

Novo-G: At the Forefront of Scalable Reconfigurable Supercomputing

Alan George, Herman Lam, and Greg Stitt
 NSF Center for High-Performance Reconfigurable Computing (CHREC)¹
 ECE Department, University of Florida, Gainesville, FL 32611-6200
 {george, hlam, gstitt}@chrec.org

In this article we present Novo-G, an innovative new supercomputer in the NSF CHREC Center at Florida whose architecture can adapt to match the unique needs (e.g. parallelism, precision) of each application, and thereby attain more performance with less energy than conventional machines for key problems in health and life sciences, signal and image processing, and more.

1. Introduction

Throughout the long history of computing, and within the many forms of computers existing today spanning hand-held smartphones to mammoth supercomputers, the one common denominator is fixed-logic processors. In this theme that is conventional computing, each application must be adapted to match the fixed structures, parallelism, functionality, and precision of the target processor (e.g. CPU, DSP, or GPU) as dictated by the device vendor. This “one size fits all” approach, while advantageous for greater uniformity, can lead to dramatic inefficiencies in speed, area, and energy when the application does not conform to the ideal case for that device. By contrast, a relatively new paradigm of computing known as reconfigurable computing (RC) takes the opposite approach, wherein the architecture adapts to match the unique needs of each application, consequently approaching the speed and energy advantages of application-specific ICs (ASICs) with the versatility of CPUs. Many RC systems have recently emerged in the research community and marketplace, addressing an increasingly broad range of applications, from sensor processing for space science to proteomics for cancer diagnosis. Most of these systems are of relatively modest scale, featuring one or several reconfigurable processors such as field-programmable gate arrays (FPGAs). However, at the extreme scale of RC in the world today is the Novo-G machine in the NSF CHREC Center at the University of Florida. Initial studies on Novo-G demonstrate that for some important applications a scalable system with 192 reconfigurable processors can rival the speed of the world’s largest supercomputers at a tiny fraction of their cost, size, power, and cooling. This article provides an overview of this novel system and its initial applications and performance breakthroughs.

2. Background

Demands for innovation in computing are growing rapidly. Technological advances are transforming many data-starved science domains into data-rich ones. For example, in genomics research in the health and life sciences, contemporary DNA sequencing instruments are capable of determining 150-200 billion nucleotide bases per run, resulting in output files routinely in excess of 1 TB per instrument run. In the near future, DNA sequence output from a single instrument run will easily exceed the size of the human genome by more than 100-fold. Thus, it is increasingly clear that the discordant trajectories growing between data production and the capacity for timely analysis are threatening to impede new scientific discoveries and progress in many scientific domains, not because we cannot generate the data, but because we cannot analyze it. To address these growing demands with a sustainable computing infrastructure in terms of power, cooling, size, weight, and cost, adaptive systems that can be dynamically tailored to the unique needs of each application are coming to the forefront. At the heart of these systems are reconfigurable-logic devices, processors such as FPGAs that under software control can adapt their hardware structures to reflect the unique operations, precision, and parallelism associated with compute-intensive, data-driven applications in fields such as health and life sciences, signal and image processing, and cryptology.

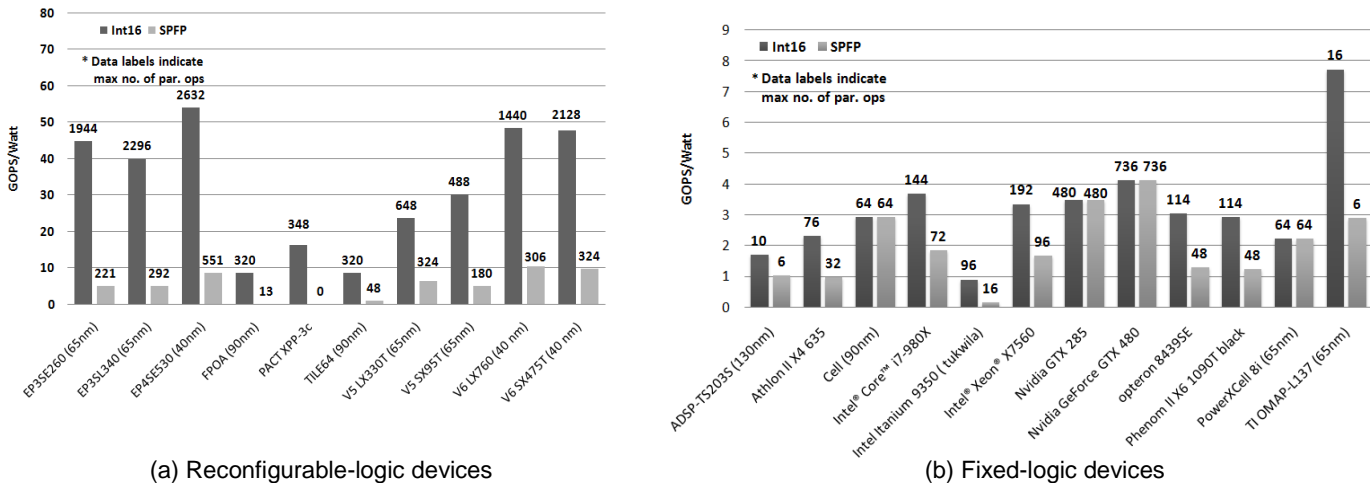


Figure 1. Computational density (in GOPS) per Watt of modern fixed- and reconfigurable-logic devices [2]

¹ This work was supported in part by the I/UCRC Program of the National Science Foundation under Grant No. EEC-0642422. The authors gratefully acknowledge contributions of numerous researchers at the University of Florida and in the Novo-G Forum, as well as providers at Altera and GiDEL.

The benefit of RC with modern FPGAs for such applications comes from their reconfigurable structure. Unlike fixed-logic processors where applications must conform to their fixed structure (for better or for worse), with RC the architecture conforms to the unique needs of each application. This adaptive nature enables FPGA devices to exploit higher degrees of parallelism while running at lower clock rates and thereby in many cases achieve better execution speed while consuming less energy. For example, from a comparative suite of device metrics [1-2], Figure 1 illustrates performance (in terms of computational density) per Watt of some of the latest reconfigurable- and fixed-logic processing devices for 16-bit integer (*Int16*) or single-precision floating-point (*SPFP*), assuming an equal number of add and multiply operations. In these charts, the peak number of sustainable parallel operations of each type is cited atop each bar. In general, FPGAs achieve more speed per unit of power as compared to CPUs, DSPs, and GPUs. For instance, the leading FPGA for Int16 operations in this study (Altera Stratix-IV EP4SE530) can support over 50 billion operations per second (GOPS) per Watt while the leading fixed-logic processor (TI OMAP-L137 DSP) attains less than 8 GOPS/Watt. With SPFP, the gap is narrower, but FPGAs continue to enable more GOPS/Watt. Similarly, not shown in the figure, the gap widens increasingly for simpler (e.g. byte, bit) operations. With RC, the simpler the task, the less chip area needed for each, and thus the more that can fit and operate concurrently in hardware.


However, this promising approach while motivating has to date mostly been limited to small systems, studies, and datasets, and to move beyond these limits some key challenges must be overcome. Chief among these challenges is parallelization, evaluation, and optimization of critical applications in data-intensive fields in a manner that is transparent, flexible, portable, and performed at a much larger scale, commensurate with the massive needs of emerging, real-world datasets. When successful, the impact can be a dramatic speedup in execution time concomitant with savings in energy and cooling. As described later in this article, our initial studies show critical applications executing at scale on Novo-G, achieving speeds rivaling the largest conventional supercomputers in existence yet at a small fraction of their size, energy, and cost. While processing speed and energy efficiency are important, the principal impact of a reconfigurable supercomputer like Novo-G is the freedom that this innovative approach to computing can give to scientists to conduct more types of analysis, examine larger datasets, ask more questions, and find better answers.

3. Novo-G Reconfigurable Supercomputer

The Novo-G machine is an experimental research testbed operating since July 2009 in the NSF CHREC Center at the University of Florida and supporting a variety of research projects on the challenges of scalable RC. The principal emphases of Novo-G are *performance* (device, subsystem, system), *productivity* (concepts, languages, tools), and *impact* (scalable applications). Figure 2 shows the Novo-G machine and one of its quad-FPGA boards. The current Novo-G configuration consists of 24 compute nodes, each a standard 4U Linux server with an Intel quad-core Xeon (E5520) processor, memory, disk, etc., housed in three racks. A single 1U server with twin quad-core Xeons functions as head node. Compute nodes communicate and synchronize via Gigabit Ethernet and a non-blocking fabric of 20 Gb/s InfiniBand. Each of the compute nodes houses two PROCStar-III boards from GiDEL in its PCIe slots. The novel computing power of Novo-G is derived from these boards, each containing four Stratix-III E260 FPGAs from Altera, resulting in a system of 48 boards and 192 FPGAs.² Concomitantly, when fully loaded, power consumption of the entire Novo-G system peaks at ~8K Watts.



(a) Novo-G machine



(b) GiDEL PROCStar-III board

Figure 2. Novo-G supercomputer (a) and one of its quad-FPGA reconfigurable processing boards (b)

While this set of FPGAs is theoretically capable of providing a massive degree of computing power for the system, memory capacity, throughput, and latency often limit performance if unbalanced. As shown in Figure 2(b), attached to each of the FPGAs is 4.25 GB of dedicated memory in three banks. Data transfer between adjacent FPGAs can be made directly through a wide, bidirectional bus at rates up to 25.6 Gb/s and latencies of a single clock cycle up to 300 MHz, and transfer between FPGAs across two boards in the same server is also supported via a high-speed cable. By supplying each FPGA with large, dedicated memory banks, as well as high bandwidth and low latency for inter-FPGA data transfer, the system strongly supports RC-centric applications. Processing engines on the FPGAs can execute with minimal involvement by the host CPU cores, enabling maximum utilization of the FPGAs.

² An impending upgrade will soon add 72 Stratix-IV E530 FPGAs to Novo-G, each with twice the reconfigurable logic of a Stratix-III E260 (yet roughly the same power consumption), thereby expanding total reconfigurable logic in Novo-G by nearly 80%.

Alongside the architecture, equally important are the design tools available and upcoming for Novo-G. The very nature of RC empowers the application developer with far more capability, control, and influence over the architecture. Instead of stipulating all architecture decisions to the device vendors, as with CPUs and GPUs, in RC the application developer specifies a custom architecture configuration, such as quantity and types of operations, numerical precision, and breadth and depth of parallelism. Consequently, RC is a more challenging environment for the development of applications, and thus productivity concepts and tools are vital. Novo-G features a broad and growing range of academic and commercial tools, in areas that include: (1) strategic design and performance prediction tools for parallel algorithm and mapping studies; (2) MPI, UPC, and SHMEM for system-level programming with C; (3) VHDL, Verilog, and an expanding list of high-level synthesis tools for FPGA-level programming; (4) an assortment of core libraries; (5) middleware and API for design abstraction, platform virtualization, and portability of apps and tools; and (6) verification and performance-optimization tools.

To help expand the applications and tools available on Novo-G, and establish and showcase advantages of RC at scale, the Novo-G Forum was formed in 2010. This forum is an international group of academic researchers and technology providers working collaboratively with a common goal of realizing the promise of reconfigurable supercomputing by demonstrating unprecedented levels of performance, productivity, and sustainability. Faculty and students in each academic research team are committed to contribute innovative applications and/or tools research on the Novo-G machine based upon their unique expertise and interests. Currently committed academic participants in the Novo-G Forum include Boston University, Clemson University, University of Florida, George Washington University, University of Glasgow (UK), Imperial College (UK), Northeastern University, Federal University of Pernambuco (Brazil), University of South Carolina, University of Tennessee, and Washington University at St. Louis. Each academic team is equipped with one or more Novo-G boards for local experiments, and supported by remote access to the large Novo-G machine at Florida for scalability studies.

4. Initial Applications Studies

Of the three principal emphases cited for Novo-G, *impact* is undoubtedly the most important. What good is a new and innovative high-performance system if the resulting applications have little impact in science and society? In this section, we overview initial performance breakthroughs on a set of bioinformatics applications for genomics, developed in collaboration with the Interdisciplinary Center for Biotechnology Research (ICBR) at Florida. Results of such breakthroughs can potentially revolutionize the processing of massive genomics datasets, which in turn may enable revolutionary discoveries for a broad range of challenges in the health, life, and agricultural sciences.

Although more than a dozen challenging application designs are underway on Novo-G, from several domains of science, here we focus upon our first case studies. These studies include two popular genomics applications for optimal sequence alignment based upon wavefront algorithms, Needleman-Wunsch (NW) and Smith-Waterman (SW) without traceback³, and a metagenomics application, Needle-Distance (ND), which is an augmentation of NW with distance calculations. Each of these applications features massive data parallelism with minimal communication and synchronization between FPGAs, and a highly optimized systolic array of processing elements (PEs) within each FPGA and optionally spanning multiple FPGAs. Using a novel method for in-stream control [3], we optimized each of the three designs to fit up to 850 PEs per FPGA for NW, 650 for SW, and 450 for ND, all operating at 125 MHz.

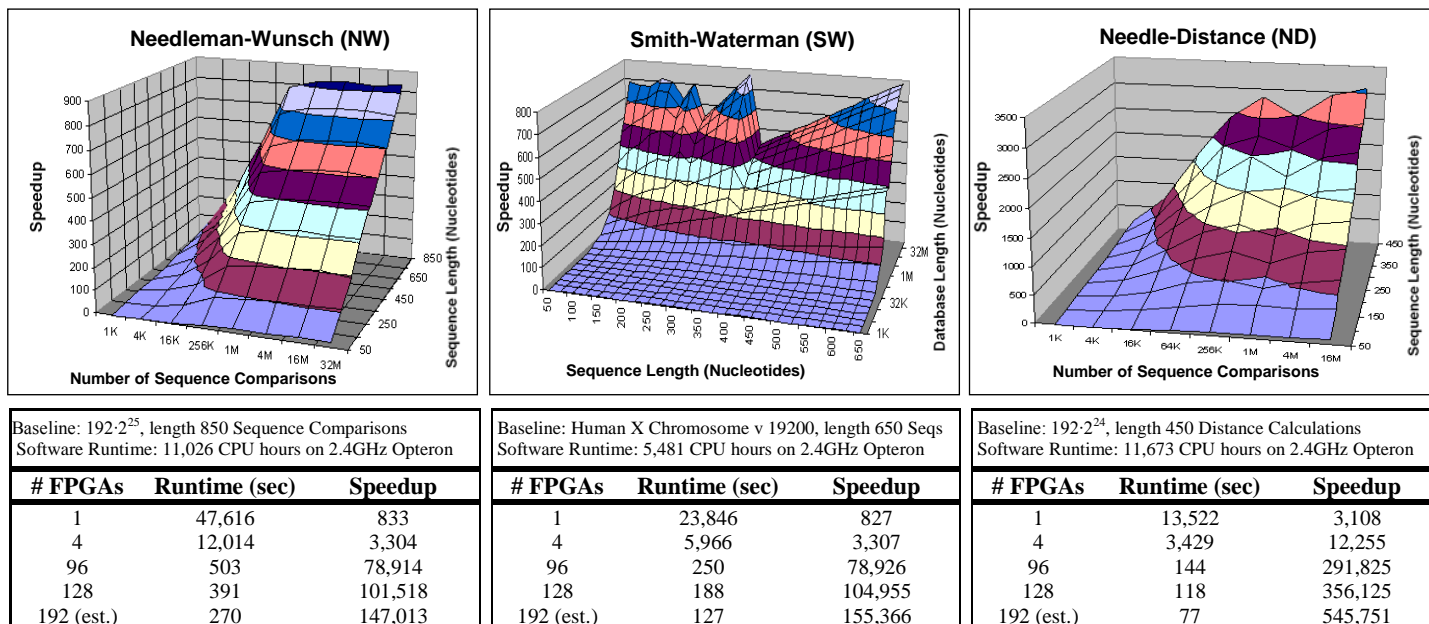


Figure 3: Results on Novo-G for NW (left), SW (center), and ND (right). Each chart illustrates performance of a single FPGA under varying input conditions. Each table shows performance with varying number of FPGAs under optimal input conditions. [3]

For each application a contour plot in Figure 3 illustrates relative design performance on one FPGA under varying input conditions. The corresponding tables show how the three designs scale when executed on multiple FPGAs in Novo-G. In all cases, speedup is defined in

³ An extended version of SW with the traceback option (SW+TB) is nearing completion, by augmenting our SW hardware design to collect and feed data for traceback to the hosts, such that FPGAs perform SW while CPU cores perform TB concurrently. Initial results indicate that, after adding TB, execution times increase less on Novo-G than on the C/Opteron baseline, and thus Novo-G speedups with SW+TB exceed those of SW.

terms of an optimized C-code software baseline running on a 2.4 GHz Opteron core in our lab. More details on these algorithms, architectures, experiments, and results are provided with [3]. All data except the final row of the tables came directly from measurements on Novo-G and include the full execution time, including data transfers to/from the FPGAs, and not merely computation time.

Speedup with one FPGA on each of these three applications peaked at approximately 830 for NW and SW and more than 3100 for ND. When ramping up from a single FPGA to a quad-FPGA board, speedups were measured and observed to grow almost linearly to about 3300 for NW and SW and more than 12K for ND. At the largest scale of our testbed experiments, with 32 boards (i.e. 128 FPGAs), speedups for NW and SW exceeded 100K and ND exceeded 356K. By extrapolating these trends (since not all of our 48 boards in Novo-G were operational at the time), we estimate speedups on all 192 FPGAs of Novo-G of about 150K for NW and SW and almost 550K for ND. Putting these numbers in context, the latter implies that a conventional supercomputer would require more than a half-million Opteron cores operating optimally to match the performance of Novo-G on the ND application. By contrast, none of the world's largest supercomputing machines (e.g. as cited at the top of the rankings at www.top500.org) has this many cores, and thus none would be able to achieve such performance on this application despite being orders of magnitude larger in cost, size, weight, power, and cooling. Although Novo-G will not provide all applications with the same speedups as these examples, they do highlight the potential advantages of RC, especially in solving problems where conventional, fixed-logic computing falls far short of achieving optimal performance.

5. Conclusions

For a growing list of important applications from a broad range of science domains, underlying computations and data-driven demands are proving to be underserved by conventional “one size fits all” processing devices. By changing the mindset of computing, from processor-centric to application-centric, reconfigurable computing can provide solutions for domain scientists in a small fraction of the time and/or cost of traditional servers or supercomputers. This article has provided an overview of the Novo-G machine, applications, research forum, and preliminary results that are helping to pave the way for scalable reconfigurable computing.

References

1. J. Williams, A. George, J. Richardson, K. Gosrani, C. Massie, and H. Lam, “Characterization of Fixed and Reconfigurable Multi-Core Devices for Application Acceleration,” *ACM Transactions on Reconfigurable Technology and Systems (TRETTS)*, Vol. 3, No. 4, Jan. 2011, to appear.
2. J. Richardson, S. Fingulin, D. Raghunathan, C. Massie, A. George, and H. Lam, “Comparative Analysis of HPC and Accelerator Devices: Computation, Memory, I/O, and Power,” *Proc. of High-Performance Reconfigurable Computing Technology and Applications Workshop (HPRCTA) at the ACM/IEEE Supercomputing Conference (SC10)*, New Orleans, LA, Nov. 14, 2010, to appear.
3. C. Pascoe, A. Lawande, H. Lam, A. George, Y. Sun, W. Farmerie, and M. Herbordt, “Reconfigurable Supercomputing with Scalable Systolic Arrays and In-Stream Control for Wavefront Genomics Processing,” *Proc. of Symposium on Application Accelerators in High-Performance Computing (SAAHPC)*, Knoxville, TN, July 13-15, 2010.