

# Algorithmic Adaptations to Extreme Scale Computing

David E. Keyes

Extreme Computing Research Center

King Abdullah University of Science and Technology

Algorithmic adaptations to use next-generation computers close to their potential are underway. Instead of squeezing out flops – the traditional goal of algorithmic optimality, which once served as a reasonable proxy for all associated costs – algorithms must now squeeze synchronizations, memory, and data transfers, while extra flops on locally cached data represent only small costs in time and energy. After decades of programming model stability with bulk synchronous processing, new programming models and new algorithmic capabilities (to make forays into, e.g., data assimilation, inverse problems, and uncertainty quantification) must be co-designed with the hardware. We briefly recap the architectural constraints and application opportunities. We then concentrate on two types of tasks each of occupies a large portion of all scientific computing cycles: large dense symmetric/Hermitian linear systems (covariances, Hamiltonians, Hessians, Schur complements) and large sparse Poisson/Helmholtz systems (solids, fluids, electromagnetism, radiation diffusion, gravitation). We examine progress in porting solvers for these tasks to the hybrid distributed-shared programming environment, including the GPU and the MIC architectures that make up the cores of the top scientific computers “on the floor” and “on the books.”