

# Architecturally Tailored Reconfigurable Platforms for Scientific Computing

by Ali Akoglu

Department of Electrical and Computer Engineering

University of Arizona

akoglu@ece.arizona.edu

High-throughput technologies have led to an exponential growth in the amount of data generated over the past several years. This data explosion is forcing scientists to search for innovative computational designs to meet the growing demands. The complexity, variety of techniques and tools, and the high computation, storage and I/O bandwidth demands associated with the applications of biological systems pose several challenges, particularly from the points of scalability, resource utilization and real-time implementation. Current technologies fall short of providing low cost and flexible solutions in responding to processing demand. These drawbacks have led us into exploration of the reconfigurable hardware systems in order to reduce the timeframe with which computational study of biological systems can be achieved. Configurable computers are capable of accelerating suitable applications by several orders of magnitude when compared to traditional processor based architectures. This significant speed advantage is due to the highly parallel nature of FPGA (field programmable gate array) hardware. Fortunately, many problems in scientific computing are inherently parallel, and benefit from concurrent computing models. New hybrid chips, those with both a general purpose CPU (central processing unit) and FPGA are now being fabricated at commodity prices. Exploiting the potential of reconfigurable systems for high performance scientific computing requires the ability to abstract the boundary between the CPU and FPGA, and allow the user (application programmers) to guide the mapping of high level function operations across all components with a minimal effort. To the best of our knowledge there isn't a well known methodology for scientific applications that performs a design space exploration of architectures that fall in between coarse grain (cluster based systems) and fine grain (instruction level) parallelism on CPU-FPGA systems in an automated manner. This presentation will lay out the need for such methodology and then show that reconfigurable platforms tailored to the computation characteristics of the target applications have the great potential of delivering affordable, flexible systems with higher number crunching capability compared to general purpose processors.