

A GPU-based Correlator X-engine Implemented on the CHIME Pathfinder

Nolan Denman,^{*†} Mandana Amiri,[§] Kevin Bandura,[‡] Jean-François Cliche,[‡] Liam Connor,^{*†¶}
 Matt Dobbs,^{‡||} Mateus Fandino,[§] Mark Halpern,[§] Adam Hincks,[§] Gary Hinshaw,[§] Carolin Höfer,[§]
 Peter Klages,^{*} Kiyoshi Masui,^{§||} Juan Mena Parra,[‡] Laura Newburgh,^{*} Andre Recnik,^{*} J. Richard Shaw,[¶]
 Kris Sigurdson,[§] Kendrick Smith,^{**} and Keith Vanderlinde^{*†}

^{*} Dunlap Institute, University of Toronto

[†] Department of Astronomy and Astrophysics, University of Toronto

[§] Department of Physics and Astronomy, University of British Columbia

[‡] Department of Physics, McGill University

[¶] Canadian Institute for Theoretical Astrophysics

^{||} Canadian Institute for Advanced Research, CIFAR Program in Cosmology and Gravity

^{**} Perimeter Institute for Theoretical Physics

Contact E-Mail: denman@astro.utoronto.ca

Abstract—We present the design and implementation of a custom GPU-based compute cluster that provides the correlation X-engine of the CHIME Pathfinder radio telescope. It is among the largest such systems in operation, correlating 32,896 baselines (256 inputs) over 400 MHz of radio bandwidth. Making heavy use of consumer-grade parts and a custom software stack, the system was developed at a small fraction of the cost of comparable installations. Unlike existing GPU backends, this system is built around OpenCL kernels running on consumer-level AMD GPUs, taking advantage of low-cost hardware and leveraging packed integer operations to double algorithmic efficiency. The system achieves the required 105 TOPS in a 10 kW power envelope, making it one of the most power-efficient X-engines in use today.

I. INTRODUCTION

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) is an interferometric radio telescope, presently under construction at the Dominion Radio Astrophysical Observatory (DRAO) in British Columbia, Canada, which will map the northern sky over a radio band from 400 to 800 MHz. With over 2000 inputs and a 400 MHz bandwidth, the correlation task (measured as the bandwidth-baselines product) on CHIME will be an order-of-magnitude larger than on any currently existing telescope array. The correlator follows an FX split design, with a first stage Field Programmable Gate Array (FPGA)-based F-engine which digitizes, Fourier transforms (channelizes), and bundles the data into independent frequency bands, followed by a second-stage Graphics Processing Unit (GPU)-based X-engine which produces a spatial correlation matrix consisting of the integrated pairwise products of all the inputs at each frequency.

The CHIME Pathfinder instrument [1] features 128 dual-polarization feeds and a reduced-scale prototype of the full CHIME correlator. This paper describes the X-engine of the Pathfinder’s 256-input hybrid FPGA/GPU FX correlator, among the largest such systems in operation. Through extensive use of off-the-shelf consumer-grade hardware and heavy optimization of custom data handling and processing software,

the system achieves high performance at a small fraction of the hardware cost of comparable installations.

This paper focuses on the system architecture and implementation, while two companion papers describe the custom software stacks, one focusing on an innovative OpenCL-based X-engine GPU kernel [2], and one on the handling of the vast data volume flowing through the system [3].

The paper is structured as follows: design considerations and constraints are discussed in §II; the hardware components of the system are described in §III, and the software in §IV; the scaling of the X-engine to the full-size CHIME telescope is described in §V, and a summary and conclusion follow in §VI.

II. DESIGN CONSIDERATIONS

While most components in CHIME scale linearly with number of inputs N , the computational cost of pairwise correlation scales as N^2 , making efficiency in the X-engine a primary concern. There are correlation techniques which rely on the redundancy of CHIME feed separations to scale as $N \log N$, but the real-time calibrations these require for precision observations remain largely unproven in an astrophysical context. Design decisions were guided by the need to produce an inexpensive system capable of scaling to support full CHIME, and which would support rapid development and deployment of new data processing algorithms. These requirements of computational power and ease of development drove the decision to build the X-engine around GPUs rather than Application-Specific Integrated Circuits (ASICs) or FPGAs.

The computational cost η of pairwise element correlation for N elements across a bandwidth of $\Delta\nu$ is

$$\eta = \Delta\nu \cdot N \cdot (N + 1)/2 \quad (1)$$

measured in complex multiply-accumulate (cMAC) operations per second; for the CHIME Pathfinder, $\eta = 13 \text{ TcMAC/s}$. For large N this dominates the cost of any other processing proposed for the X-engine. Top-end GPUs in 2014 provided of

order 4 TFLOPS of processing power per chip, equivalent to 0.5 TcMAC/s; a naïve but efficient X-engine implementation for the CHIME Pathfinder would require of order 26 GPUs, although network topology considerations favor a baseline target of 32 GPUs in the cluster.

Several factors favor a densely-packed configuration for the X-engine. The proximity of the Pathfinder correlator system to the telescope itself, in a national radio-quiet zone, requires substantial Faraday shielding around the entire system, the cost of which increases rapidly with size. The 10 GbE 4xSFP+ \leftrightarrow QSFP+ cabling chosen for the Pathfinder system shows a steep increase in cost above a length of 7 m, where active cabling becomes necessary. Together, these apply substantial pressure to minimize the physical dimensions of the correlator system. The nodes hosting the GPUs also form the bulk of the hardware cost, so maximizing the density of GPUs within each node is an efficient cost control measure.

A. I/O Requirements

Data arrive from the F-engine as 4+4-bit offset-encoded¹ complex numbers, arranged into 1024 independent frequency sub-bands. Assuming the baseline 32-GPU layout, each GPU is independently processing 32 of these bands, corresponding to a 12.5 MHz band of radio data. In this layout each GPU requires a continuous 25.6 Gbps input stream of sky data drawn from four 10 GbE connections.

If limited to standard bus widths, 8-lane PCI Express Revision 3 becomes the baseline interconnect, meaning 16 lanes of PCIe3 are required for each of the 32 GPUs – 8 to feed the GPU, and 8 to receive data from a NIC. While PCIe3x4 or PCIe2x8 links could theoretically support this transfer rate, tests showed it to be unreliable and prone to lost packets and bottlenecking. The use of PCIe3 interconnects restricts the choice of CPU to Intel varieties, and the Ivy Bridge line of consumer processors provide up to 40 lanes of PCIe3, allowing each host system to support two GPUs, fed from eight 10 GbE links.

The rate at which correlated output data are produced depends linearly on the choice of accumulation period, but is not a significant driver in the design; for the Pathfinder system’s default 21 s integration, the output is a modest ~ 8 Mb/s per node.

B. GPU Selection

Consumer-grade hardware was generally preferred to professional for its increased availability, interoperability of components, and drastically reduced cost. Dedicated scientific GPUs are differentiated by their use of error-correcting memory and the availability of native double-precision functions, neither of which are significant considerations in a radio correlator. The AMD offerings feature higher computational throughput per unit cost than the comparable NVIDIA GPUs, which led to an early preference for AMD-brand consumer-grade GPUs. This is a significant departure from prior GPU X-engines, which typically use the xGPU [4] software package, a highly optimized X-engine implementation built on NVIDIA’s

TABLE I. OPERATIONAL THERMAL PARAMETERS OF DENSE-PACKED AMD R9 280X GPUS FROM DIFFERENT MANUFACTURERS

Manufacturer	Max T (°C)	Power (W)	BW (MHz)
MSI OC Edition	>99	825	13.6
XFX	>99	800	13.0
Club3D RoyalQueen	96	875	14.2
TurboDuo	96	800	13.6
HIS IceQ	89	730	14.5
Sapphire Dual-X OC	85	715	13.5
Sapphire, Watercooled	$T_{\text{water}} + 10^\circ\text{C}$	510	14.5

proprietary CUDA programming framework. A custom software X-engine was developed in its place (§IV-C1, [2]), built on the vendor-independent OpenCL standard. [5]

Proof-of-concept software was developed on 2012-era AMD Southern Islands GPUs similar to those later adopted for the CHIME Pathfinder. This software was able to process all 32,896 baselines of simulated Pathfinder data at the required 12.5 MHz/GPU rate. Further software optimization brought the per-GPU throughput up to ~ 19 MHz, leaving ample headroom for additional processing (§IV-C).

Consumer-grade GPU boards are manufactured by a variety of third parties, with proprietary modifications to the reference design which can result in higher or lower power draw, faster or slower default clocking, variable power-saving or high-temperature operation; and which vary dramatically in their solutions for heat dissipation. The cooling systems are generally not designed for the close packing required for the CHIME Pathfinder: of the 6 varieties of R9 280X boards tested in a packed 4U case, many reached die temperatures close to 100°C. Operation at these temperatures lowers the GPUs’ efficiency and increases their failure rate while drawing additional power for cooling fans. Table I shows the maximum die temperature, mean power draw, and processing throughput of GPUs running an unoptimized ‘stress test’ kernel in a closely-packed case similar to those used in the Pathfinder X-engine. These measurements were performed at an ambient air temperature of $\sim 25^\circ\text{C}$, and differences in power draw and throughput are attributable to a combination of changes in clocking, efficiency, and variable fan speeds.

The best performers were Sapphire-branded “Dual-X OC” boards, which maintained average die temperatures of $\sim 75^\circ\text{C}$ in the test configuration, while matching the computational performance of the other boards and consuming less power. Other strong performers (e.g., the HIS-branded boards) showed higher pricing markups and proved more difficult to source.

III. IMPLEMENTATION - HARDWARE

The X-engine is a cluster of 16 nodes, each housed in a 4U case. It fits in 64U spread across two 42U racks, dubbed East and West, containing 6 and 10 nodes respectively. Each rack is powered by a Tripp Lite PDU3VSR10H50 managed PDU, allowing remote monitoring and power cycling of all nodes. The remaining space in the East rack contains the FPGA-based F-engine as well the 1GbE interchange switch, while the West rack additionally houses a 1U control system dubbed *gamelan*.² A diagram of the X-engine computational cluster

¹That is, a range of $[-8,7]$ is set to $[0,15]$ by adding 8 to the values.

²An Indonesian percussion ensemble featuring chimes; see <http://en.wikipedia.org/wiki/Gamelan>

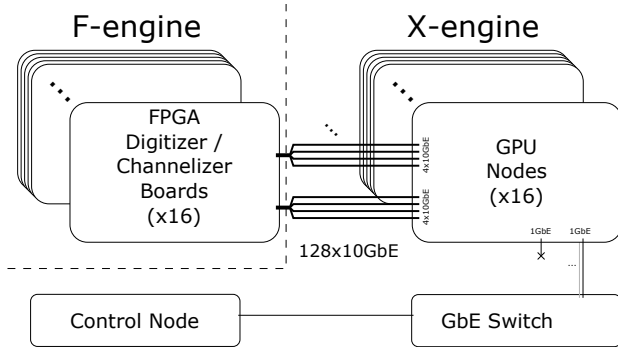


Fig. 1. A schematic of the data flow between the F- and X-engine portions of the correlator. The interchange switch and the gamelan control node are also shown.

is shown in Figure 1.

A. Node Description

A schematic of the hardware layout in each GPU node is shown in Figure 2. The processing nodes are built on EVGA X79 Dark motherboards, Intel i7 4820k CPUs, and 4x4GB DDR3 RAM overclocked to 2133 MHz. They are powered by 1000 W high-efficiency³ power supplies and housed in 4U Chenbro RM41300-FS81 cases. Each node is fed by eight 10GbE lines connected to a pair of Silicom PE310G4SPi9 quad-10GbE network interface cards.

Three AMD R9-series GPUs – two 280X and one 270X – receive and process a total of 25 MHz of radio bandwidth from all the spatial inputs. The 270X is not required for the pairwise correlation, but was added to ensure surplus computational power was available to explore additional real-time processing and alternate correlation algorithms (§IV-C).

The nodes are diskless, both booting off and streaming data back to the gamelan control node via onboard GbE, through a Cisco SRW2048-K9-NA switch. A second on-board GbE connection allows for future expansion and data exchange between nodes, not required in standard operation but available for planned algorithmic upgrades.

The average per-node power usage is ~ 680 W for the air-cooled and ~ 630 W for the liquid-cooled; the total system exceeds the 105 TOPS requirement with a direct power consumption of ≈ 10 kW. The X-engine therefore achieves ~ 11 GOPS/W, though we note that these are 4-bit arithmetic operations, not directly comparable to single-precision floating-point operations.

B. Liquid Cooling Upgrade

In the summer of 2014, 10 of the GPU nodes were retrofitted to cool their CPU and GPUs via direct-to-chip liquid cooling systems. This upgrade significantly eases the strain upon the existing A/C system, and is intended as a proving ground for the full CHIME correlator, where traditional HVAC solutions become prohibitively expensive.

Liquid cooling was implemented using aftermarket heatsinks attached to each GPU, and a Swiftech Apogee II

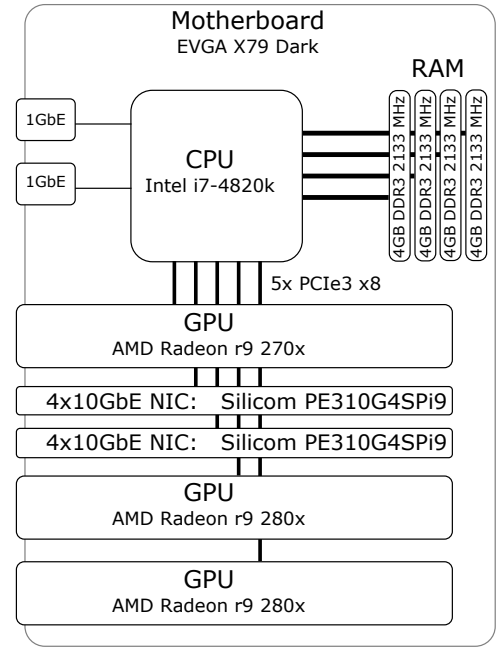


Fig. 2. The internal components of the GPU nodes, including network interfaces, as described in §III-A.

combination pump/heatsink attached to the CPU. Lab testing showed a reduction in GPU die temperatures to 10°C above the input water temperature, with minimal dependence on the specific variety of heatsink. Varieties from several vendors were ultimately deployed, with similar performance. In operation, the watercooled GPUs maintained stable temperatures of 40 – 60°C , compared to $\sim 85^\circ\text{C}$ for the air-cooled GPUs. The cooler operating point and the removal of the low-efficiency GPU cooling fans significantly dropped the power draw per node with no impact on performance, and is expected to reduce the long-term failure rates.

Figure 3 shows the overall structure of the watercooling system as currently deployed. A sealed loop circulates a coolant consisting of 50% water and 50% ethylene glycol through the GPU nodes in parallel. Each node has a small pump contained in the CPU heat sink and passive water blocks on each of the GPUs, while a constant pressure differential is held across the nodes by a larger external pump. The coolant runs through a heat exchanger which transfers heat from the node-cooling loop to an external liquid loop. That loop travels out of the RFI enclosure and is then cooled to the ambient external air temperature in a liquid-air heat exchanger. A temperature-dependent mixing valve permits some fraction of the hot coolant to immediately recirculate, to avoid cooling the system below the dew point inside the RFI enclosure during winter operation.

IV. IMPLEMENTATION - SOFTWARE

A. Data Flow

Channelized data, flags, and metadata from the F-engine arrive at the GPU nodes on eight 10GbE lines. Each 10GbE link carries all spatial inputs for 8 of the 1024 total frequency channels. With eight links per node, each node processes data covering $1/16$ of the full CHIME frequency band. The

³80 Plus Platinum, corresponding to $\geq 92\%$ efficiency at 50% load




Fig. 3. A diagram showing the liquid-cooling structure in the CHIME Pathfinder, with red and blue indicating hot and cold coolant, respectively. The heat exchanger uses a large fan to cool the liquid using ambient outside air. The object marked ‘M’ is a temperature-controlled mixing valve, which regulates the temperature of the ‘cold’ sections of the loop.

*kotekan*⁴ software pipeline manages the data flow and processing within GPU nodes. Due to the high I/O demands (~ 820 Gb/s in total), the system must make maximally efficient use of the available bandwidth at each stage. A packetized and loss-tolerant data handling system, similar to that in operation in the PAPER [6] correlator [7] ensures that momentary faults do not impede long-term data gathering. Recnik et. al. [3] discuss the data handling in detail; a brief description follows.

Data arrive as UDP packets and are buffered by the host CPU in system memory for inspection and staging prior to transfer into the GPUs for processing. Packet loss, though rare, is tracked along with other flags from the F-engine from e.g. saturation of the ADCs or from FFT channelizing. The count of missing or saturated data is used to renormalize the post-integration correlation matrices.

A series of OpenCL kernels are executed on the GPUs; these pre-condition the data, compute and integrate correlation matrices, and post-process the data if necessary (see [2] for more details). Computed correlation matrices are assembled by the CPU and forwarded to *gamelan*, which stitches the full 400 MHz band back together using data from all active nodes. This reassembly is robust against individual node failures or outages; they simply result in loss of data from the inactive nodes.

Integrated correlation matrices are recorded onto an array of disks in *gamelan*, and these data are asynchronously copied to a remote archive server hosting a much larger array of drives, and then copied off-site for scientific analyses.

⁴A style of playing fast interlocking parts in Balinese Gamelan music; see <http://en.wikipedia.org/wiki/Kotekan>

FPGA/Node Name:	Correlator State:	FPGA State:	FPGA Temp:	Node State:	GPU1:	GPU2:	GPU3:	Node Power:
slot_1_iceboard / chi01	Kotekan Running	UP	78.0 °C	UP	61.0 °C	60.0 °C	60.0 °C	657 W
slot_2_iceboard / chi02	Kotekan Stopped	UP	76.2 °C	UP	55.0 °C	53.0 °C	53.6 °C	204 W
slot_3_iceboard / chi03	Kotekan Running	UP	80.1 °C	UP	55.0 °C	53.7 °C	56.0 °C	649 W
slot_4_iceboard / chi04	Kotekan Running	UP	80.3 °C	UP	38.0 °C	37.0 °C	41.0 °C	625 W
slot_5_iceboard / chi05	Kotekan Running	UP	75.9 °C	UP	36.0 °C	37.0 °C	36.0 °C	643 W
slot_6_iceboard / chi06	Kotekan Running	UP	76.2 °C	UP	39.0 °C	37.0 °C	39.0 °C	597 W
slot_7_iceboard / chi07	Kotekan Running	UP	87.6 °C	UP	40.0 °C	38.0 °C	42.0 °C	617 W
slot_8_iceboard / chi08	Kotekan Running	UP	79.9 °C	UP	37.0 °C	37.0 °C	37.0 °C	614 W
slot_9_iceboard / chi09	Kotekan Running	UP	77.6 °C	UP	40.0 °C	34.0 °C	40.0 °C	630 W
slot_10_iceboard / chi10	Kotekan Running	UP	82.9 °C	UP	39.0 °C	48.0 °C	38.0 °C	611 W
slot_11_iceboard / chi11	Kotekan Running	UP	83.1 °C	UP	70.0 °C	71.0 °C	79.0 °C	701 W
slot_12_iceboard / chi12	Kotekan Running	UP	71.5 °C	UP	73.0 °C	73.0 °C	79.0 °C	682 W
slot_13_iceboard / chi13	Kotekan Running	UP	79.9 °C	UP	71.0 °C	70.0 °C	78.0 °C	667 W
slot_14_iceboard / chi14	Kotekan Stopped	UP	77.7 °C	UP	28.0 °C	28.0 °C	24.0 °C	182 W
slot_15_iceboard / chi15	Kotekan Stopped	UP	0.0 °C	UP	59.0 °C	27.9 °C	32.0 °C	286 W
slot_16_iceboard / chi16	Kotekan Running	UP	79.3 °C	UP	70.0 °C	70.0 °C	78.0 °C	681 W

Fig. 4. Screenshot showing the correlator status webpage. The monitoring system displays the status of each FPGA and GPU node, and of the correlator software; per-GPU temperatures and per-node power consumption are also available.

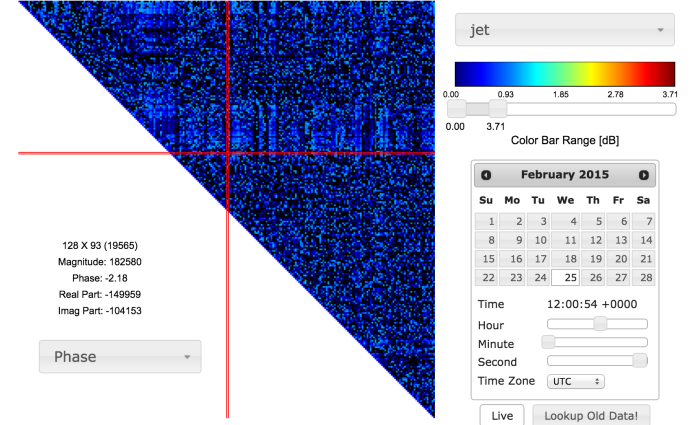


Fig. 5. Screenshot showing the live view webpage. The triangle is the full correlation matrix for a particular frequency, with colour indicating the complex value’s phase; the website may be queried for any of the associated data.

B. Monitoring and Control

The X-engine software pipeline (composed of the *kotekan* instances on each node, along with the collection server software) is launched and controlled through scripts run on *gamelan*. The nodes run CentOS Linux 6.5 and can be accessed by remote shell login, while the PDUs allow remote power cycling to aid in recovery of crashed systems.

The status of each of the GPU nodes is tracked by the *gamelan* control node, and made available via a web interface; see Figure 4 for an example of the tracking display.

In addition, the last few hours of data are streamed over TCP to a second server, where it is available for live analysis and monitoring. An example of the live-monitoring webpage is shown in Figure 5.

C. GPU Data Processing Tasks

The most computationally expensive operation performed on the GPU nodes is the mission-critical pairwise feed correlation. In parallel with this, the Pathfinder correlator will explore alternate correlation methods which leverage the redundant layout of the CHIME baselines. Supplemental tasks include beamforming, gating, time-shifting of inputs, and RFI excision. Brief descriptions of these tasks follow.

1) *Full Correlation*: The primary responsibility of the X-engine is to calculate and integrate the correlation matrix of all the spatial inputs. This involves accumulation of 32,896 pairwise products for each of 1024 frequency bands. The default integration period is 2^{23} samples, corresponding to 21.47 s, much faster than the $\gtrsim 2.5$ minute beam-crossing time from sky rotation. The computational requirements of the full X-engine system are dominated by this correlation operation, such that all other processes constitute an insignificant additional burden. The current implementation achieves near-maximum-theoretical throughput; for details, see Klages et. al. [2]

2) *Alternate Correlation Techniques*: Interferometric arrays with highly redundant baselines can take advantage of correlation techniques that are more efficient than the naïve pairwise method. In the case of feeds which are evenly spaced, FFT-based transformations can be used to increase the efficiency of the correlation to $N \log N$, at the cost of strict calibration requirements. [8][9][10] These correlation strategies will be tested in parallel with the pairwise N^2 correlation; additionally, they may be used in hybrid form with some N^2 and some $N \log N$ stages.

3) *Discrete Beamforming*: The CHIME Pathfinder is a stationary telescope that cannot physically point at a specific source or location on the sky. When observing localized sources, it is desirable to form one or more beams, ‘pointing’ the telescope digitally to an arbitrary location within the main beam. This is accomplished in the GPUs by phase-shifting and summing the data from all antennas, so that signals originating in one region of the sky interfere constructively. This signal is then written out at very high cadence, allowing examination of a localized source with very fine time resolution.

4) *Output Gating*: The CHIME Pathfinder will observe periodic sources such as astronomical pulsars and injected calibration signals. These sources generally vary faster than the default ~ 21 s integration period, but high-cadence gating may be used to observe sub-integration signal structure. Gating consists of partitioning the output into a set of sub-buffers based on the time relative to the period of the source, so that independent ‘on’ and ‘off’ signals may be constructed.

5) *Time Shifting*: Signals from outlying telescope stations can be fed into the correlator. Large spatial separations introduce decorrelation between inputs, which can be corrected for by time-shifting samples within the GPUs. The current implementation permits the correction of any input by up to 168 ms, and has been tested with the nearby John A. Galt 26m radio telescope at DRAO.

6) *RFI Cleaning*: Anthropogenic radio frequency interference (RFI) introduces a significant source of additional noise to the astronomical signal. These signals are generally narrow-band and intermittent, coming and going on timescales much shorter than the default 21 s integration period, but with relatively low duty cycles. High-cadence identification and excision of RFI can be performed within the GPUs, and a variety of algorithms are under development including robust outlier and higher-moment statistical tests. [11][12]

V. X-ENGINE SCALABILITY

The X-engine described here was designed for the CHIME Pathfinder, and must be scaled up significantly for the full CHIME instrument. Given a scaled F-engine providing channelized data, the X-engine’s design allows it to scale straightforwardly to a broader band or larger- N arrays.

Additional radio bandwidth is trivially added through additional nodes; increasing the number of inputs adds to the computational demand on each node, and can be addressed through newer, more powerful GPUs. At the time of writing, the computational power per node could be roughly tripled by simply replacing the GPUs.

To support larger N^2 requirements, the bandwidth handled in each node can be reduced, in exchange for proportionally more nodes. The bandwidth fed to each GPU can similarly be reduced, and for very large N , when even a single frequency band is beyond the capacity of a single processing node, data can be time-multiplexed across multiple GPUs.

The expansion to full CHIME ($N = 2048$) yields an N^2 computational requirement η an order-of-magnitude greater than any system currently in existence. Using *current* technology, a straightforward scaling of the current system — 256 nodes each containing 2 dual-chip R9 295X2 GPUs — could handle the entire pairwise correlation task, without additional software development. This density of processors is easily achievable with the liquid cooling demonstrated, and would occupy a modest physical footprint, at a very low hardware cost of $\sim \$1\text{M}$. However, it is not expected that the full CHIME instrument will rely on a complete N^2 correlation, instead pursuing a fast alternate correlation technique as discussed in §IV-C2.

VI. CONCLUSION

We have implemented a low-cost, high-efficiency GPU-based correlator X-engine for the CHIME Pathfinder. Capable of correlating 32,896 baselines over 400 MHz of radio bandwidth, it makes efficient use of consumer-grade parts and executes a highly optimized software stack. Measured by the computational requirement of a naïve N^2 correlation — the bandwidth-baseline product η defined by Equation 1 — the CHIME Pathfinder correlator is among the largest in the world. Aspects of the system such as the cooling systems have been substantially modified, optimizing the X-engine’s efficiency and ensuring economical scaling to the full-size CHIME instrument.

ACKNOWLEDGEMENTS

We are very grateful for the warm reception and skillful help we have received from the staff of the Dominion Radio Astrophysical Observatory, which is operated by the National Research Council of Canada.

We acknowledge support from the Canada Foundation for Innovation, the Natural Sciences and Engineering Research Council of Canada, the B.C. Knowledge Development Fund, le Cofinancement gouvernement du Québec-FCI, the Ontario Research Fund, the CIFAR Cosmology and Gravity program, the Canada Research Chairs program, and the National Research Council of Canada. PK thanks IBM Canada for funding

his research and work through the Southern Ontario Smart Computing Innovation Platform (SOSCIIP).

We thank Xilinx University Programs for their generous support of the CHIME project, and AMD for donation of test units.

REFERENCES

- [1] K. Bandura *et al.*, “Canadian Hydrogen Intensity Mapping Experiment (CHIME) Pathfinder,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9145, Jul. 2014, p. 22.
- [2] P. Klages *et al.*, “Data Packing for High-Speed 4-Bit GPU Correlators,” 2015, In press; accepted to IEEE ASAP 2015.
- [3] A. Recnik *et al.*, “An Efficient Real-time Data Pipeline for the CHIME Pathfinder Radio Telescope X-Engine,” 2015, In press; accepted to IEEE ASAP 2015.
- [4] M. A. Clark, P. C. La Plante, and L. J. Greenhill, “Accelerating Radio Astronomy Cross-Correlation with Graphics Processing Units,” *ArXiv e-prints*, Jul. 2011.
- [5] Khronos Group: The OpenCL Specification. [Online]. Available: khronos.org/opencl
- [6] A. R. Parsons *et al.*, “New Limits on 21 cm Epoch of Reionization from PAPER-32 Consistent with an X-Ray Heated Intergalactic Medium at $z = 7.7$,” *The Astrophysical Journal*, vol. 788, p. 106, Jun. 2014.
- [7] A. Parsons *et al.*, “A Scalable Correlator Architecture Based on Modular FPGA Hardware, Reuseable Gateway, and Data Packetization,” *Publications of the Astronomical Society of the Pacific*, vol. 120, pp. 1207–1221, Nov. 2008.
- [8] J. D. Bunton, “Antenna Array Geometries to Reduce the Compute Load in Radio Telescopes,” *IEEE Transactions on Antennas and Propagation*, vol. 59, pp. 2041–2046, Jun. 2011.
- [9] M. Tegmark and M. Zaldarriaga, “Fast Fourier transform telescope,” *Physical Review D*, vol. 79, no. 8, p. 083530, Apr. 2009.
- [10] —, “Omniscopes: Large area telescope arrays with only $N \log N$ computational cost,” *Physical Review D*, vol. 82, no. 10, p. 103501, Nov. 2010.
- [11] G. M. Nita, D. E. Gary, Z. Liu, G. J. Hurford, and S. M. White, “Radio Frequency Interference Excision Using Spectral-Domain Statistics,” *Publications of the Astronomical Society of the Pacific*, vol. 119, pp. 805–827, Jul. 2007.
- [12] G. M. Nita and D. E. Gary, “Statistics of the Spectral Kurtosis Estimator,” *Publications of the Astronomical Society of the Pacific*, vol. 122, pp. 595–607, May 2010.