

The Case for Repeated Research In Operating Systems

Jeanna Neefe Matthews
Clarkson University
jnm@clarkson.edu

1. Introduction

Repeated research is a well-respected model of investigation in many sciences. Independent tests of published research are valued because they document the general applicability of results. In addition, repeated research often sheds new light on aspects of a work not fully explored in the original publication and exposes unreported limitations. In computer science, researchers typically report results from testing software that they themselves have implemented. It is natural to wonder why independent tests of published results are so rare.

One argument might be that it is difficult to transfer fragile research software to a new environment, especially given the wide variety of hardware and software platforms. However, many computer science researchers, especially in operating systems, work with commodity hardware and build on open source operating systems. Many researchers even make their own source code available. This provides a readily available platform for repeated research.

Another argument might be that it is the job of industry to validate and extend research ideas. However, researchers often note with dismay the number of great ideas that are not incorporated into production computer system. Repeated research would establish reported experience as repeatable fact and be an important step towards this transfer of technology.

In addition to providing independent validation of published research results, I have found repeated research to be a highly effective model for directing student research projects. I have used it as the basis for individual student research projects and for class projects. It is an excellent model for transferring research results between universities and into course materials.

In the sections that follow, I describe a model for repeated research in computer science. I also discuss my experiences using this model in more detail. I argue

the benefits both to the research community and to students. Finally, I propose some concrete steps that could encourage repeated research – including encouraging the publication of repeated research results in *Operating Systems Review*.

2. A Repeated Research Project

A repeated research project begins with learning the intimate details of a published work – understanding the motivation for the work, examining each detail of the figures, reading the related work, etc.

The next step is to reproduce an experimental set-up similar to that used in the paper. This frequently involves downloading, configuring and using publicly available software. Alternatively, more advanced projects could involve reimplementing software based on the published description.

A key milestone in repeated research is determining if the results from the published work can be reproduced as described. Key experimental details that affect the results may not be highlighted or even mentioned in the original paper.

Repeated research also typically extends the original work by performing additional experiments. This step provides important information about whether the results hold only in a narrow set of conditions or whether they are more generally applicable.

To provide one concrete example, I recently directed a repeated research based on the paper “Xen and Art of Virtualization” from SOSP 2003 [1]. Xen is an x86 virtual machine monitor produced by the University of Cambridge Computer Laboratory and re-leased under the GNU General Public License. Students in an advanced operating systems class reproduced the experiments described in the SOSP03 paper comparing Xen to native Linux and to other virtualization techniques like User Mode Linux. The tests described in the paper

used a new Dell Xeon server. In addition to reproducing the tests on a nearly identical hardware platform, we also evaluated Xen on a 3-year-old x86. We also evaluated the results of similar tests running on an IBM zServer. The results of these tests will be appear in the FREENIX Track of the 2004 USENIX Technical conference [2].

3. Benefits of Repeated Research

- **Provides Independent Validation of Published Results**

When researchers test software that they have implemented, it is natural that they strive to show it in the best light. Repeated research provides an opportunity for more impartial evaluation. Repeating research requires a more detailed review of research results than is possible with even the most detailed reading. It is well suited for determining if all the relevant aspects of the experiments have been described.

- **Question and Extend the Range of Tests Performed**

It is natural while repeating a set of experiments to consider other experiments that could be performed. Once an independent group has duplicated a similar experimental platform, they can make targeted variations to the test conditions. This can yield confidence that the results are robust to logical perturbations in platform or may highlight some surprising limitations.

Completing a repeated research project lays a foundation for exploring some of the future work described in the paper itself. It also places the group repeating the research in a good position to make significant enhancements of their own.

- **Document and Generalize Research Software**

Repeating research once, can make it even easier for the next repetition by documenting the process and removing dependencies on the original environment. In my experience, repeating published research is difficult enough that it should not be left as an “exercise to the reader”. For example, in repeating the SOSP03 Xen results, we ported three device drivers, wrote a dozen testing scripts and documented some important configuration requirements for achieving the reported results.

- **Speed Transfer Into Production Systems**

Documented success with recreating the experimental system and independent validation of results would significantly reduce the risk involved in transferring the ideas into production systems.

- **Involve Students In Research**

Involving students in research, especially early in their career, is a high priority of federal research funding agencies. A positive initial experience with research can be the impetus for a research career.

Often, a student’s first research experience involves a small task within a larger project. In many cases, students emerge without a good sense of how their task related to the entire project and without the passion that comes from seeing the big picture. In addition, the number of useful small tasks available in ongoing research projects is limited even at large research universities and students at teaching universities may have even fewer opportunities.

Students engaged in repeated research leave with a clear understanding of a published work – from its high level motivation to its experimental details. They gain important practical skills including setting up an experimental platform, collecting and analyzing results, writing a research paper and giving a research presentation. Further, they learn how to build upon and critically evaluate others work and how to put their own work in the context of prior work. With this training in place, students are poised to pursue their own original research ideas.

Following a previously published work allows students to emulate a successful example. I have seen this translate directly into tangible external accomplishments. Students I have mentored in this model have twice won the IBM Linux Challenge (3 winners in 2001 including the institutional prize and 2 winners in 2002) and have presented at the 2003 National Conference on Undergraduate Research. In addition, many students mentored in this model as undergraduates have gone on to graduate school.

- **Speed Transfer Into Curriculum**

Research that has been repeated enough times to develop a clear set of instructions can be used as the basis for classwork. I have also used this model as the basis

for assignments in both undergraduate and graduate courses.

5. Concrete Steps To Encourage Repeated Research In Operating Systems

Encouraging publication of repeated research results in *Operating Systems Review* and other venues would be an excellent way to motivate such work. It would be interesting to invite original researchers to comment on the repeated results as well.

It would also help to establish an online repository of repeated research suggestions and materials. As I have mentored students in this model, I have made it a habit to look specifically for research that can be readily repeated while reading publications or attending conferences. Certainly some research can be more readily repeated and it would help to steer students towards these works. The tradition of open source software in computer science results in many research projects that use readily available software and hardware. In addition, many original researchers make their software and data sets available online. I have found a set of examples accessible to students at various stages of their education.

I would love to hear directly from others interested in encouraging repeated research and in contributing to a repository of repeated research suggestions.

5. References

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt and A. Warfield. Xen and the Art of Virtualization. Proceedings of the nineteenth ACM symposium on Operating systems principles, pp 164-177, Bolton Landing, NY, USA, 2003.

[2] B. Clark, T. Deshane, E. Dow, S. Evanchik, M. Finlayson, J. Herne and J. Matthews. Xen and the Art of Repeated Research. To appear in the 2004 USENIX Technical Conference FREENIX Track, June 2004.