

Nullius in Verba: Reproducibility for Database Systems Research, Revisited

Wolfgang Mauerer^{1,2}
Stefanie Scherzinger³

¹Technical University of Applied Sciences Regensburg, Regensburg, Germany

²Siemens AG, Corporate Research, Munich

³University of Passau, Germany

April 21, 2021



The Royal Society motto [...]:
The determination [...] to verify all statements by an appeal
to facts determined by experiment. (Wikipedia)

Wolfgang Mauerer

- ▶ Industry: Siemens Corporate Research
- ▶ Academia: Professor at Technical University of Applied Sciences Regensburg
- ▶ Favourite reproducibility badge: ICSE 2019

Stefanie Scherzinger

- ▶ Industry (earlier): Software developer (IBM, Google)
- ▶ Academia: Professor at University of Passau
- ▶ Favourite reproducibility badge: BTW 2021

The Repeatability Experiment of SIGMOD 2008

I. Manolescu¹, L. Alraissini², A. Arion³, J. Ditrich⁴, S. Manegold⁵
 N. Polyzotis⁶, K. Schmitter⁷, P. Senellart⁸, S. Zoupanos⁹
 D. Shasha¹⁰

¹ INRIA Saclay - Île-de-France, France ² Informatik Institut für Informatik
³ University of Amsterdam, Netherlands ⁴ jens.ditrich@inf.ethz.ch
⁵ ETH Zurich, Switzerland ⁶ jens.manolescu@ethz.ch
⁷ CNRS, Valenciennes, France ⁸ stefan.senellart@epfl.ch
⁹ U. California, Santa Cruz, USA ¹⁰ (shasha@cs.cornell.edu)
¹¹ Courant Institute, New York, USA shasha@courant.nyu.edu

ABSTRACT
 SIGMOD 2008 was the first database conference that allowed to test submitted programs against their data to verify the experimental conditions. This paper discusses the conditions for this effort, the community's reaction, our experience, and advice for future similar efforts.

1. MOTIVATION
 Repeatability has been a fundamental driver of progress in science since the time of Francis Bacon in the 16th century. In natural sciences, repeatability allows one scientist to verify the assertions of another, occasionally exposing fraud, but more often simply providing a check against inadvertent errors.

Natural science papers continue to be the responsibility requirement by providing a complete description of the protocol used in an experiment (reagents, equipment used, data to the world number, times, temperatures etc.). The protocol used here is described in sufficient detail for another lab to replicate the experiment. Computer science papers can't generally do this, because software is far more complex than laboratory procedures.

Fortunately for computer science, however, a computerized paper could, in the experiment to build data to enable repeat paper or others. The community is that the need for data helps help experimental results, that there are no field in a computer setting also, fortunately it better can easily share

Philippe Bonnet, Juliana Freire, Stratos Idreos, Stefan Manegold, Ioana Manolescu, and Dennis Shasha — 2020 SIGMOD Contributions Award

SIGMOD AWARDS

Philippe Bonnet, Juliana Freire, Stratos Idreos, Stefan Manegold, Ioana Manolescu, and Dennis Shasha
 2020 SIGMOD Contributions Award



Long-Term Maintenance in Industry

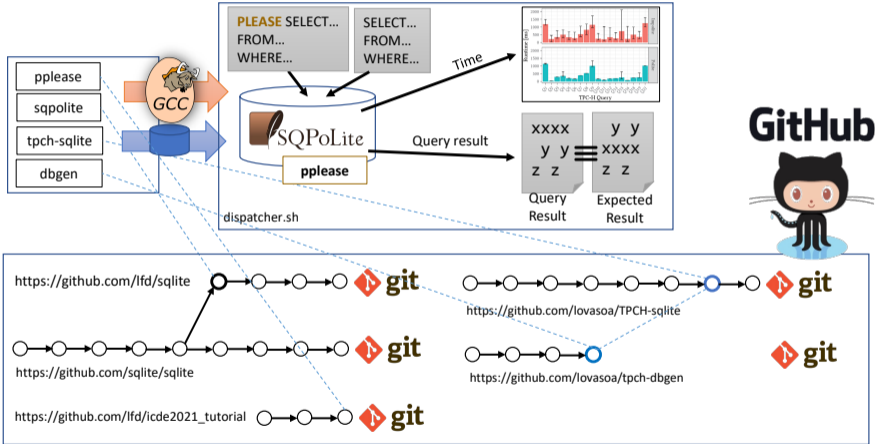
- ▶ Boeing 747 aircraft
 - ▶ Development started in 1966
 - ▶ Last machines produced in 2022 will be in service until about 2050
- ▶ Bitcoin
- ▶ Tor Browser
- ▶ [Civil Infrastructure Platform Initiative](#): Linux Kernels

Piggyback Strategy for Research

Put your money on open source tools that are massively employed:
as industry has a strong incentive (and the resources) to maintain them.

1. Working Example
2. Reproduction 101
 - 2.1 Goals
 - 2.2 Building and Comparing Artefacts
3. Building Docker Reproduction Images
 - 3.1 Self-Contained Execution Packages
4. Designing for Reproducibility
 - 4.1 Standing on the Shoulders of Giants
 - 4.2 Making History
 - 4.3 Licensing Considerations
5. End-to-end Reproduction
6. Summary

Play-along docker recipe: https://github.com/lfd/icde2021_tutorial.



We measure

- ▶ Query runtime
- ▶ Throughput
- ▶ Latency

We encounter:

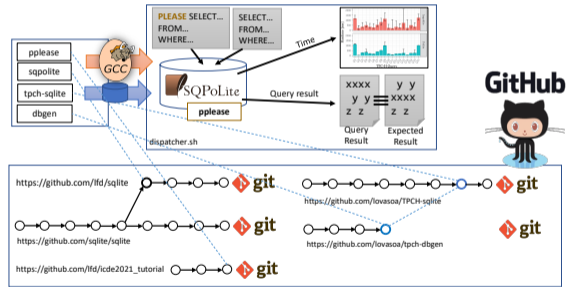
- ▶ Noise
- ▶ Variability


```
1 void measure_query(char *dbfile , char *basepath, char *query, int iterations) {
2     // ... pre-allocate mm buffer to hold time stamps ...
3     // ... open DB connection, open dbfile, read in query string ...
4
5     tstamps[0] = get_tstamp(); // Record start timestamp.
6     for (int count = 1; count < iterations; count++) {
7         rc = sqlite3_exec(db, sql, 0, 0, &err_msg);
8         // ... handle potential errors ...
9
10        tstamps[count] = get_tstamp(); // Record end timestamp.
11    }
12
13    // ... write durations from mm buffer to to stdout ...
14    // ... free resources ...
15 }
```

latency.c

```
1  sqpolite> .load ./pplease
2  sqpolite>
3  sqpolite> select pplease(o_comment) from orders limit 3;
4  What was the magic word again?
5  nstructions sleep furiously among
6   foxes. pending accounts at the pending, silent asymptot
7  sly final accounts boost, pplease. carefully regular
8  sqpolite>
9  sqpolite>
10 sqpolite> select pplease(o_comment) from orders limit 3;
11 nstructions sleep furiously among
12 foxes, pplease. pending accounts at the pending, silent asymptot
13 sly final accounts boost. carefully regular
```

- ▶ Artefacts (data, code, SW tools, scripts)
- ▶ Build process
- ▶ Need to compare both the artefacts *and* the results.



Various choices for target dimensions:

- ▶ Artefacts: bitwise identity, functional identity
- ▶ Observables
 - ▶ physical quantities (time, energy consumption, ...)
 - ▶ functional results (bitwise-identical, permutation-identical, ...)
 - ▶ stochastic results (numerical noise, ...)
- ▶ Build the toolchain itself (y/n)
- ▶ Unspecified dependencies (y/n)
- ▶ External sources (y/n)

Hands-on examples: Building & Comparing

- ▶ Binaries (C/C++)

Further Challenges

- ▶ Linking
- ▶ Libraries
- ▶ Tarballs

Tool Support

- ▶ Reptest (pypi.org/project/reptest/)
- ▶ Diffoscope (diffoscope.org/)
- ▶ See the accompanying website for further links

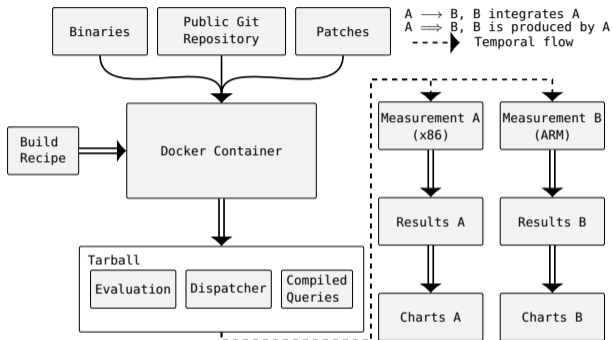
Create “an island of stability in an ocean of activity”

Be independent of updates to external components, because...

- ▶ Base system details change (package versions, etc.)
- ▶ System runtime configuration (beyond distro and kernel versions) changes
- ▶ Github repositories disappear
- ▶ Projects move between hosts (Sourceforge, GitHub, ...)
- ▶ External software is no longer maintained, download links disappear
- ▶ ...


Goal: Build self-contained, complete environments

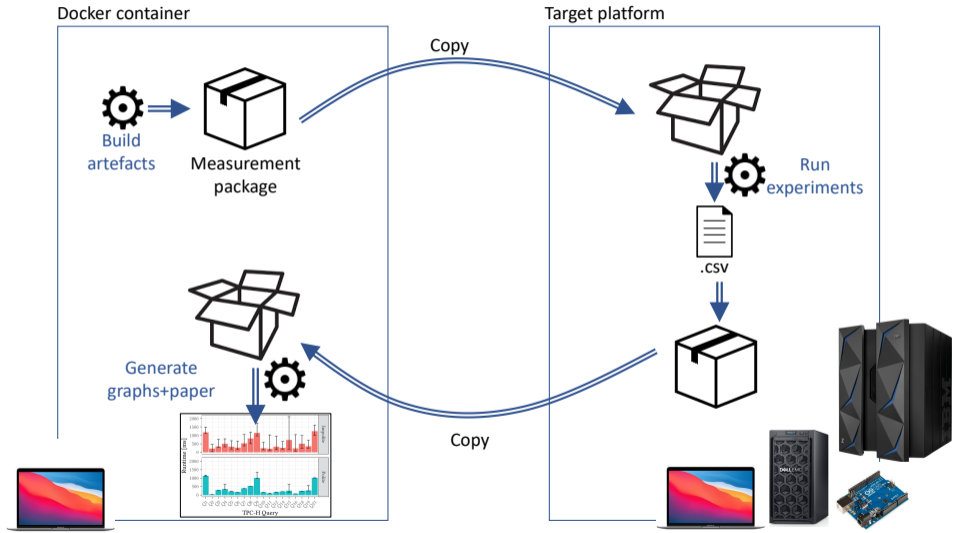
Be ready to build your stuff even when you trapped on an island without internet access, or 20 years after all the repositories have gone.




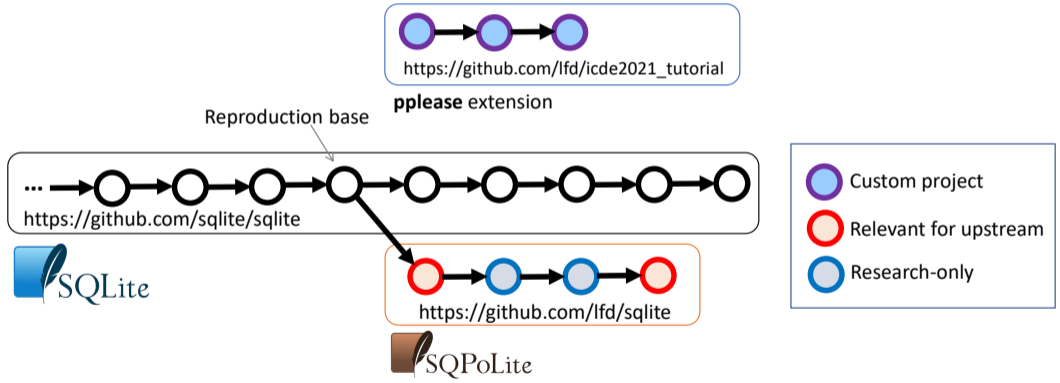
Wolfgang Mauerer, Ralf Ramsauer, Edson Ramiro Lucas Filho, Daniel Lohmann, Stefanie Scherzinger:
Silentium! Run-Analyse-Eradicate the Noise out of the DB/OS Stack. BTW 2021.











Hands-On Example: Patch Stacks

- ▶ Presenting thoughts versus presenting results
- ▶ Rebasing, squashing, rewriting and all of that
- ▶ Self-documenting patches
- ▶ Trail of responsibility/lineage and provenance
- ▶ Upstream, integrate, externalise?



Steps

- ▶ Choosing a license
- ▶ Specifying a license
- ▶ SPDX

Proprietary, closed-source components

~~Just don't!!~~ Try not to use them.

We have solved two problems:

- ▶ Our build is reproducible.
- ▶ Our results are reproducible.

Two more to go:

- ▶ We need to write a paper.
- ▶ We need to make our artefacts available.

The screenshot shows a Zenodo repository page for a software package. The page is titled "Reproduction package for 'Silentium! Run-Analyse-Eradicate the Noise out of the DB/OS Stack'". It includes a search bar, a navigation menu with "Upload" and "Communities", and a "Log in" / "Sign up" button. The main content area displays the package title, authors (Maurer, Wolfgang; Ramsauer, Ralf; Lucas, Edson; Lohmann, Denise; Scherzinger, Stefanie), and a list of files for download. The files table has columns for Name, Size, and a Download button. The files listed are:

Name	Size	Download
btw2021-data.tar.bz2	5.0 GB	Download
md5fd197bb932064d7a3fb810536cd9670		
btw2021.img.bz2	1.7 GB	Download
md5d1396f4375be2591382d993798599		
btw2021.tar.bz2	24.8 MB	Download
md569bc5840cdcafd17a43efc9f6c8453e		

Below the files table, there is a "Citations" section with a search bar and a "Show only" filter. The filter options are Literature (0), Dataset (0), Software (0), and Unknown (0). There are no citations listed.

On the right side of the page, there are several summary boxes:

- Views and Downloads:** 6 views, 5 downloads. A link to "See more details..." is provided.
- Indexed in:** OpenAIRE
- Publication date:** March 12, 2021
- DOI:** 10.5281/zenodo.4602296
- Grants:** European Commission, IDEV40 - Integrated Development 4.0 (783163)
- License (for files):** Creative Commons Attribution 4.0 International
- Versions:** Version 1 (Mar 12, 2021) with DOI 10.5281/zenodo.4602296
- Cite all versions:** A note explaining that the DOI 10.5281/zenodo.4602295 represents all versions and will always resolve to the latest one.
- Share:** A button to share the page.

Benefits vs. Costs

- ▶ We won't lie to you: Reproducibility engineering is a lot of work.
- ▶ Enablers to reproducibility: Testability & automation
- ▶ However, there are clear benefits for your team, just imagine:
 - ▶ Long-term, automation will save you time.
 - ▶ You can build your stuff even after your PhD student has graduated.
 - ▶ You can switch to a new notebook without risking your progress in research.
 - ▶ You will no longer have to negotiate within the team why s.th. works for you, but only you.

- ▶ Stefanie Scherzinger's contribution was supported by the *Deutsche Forschungsgemeinschaft* (DFG, German Research Foundation) – 385808805.