



ABSTRACTS PMACS

SPARSE GRID REGRESSION FOR PERFORMANCE PREDICTION USING HIGH-DIMENSIONAL RUN TIME DATA

Philipp Neumann

We employ sparse grid regression to predict the run time in three types of numerical simulation: molecular dynamics (MD), weather and climate simulation. The impact of algorithmic, OpenMP/MPI and hardware-aware optimization parameters on performance is studied. We show that normalization of run time data via algorithmic complexity arguments significantly improves prediction accuracy. Mean relative prediction errors are in the range of few percent; in MD, a five-dimensional parameter space exploration results in mean relative prediction errors of ca. 20% using ca. 178 run time samples.

MPCDF HPC PERFORMANCE MONITORING SYSTEM: ENABLING INSIGHT VIA JOB-SPECIFIC ANALYSIS

Luka Stanisic and Klaus Reuter

This paper reports on the design and implementation of the HPC performance monitoring system, deployed to continuously monitor performance metrics of all jobs on HPC systems at the Max Planck Computing and Data Facility (MPCDF). Thereby it reveals important information to various stakeholders, in particular to the users, application support, system administrators, and to management. On each compute node, hardware and software performance monitoring data is collected by our newly developed lightweight open-source hpcmd middleware which builds upon standard Linux tools. The data is transported via rsyslog, and aggregated and processed by a Splunk system, enabling detailed per-cluster and per-job interactive analysis in a web browser. Additionally, performance reports are provided to the users as PDF files. Finally, we report on practical experience and benefits from large-scale deployments on MPCDF HPC systems, demonstrating how our solution can be useful to any HPC center.

TOWARDS A PREDICTIVE ENERGY MODEL FOR HPC RUNTIME SYSTEMS USING SUPERVISED LEARNING

Gence Ozer, Sarthak Garg, Neda Davoudi, Gabrielle Poerwawinata, Matthias Maiterth, Alessio Netti, and Daniele Tafani

High-Performance Computing systems collect vast amounts of operational data with the employment of monitoring frameworks, often augmented with additional information from schedulers and runtime systems. This amount of data can be used and turned into a benefit for operational requirements, rather than being a data pool for post-mortem analysis. This work focuses on deriving a model with supervised learning which enables optimal selection of CPU frequency during the execution of a job, with the objective of minimizing the energy consumption of a HPC system. Our model is trained utilizing sensor data and performance metrics collected with two distinct open-source frameworks for monitoring and runtime optimization. Our results show good prediction of CPU power draw and number of instructions retired on realistic dynamic runtime settings within a relatively low error margin.

RESOURCE AWARE SCHEDULING FOR EDA REGRESSION JOBS

Saurav Nanda, Ganapathy Parthasarathy, Parivesh Choudhary, and Arun Venkatachar

Typical Integrated Circuit (IC) design projects use Electronic Design Automation (EDA) tool flows to launch thousands of regressions every day on shared compute grids to complete the IC design verification process. These regressions in turn launch compute jobs with varied resource requirements and inter-job dependency constraints. Traditional grid schedulers, such as the Univa Grid Engine (UGE) [1] prioritize fairness over performance to maximize the number of jobs run with equal distribution of resources at any time. A constant challenge in day-to-day operations is to schedule these jobs for minimum overall job completion time so that developers can expect predictable regression turn-around time (TAT).

We propose a resource-aware scheduling mechanism that balances performance and fairness for real-word EDA-centric workloads. We present an analysis of historical profile information from a set of regressions with complex inter-job dependencies and highly variable resource requirements to show that many of these regression jobs are well suited for efficient packing on grid machines.

We formulate the regression scheduling problem as a variant of the *bin packing problem*, where the size of bins and balls may vary according to job-resource requirements and differing server configurations on the grid. We propose using two analytic techniques – namely K-means Clustering [2] and adaptive histogram-based binning [3], to solve this problem. We then evaluate the performance of our proposed solution using real workloads from daily regressions on an enterprise compute grid.



Submission deadline for the Euro-Pas PhD Symposium has been extended to 27 May, 2023. Click here for more information - <https://t.co/wWxislCJSC>

17.05.2023 - 11:51

The Euro-Par PhD Symposium is a welcoming and supportive forum for PhD students to present their work. Click here for more information: <https://t.co/wWxislCJSC>

04.04.2023 - 09:25

Submit your paper for EURO-PAR 2023 Workshops and Minisymposia! Click here for more information. <https://t.co/UEseXWb3Dz>

07.03.2023 - 08:18

Abstract submission deadline: 24 Feb, 2023. Click here for more information: <https://t.co/UEseXWb3Dz>



CONTACT US

HOSTS



SPONSORS



SHARE ON:



