

Reducing Technical Debt with Reproducible Containers

(the slides are available under "Presentation Materials" in the above URL)

Date: November 4, 2020

Presented by: Tanu Malik (DePaul University)

Q. On your description of tested languages I noticed that CUDA wasn't present. Does this work with cuda compiled pytorch, tensorflow, etc?

A. Currently we can audit and create a sciunit. We can repeat that Sciunit in a Docker container environment, but exact repeat (using Sciunit repeat) is not possible due to GPU specific system calls.

Q. Who "pays down" the technical debt of the Sciunit system itself?

A. Sciunit is a software and so its technical debt is not different than technical debt of other softwares. We the developers of Sciunit have to pay its Technical Debt.

B. Sciunit's technical debt is intentionally kept low by keeping it dependent on only Python \geq 2.7 and glibc \geq 2.17

That being said, Sciunit automates repayment of technical debt for other softwares.

Q. Does sciunit work similarly with Singularity?

A. Yes.

Q. Is there a Sciunit to Docker converter? I would imagine for wider deployment, that might be useful

A. Yes. Sciunit to Dockerfile convertor is the basis for the following publication:
J. Chuah, M. Deeds, T. Malik, Y. Choi, J. Goodall, "Documenting Computing Environments for Reproducible Experiments", In *Parallel Computing: Technology Trends*, 756-765, 2020, *doi*: 10.3233/APC200106.
The convertor is yet to be publicly released.

Q. How do you deal with host cpu/system-dependent system calls such as AVX512 intrinsics, MKL presence or absence etc?

A. We don't deal with them. Note we only need to deal with data-dependent system calls--not all system calls.

Q. How does Sciunit handle interfacing with fabric libraries (verbs) on various machines for MPI applications, which are typically installed/managed by the site systems administrators?

A. Sciunit does not require root privileges. So if the application running in userspace accesses them the resulting container will have it. I am not exactly sure what the audience member refers to by "interfacing".

Q. Singularity MPI applications can scale to 100s or 1000's of nodes at near the same timing as bare metal, can one do the same with an MPI application that is packaged with Sciunit?

A. This is part of our current work.