



Simulation Update

Paul Nilsson, Raees Ahmad Khan, Sairam Sri Vatsavai

October 1, 2024

Redwood Meeting



@BrookhavenLab

XML Generation for SimGrid/WRENCH

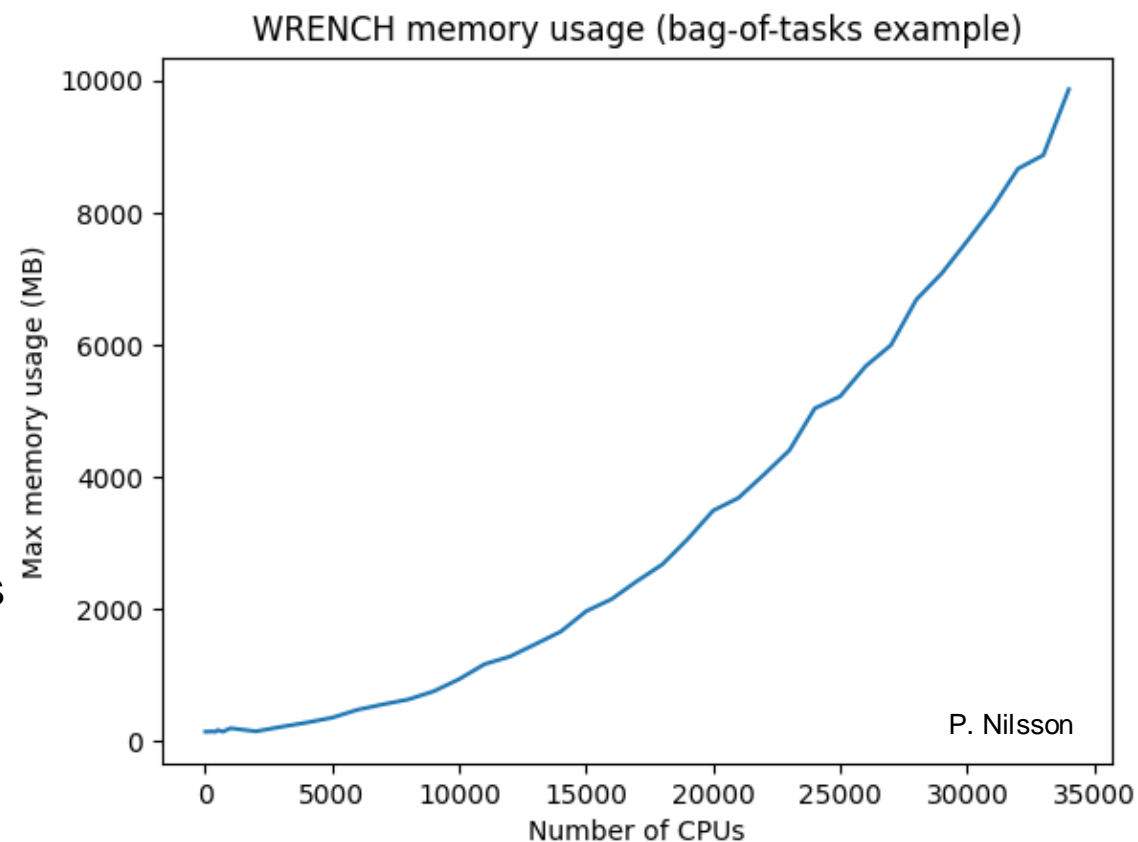
- Topology for individual grid site as well as network of sites can be defined in so-called platform.xml, used as input to SimGrid/WRENCH
 - Contains info about all the hardware
- Platform files can be generated using information we already have collected (JSON)
- XML gets rather large very quickly – a site with 30k CPUs leads to a platform file over 27MB
- Better to get all necessary info from JSON files created earlier and set up infrastructure in the simulator programmatically
- Scripts for creating JSON files and for XML generation in GitHub
 - <https://github.com/PalNilsson/redwood>

Memory usage in WRENCH example

- Modified standard “bag-of-tasks” example using WRENCH version 2.2 (based on SimGrid version 3.34)
- Dynamically creates simple geometry for single computing resource (“site” with many compute hosts)
 - Essentially loop over create_hosts() using some default values

```
// Create the ComputeHosts and their network links
for (int i = 0; i < num_hosts; ++i) {
    std::string host_name = "ComputeHost_" + std::to_string(i);
    auto compute_host = zone->create_host(host_name, "1Gf");
    compute_host->set_core_count(8);
    [plus network routes etc]
```

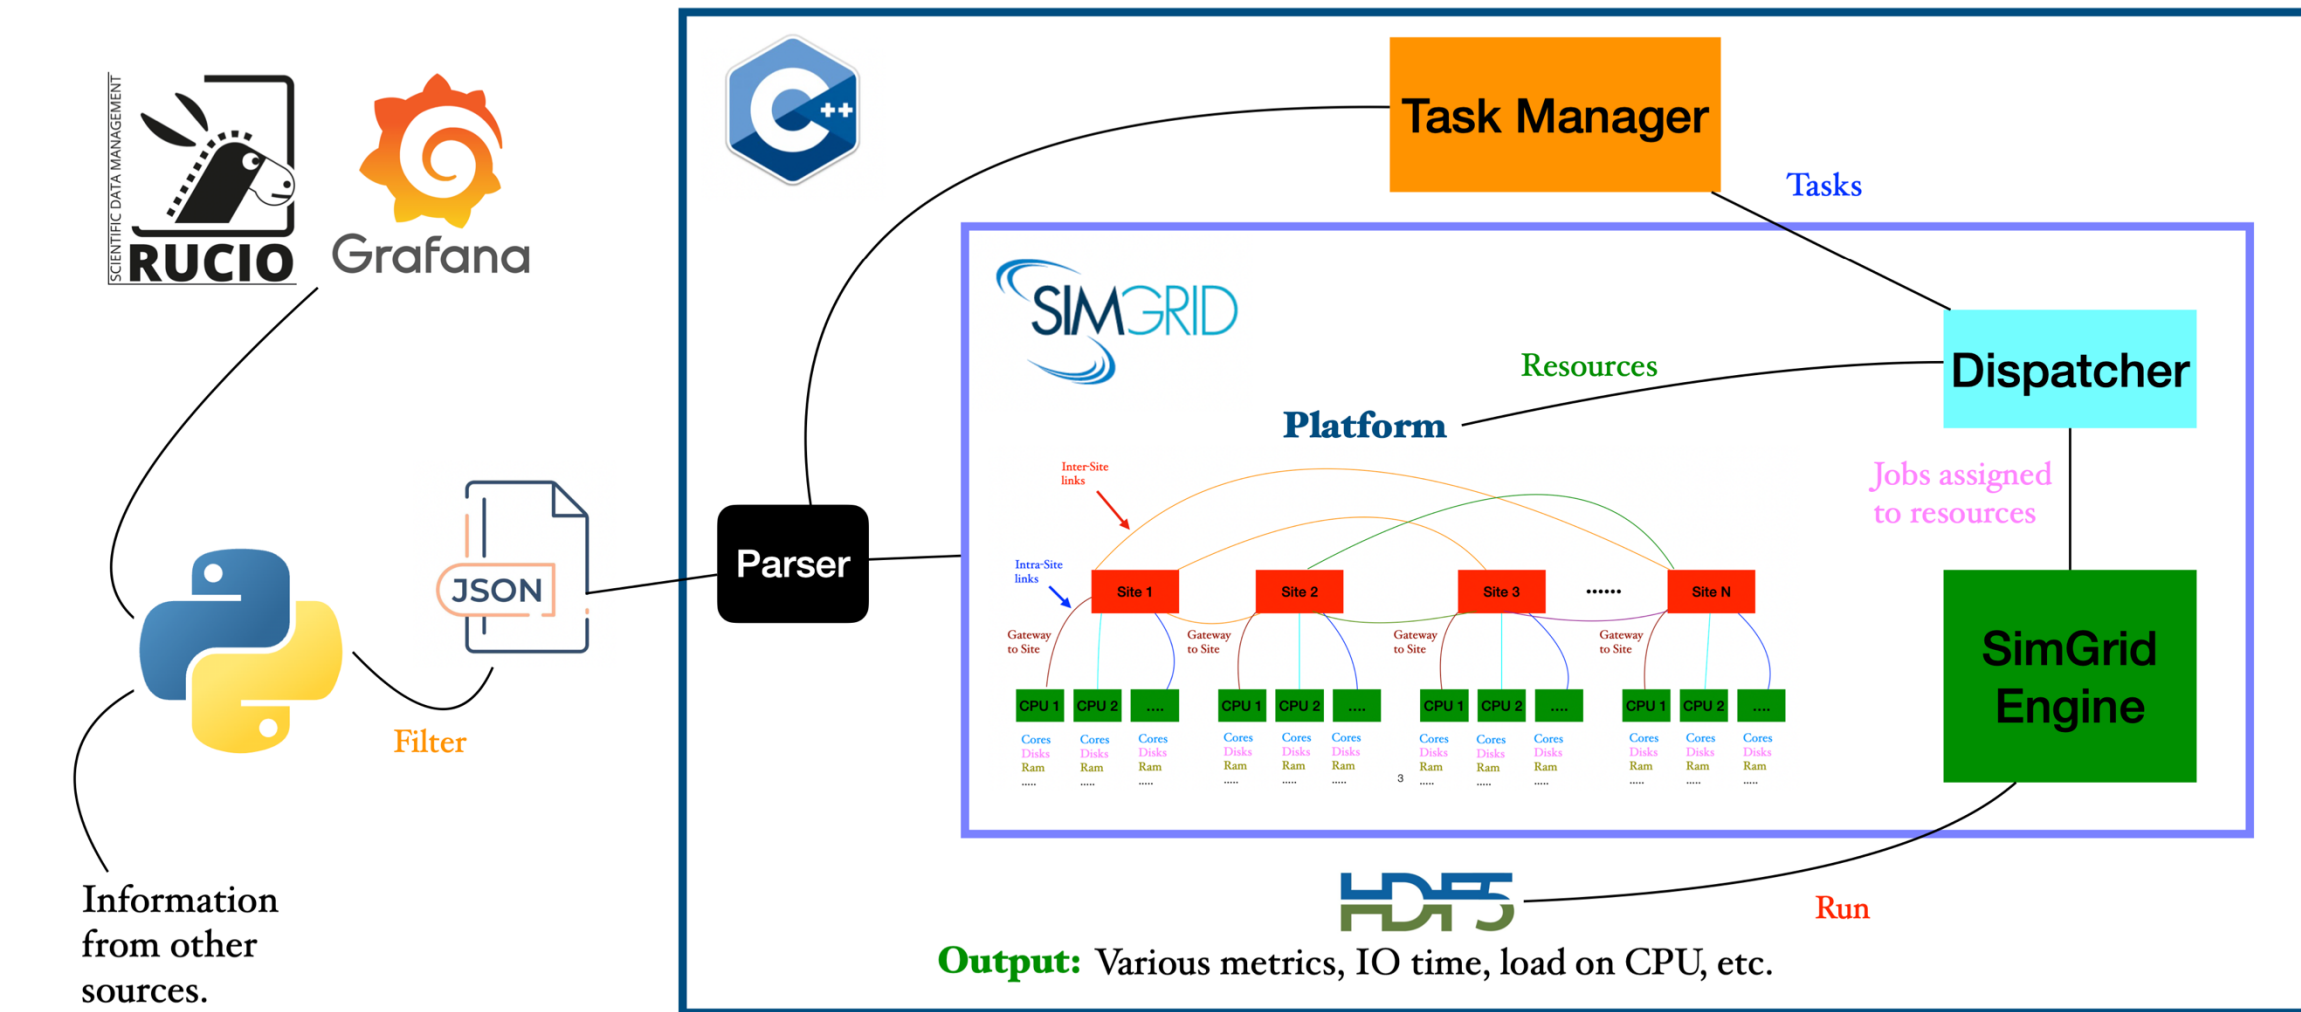
- Dramatic increase of used memory as number of hosts increase
 - Site with 30k CPUs uses around 8 GB during simulation
 - I.e. simulating 150 grid sites is not going to work with current version
 - Most of the memory is allocated while setting up geometry, another 10-20% increase in memory usage when simulation starts (10,000 jobs)



Going Forward – Switched to SimGrid

- Raees has developed first version of ATLAS grid simulation based on SimGrid 3.36 (latest)
 - Installed and tested by Paul and Sairam
- Simulation is a preliminary version of the ATLAS grid with disk space and connections taken from previously created JSON file
- Code repository <https://github.com/REDWOOD24/ATLAS-GRID-SIMULATION>
 - Repo is now "locked", development is done in private repos
 - New code will only be added via pull requests
 - Versioning just introduced to facilitate distributed team development and testing
 - "1.0.0"-style, [VERSION].[REVISION].[BUILD]
- Code validation to be added and used with GitHub Actions (a la PanDA Pilot)
 - E.g. run static analysis tool and linter upon a pull request
 - Practical to start with manual testing (use GitHub Actions later), but start soon with using validation tools before we have too much code (as soon as we have agreed on which tools to rely on, which coding convention to use, ..)
- Preliminary study of memory usage as function of number of CPUs looks good
 - Memory usage did not pass 3 GB while running 5k tasks with "full" ATLAS grid geometry

Basic Workflow



Status and Plans

- We have tested both WRENCH and SimGrid, and decided on building the simulation around SimGrid which seems to avoid memory issues
 - Developed a preliminary software description of the ATLAS GRID and tested submitting jobs on it and collecting the output statistics
- Need to make things more accurate (CPU, DISK, etc)
 - Currently, the number of CPUs are randomly distributed (to realistic order), disk sizes are preliminary
 - Investigate if disk info is available for all queues, chase it down if missing
- Investigate running time of simulation
 - Processing 5k tasks (152k jobs) uses 2.7 GB memory but takes 2h 41 minutes to complete on M1 Mac
 - Can this be optimized?
- Move from simple to realistic brokering (longer term goal)
 - Need to discuss how exactly to implement this (the algorithm used by PanDA is rather complex)
 - Implement in steps of increasing complexity