

# Gigabyte-Scale Alignment Acceleration of Biological Sequences via Ethernet Streaming


Theepan Moorthy and Sathish Gopalakrishnan  
The University of British Columbia




a place of mind

## The DIALIGN Alignment Algorithm

- The goal is to find the best possible match between a Query Sequence and a much longer Reference Sequence.
- The Query Sequence can be split into multiple segments if this results in a better match being obtained.
- For hardware acceleration purposes, this problem can be viewed as a simple string-matching problem, with a unique scoring algorithm for each character match being detected.



## Pipeline Processing Elements




- Within the Systolic Array, an entire diagonal-line of match values are calculated in parallel.
- Input to this FPGA based daisy chain of Processing Elements is achieved via an Ethernet connection to a PC, which has the Query Sequence that is to be aligned against a desired Reference Sequence.
- The Ethernet channel is used to first load all of the PEs locally with each Query Sequence character, then stream an entire Reference Sequence once accelerator operation starts.

## Pipeline Limitations

- Given that Query Sequence lengths can be 1 megabyte or more, the number of Processing Elements that can be daisy chained together to make up the Query Sequence is limited by FPGA logic.
- A solution to this problem is to time multiplex the Processing Elements across sequential partition passes over the entire Query Sequence

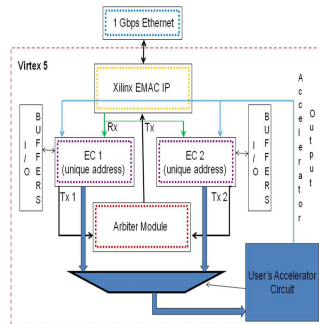
## External Data Requirement

- For each partition pass of the Query Sequence, the Reference Sequence must be completely re-streamed each time.
- This is accomplished via a Simple Interface for Reconfigurable Computing (SIRC) based Ethernet Channel to a host PC.

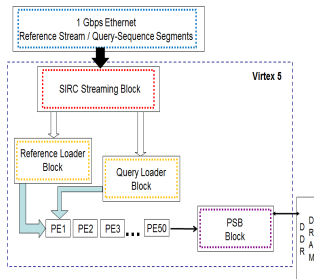


## SIRC Modified for Ethernet Streaming

- A single physical Ethernet link is multiplexed with two logical controller cores (i.e. SIRC APIs) to achieve streaming of the Reference Sequence from the PC to the FPGA Accelerator



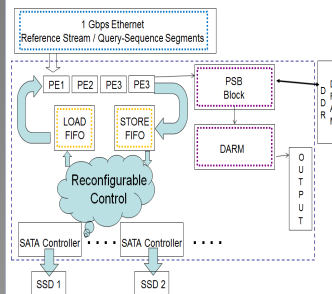
## Streaming Reference/Loading Query Sequences



- The Partition State Bank (PSB) Block stores the key alignment information produced by each Processing Element once an entire partition has been processed.
- Once all partitions have been processed, this information is then retrieved to stitch together the best alignment calculated by each partition's worth of processing.

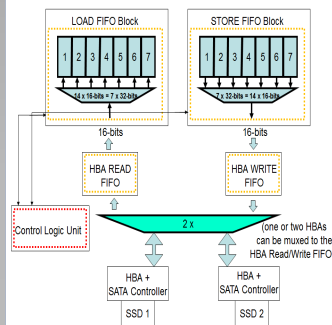
## Stitching Partitions Together

- The terminal Processing Element in the daisy chain produces 28-bytes of intermediate data that the Alignment algorithm needs to compute the best matching scores.
- These 28-bytes are outputted per clock cycle for the entire length of the Reference Sequence that is being streamed.
- Thus even for a single gigabyte length Reference Sequence, 28 Gigabytes of data will be generated. This large amount of data must be stored and retrieved with as much minimal latency as possible.
- The data rate of the intermediate data will be equal to 28-bytes per clock cycle x maximum circuit frequency (67.689 MHz) = 1896 MB/s




- Even with the integration of a single SATA controller and large storage Solid State Drive, a deep enough FIFO can be built to stitch together a small physical chain of PEs into a much larger logical daisy chain.
- However, SATA write bandwidth will be the bottleneck to running the accelerator at its maximum synthesized clock rate.
- A dual-SATA controller system can also be built to concurrently overlap data reads with that of the slower data write speeds.

## Exact Implementation of SATA data Flow



- The SIRC based Groundhog SATA controller core that we employ achieves an average write rate of 66.324 MB/s, allowing for all 28 Gigabytes of required intermediate data between partitions to be stored within roughly 7.2 minutes.
- Since the 66.324 MB/s write rate is clearly much less than the 1896 MB/s data rate of the accelerator, the accelerator itself is not functioning at full potential-throughput levels

## Single SATA-SSD Platform



## Experimental Results

- 50 Processing Elements were synthesizable within our XC5VLX10T Virtex 5 FPGA, with 88% logic utilization, including all SIRC Ethernet and SATA controllers.
- Maximum clock frequency of Alignment Acceleration logic was 67.689 MHz
- A 200-character long Query Sequence, requiring a total of 4 Partition Executions, was aligned a 1 Gigabyte Reference Sequence
- Total wall-clock execution time on average was 28.61 minutes, or just under a half an hour
- Previous single PC, non-multithreaded, execution of the DIALIGN alignment algorithm requires roughly 18.36 hours. Modern multicore parallel software implementations would improve that time.

## Conclusions

- Relative to single PC execution a speedup of 38x was demonstrated.
- This is a low-cost (open-source IP) solution to work around limited FPGA logic capacity to process large Query and Reference Sequences
- The platform easily accommodates next generation SATA line rates

## Contact Information

Theepan Moorthy  
PhD Candidate  
Department of Electrical and Computer Engineering  
The University of British Columbia  
tmoorthy@ece.ubc.ca