other environments, such experiments aid in developing approaches usable elsewhere. However, conducting larger-scale software engineering research in classroom environments can be quite complex and time-consuming without proper tool support.

A single environment (e.g., a single class) is often insufficient for obtaining significant results. Therefore, multiple replications of a given experiment must be carried out by different faculty at different universities in order to provide sufficient data to make conclusions. This means each experiment must be handled in a similar manner to allow for combining partial results. The complexity of providing consistent data across various experimental protocols has been overwhelming so far.

In this paper, we describe an environment we are developing to simplify the process of conducting software engineering experiments that involve development effort and workflow and ensure consistency in data collection across experiments in classroom environments. We have used this environment to carry out research to identify the effect of parallel programming languages on novice programmer productivity in the domain of high-performance computing (e.g., MPI [DO96], OpenMP [DA98], UPC [CA99], Matlab*P [CH03]). Although there are often issues regarding the external validity of students as subjects, this is not a major concern here since we are explicitly interested in studying student programmers.

This work was carried out in multiple universities across the United States in courses where the professors were not software engineering researchers and were, therefore, not experienced with conducting experiments that involved human subjects.

2. Classroom as software engineering lab

The classroom is an appealing environment for conducting software engineering experiments, for several reasons:

- Most researchers are located at universities.

- Training can be integrated into the course.
- Tasks required can be integrated into the course.
- All subjects are performing identical programming tasks, which is not generally true in industry.

In addition to the results that are obtained directly from these studies, such experiments are also useful for piloting experimental designs and protocols which can later be applied to industry subjects, an approach which has been used successfully elsewhere (e.g., [BA94] [BA97] [BA99]).

Aside from the threats to validity in using students as subjects, there are additional complexities that are specific to research in this type of environment. We encountered each of these issues when conducting research on the effects of parallel programming model on effort in high-performance computing [HO05b].

Complexity. Conducting an experiment in a classroom environment is a complex process that requires many different activities (e.g., planning the experimental design, identifying appropriate artifacts and treatments, enrolling students, providing for data collection, checking for process compliance, sanitizing data for privacy, analyzing data). Each such activity identifies multiple points of failure, thus requiring a large effort to organize and run multiple studies. In order to collect data from multiple universities, the professor may also be a novice in organizing and conducting such experiments.

Research vs. pedagogy. When the experiment is integrated into a course, the experimentalist must take care to balance research and pedagogy [CA03]. Studies must have minimal interference with the course. If the students in one class are divided up into treatment groups and the task is part of an assignment, then care must be taken to ensure that the assignment is of equivalent difficulty across groups. Students who consent to participate must not have any advantage or disadvantage over students who do not consent to participate, which limits additional overhead required by the experiment. In fact, each university's Institutional Review Board (IRB), required in all USA universities performing experiments with human subjects, insists that participation (or non-participation) must have no effect on the student's grade in the course.

Consistent replication across classes. To build empirical knowledge with confidence, researchers replicate studies in different environments. If studies are to be replicated in different classes, then care must be taken to ensure that the artifacts and data collection protocols are consistent. This can be quite challenging because professors have their own style of giving assignments. Common projects across multiple locations often differ in crucial ways making meta-analysis of the combined results impossible [MI00].

Participation overhead for professors. In our experience, many professors are quite willing to integrate software engineering studies into their classroom environment. However, for professors who are unfamiliar with experimental protocols, the more effort required of them to conduct a study, the less likely it will be a success. Therefore, empirical researchers must try to minimize any additional effort required to run an empirical study in the course.

The required IRB approval, when attempted for the first time, seems like a formidable task. Help in understanding IRB approval would greatly aid the ability of conducting such research experiments.

Participation overhead for students. An advantage of integrating a study into a classroom environment is that the students are already required to perform the assigned task as part of the course, so the additional effort involved in participating in the study is much lower than if subjects were recruited from elsewhere. However, while the additional overhead is low, it is not zero. The motivation to conform

to the data collection process is, in general, much lower than the motivation to perform the task, because process conformance cannot be graded.

This can be particularly problematic when trying to collect process data from subjects (e.g. effort, activities, defects), especially for assignments that take several weeks (e.g., we saw a reduction in process conformance over time when subjects had to fill out effort logs over the course of multiple assignments).

Automatic data collection of software process. To reduce subject overhead and increase data accuracy, it is possible to collect data automatically from the programmer's environment. Capturing data at the right level of granularity is difficult. All user-generated events can be captured (keyboard events, mouse events), but this produces an enormous volume of data that may not abstract to useful information. Allowing this raw data to be used can create privacy issues, such as revealing account names, with the ability to then determine how long specific users took to build a product or how many defects they made.

All development activities taking place within a particular environment (e.g., Eclipse) simplifies the task of data collection, and tools exist to support such cases (e.g. Marmoset [SP05]). However, in many domains development will involve a wide range of tools and possibly even multiple machines. For example in the domain of high-performance computing, preliminary programs may be compiled on a home PC, final programs are developed on the university multiprocessor and are ultimately run on remote machines. Programmers typically use a wide variety of tools.

Data management. Conducting multiple studies generates an enormous volume of heterogeneous data. Along with automatically collected data and manually reported data, additional data includes versions of the programs, pre- and post-questionnaires, and various quality outcome measures (e.g. grades, code performance, defects). Because of privacy issues, and to conform to IRB regulations, all data must be stored with appropriate access controls, and any exported data must be appropriately sanitized. Managing this data manually is labor-intensive and error-prone, especially when conducting studies at multiple sites.

3. The Experiment Manager framework

We evolved the Experiment Manager framework (Figure 1) to mitigate the complexities described in the previous section. The framework is an integrated set of tools to support software engineering experiments in HPC classroom environments. While aspects of the framework have been studied by others, the integration of all features allows for a uniform environment that has been used in over 20 classroom studies over the past 3 years. The framework supports the following.

Minimal disruption of the typical programming process. Study participants solve programming tasks under investigation using their typical work habits, spreading out programming tasks over several days. The only additional activity required is filling out some online forms. Since we do not require them to complete the task in an alien environment or work for a fixed, uninterrupted length of time, we minimize any negative impact on pedagogy or subject overhead.

Consistent instruments and artifacts. Use of the framework ensures that the same types of data will be collected and the same types of problems will be solved, which increases confidence in meta-analysis across studies.

Centralized data repository with web interface. The framework provides a simple, consistent interface to the experimental data for experimentalists, subjects, and collaborating professors. This reduces overhead for all stakeholders and ensures that data is consistently collected across studies.

Instrumentation package. Our instrumentation package called *umdist* supports automatic collection of software process data in a Unix-based, command-line development environment, which is commonly used in high-performance computing. The package is designed to be installed in a master account on the local server and then enabled in the accounts of each subject by executing a setup script, or be installed in the account of individual subjects. The appropriate installation mode depends on the need of a specific experiment and the configuration of the machine to be instrumented. In either case, the package can be used without the intervention of system administrators.

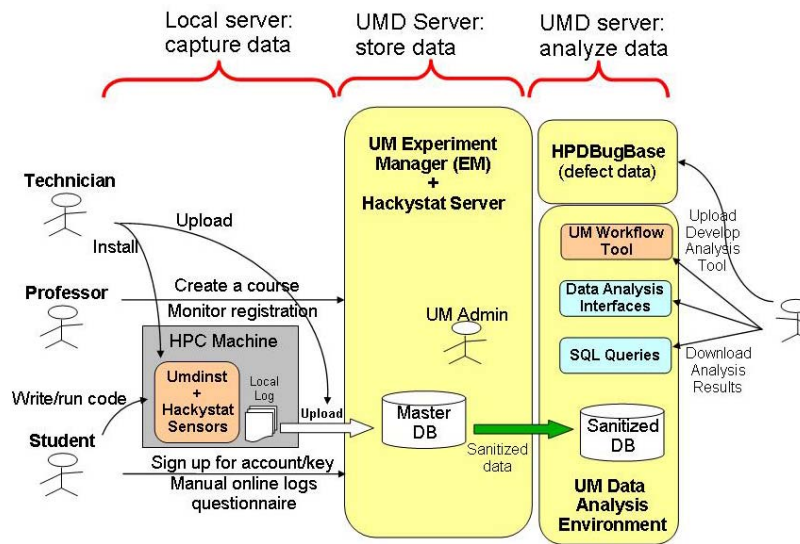


Figure 1 Experiment Manager structure

Umdinst package instruments programs that are involved in the software development process by replacing each tool (e.g., compiler) with a script that first collects the desired information and then calls the original tool. It is used for instrumenting compilers, although it is also designed to support job schedulers (common in high-performance computing environments), debuggers and profilers. For each compile the following data is captured:

- a timestamp
- contents of the source file that were compiled
- contents of local header files referenced in the source file
- the command used to invoke the compiler
- the return code of the compiler
- the time to compile

The umdist package includes Hackstat sensors [JO04] to instrument supported editors such Emacs and vi, and to capture shell commands and timestamps. The collected data is used in studies to estimate total effort as well as to infer development activities (e.g., debugging, parallelizing, tuning).

Web portal. The heart of the Experiment Manager framework is the web portal, which serves as a front-end to the database that contains all of the raw data, along with metadata about individual

experiments. Multiple stakeholders use the web interface: experimenters, subjects, and data analysts. For example, experimenters specify treatments (in our case, parallel programming models), assignment problem, participation rate, and grades. They also upload data captured automatically from *umdist*. Subjects fill in questionnaires, and report on process data such as time worked on different activities and defects. Analysts would export data of interest, such as total effort [Ho05] for hypothesis testing, or a stream of timestamped events for workflow modeling.

We have divided the functionality of the Experiment Manager into four roles. For each role, we developed a few use cases that describe its functionality, thus simplifying the design of the software.

1. **Technician:** The technician sets up the environment on the local server. This will be someone at a university with access to the machine the students use for the class. Often it is the Teaching Assistant in the course the software will be used in. The tasks for the technician are to install *umdist* so that students can use them. At the end of the semester the technician also sends the collected data to the University of Maryland server in case it was collected locally.
2. **Professor:** A database provides the professor with sample IRB questionnaires for submittal. This cannot be fully automated since each university has its own guidelines for submitting the IRB approval form. But experience with many universities over the last 3 years allows us to help in answering the most common questions on these forms.

The instructor registers each class with the Experiment Manager to set up a classroom experiment. For each such class, the professor can assign several programming projects from our collected database of assignments or assign one of his own. The system allows the professor to see if students have completed their assignments, but does not allow access to any of the collected data until the grades for the assignment are completed. In reality, the Teaching Assistant may actually perform this task, but in the role of the professor.

3. **Student:** A student who takes part in the experiment sends in data on HPC development. This requires the student to: (1) Register with the Experiment Manager by filling out a background questionnaire on courses taken and experiences in computer science in general and HPC programming in particular. Although this registration process can take up to 15 minutes, it is required only once during the semester. (2) Run the script to set up the wrappers for commands to edit, compile and run programs. Once an assignment is underway, the data collection process is mostly automatic and painless to provide.
4. **Analyst:** An analyst accesses the collected data for evaluating some hypothesis about HPC development. At the present time the analysis tools are relatively simple. Analysts can see the total effort and defects made by each student and collect workflow data.

Many tools exist in prototypes to support the various types of studies. For example, the HPC community is developing concepts of what productivity means in the HPC environment [ZE05] and we have been looking at developing workflow models (e.g., how much time is spent in various activities, such as developing code, testing, parallelizing the code, tuning the code for better performance). To support the study of workflows and productivity, we have developed a tool to allow the experimenter to apply various heuristics to the base data to see if we can automatically deduce the developer's programming activity, e.g., testing and debugging vs. development. This tool takes raw data, collected from online tools such as Hackystat and manual logs generated by students, and we have been developing

algorithms for automatically inferring the workflow cycle [VO06]. Results of this work are described in Section 5. Current activities are looking at extending these tools.




Figure 2(a). Effort capture tool.

Data collection. The actual data collection activity was designed to present minimal complexity to the student (Figures 2(a) and 2(b)). Within the Experiment Manager, the student has two options. If data was not collected automatically, the student can enter a set of activities, with the times each activity started and ended (Figure 2(b)). However, the Effort tool simplifies the process greatly. If the student clicks to start the tool (small oval near bottom of Figure 2(b)), then a small window opens on the top left corner of the screen (Large oval in Figure 2(b)). Each time the student starts a different activity, the student only needs to pull down the menu in the Effort tool and set the new activity type (Figure 2(a)).

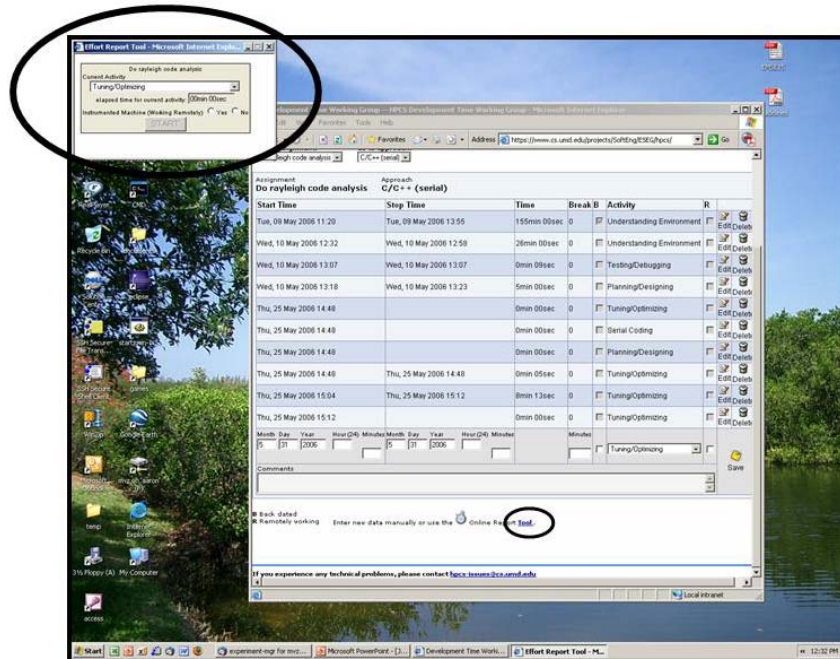


Figure 2(b). Effort collection screen.

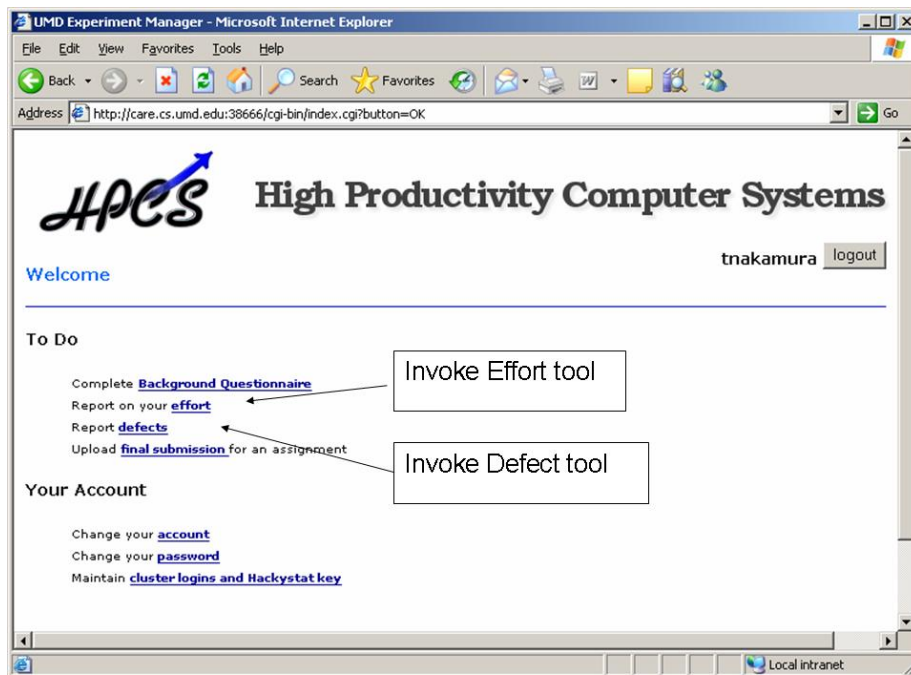


Figure 3. Student view of Experiment Manager.

The tool automatically computes elapsed time and saves the data in the database. If the student stops for a period of time (e.g., goes to lunch, surfs the web), there is a *stop* button on the tool. Upon returning, the user simply clicks on *start* to resume timing.

For most HPC development, the student simply has to:

1. Log into Experiment Manager to go to Effort page (Figure 3). Then click on Effort tool (Figure 2(b)).
2. Develop program as usual.
3. Each time a new activity starts, click on the new activity in the Effort tool (Figure 2(a)).
4. If any errors are found, the student records that defect by invoking the Defect tool (Figure 3) to explain the defect on a separate page (Figure 4).

Only step 3 and 4 involves any separate activity for participating in these experiments, and such activity is minimal.

Data sanitization. While personal data collected by the experiments must be kept private, we would like to provide as much data as possible to the community as part of our analysis. The sanitization process exports “safe” data into a database that can be made accessible to other researchers, running on a separate machine.

The sanitization process is briefly described in Figure 5. Each data object we obtained in an experiment is classified as one of:

1. *Prohibited.* Data contains personal data we cannot reveal (e.g., name or other personal identifiers)
2. *Clean.* Data we can reveal (e.g., effort data for development of another clean object).
3. *Modified.* Data we can modify to make it clean (e.g., removing all personal identification in the source program)

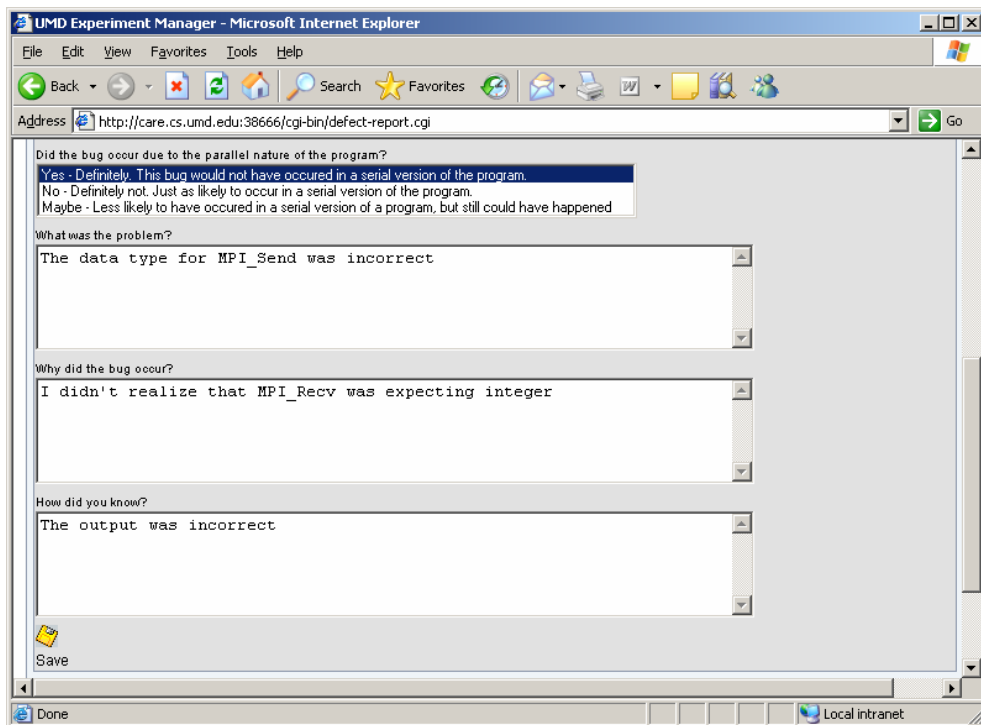


Figure 4. Defect reporting tool.

Clean data can be moved to the analysis server and modified data can also be moved. Only prohibited data cannot be exported to others desiring to look at our collected database. Our sanitization process on data consists of the following 4 functions:

1. *Normalization* - Normalize the timestamps for each class on a common basis. By making each time stamp relative to 0 from the beginning of that experiment, information about in which semester it was collected (and hence from which school the data was collected) is hidden.
2. *Discretization* – Since grades are considered private data, we define a mapping table that maps grades on a small set, such as {good,bad}. Converting other interval or ratio data into less specific ordinal sets, while it loses granularity, it helps to present anonymity.
3. *Summarization* – With some of the universities, we can give out source code if we remove all personal identifiers of the students who wrote the code. But in some cases, we are prohibited from doing even that. If we cannot give out source code, we can collect patterns and counts of their occurrence in the source code. For example, we can count lines of code, or provide analyses of “diffs” of successive versions of code.
4. *Anonymization* - We can create hash values for dates, school names and other personal identifiers.

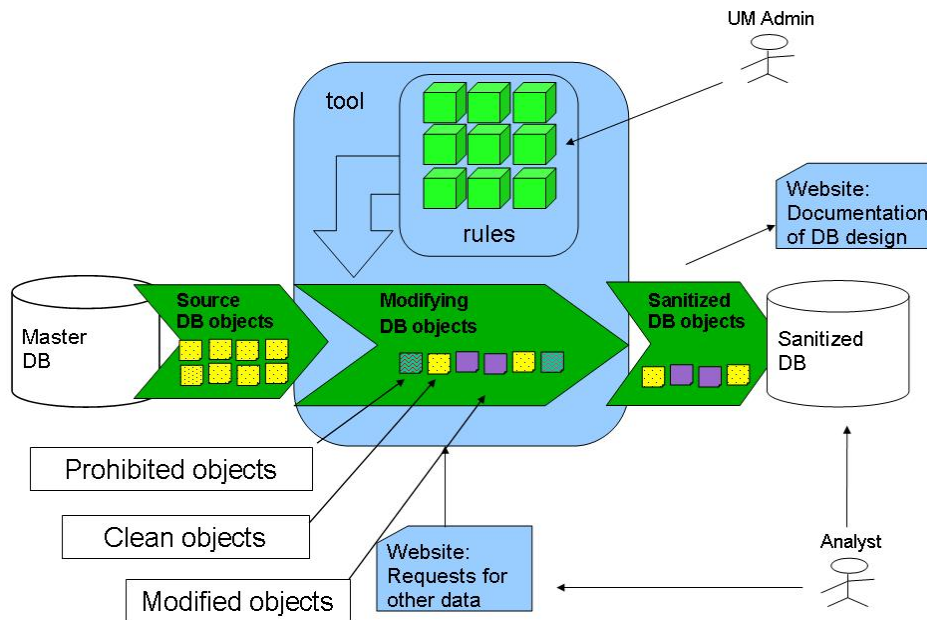


Figure 5. Basic sanitization process.

4. Current status

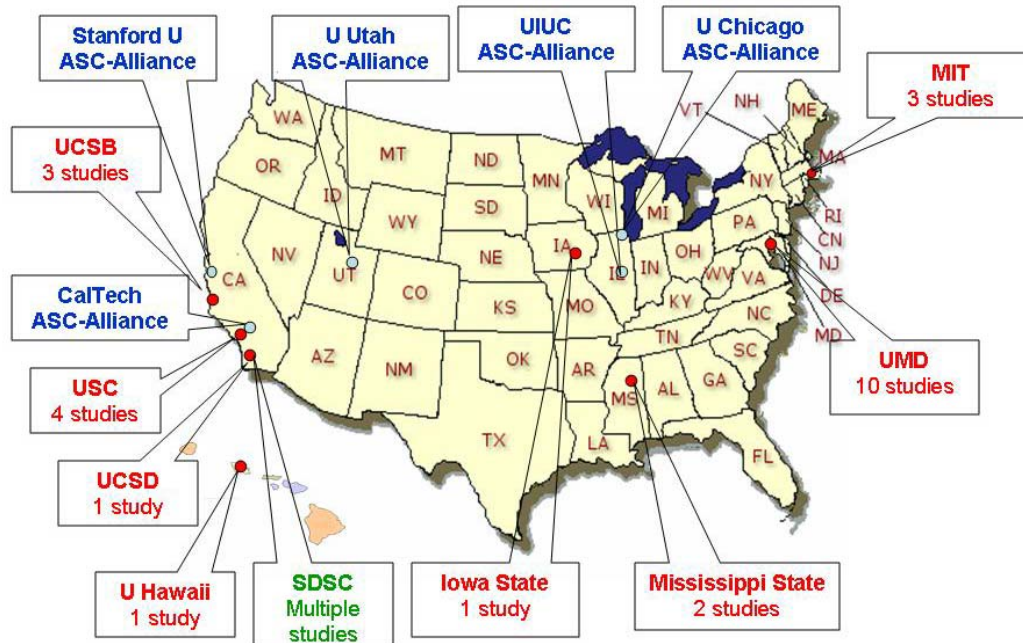


Figure 6 Studies conducted

Our Experiment Manager framework currently contains data from over 20 classroom experiments conducted at various universities in the U.S. (Figure 6). While some of the experiments preceded the Experiment Manager (and motivated its development) and their data was imported into the system, perhaps half the experiments used parts or all of the system.

Experiment Manager Effectiveness. Before using the experiment manager, we discovered many discrepancies between successive classroom projects which prevented merging the results. Some of these were:

1. It was not often obvious for how many processors the final programs were written.
2. Related to the previous problem, the projects all had differing requirements for final program complexity (e.g., the number of replicated cells needing to be computed).
3. Grading requirements differed. Was performance on an HPC machine important? Sometimes just getting a valid solution mattered. Maximum speedup, or the decrease in execution time of an HPC machine over a serial implementation, was the major goal.

By using our collected database of potential assignments, as well as our checklist of project attributes, this problem has lessened across multiple classes recently, allowing for the combination of results across different universities.

The IRB process seems like a formidable roadblock the first time any new professor encounters it. Often, in contacting faculty at a new university we would lose a semester's activity simply because the IRB approval process was too onerous the first time it was attempted. With our database of common IRB issues, this no longer seems a problem.

A related problem was the installation of software on the host computer for the collection of data. Again, this often meant the delay by a semester since the installation process was too complex. This was a major driving force to host much of this software as a web server at the University of Maryland, with a relatively simple *umdist* package that needed to be installed at each university's site.

The Effort tool (pictured earlier as Figures 2(a) and 2(b)) also solved some of our data collection problems. We can collect effort data 3 ways (Hackystat at the operating system level, manual data via programmer filled-in forms, and command line time stamps via *umdist*). All give different results [HO05]. The use of the effort tool greatly eases the data collection problem, which we believe increases the reliability of such data.

Most of our results, so far, are anecdotal. But we have been able to address new universities and additional classes in a more methodical manner at present and believe the Experiment Manager software is a major part of this improvement.

Experiment Manager Evolution. The system is continuing to evolve. Current efforts focus on the following tasks:

- We are evolving the user interface to the Experiment Manager web-based tool. The goal is to minimize the workload of various stakeholders (i.e., roles) for setting up the experiment environment, registering with the system, entering the data and conducting an analysis. We would like to develop small native applications that provide a more integrated interface to the operating system, making it less disruptive to users.
- We want to continue our experimentation with organizations that are often behind firewalls. Although we are currently studying professionals in an open environment, we have only begun to use the Experiment Manager in this environment. Although the *umdist* instrumentation package can be set up in secure environments, the collected data cannot be directly uploaded to the University of Maryland servers. We are enhancing the Experiment Manager architecture to better support the experimentations with these organizations. Working in these environments is necessary to see how professionals compare to the students.
- We will continue to evolve our analysis tools. For example, our prototype experience bases for evolving hypotheses and high end computing defects will continue to evolve both in content and usability.
- We will evolve the problem-specific harnesses that automatically capture information about correctness and performance of intermediate versions of the code during development.

This also requires us to evolve our experience bases to generate performance measures for each program submitted in order to have a consistent performance and speedup measure for use in our workflow and time to solution studies.

Our long-range goal is to allow the community access to our collected data. This requires additional work on a sanitized database that removes personal indicators and fulfills legal privacy requirements for use of such data.

5. Supported Analyses

The Experiment Manager was designed to ease data analysis, in order to support the investigation of a range of different research questions. Some of these analyses are focused in detail on a single developer being studied, while others aggregate data over several classes, allowing us to look across experimental data sets to discover influencing factors on effective HPCS development.

Views of a Single Subject. One view of a subject's work patterns is provided directly by the instrumentation. We refer to this view as the *physical level view* since it objectively reports incontrovertible occurrences at the operating system level, such as the timestamp of each compilation.

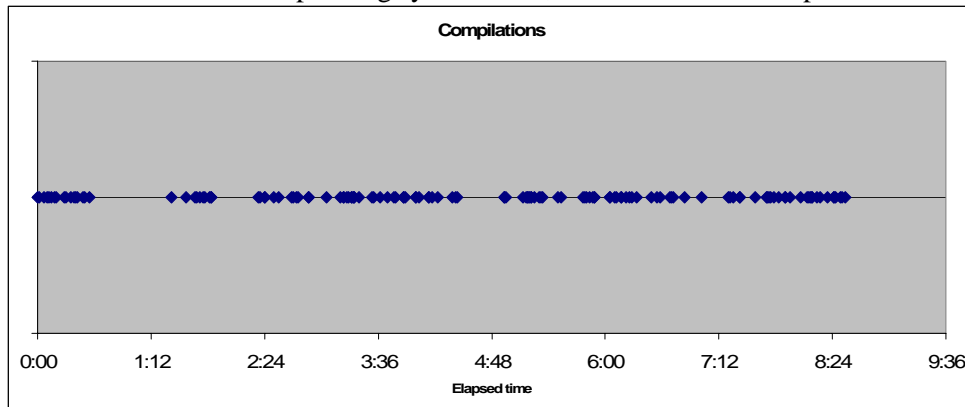


Figure 7: Compilation events for one subject

Figure 7 shows such a physical view for a 9-hour segment of work done by a given subject. The x-axis represents time and each dot on the graph represents a compile event. Although we have the tools to measure physical activities with a high degree of accuracy, this type of analysis does not yield much insight. For example, although we know how often the compiler was invoked, we do not know why: We cannot distinguish compilations which add new functionality from compilations which correct problems or defects that could have been avoided. This is an important distinction to make if we want to know the cause of unnecessary rework so it can be avoided.

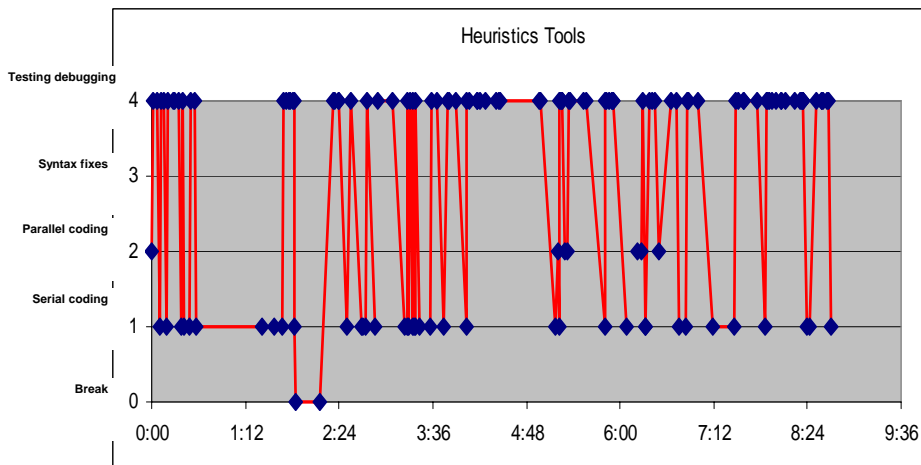


Figure 8: Types of development activities for one subject

A second view built on top of the logged compile times along with a snapshot of the code at each compile, applies a set of heuristics to guess at the semantics of the activities. There is a tool in the Experiment Manager that allows us to use various algorithms to manipulate these heuristics. The purpose of determining the semantic activities is to build baselines for predicting and evaluating the impact of new tools and languages.

Figure 8 uses the same data as Figure 7 to illustrate this view. Again, each dot represents one compile but in this case the activity preceding each compile has been classified as one of:

- *Serial coding*: The developer is primarily focused on adding functionality through serial code.
- *Parallel coding*: The developer is adding code to take advantage of multiple processors, not just adding function calls to the parallel library. We decided to separate out this activity since the amount of effort spent in this activity is indicative of how difficult it is to take advantage of the parallel architecture for solving the problem.
- *Testing and debugging*: The developer is focused on finding and fixing a problem, not adding new functionality. This activity can be identified via some typical and recognizable testing strategies, such as when a high percentage of the code added before a compile were output statements (so that variable values can be checked at runtime); creating/modifying test data files instead of the main code block; or removing test code back out of the system at the end of a debugging session. Our hypothesis is that effort spent on these activities can come from misunderstanding of the problem or the proposed solution and so could be addressed with more sophisticated aides given to developers.

Such a view helps us understand better the approach used by the subject, and how much of his/her time was spent on rework as opposed to adding new functionality. The duration data associated with each activity also helps us identify interesting events during the development: For example, when a large amount of time is spent debugging, analysts can focus on events preceding the debugging activity to understand what type of bug entered the system and how it was detected. This type of information can be used to understand how hard or easy it is for the developer to program in a given environment, and allows us to reason about what could be changed to improve the situation.

Validation of workflow heuristics. The heuristics to date have been developed by having researchers examine the collected data in detail (e.g. examining successive changes in source code versions). However, the accuracy of these heuristics is not generally known. We have developed a tool that provides a programmer with information about the inferred activities in real-time. Using the tool the developer then provides feedback about whether the heuristic has correctly classified the activity. This allows us to evaluate how well the heuristics agree with the programmer's belief about the current development activity.

Views of Multiple Subjects across several classes. From the same data, total effort can be calculated for each developer and examined across problems and across classes to understand the range of variation and whether causal factors can be related to changes in other measures of outcome, such as the performance of the code produced.

We note that for our experimental paradigm to support effective analyses, subjects in different classes who tackle the same problem using the same HPC approach should exhibit similar results regarding effort and the performance achieved. Moreover, we must be able to find measurable differences between subjects who, for example, applied different approaches to the same problem. In a previous paper [SH05], we presented some initial analyses of the data showing that both conditions hold.

Such analyses have been instrumental in developing an understanding of the effectiveness of different HPC approaches in different contexts. For example, Figure 9 shows a comparison of effort data for two HPC approaches, "OpenMP" and "MPI." The percentage of effort saved by using OpenMP instead of MPI is shown for each of 5 parallel programming problems ("Buffon", "Matvec," etc.) representing

different classes of parallel programs. Thus, 50 on the y-axis represents 50% less effort for OpenMP; -50% would indicate that OpenMP required 50% more effort. The height of each bar represents the range of values from across an entire dataset of subjects. As can be seen, in 2 cases OpenMP yielded better results than MPI as all subjects required less effort; for 2 other cases although there were some who required less effort for MPI the majority of data points indicated an effort savings associated with OpenMP. In only one case, for the Buffon problem, did MPI appear to give most subjects a savings in effort. As we gather more datasets using the Experiment Manager tool suite, we will continue this type of analysis to understand what other problems are in the set for which MPI requires less effort, and what it is about these situations that sets them apart from the ones where OpenMP was the less effort-intensive approach. (Interested readers can find a description of these approaches and programming problems in other publications [SH05].)

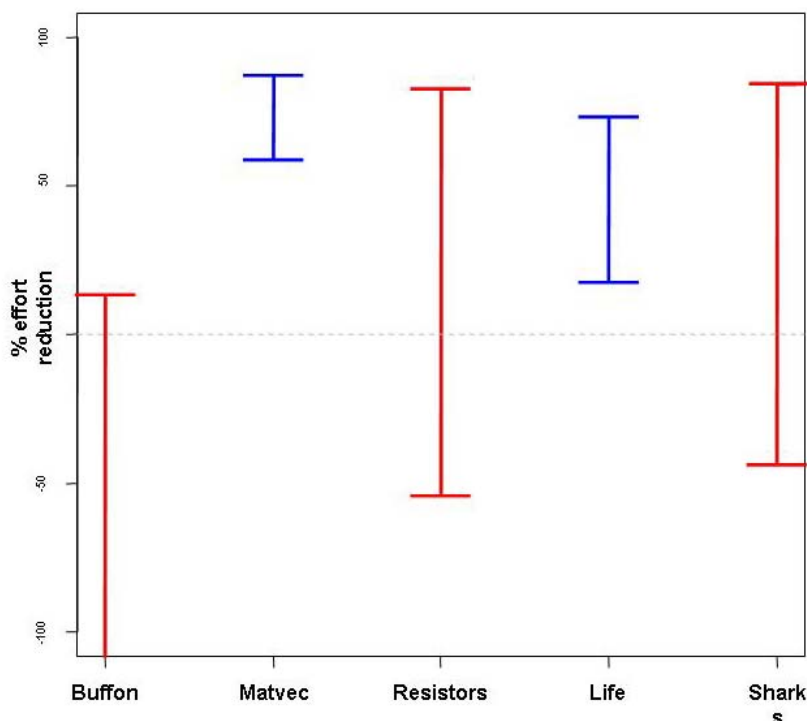


Figure 9: % effort reduction for OpenMP over MPI

6. Evaluation

We have been performing classroom experiments in the HPC domain since early 2003. While we have not performed a careful controlled experiment of its effectiveness, we have observed anecdotally that the Experiment Manager avoids many of the problems others (including ourselves) have observed in running experiments. Many of these problems have already been reported in this paper. We have been able to run the same experiment across multiple classes in multiple universities and combine the results. Data is collected reliably and consistently across multiple development platforms. We have been able to obtain data to install into our database effortlessly without the need for students to perform any post-development activity. Faculty, who are not experimental researchers, have been able to run their own experiments with only minimal help from us. And finally, the ability to sanitize data allows us to provide

copies of datasets to others wanting to perform their own analysis without running into IRB and privacy restrictions.

7. Related work

There are various other projects that either support software engineering experiments, or support automatic data collection during development, but not both.

The SESE system [AR02] has many similarities: it is web-based and supports features such as managing subjects, supporting multiple roles, administering questionnaires, capturing time spent during the experiment, collection of work products, and monitoring of subject activity. By comparison, Experiment Manager supports additional data capture (e.g. intermediate source files, defects) and data analysis (e.g. sanitization, workflow analysis).

PLUM, back in 1976, was one of the first systems to automatically collect development data [ZE76]. It, along with Hackystat [JO04], Ginger2 [TO99], and Marmoset [SP05] are examples of systems which are designed to collect data during the development process, but do not have data management facilities that are specifically oriented towards running multiple experiments. Hackystat, which we are using in the Experiment Manager, can collect data from several different types of applications (e.g. vi, Emacs, Eclipse, junit, Microsoft Word) via sensors. It was originally designed for project monitoring rather than running experiments. We have adopted the use of some of the Hackystat sensors into our data collection system. Ginger2 is an environment for collecting an enormous amount of low-level detail during software development, including eye-tracking and skin resistance. Marmoset is an Eclipse-specific system which captures source code snapshots at each compile, and is designed for computer science education research.

8. Conclusions

The classroom provides an excellent opportunity for conducting software engineering experiments, but the complexities inherent in this environment makes such research difficult to perform across multiple classes and at multiple sites. The Experiment Manager framework supports the end-to-end process of conducting software engineering experiments in the classroom environment. This allows many others to run such experiments on their own in a way that allows for the appropriate controls of the experiment so that results across classes and organization at geographically diverse locations can be compared. The Experiment Manager significantly reduces the effort on behalf of the experimentalists who are managing the family of studies, and on the subjects themselves, by applying heuristics to infer programmer activities.

We have successfully applied the Experiment manager framework and with each application are learning and improving the interface, simplifying the use by students, making its use of value in shrinking the overall problem solving process by students, e.g., various forms of harnesses, the support for analysis, in order to get a thorough understanding of the HPC development model.

Acknowledgement

This research was supported in part by Department of Energy contract DE-FG02-04ER25633 and Air Force grant FA8750-05-1-0100 to the University of Maryland. Several students worked on the Experiment manager including Thiago Escudeiro Craveiro and Patrick R. Borek.

References

- [AR02] E. Arisholm, D.I.K. Sjoberg, G.J. Carelius, Y. Lindsjom, A web-based support environment for software engineering experiments, *Nordic Journal of Computing* 9(3) (September, 2002), 231-247.
- [BA94] V. Basili, S. Green, Software Process Evolution at the SEL, *IEEE Software* 11(4), (July 1994), 58-66.
- [BA97] V. Basili, "Evolving and Packaging Reading Technologies," *Journal of Systems and Software*, vol. 38 (1): 3-12, July 1997.
- [BA99] V. Basili, F. Shull, and F. Lanubile, "Building Knowledge through Families of Experiments," *IEEE Transactions on Software Engineering*, vol. 25(4): 456-473, July 1999.
- [BA02] V. Basili, F. McGarry, R. Pajerski, M. Zelkowitz, Lessons learned from 25 years of process improvement: The rise and fall of the NASA Software Engineering Laboratory, IEEE Computer Society and ACM International Conf. on Soft. Eng., Orlando FL, (May 2002) 69-79.
- [CA99] W. Carlson, D. Culler, K. Yellick, E. Brooks, K Warren, Introduction to UPC and language specification (CCS-TR-99-157), Technical report, Center for Computing Sciences, 1999.
- [CA03] J. Carver, L. Jaccheri, S. Morasca, F. Shull, Issues in using students in empirical studies in software engineering education, International Symposium on Software Metrics, Sydney, Australia, (2003), 239-249.
- [CH03] R. Choy, Al Edelman, MATLAB*P 2.0: A unified parallel MATLAB, Singapore MIT Alliance Symposium, 2003.
- [DA98] L. Dagum, R. Memon, OpenMP: an industry-standard API for shared-memory programming, *IEEE Computational Science & Engineering* 5(1), (January 1998), 46-55.
- [DO96] J.J. Dongarra, S.W. Otta, M. Snir, D. Walker, A message passing standard for MPP and workstations, *Communications of the ACM* 39(7), (July 1996), 84-90.
- [HO05] L. Hochstein, V. Basili, M. Zelkowitz, J. Hollingsworth and J. Carver, Combining self-reported and automatic data to improve effort measurement, Joint 10th European Software Engineering Conference and 13th ACM SIGSOFT Symposium on the Foundations of Software Engineering , Lisbon, Portugal, (September 2005) 356-365.
- [HO05b] L. Hochstein, J. Carver, F. Shull, S. Asgari, V. Basili, J. K. Hollingsworth, M. Zelkowitz, HPC Programmer Productivity: A Case Study of Novice HPC Programmers, Supercomputing 2005, Seattle, WA, November 2005.
- [JO04] P. M. Johnson, H. Kou, J. M. Agustin, Q. Zhang, A. Kagawa and T. Yamashita, Practical automated process and product metric collection and analysis in a classroom setting: Lessons learned from Hackystat-UH, International Symposium on Empirical Software Engineering, Los Angeles, California, (August, 2004).
- [MI00] J. Miller, Applying meta-analytical procedures to software engineering experiments, *Journal of Systems and Software* 54, 1, (September, 2000) 29-39.
- [SH05] F. Shull, J. Carver, L. Hochstein, V. Basili, Empirical Study Design in the Area of High Performance Computing (HPC), International Symposium on Empirical Software Engineering, Noosa Heads, Australia, (2005).
- [SJ05] D. Sjoberg, J. Hannay, O. Hansen, V. Kampenes, A. Karahasanovic, N. Liborg, A.C. Rekdal, A Survey of Controlled Experiments in Software Engineering, *IEEE Transactions on Software Engineering* 31(9), (September, 2005), 733-753.

- [SP05] J. Spacco, J. Strecker, D. Hovemeyer, W. Pugh, Software repository mining with Marmoset: an automated programming project snapshot and testing system, Proceedings of the 2005 International Workshop on Mining Software Repositories, St. Louis, Missouri, (2005), 1-5.
- [TO99] K. Torii, K. Mastumoto, K. Nakakoji, Y. Takada, S. Takada, K. Shima, Ginger2: an environment for computer-aided empirical software engineering, *IEEE Transactions on Software Engineering*, 25 (4), (July, 1999).
- [VO06] M. Voelp, Diploma Thesis, Computer Science, University of Applied Sciences, Mannheim, Germany (August, 2006).
- [ZE76] Zelkowitz M. V., Automatic program analysis and evaluation, International Conf. on Software Engineering, San Francisco, CA (October, 1976) 158-163.
- [ZE98] M. V. Zelkowitz and D. Wallace, Experimental models for validating computer technology, *IEEE Computer* 31, 5 (May, 1998) 23-31.
- [ZE05] M. V. Zelkowitz, V. Basili, S. Asgari, L. Hochstein, J. Hollingsworth, and T. Nakamura, Productivity measures for high performance computers, International Symposium on Software Metrics, Como, Italy, September, 2005.