Editorial: Need for Artifact Verified Articles in ACM Transactions

One of the complaints by many researchers in the field of design automation, and embedded systems including myself, has been that the results in most papers are hard to reproduce. Often, we find papers with astounding performance numbers—speed, latency, power consumption, accuracy, and many other performance characteristics. Unfortunately, when one contacts the authors for the code, or the models, they are unavailable—because the graduate student working on it has graduated by the time of publication and has not archived the artifacts properly, or because the system built has dependence on licensed software that may not be available to the one asking for the code, or may be the authors are unwilling to provide the artifacts.

The result of this unavailability of research artifacts are multiple folds—when comparing new results against reported results in the literature, one must quote numbers from the papers that were obtained in very different circumstances, with different—possibly better—implementation of algorithms or worse implementations, or the dataset on which results were obtained are drastically of different characteristics, and so on. Thus, the comparison becomes between apples and oranges. In many cases, the authors may have made subtle mistakes in their implementation, modeling, or pre-processing of datasets, getting results that are not reproducible under correct implementation. In some cases, the suspicion remains that the results are manufactured, which hopefully is rare. Finally, because of the ownership changes (graduate students moving on, project scientists and engineers taking up new positions), these artifacts are not archived anywhere, and the original authors themselves are unable to demonstrate the correctness of the results reported.

In other fields of science, there are laboratory notes kept with specific protocols, and those can be recalled by editors while considering the veracity of the results, and all experiments are presented with the exact protocols in order that other scientists can reproduce the results in their laboratories.

Somehow, until very recently, the field of Computer Science had not adopted these practices for many reasons, perhaps because the computer science experiments are very different in nature than typical laboratory experiments in other sciences. The assumption is that if the code is available, one can create the results – given the right dataset, possibly benchmark dataset, and the similar hardware/software configurations to run the code on.

As I was complaining about it, Ms. Laura Lander, the ACM journals manager at the ACM headquarters, as helpfully as ever, pointed me to the ACM website on “Artifact Review and Badging” (<https://www.acm.org/publications/policies/artifact-review-badging>), which describes the various types of badges that can be stamped on papers whose artifacts have been reviewed, verified, made available, replicated, or reproduced. There are separate badges for each category. I knew that several conferences—and I believe at least two ACM transactions (TOMS and TOMACS)—have started this special category of articles, namely, those stamped with “artifact verified.” In fact, my first experience with artifact verification was for a paper we presented at the IEEE Real-Time Systems Symposium last year. We were asked to submit the artifacts in the form of code, and it was verified by a committee formed by the program chairs. It seems that authors are quite enthusiastic about this process, as the artifact-verified paper is considered to be of more value than those without.

2018 Copyright is held by the owner/author(s).

1539-9087/2018/11-ART82

<https://doi.org/10.1145/3282437>

So, I propose that we start a special category for “artifact-verified” papers in *ACM Transactions on Embedded Computing* from Volume 18. We should start with moderate aim to have 25% papers in a year to be of the “artifact-verified” variety. A subcommittee of the editorial board will take the responsibility to verify artifact of papers that are declared at the submission time as candidates for artifact verification. We will also request ACM to provide mechanism to store these artifacts in the ACM TECS webserver, and articles associated with the artifacts should have links to download the artifacts with instructions on how to compile, run, and reproduce the results.

I believe that this will bring about a change in the way we do our comparisons with the past results, and it will help researchers and students to develop these artifacts further to improve, extend, and possibly even make larger systems including those artifacts. ACM must decide on license terms for using artifacts made available through ACM.

I request all future authors to consider submitting artifacts associated with their work along with their paper as a zip file containing detailed instructions on how to reproduce their results. This will only work for certain types of work—anything that will require specialized hardware, or not so commonly available experimental set up will not be qualified for this—at least for now.

Finally, this issue is the fifth issue of the Volume 17 with papers that are regular submissions, and no special issue or section has been included in this volume. I hope, as usual, that the readers enjoy the diversity of topics in the realm of embedded computing in this issue.

Sandeep K. Shukla
Indian Institute of Technology Kanpur

Editor-in-Chief