
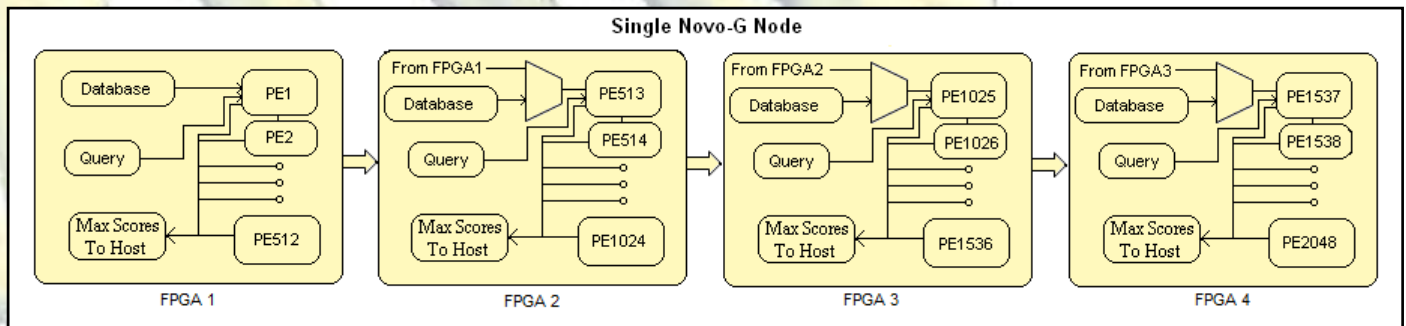


# Application Acceleration on CHREC Novo-G: Breaking New Ground for Bioinformatics




**PROBLEM:** The act of aligning DNA, RNA, or protein sequences in order to identify regions of similarity is one of the most common and important tasks in computational biology. The search for similarities is performed under the assumption that they imply functional, structural, or evolutionary relationships between sequences and their sources. Contemporary software implementations of sequence alignment (whether global, local, or anything in-between) are based upon a computation-intensive dynamic programming algorithm that breaks down the process of sequence alignment into a set of recursive computations, calculating alignment scores for all possible subsequences, and then choosing the sequence alignment based upon the set of sub-alignment scores. Alignments can be computed optimally through a series of rigorous calculations or else approximated with a set of faster heuristics. For two sequences of length  $A$  and  $B$ , optimum alignment requires the calculation of  $A \cdot B$  scores, with serial implementations operating in  $O(A \cdot B)$  execution time and  $O(\min\{A, B\})$  space complexity. Rigorous, optimal alignments in bioinformatics applications are the ideal, most accurate, but in practice are often replaced with heuristic estimates due to prohibitively large execution times of all-software implementations. However, on-going research by the NSF Center for High-Performance Reconfigurable Computing (CHREC, chrec.org) at the University of Florida is now making it possible for scientists to have the best of both worlds, optimal results computed with high efficiency, by exploiting reconfigurable computing and the Novo-G machine, the most powerful reconfigurable supercomputer ever fielded in academia.



**SOLUTION:** Adaptation of the dynamic programming algorithm to produce optimum, local sequence alignment is commonly called the Smith-Waterman (S-W) algorithm. CHREC students have succeeded in accelerating S-W targeted for the Novo-G machine at CHREC by designing novel hardware cores to exploit massive parallelism that allows up to  $A$  of the  $A \cdot B$  scores of a single alignment to be calculated simultaneously in a single 125MHz clock cycle, reducing time complexity of this algorithm to  $O(A+B)$ . Novo-G is a cluster of 24 Linux servers (i.e. nodes), each housing a quad-FPGA ProcStar-III accelerator board from GiDEL, altogether featuring 96 high-end, Stratix-III E260 FPGAs from Altera. Our basic design for S-W consists of a long systolic array of 512 processing elements (PEs) per FPGA with input FIFOs for streaming of database and query sequences, and various score reporting registers for feedback to the host CPU in each node. Each PE simultaneously performs all of the calculations required to produce a single score in a single clock cycle and makes extensive use of dedicated registers to store intermediate data. Because of the high bandwidth and low latency of the interconnect between FPGAs on each accelerator board, this systolic array can be extended across all four FPGAs with only a single clock cycle of delay from one FPGA to the next. Thus, with an array of 2048 PEs per node and 24 total nodes in Novo-G, an unprecedented 49152 scores can be simultaneously computed per clock cycle by exploiting all of Novo-G's FPGA hardware. In order to maximize hardware utilization, each FPGA can be configured to perform alignments on multiple query sequences simultaneously; with a granularity of 32 PEs, if the first query sequence does not require all 2048 PEs per node to compute the alignment, additional queries can be added to exploit unallocated PEs.

**RESULTS:** The chart and table provided depict execution times for a 34MB chromosome sequence aligned with 16K 128-character sequences randomly extracted from a different chromosome. Novo-G performance was compared with an optimized software implementation executed on a single 64-bit, 2.4GHz AMD Opteron core. After validation tests to confirm matching alignment scores, a speedup of **2665** was measured for this dataset (though different datasets produce similar results) on a single Novo-G node. Thus, on a conventional HPC machine, 2665 of these Opteron cores working in parallel with no overhead would be required to perform the same amount of work in the same period of time as a single Novo-G node. The same dataset and program were partitioned across 16 Novo-G nodes, using MPI for coordination, where a speedup of **40849** was observed. If application behavior is extrapolated to all 24 Novo-G nodes, a speedup of about **60053** can be expected. Thus, Novo-G can achieve in about 12 seconds what would require a processor core nearly 9 days! When



compared to two of the newest machines on the national TeraGrid, Kraken at UTK/ORNL (66,000 2.3GHz Opteron cores) or Ranger at UTA/TACC (62,976 2.3GHz Opteron cores), and allowing for ~4% slower clocks in these machines (vs. our baseline), it is noted that for this application the Novo-G reconfigurable supercomputer achieves comparable or better performance while being ~100 times less expensive in cost, power, cooling, etc.

Execution Time of Serial Baseline on Single 2.4 GHz Opteron Core = 743,460 Seconds (~8.6 Days)	Number of Novo-G Nodes in Execution					
	1	4	8	12	16	24(E)
Execution Time (Sec) of Novo-G	279	70.4	35.6	24.2	18.2	12.38
Novo-G Speedup vs. Single Core	2665	10561	20884	30721	40849	60053