WRENCH: Cyberinfrastructure Simulation Workbench

Henri Casanova[1], Rafael Ferreira da Silva[2], Ryan Tanaka[1,2], Gautam Jethwani[2], William Koch[1], Tongyu Zhu[2], Frédéric Suter[3]

[1]University of Hawai`i at Mānoa – Computer Science Department
[2]University of Southern California – Information Sciences Institute
[3]Centre National de la Recherche Scientifique – Institut National de Physique Nucléaire et de Physique des Particules

MOTIVATION

✦Disconnect between CI theory and practice

✦Many theoretical results are not useful to practitioners
✦One well-known reason is that theoretical results are obtained with models that, to be tractable, are often unrealistic or unattainable in practice
✦As a result, practical work must be experimental

✦Real-World Experiments are Limited

✦One is limited to particular platform configurations (and sub-configurations)
✦One is limited by specifics of the software infrastructure that impose constraints on CI application executions
✦Limited Experimental Scope impedes progress / discovery

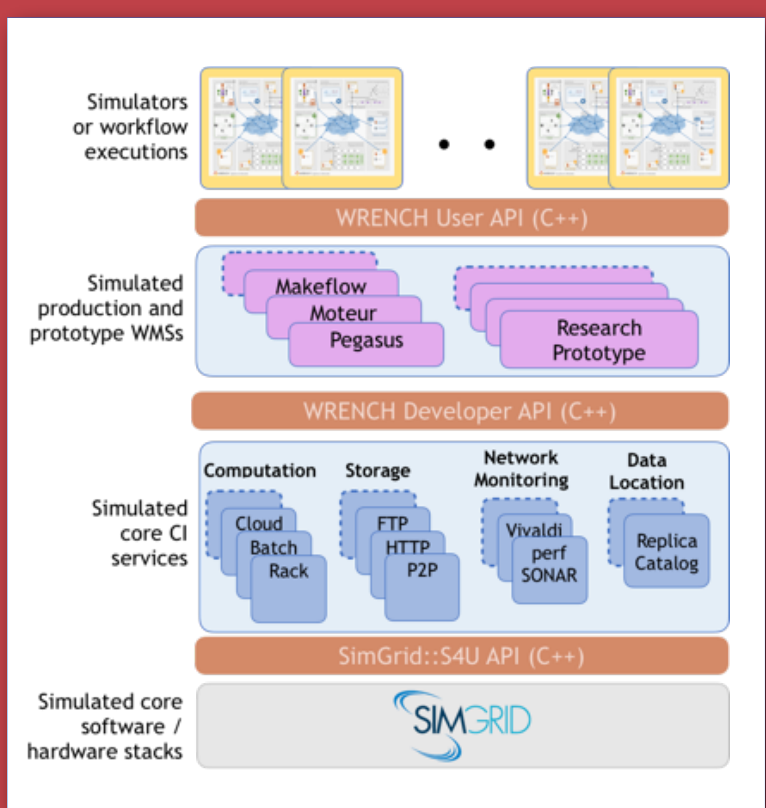

OBJECTIVES

✦Simulation

✦When one works in an experimental field in which experiments are problematic, one resorts to simulation
✦In some fields of Computer Science simulation is a standard research and development methodology

✦Simulation-driven engineering life cycle

✦The ability to easily develop accurate CI simulators, from which research products evaluated via experimental simulation could be seamlessly integrated into actual CI platforms

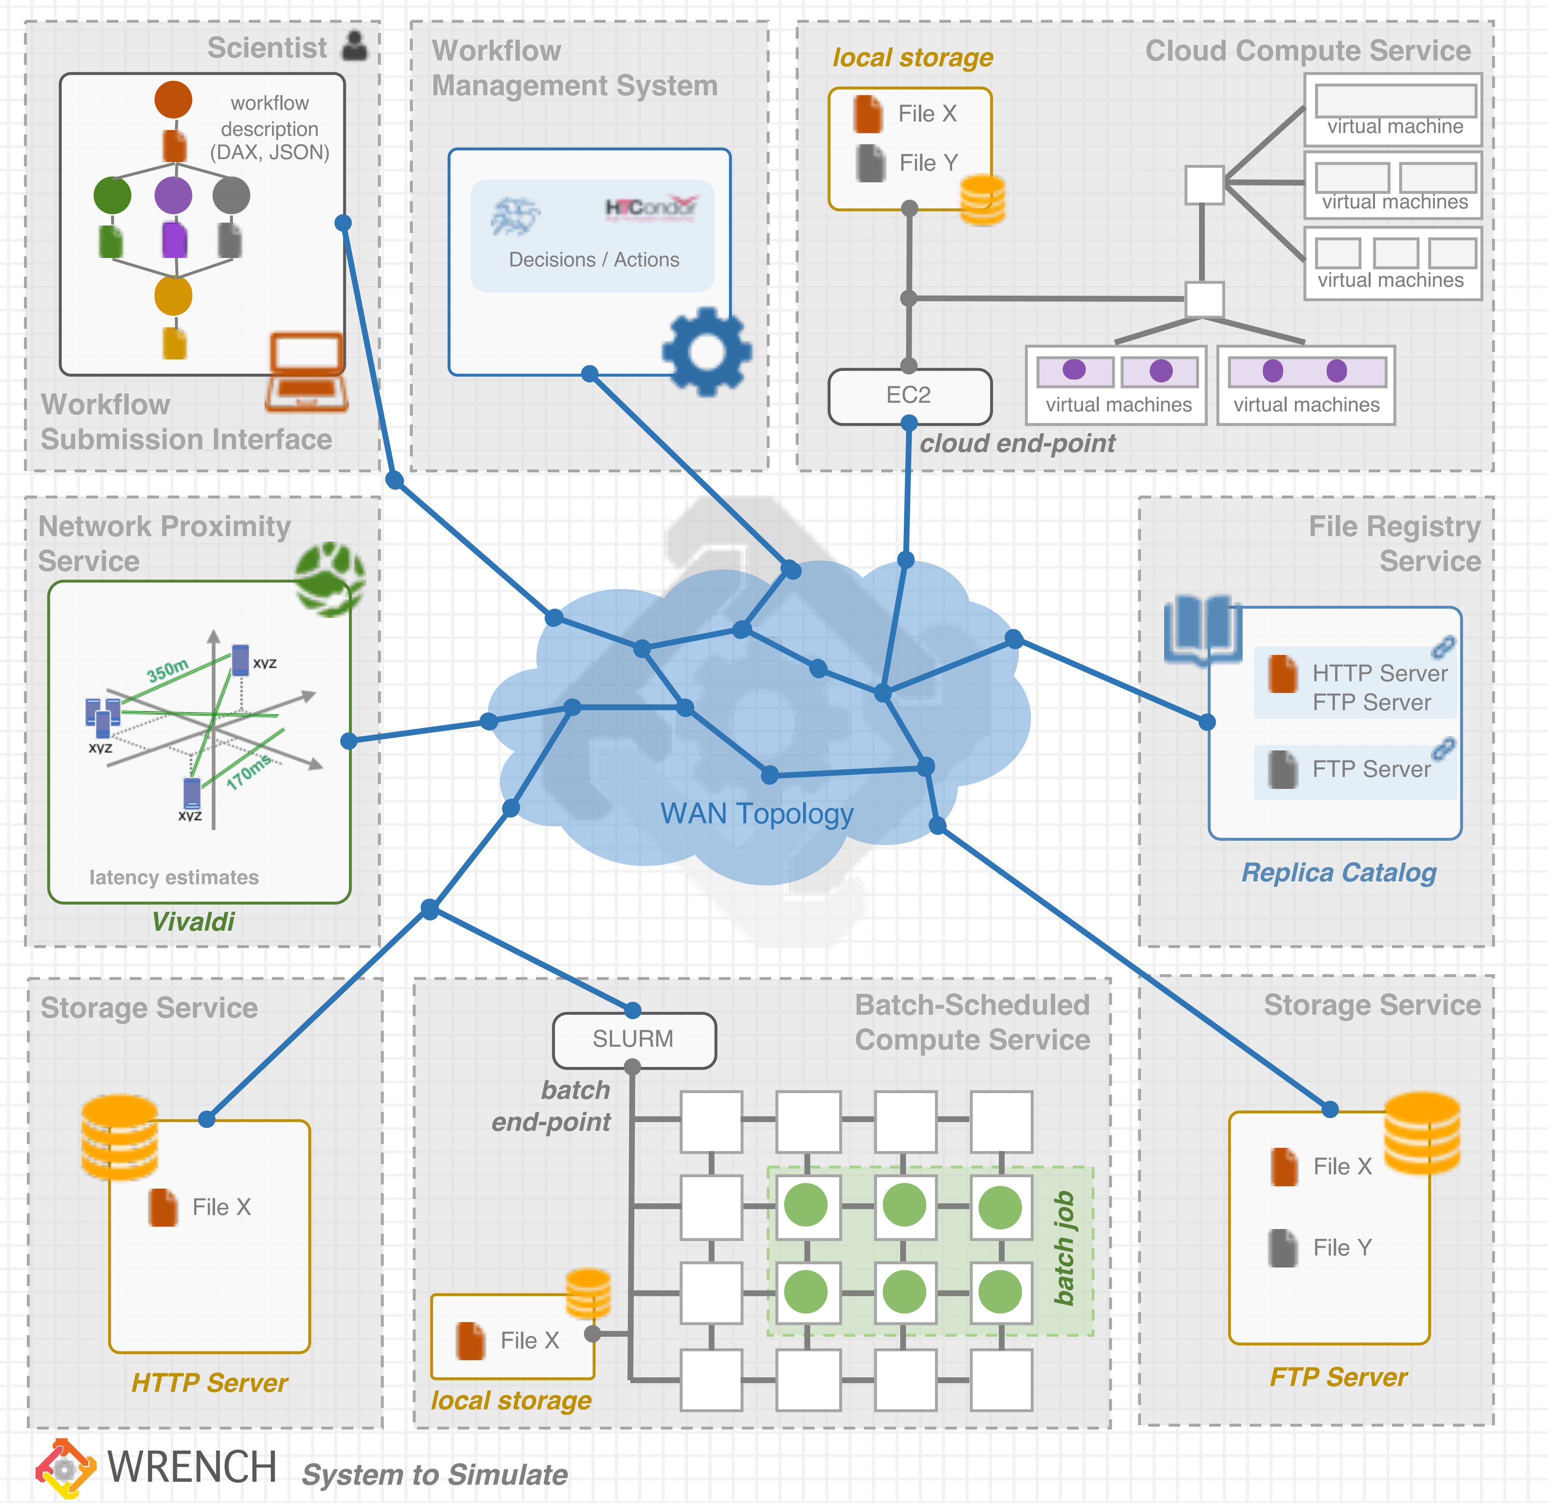

SIMULATION ACCURACY

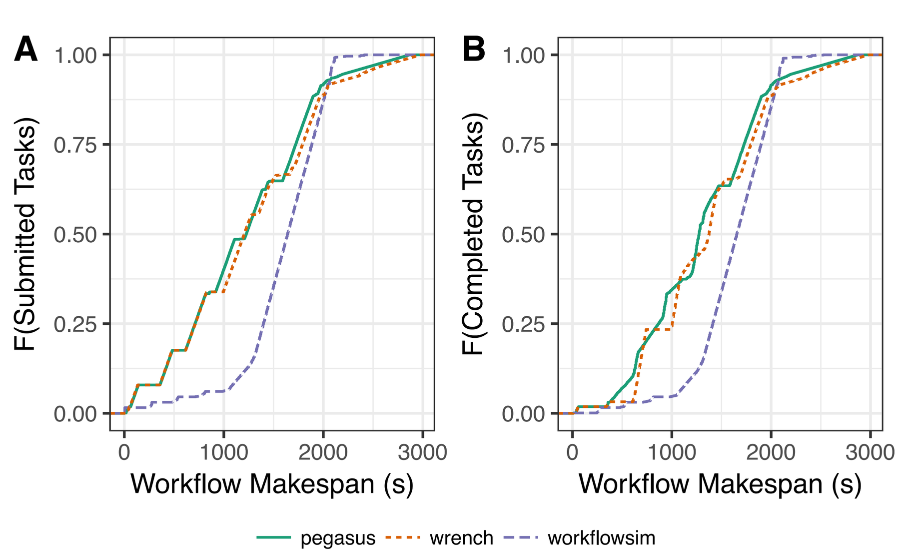

Accuracy: the ability to capture the **behavior** of a real-world system with as little bias as possible

SIMULATION SCALABILITY

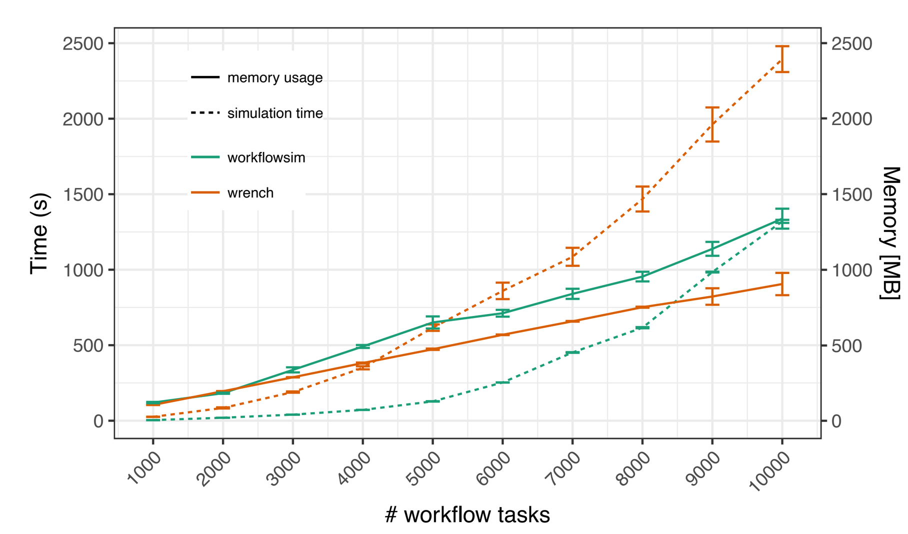

Scalability: the ability to simulate **large systems** with as few CPU cycles and bytes of RAM as possible

WRENCH's IMPACT ON CI RESEARCH

Investigated the impact of resource utilization and **I/O operations** on the **energy usage**, as well as the impact of executing multiple tasks concurrently on multi-socket, multi-core compute nodes

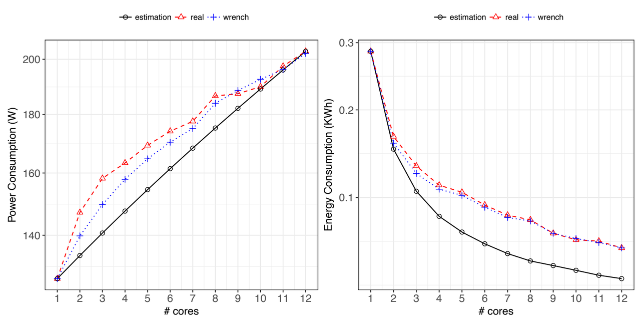

Data-aware planning algorithm that leverages two characteristics of a family of **virtual machines** instances, i.e., a large number of cores and a dedicated storage space on fast SSD drives, to improve data locality, hence **reducing the amount of data transfers** over the network during the execution of a workflow

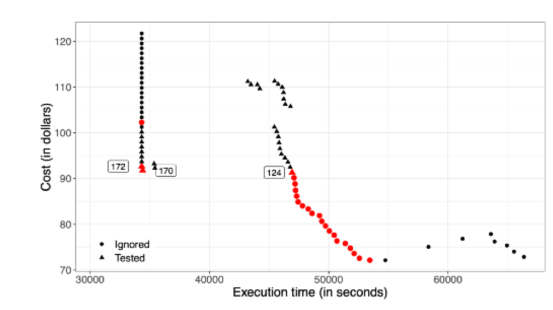

Using **accurate simulations** and real-world experiments, they study how to best use new storage layers when executing scientific workflows

They characterize the I/O behaviors of several real-world workflows when using several data placements strategies, and use these characterizations to calibrate a simulator of workflow executions on **HPC platforms with burst buffers** to estimate potential gains

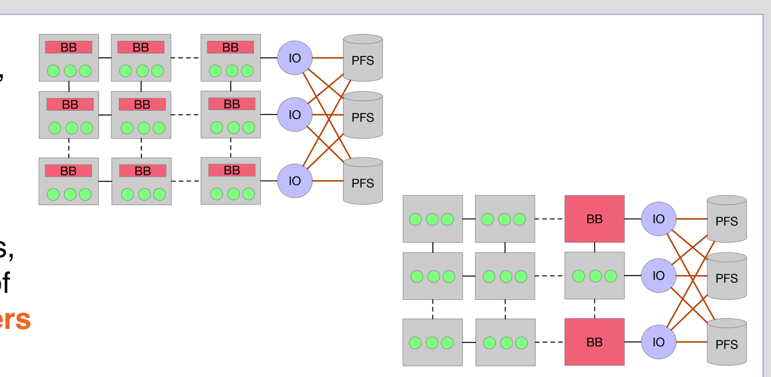

WRENCH PEDAGOGICAL MODULES

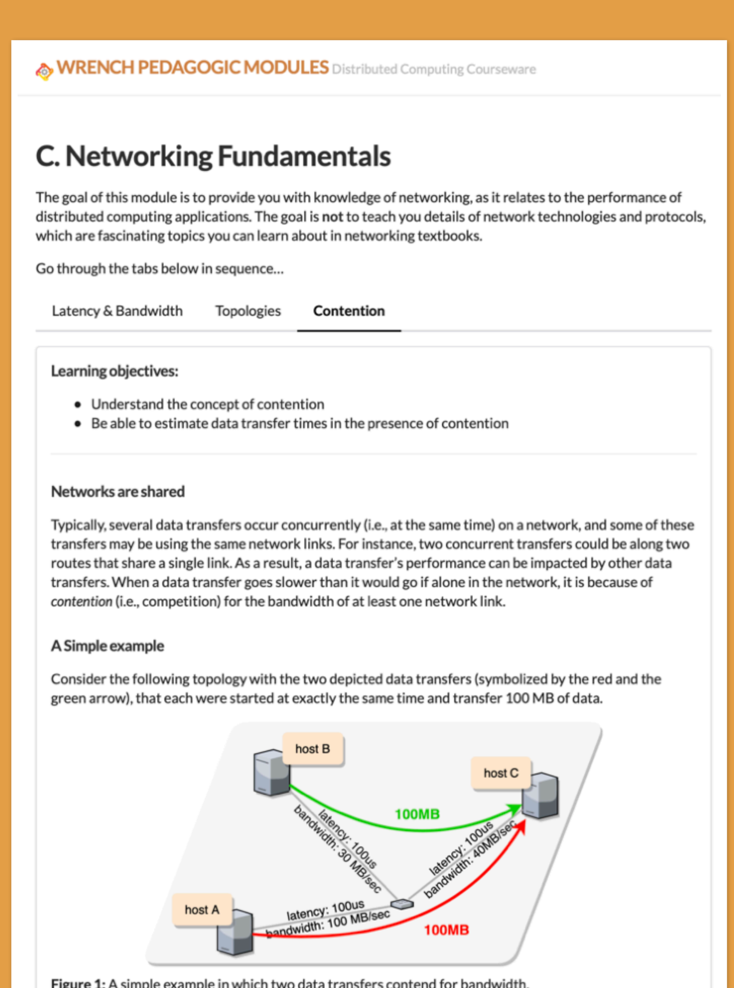

https://wrench-project.org/wrench-pedagogic-modules

Simulation-driven self-contained pedagogic modules supported by WRENCH-based (accurate and scalable) simulators

Activities entail running, through a **Web application**, a simulator with different input parameters

Principles of Computing and Distributed Computing

Single-core computing speed, work, RAM	Networking latencies, bandwidth, topologies, contention
Multi-core computing speedup, efficiency, idle time	I/O HDD/SDD, data rates, overlap with computation

Applying Principles To Workflows

Scientific workflows basic concepts	Workflows and parallelism multi-core, multi-node
Workflows and data locality network proximity of data	Provisioning resources meeting performance goals within budget

SIMULATED CORE CI SERVICES

Abstractions for simulated CI components to execute computational workloads

Compute Services

Provide mechanisms for executing application tasks, which entail I/O and computation

bare-metal · virtualized cluster · cloud · batch-scheduled cluster

Storage Services

Store application files, which can then be accessed in reading/writing by the compute services when executing tasks that read/write files

File Registry Services

Databases of key-value pairs of storage services and files replicas

Network Proximity Services

Monitor the network and maintain a database of host-to-host network distances

Workflow Management System

Provides the mechanisms for executing workflow applications, including decision-making for optimizing various objectives

WRENCH in numbers (since 2018)

- **8** stable software releases
- **3400+** unique users visiting the website from **100+** countries
- **4** peer-reviewed research papers
- recruited undergraduate students **6**
- **7** simulators for pedagogic modules
- WRENCH-enabled simulators for research **5**
- **8** pedagogic modules
- institutions from **7**
- **4** and graduate students
- **3** different levels of documentation
- **5** countries using WRENCH for CI-related research and education
- **100+** software downloads from GitHub
- contributors for WRENCH core software **11**
- **>90%** code coverage
- **19K** lines of code of WRENCH core software
- **2000+** continuous integration builds
- **2** research papers from WRENCH users*
- **250+** pulls of Docker containers

SIMGRID AS A CORE SIMULATION TECHNOLOGY

https://simgrid.org

✦ SimGrid is a **research project**
 ✦ Development of simulation models of hardware/software stacks
 ✦ Models are accurate (validated/invalidated) and scalable (low computational complexity, low memory footprint)
✦ SimGrid is **open source usable software**
 ✦ Provides different APIs for a range of simulation needs, e.g.:
  ✦ S4U: General simulation of Concurrent Sequential Processes
  ✦ SMPI: Fine-grained simulation of MPI applications
✦ SimGrid is a **versatile scientific Instrument**
 ✦ Used for (combinations of) Grid, HPC, Peer-to-Peer, Cloud simulation projects
✦ First developed in 2000, latest release: v3.24 (October 2019)

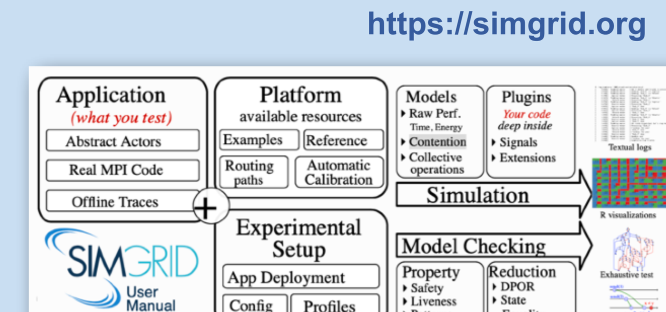

LEARN MORE

Get in Touch

https://wrench-project.org – support@wrench-project.org

WRENCH is funded by the National Science Foundation (NSF) under grants number 1642369, 1642335, 1923539, and 1923621; and the National Center for Scientific Research (CNRS) under grant number PICS07239.