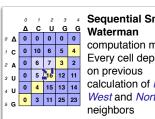# GATOR and SWAMP: GPU Computing for Sequence Alignment

SWAMP: Smith Waterman using Associative Massive Parallelism

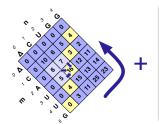**Shannon Steinfadt & Johnnie W. Baker**
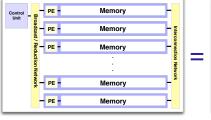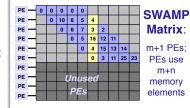Dept. of Computer Science, Kent State University, Kent, Ohio 44242

Sequential Smith-Waterman computation matrix. Every cell depends on previous calculation of *North*, *West* and *Northwest* neighbors

ASC Model with a single Control Unit (Instruction Stream)

SWAMP Matrix: m+1 PEs; PEs use m+n memory elements

"Shifted" SWAMP matrix

Sequence Alignment

similar characters → similar structure similar function

Ancestral Relationships
Gene Functionality
Aid in Drug Discovery

SLOW
Rigorous Sequential Alignment
O(m*n) time
Smith-Waterman

FAST
Heuristic Sequential Alignment Quicker
Can Miss Alignments
Lose Sensitivity
BLAST, FASTA

SWAMP
Fast & Rigorous Sequential Alignment
O(m+n) time using *m* simple processors (PEs)
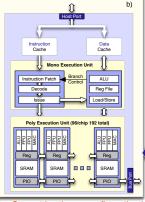
SWAMP+

Sequence 1
Sequence 2

SWAMP finds the best local-alignment. SWAMP+ will return at most the top *k* non-intersecting, non-overlapping segments (here *k* = 3).

ASC on Metal: Clear*Speed* X620

- **"SIMD-like"** model on a parallel accelerator
- 1 board, 2 chips, each with 96 processing elements for a total of **192 Processing Elements** *or* Poly Execution Units **(PEs)**
- Each PE has 6KB **local** SRAM **memory** and **dual 64-bit FPU's**
- 25 W average power consumption
- 66+ GFLOP sustained performance on DGEMM
- Cn Language, SDK, and Visual Profiler

a) Clear*Speed* Advance X620 SIMD board with 2 chips and 192 PEs
b) Single Chip Execution Units

ASC - a SIMD with Additional Features:

- Locates data (**searches**) **based on content**, not memory location
- **Processing elements** (PEs) have **own local memory**
- Existing algorithms, programming language, compiler, and emulator
- **Data** configuration is **tabular**
- **Fast communication**; Tightly coupled, local raw data movement (NOT over LAN/WAN)
- **Constant time search/respond operations** and constant time global reduction operators:
  - **Maximum / Minimum**
  - **"Any Responders?"**
  - **"Pick One"**

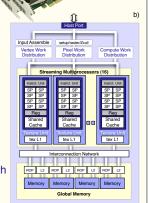**NEED ASSOCIATIVE OPERATORS FOR EACH "ASC ON METAL" ARCHITECTURE**

ASC on Metal: NVIDIA C870 Tesla

GATOR: Genetic Alignment on sTream processORs

- **"SPMD"** model on a GPGPU (General Purpose Graphical Processing Unit) platform
- 1 GPU/board, 16 streaming multiprocessors (MP), each with 8 streaming processor cores (SPs) for **128** 32-bit FP **SPs**
- Each MP has 16KB of **shared** memory
- 120 W typical power consumption
- 518 Peak GFLOPs on Tesla series
- CUDA - Compute Unified Device Architecture - C-like language and virtual platform
- Multiple levels of parallelism:
  - Up to 512 threads per block
  - Grids of thread blocks

a) Dual NVIDIA Tesla C870 GPGPU's each with 128 stream processors
b) Tesla/GPGPU organization

Note: Current hardware configuration is that the three acceleration cards are housed in a single machine for comparative timings

Grace Hopper 2008