Beyond the Standard Model

-- Towards an integrated modeling methodology for performance and power --

Adolfy Hoisie[1] (PI), Kevin J. Barker[1], Greg Bronevetsky[2], Laura Carrington[3], Marc Casas[2], Daniel Chavarria[1], Roberto Gioiosa[1], Darren J. Kerbyson[1], Gokcen Kestor[1], Nathan R. Tallent[1], Ananta Tiwari[3]

[1]Pacific Northwest National Lab, [2]Lawrence Livermore National Lab, [3]UC San Diego/SDSC

Multiple significant technical challenges punctuate the path to Exascale. Millions of threads of parallelism need be tractable and manageable. New algorithms may be required that can utilize these extreme levels of parallelism and can be efficiently mapped to system resources. Modeling is an ideal vehicle to explore the needs in advance of implementation and system availability as well as for use in dynamic decision-making within intelligent runtime systems.

Within Beyond the Standard Model (BSM) new approaches to predict, assess and model performance and power both statically and dynamically are being developed. This includes methodologies, tools, and new ways to extract insight. Required is a rational system design: co-design of systems (hardware, software, run-time systems, execution models) and applications through an accurate and quantitative methodology. We envision models will be part of the lifecycle of systems and applications from their design phases, to optimization, to production. Overall, models in all facets, area of applicability and use, and methodology base will become ubiquitous.

Our high level view of *dynamic modeling* is shown in Figure 1. The idea is that models will sit in the middle of and help all the other layers of the Exascale stack make intelligent decisions in real time. For example, the goal of minimizing data movement requires models to effectively guide runtime decisions made by the runtime software to determine when and where data should be moved/copied/recomputed. Only then can the data motion challenges be conquered and the door to Exascale be opened.

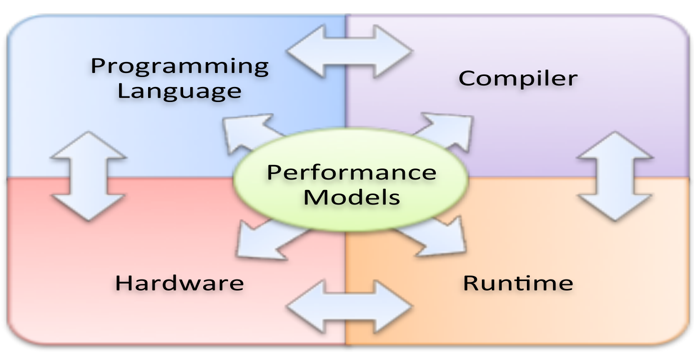

Figure 1: Performance models enable real-time decision-making between system levels.

Within BSM we set the stage for establishing an integrated approach to modeling of performance and power, which will take us "beyond the standard" modeling and simulation approaches that represent today's state-of-the-art. Co-design *is* modeling, in its broadest definition—so a frontal attack on the way in which we do modeling will enable practical co-design solutions among architectures, applications/algorithms, programming models and runtimes. Research in BSM is structured along four areas of investigation to which we believe will influence Exascale systems and applications.

Modeling of Performance and Power – we are establishing the modeling of performance and power *in concert* as the ultimate goal, beyond the current state-of-the-art in which (except for limited instances) performance only is the modeling target.

Modeling at different scales – From definition of metrics, to application models, to detailed architectural descriptions, models capture the performance and power characteristics at the various boundaries of the hardware/software stack with the desired accuracy and predictive capability needed to make the decision at hand. The integrated nature of this approach goes beyond the current state-of-the-art.

Dynamic Modeling of Performance, Power and Data Movement – This is at the heart of modeling performance and power together, and aims at going beyond the current practice that regardless of the methodology employed is static (off-line) in nature. We envision models operating in the entire spectrum from static to dynamic, the latter models serving as the engine of intelligent runtime systems, among others.

Techniques for Model Generation – This research direction aims at simplifying static model generation, including through compiler based approaches, and at coming up with methodologies for generating models dynamically based on monitoring of systems and application behavior at runtime.

Modeling methodology, in the context of this project, is the generation of quantitative, accurate mappings of application characteristics onto system resources, whose goal is to explain and predict the power and time consumed by applications running on High Performance Computing (HPC) architectures. Central to this work is the concept that future applications, HPC architectures, and all the system software layers in-between ought to be designed by taking into account the characteristics of each other, and optimized for performance under power/energy constraints.

Our initial focus in BSM has been to: develop accurate power models to predict system-level power draw; to consider computational phases within an application; to asses the dependence of an application performance on the availability of cache storage and bandwidth; a prototype of a highly-accurate and light-weight monitor that will serve as a basis for dynamic modeling; and a preliminary tool to generate performance models from application code.

This project is a collaboration between researchers at the Pacific Northwest National Laboratory, Lawrence Livermore National Laboratory, and San Diego Supercomputing Center that was initiated in late 2012.