CASPER 2021 Conference Program

Monday May 17

7:15 a.m. Pacific Time, 16:15 SA Time
Welcome remarks and Casper Overview -- Kathryn Plant and Jonathon Kocz

7:30-8:00 a.m. Pacific Time, 16:30-17:00 SA Time
Invited talk: Event Horizon Telescope
Title: TBD
Abstract: TBD

8:00-8:30 a.m. Pacific Time, 17:00-17:30 SA Time

Invited talk: Mitchell Burnett
Brigham Young University
mitch.burnett@byu.edu
Bringing RFSoC to CASPER
The Xilinx Zynq UltraScale+ RFSoC integrates programmable logic with the Zynq ARM (A53) processor, high speed serial transceivers, and several multi-gigasample per second RF Data Converters (RFDC) capable of digital down conversion and direct sampling for RF signals up to 6 GHz. These features are ideal for small form factor and low power digitizers streaming raw voltages over 2x100GbE, or suitable as modest sized channelizers using the available fabric resources. This presentation will provide a brief technology overview of the RFSoC, demonstrate the new RFDC yellow block and six newly supported RFSoC CASPER platforms: zcu216, zcu208, zcu111, zrf16 (both 29/49DR rev.), and PYNQ 2x2. The current capabilities of each platform, how the RFDC integrates with the tool flow, current software (i.e., casperfpga) support, available tutorials and resources, and future development roadmap will also be shown. The addition of these new CASPER platforms and the extent to which the RFDC yellow block integrates with the hardware to provide flexible design configurations will continue to proliferate the design philosophy of CASPER of decreasing the time-to-science metric and provide a way of bringing the needed capabilities to next generation instruments.

8:30-9:15 a.m. Pacific Time, 17:30-18:15 SA Time

Session I

Session Chair: Dan Werthimer

GNU Radio X-Engine at the ATA

Michael Piscopo
DeepSIG and GNU Radio
mike.piscopo@gmail.com

Abstract

A High-Throughput Oversampled Polyphase Filter Bank using Vivado HLS and PYNQ on a RFSoC

Jenny Smith
University of California, Santa Barbara
jennifer_smith@ucsb.edu

Many digital signal processing applications require a channelizer capable of moving sections of the incoming spectrum to baseband quickly and efficiently with minimal spectral leakage and signal distortion. We report the design and implementation of a 4 GHz, 4096-branch, 8-tap, 2/1 oversampled polyphase channelizer implemented on a Xilinx RFSoC. The open-source design consists of only IP cores created using Vivado HLS (C++) and IP cores available in Vivado HLx and System Generator versions 2019.1+. The channelizer was tested using a PYNQ overlay and Jupyter Notebook (Python) hosted on the RFSoC embedded CPU. The design uses 24% of the LUTs, 9% of the DSP48s, and 11% of the BRAMs on the ZCU111 RFSoC evaluation board and meets timing constraints at 512 MHz. The oversampled polyphase channelizer suppresses spectral image components below -60 dB. This design provides the first example of an oversampled polyphase channelizer running on a system on a chip architecture created without direct use of hardware description language. The presented approach leverages high-level design tools and includes source code which can be readily adapted by other researchers and development teams without the need for specialist knowledge in high-performance FPGA design.

SWRDS: SMT Wideband Radio Digital Spectrometer

Arash Roshanineshat
University of Arizona
aroshanineshat@email.arizona.edu

We introduce a new ultra-wideband backend, the SMT Wideband Radio Digital Spectrometer (SWRDS). The spectrometer uses the Xilinx ZCU111 evaluation board for the 1st generation Zynq UltraScale+ RF System on Chip. This device combines an FPGA with eight 12-bit analog-to-digital converters (ADCs) and eight 14-bit digital-to-analog converters (DACs). Each ADC has 4GHz analog bandwith and can be clocked up to 4.096 Giga-samples-per-second (Gsps), a Nyquist bandwidth of 2GHz. We interleave four pairs of ADCs to produce four channels of 4GHz bandwidth. To optimize this interleaving, we have designed and built an expansion board for the

ZCU111 that allows fine clock and gain adjustments for each channel. A robust embedded software structure is implemented that provides an extensive API for controlling and designing a custom user interface. A python command line interface and a web-based GUI are used to communicate with the board and visualize data.

9:15 a.m.--9:30a.m. Pacific Time, 18:15--18:30 SA Time
Break
9:30-10:15 a.m. Pacific Time, 18:30-19:15 SA Time

Session II

Session Chair: Amish Patel

Moving from SKARAB to Alveo

Adam Isaacson
South African Radio Astronomy Observatory
aisaacson@ska.ac.za

The MeerKAT radio telescope is now available for scientific observations. The toolflow development for MeerKAT is now complete and we are entering a support phase for the SKARAB. This presentation briefly covers the recent functionality added to the SKARAB toolflow for MeerKAT and our latest SKARAB ADC work.

The presentation will then cover the next generation COTS hardware: the Alveo. It will explain why SARAO selected the Alveo hardware, the initial efforts of trying to utilise an RTL design using our 100GbE design with the Alveo using the existing toolflow and why SARAO finally decided to develop the way the data accelerator cards were intended for. This presentation will cover the PROs and CONs of using the Alveo based on our early experiences with the Alveo.

MeerKAT Correlator Beamformer and the Next Generation (Correlator)

Tyrone van Balla
SARAO
tvanballa@ska.ac.za

The DSP Team at SARAO is currently in a state of transition as we complete work on the MeerKAT Radio Telescope's Correlator Beamformer and begin designing and developing the next-generation correlator for the MeerKAT Extension (or MeerKAT+) project. This talk will provide a brief status update on the final stages of development of the MeerKAT Radio Telescope's Correlator Beamformer, introduce the MeerKAT+ project and outline the plans and architecture for the entirely GPU-based next-generation correlator (NGC). (More detailed presentations on the GPU F-Engine and GPU X-Engine will be delivered by other SARAO DSP Team Members.)

Integrating the Alveo into the CASPER Ecosystem

Wesley New

SARAO
wesley@ska.ac.za
With FPGA companies pushing into the datacenter markets, we have been assessing the viability of using the PCIe based Alveo board for radio astronomy and streaming data processing. These cards do not lend themselves to the traditional CASPER graphical design methodology, but rather to design with high-level languages. This means that the collaborative focus will target primarily the libraries that can be used with this hardware. This presentation will discuss how we came to these conclusions and how we plan to feedback the work into the CASPER community.

10:15 a.m.--10:30a.m. Pacific Time, 19:15--19:30 SA Time

Sparkler Session (Many Short Talks)

16G Spectrometer Development

Homin Jiang
Institute of Astronomy and Astrophysics, Academia Sinica
homin@asiaa.sinica.edu.tw

Radio Astronomy Spectrometer Deployment and Initial Results at NASA's Canberra Deep Space Communications Complex

Kristen Virkler
Jet Propulsion Laboratory, California Institute of Technology
kristen.virkler@jpl.nasa.gov

Kristen Virkler (Jet Propulsion Laboratory, California Institute of Technology), Melissa Soriano (Jet Propulsion Laboratory, California Institute of Technology), Jonathon Kocz (UC Berkeley), Jorge Pineda (Jet Propulsion Laboratory, California Institute of Technology), Tyrone McNichols (California Institute of Technology), Shinji Horiuchi (Canberra Deep Space Communications Complex), Brian Bradford (Jet Propulsion Laboratory, California Institute of Technology)

A high-resolution spectrometer was installed and commissioned at NASA's Canberra Deep Space Communications Complex (CDSCC) in November 2019 to enable spectroscopic capabilities. The spectrometer is comprised of four Reconfigurable Open Architecture Computing Hardware (ROACH)-2 boards that process a total bandwidth of 16 GHz. The spectrometer samples and processes data from CDSCC's 70-m antenna and K-band front end, covering frequencies of 17-27 GHz. The large total bandwidth enables the simultaneous observations, which can be combined together to significantly improve the sensitivity of these observations.

The system has two different firmware modes: 1) A 65k-pt FFT that provides 32,768 spectral channels at ~30.5 kHz and 2) A 16k-pt polyphase filterbank (PFB) that provides 8,192 spectral channels with ~122 kHz resolution. The firmware utilizes the tool flow developed by the Collaboration for Astronomy Signal processing and Electronics Research (CASPER).

To date, the ROACH-2 spectrometer has performed over 50 observations at CDSCC. Software modules were developed to control the spectrometer observations and to enable its operation within the Deep Space Network's infrastructure. The spectrometer is autonomously controlled through a configuration file, which includes information such as observation schedule, antenna pointing information, and firmware mode selection. The output is recorded as binary data, which is then converted to Single Dish FITS (SDFITS) format for storage and processing. There are two levels of data products: level 0 provides the raw data and all the needed metadata for a Principle Investigator (PI) to process the data themselves and level 1 corrects the data for antenna gain variations as a function of source elevation.

The spectrometer is currently being used for an extensive JPL RRL mapping project as well as for CDSCC radio astronomy projects. In the future, additional software development will enable further automation of spectrometer operations and use by external PIs.

Applying video encoders to filterbanks

Suryarao Bethapudi
Max-Planck-Institut fur Radioastronomie
shining.surya.d8@gmail.com

Total Intensity filterbank is the staple diet for all of time-domain radio astronomy. We propose making use of video encoding/decoding techniques to filterbanks (with 8bit data) as a way to reduce memory footprint and investigate its effect on data. We take 2560-samples of filterbank at once, see it as a luminance frame ("y" of "yuv") and encode it into an mp4 using FFMPEG. Decoding is simply the inverse process. Although we are encoding the mp4 as "yuv" when our data only lies in "y", the encoder encodes the "uv" as purely zeros. Across various backends ranging from GMRT band4, band5, Effelsberg Phased Array Feed, C-band, and VLA/VLA Low-frequency Ionospheric and Transient Experiment (VLITE), we achieve a minimum compression ratio of 2.54, while the mean-squared-error is not constant across the backends, the average MSE is 42.33 +/- 5.21. While the mean-squared-error (MSE) is high and not constant across the backends, the MSEs are constant in each filterbank implying a systematic in the encoding/decoding pipeline is not being accounted for. Since video compression techniques are unlike loss-less, entropy-based coding, we see that even when sending Gaussian white noise-like data, we achieve similar compression ratios. As a proof of concept, we uploaded a 7.5 minute observation as an MP4 on youtube: https://youtu.be/UmwE2tWI6mI. This technique can be applied to use cases such as long term storage, rapid transport, and streaming realtime filterbank data from geographically separated stations for realtime searches. Lastly, we would like to note that we have merely just scratched the surface of the video compression techniques, and while there is an abundance of models/ideas that can be borrowed from compression techniques to apply in radio astronomy, the effect of doing so on the data should be more carefully examined.

FPGA Based Radiometer with four Spectral Windows

Heystek Grobler
SARAO - Hartebeesthoek Site (HartRAO)
heystek@hartrao.ac.za

Radio astronomy is a subfield of the field of astronomy that studies the radio frequency part of the electromagnetic spectrum with the means of radio telescopes. To record and analyse signals, high-performance Digital Signal Processing (DSP) equipment and techniques are used. This technology makes use of customised complicated hardware and is very specialised. In radio astronomy, objects are normally studied either with total power (continuum) or by means of spectra. Thus, there is a need to provide simultaneous observations of both. This project investigates the combination of a radiometer and spectrometer in one device where the total power of the signal can be determined with conventional radiometer techniques and four spectral windows added to do spectroscopy as well. This will be investigated on the CASPER (Collaboration for Astronomy Signal Processing and Electronics Research) framework with the ROACH-2 (Reconfigurable Open Source Architecture Computing Hardware) FPGA (Field Programmable Gate Array). CASPER is a research group consisting of members from across the globe that provides new methods and equipment for radio astronomy focusing on flexible reusable open source systems.

10:30 a.m.--10:45a.m. Pacific Time, 19:30--19:45 SA Time
Break

10:45a.m.--11:15a.m. Pacific Time, 19:45--20:15 SA Time

Invited talk: Patrick Lysaght

Xilinx
plysaght@xilinx.com
Abstract

11:15a.m.--12:15p.m. Pacific Time, 20:15--21:15 SA Time

Discussion: New Hardware (ADCs, FPGAs, Alveo boards, GPUs ...)

Tuesday May 18

7:15 a.m. Pacific Time, 16:15 SA Time

Applications Based on SNAP2
Lin Shu and Liangtian Zhao
Some new instruments based on SNAP2 and other boards or servers developed by our team will be introduced. (1)Beamformer for Tianlai Cylinder Array(96 antennas, each with two polarizations) in Xinjiang, China. (2) Correlator for Heliograph Array(116 antennas, each with two polarizations) in Inner Mongolia, China.

7:30-8:00 a.m. Pacific Time, 16:30-17:00 SA Time

Invited talk: Bastian Bloessl and Joshua Morman

Perspecta Labs
mail@bastibl.net, jmorman@perspectalabs.com,
GNU Radio Framework for Heterogeneous Systems
GNU Radio is an Open Source real-time signal processing framework and the de-facto standard for CPU-based SDR applications. Yet, its current runtime has several inherent limitations that hinder proper integration of accelerators and custom schedulers. The GNU Radio project is, therefore, working on a major rewrite of its runtime system to better support heterogeneous architectures and application-specific scheduling algorithms. Ideally, these efforts will form the base for GNU Radio 4.0, a new major release that sets the project up for the next decade.
In this talk, we will give a high-level overview of the long-term goals and core concepts, as well as go into detail about our recent efforts, which have focused on the integration of GPU-accelerated blocks using the CUDA API. To this end, we added support for custom accelerator buffers that allow CPU blocks to write directly into accelerator memory. This is not possible with current GNU Radio, introducing unnecessary copies which are problematic for high-performance applications. With full-integration and support of accelerator buffers, we are, furthermore, able to stay in accelerator memory, if consecutive operations are able utilize the accelerator. Compared to the current state-of-the-art, this provides significant performance gains which we will show for various scenarios.

8:00-9:00 a.m. Pacific Time, 17:00-18:00 SA Time

Session I

Session Chair: Tyrone van Balla

Ponderings, plans and progress: The MeerKAT+ GPU based F-Engine
Andrew van der Byl
SARAO
avanderbyl@ska.ac.za

This presentation will focus on the design considerations for a GPU based F-Engine for the next phase of the MeerKAT radio telescope correlator-beamformer. The presentation will provide a high-level overview of the MeerKAT+ F-Engine covering design thoughts and considerations for building a GPU based F-Engine. Design plans and current progress will be discussed.

The Tensor-Core Correlator

John Romein
ASTRON - Netherlands Institute for Radio Astronomy
romein@astron.nl

New GPU technology, called tensor cores, performs limited-precision matrix multiplications much more efficiently than regular GPU cores. In this talk, I will present the tensor-core correlator, a GPU library that exploits this technology, and allows a GPU to correlate signals 5-10 times faster and more energy efficient than traditional state-of-the-art GPU correlators. The library hides the complexity of the use of tensor cores, and can be easily integrated into the processing pipelines of existing and future instruments.

A python-based GPU X-Engine for MeerKAT+ that uses Nvidia's Tensor Cores for correlation.

Amish Patel
South African Radio Astronomy Observatory (SARAO)
apatel@ska.ac.za

Based on the ASTRON Tensor Core Correlation Kernel, this highly-configurable software is able to process data for array sizes from four to >256 antennas. The asyncio python library is central to this software's operation, utilising SPEAD2's UDP ibverbs transport to facilitate high-speed data reception and transmission. This presentation will outline the reasoning for this design choice (and details therein), as well as noteworthy observations made during this software's development.

MeerKAT+ GPU correlator

An FPGA-GPU digital back-end for the pathfinder of the Deuterium Experiment Telescope

Vinand Prayag and Nitish Ragoomundun
University of Mauritius
vinand28@gmail.com, nitish.ragoomundun@gmail.com

The Deuterium Experiment Telescope (DET) is a radio telescope being conceived to make sensitive measurements of the hyperfine ground-state spin-flip transition of Deuterium radio spectral line at 327 MHz in certain regions of theLocal Galaxy and in the Large Magellanic Cloud. A prototype of such an instru-ment was developed at the University of Mauritius to test the feasibility of the system. This prototype constitutes a pathfinder station with 19 dual-polarizedYagi-Uda antennas. The DET uses CASPER SNAP boards and an NVIDIAGTX Titan Xp graphics card as its FPGA-GPU digital back-end infrastructure.The conditioned analogue signals (amplified and band limited) from a stationare fed to 4 networked SNAP boards. The SNAP FPGAs perform digital filter-ing and packetize the data. These packets are then sent to the

GPU processing node via network. A custom bandpass sampling scheme is programmed on theFPGA. The analogue inputs are sampled at 250 MSps and the FPGA performs quadrature digital down-conversion to bring the 250 kHz Deuterium line to the baseband and output a complex signal. Then, through five successive stages ofFIR filtering and decimation, on all the signals in parallel, the sampling rate of each channel is reduced to 250 kSps. The net digital data rate at the output of the FPGA acquisition system for the pathfinder is 72.6 MiB/s.The core of the processing is done on a gaming NVIDIA GTX Titan Xpgraphics card. A parallel MPI server program running on the host CPU of the workstation receives the data through the network, performs the depacketiza-tion, and copies the data to the GPU. A parallel FX correlator was developed for the GPU using CUDA C. The F-engine was implemented as a polyphase filterbank channeliser. A custom kernel applies the windowed-sinc filter and thecuFFT library is used to perform the Fourier transform. The process of correla-tion was reformulated to be able to perform all the complex-multiply-accumulate operations independently in every frequency channel of the spectrum naturally by matrix multiplication. Then, the X-engine was implemented using a cuBLASlibrary routine. The latter is very efficient at multiplying matrices in batch. Acustom data re-ordering kernel runs between the F and X engines to ensure that data output by the F-engine is in the correct arrangement before input to thecuBLAS routine. The GPU pipeline achieves throughputs in the GiB/s regime,which is sufficient to process the data from the SNAP boards in real-time.The pathfinder digital pipeline works with the four readily available FPGAboards, a desktop PC with a high-end CPU, and a general purpose gaming graphics card. Since the processing schemes are mostly implemented in soft-ware, the system is highly flexible. Parallel computing platforms offer so much processing power that the scalability of the system depends mostly on how the components are networked. The use of libraries provided by CUDA further cuts down the development time for the telescope, thus reducing the time-to-science of the instrument. This design also makes maximum use of the GPU performance capabilities as these libraries are highly optimised for the NVIDIAhardware. We present test runs with simulated signals and we also present data from test observations with a few antennas built for the pathfinder.

9:00 a.m.--9:15a.m. Pacific Time, 18:00--18:15 SA Time
Break

9:15a.m.--10:15a.m. Pacific Time, 18:15--19:15 SA Time

Invited talk: NVIDIA

Speaker: TBD
Afiliation: NVIDIA
Email address
Abstract

Invited talk: NVIDIA II

Speaker: TBD
Afiliation: NVIDIA
Email address
Abstract

10:15-11:15 a.m. Pacific Time, 19:15-20:15 SA Time

Session II

Session Chair: Jack Hickish

The U-Board – a heterogeneous platform for FPGA and GPU based signal processing

Gerrit Grutzeck
Max-Planck-Institut für Radioastronomie Bonn
ggrutzeck@mpifr-bonn.mpg.de
Nvidia is developing embedded modules, called Jetson, which provide a platform for artificial intelligence at the edge, e.g. for self-driving cars. These embedded systems are powerful combinations of ARM CPU cores and modern CUDA enabled GPUs accompanied by a fast and large DDR memory.

We introduce a new concept where an FPGA is coupled closely via PCIe with such a Jetson module. This heterogeneous platform combines the flexible interfacing and integer processing power of the FPGA with the float point processing power of the GPU.
Not only unifies this concept the strengths of FPGAs and GPUs but also opens new possibilities to extend instruments with machine-learning algorithms. At the same time, this enables fast and easy development of algorithms within the Python/CUDA framework.

Based on this idea we developed a prototype including an analog frontend with a broad range of applications in mind. Hereby the analog frontend consists of a transmitting part and a receiving part, which provides an instantaneous bandwidth of over 5.2 GHz. The use-cases could be MKID readout, laboratory spectroscopy or high resolution spectroscopy for instances.

A Novel GPU-based Real-Time Imager for Compact Radio Arrays

Hariharan Krishnan
Arizona State University (ASU)
hari.krish@asu.edu
In this talk, I will describe a GPU-based real-time imaging correlator, called EPIC, that is being commissioned as a commensal back-end for the Long-Wavelength Array station at Sevilleta, New Mexico. The E-field Parallel Imaging Correlator (EPIC) is based on the Modular Optimal Frequency Fourier (MOFF) mathematical formalism for direct Fourier imaging, which images directly from the voltages measured by individual antennas without the expensive cross-correlation operation in the aperture-plane. It is now implemented on a GPU-accelerated architecture and integrated with

a python/C++ based high-performance streaming framework, Bifrost. I will briefly discuss the direct-imaging algorithm and present details of the architecture of the GPU implementation. I will also discuss the ongoing optimization and development efforts in upgrading the EPIC highlighting some of the recent observational results.

Building Hybrid FPGA/GPU Systems with Bifrost
Jayce Dowell on behalf of the Bifrost Collaboration
University of New Mexico
jdowell@unm.edu

Thanks to advances driven by industry, FPGAs and GPUs offer relatively inexpensive ways to build highly parallel processing systems for scientific data acquisition and analysis. Until recently, however, building these systems has required highly specialized programming techniques that has slowed their adoption. The CASPER community has made significant progress toward democratizing FPGA programming, and Bifrost seeks to do the same for GPUs. Bifrost is an open-source, modular C++/CUDA/Python framework that is intended to make building reliable, high-performance data capture and analysis pipelines using CPUs and GPUs easier. In this talk I will discuss Bifrost and its core concepts of ring memory spaces and blocks. I will also discuss use cases for Bifrost by providing examples of instruments that are currently using or investigating using the framework. Finally, future development plans and directions for the framework will be presented.

Data Acquisition with Hashpipe: Past, Present, Future
Dave MacMahon
UC Berkeley Radio Astronomy Lab
davidm@berkeley.edu
Radio astronomy DSP has become increasingly hybridized with high speed digitizers interfacing to FPGAs which send the data over high speed Ethernet links to GPU/CPU clusters. The "High Availability SHared PIPline Engine" (Hashpipe) software package has enabled this model to be used successfully by many projects. This presentation will cover some of the basics of Hashpipe and discuss recent performance enhancing developments including support for Infiniband Verbs (zero-copy, kernel-bypass packet reception) as well as integration with Julia that allows Hashpipe threads written in Julia to interact with Hashpipe threads written in C.

10:45 a.m.--11:00a.m. Pacific Time, 19:45--20:00 SA Time
Break

11:30a.m.--12:30p.m. Pacific Time, 20:30--21:30 SA Time

Discussion: Backend Pipelines

Wednesday May 19

7:30-8:00 a.m. Pacific Time, 16:30-17:00 SA Time

Invited talk: John Romein

ASTRON - Netherlands Institute for Radio Astronomy
romein@astron.nl
Title: Experiences with Intel's OpenCL/FPGA programming environment
Abstract: Programming FPGAs in a Hardware Description Language like Verilog or VHDL is difficult, time consuming, and error prone. We explored Intel's OpenCL/FPGA toolkit, which promised to dramatically reduce programming effort, and showed that implementing complex applications like an imager is now possible. We also modified an existing OpenCL Boad Support Package to read out an ADC and to allow sending and receiving 40 GbE network packets from an OpenCL program, allowing rapid development of a signal-processing pipeline demo. In this talk, I will compare CPU, GPU, and FPGA imaging, and will discuss benefits and limitations of the OpenCL/FPGA programming environment.

8:00-9:00 a.m. Pacific Time, 17:00-18:00 SA Time

Session I

Session Chair: Jenny Smith

CASPER's HDL Development

Talon Myburgh

Prompted by the inability to fully verify CASPER's DSP IP and the tether to the proprietary software's used within the CASPER tool flow, work has been done over the past year to develop DSP IP in HDL for the community. ASTRON open-sourced its   HDL DSP IP Core library. After extending and adapting these cores for use in the CASPER tool flow, CASPER may explore and adopt new software engineering practices for the management of these cores. There is also the potential to develop using alternative tools (both open source and proprietary) as well as deploy to alternative vendor products. This talk will introduce these DSP cores, discuss the research and testing conducted in continuous integration and continuous documentation and explore the possibilities of an expanded tool flow for the CASPER user.

Comparing FPGA-based FFT Implementations on a SKARAB:

Morag Brown
SARAO
mbrown@ska.ac.za

While a number of FFT modules for use on FPGA boards exist in the wild, only a small subset of these are readily available for use with the CASPER toolflow.

Of these, the CASPER FFTs are likely among the most commonly used within the collaboration. Naturally, Xilinx also provides FFT cores with their tools, including system generator blocks compatible with the toolflow, but straightforward wideband solutions are somewhat lacking. ASTRON's HDL FFT cores were recently integrated into the toolflow too, providing the community with even more resources with which to be frustrated over spectral analysis.

This talk will present preliminary results on the comparison of CASPER, ASTRON and Xilinx SSR FFT cores, for use with SKARAB hardware specifically. Metrics include resource utilization, performance at large N, quantisation losses, and spectral leakage.

He6CRES experiment

Brent Graner
University of Washington CENPA
bgraner@uw.edu

The He6CRES experiment is a new project at the University of Washington looking for evidence of physics outside the standard model by measuring the energy of helium-6 beta decay electrons using the novel technique of cyclotron radiation emission spectroscopy (CRES). The CRES technique offers unprecedented precision in energy resolution by extracting kinetic energy information from the frequency of microwaves emitted by individual electrons orbiting inside a magnetic trap. This talk will briefly introduce the CRES technique itself and cover the data acquisition methods we have implemented using the CASPER ROACH2 system for identifying sub-femtowatt signals with maximum frequency resolution and optimum electron detection efficiency.

A SNAP2 cosmic ray detector with FPGA-based RFI flagging at the OVRO-LWA: update from early commissioning

Kathryn Plant
Caltech
kp@astro.caltech.edu

The Owens Valley Radio Observatory- Long Wavelength Array (OVRO-LWA) in Eastern California is currently undergoing an expansion to 352 dual-polarization antennas and new signal processing infrastructure, which includes an FPGA-based cosmic ray detector. High energy cosmic rays produce ~5-nanosecond radio pulses from a cascade of particle interactions in the atmosphere. The capability to identify cosmic ray candidates in the high-time-resolution raw ADC output allows cosmic ray detection using radio data alone, avoiding the need for external particle detectors. This approach requires flagging RFI and identifying candidate cosmic rays within the same SNAP2 boards that simultaneously run the F-engine for the array's main astronomy functions. When fully commissioned, the OVRO-LWA will observe thousands of cosmic rays per year at energies 10^17-10^18 eV and will place new constraints on the proportion of light and heavy elements in the cosmic ray population at energies at the limits of Galactic particle accelerators. I will present the cosmic ray signal processing design, RFI flagging strategy, and a progress update from early commissioning.

9:00 a.m.--9:15a.m. Pacific Time, 18:15--18:45 SA Time
Break

9:15a.m.--9:45a.m. Pacific Time, 18:00--18:45 SA Time

Invited talk: Robert Anderson

The MathWorks Team
robert.anderson@mathworks.com

MathWorks Model-Based workflows are transforming the way engineers and scientists realize systems targeted to ASIC, FPGA and SoC. Starting with abstract systems modeled using Simulink and MATLAB, one can assess performance, and subsequently traverse to an efficient realization on custom as well as COTS platforms, including RFSoC.

This presentation will show how to start with Simulink and MATLAB models, and convert these from floating point to a fixed-point representation. Once fixed-point performance is achieved, we will cover how to generate of register transfer language (RTL) Verilog or VHDL using some of MathWorks' best practices. It will also show how to target a Simulink model to the ZCU111 RFSoC Platform for using a MathWorks developed underlying reference design.

9:45-10:15 a.m. Pacific Time, 18:45-19:15 SA Time

Session II

Session Chair: Francois Kapp

Title: Application of the AXI interface in the CASPER toolflow

Mathews Chirindo
SARAO
mchirindo@ska.ac.za

There have been different notions on whether the ARM eXtensible Interface (AXI) can be adopted in the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) toolflow going forward and/or to retain the wishbone interface. Whilst it is generally perceived that the use of the AXI interface is not fully understood and/or not properly documented, this presentation is intended to clarify/justify the AXI interface use case in the CASPER toolflow and to provide pointers to recent documentation.
The presentation will introduce basic concepts of the AXI interface and explain how it is used, particularly in the CASPER toolflow. The concept of raw AXI interface will be explained and how that has also been adopted in the toolflow. Reference to recent documentation on the AXI interface will be furnished.

PRP

John Graham
UC San Diego
jjgraham@eng.ucsd.edu

The Pacific Research Platform / Nautilus Kubernetes Federation

Abstract:
Containerized remote desktop development environments for GPU and FPGA workflows.
And much more...

Resource-optimized FPA Beamformer Prototype using CASPER Toolflow for the Expanded GMRT

Bela Dixit
The Giant Metrewave Radio Telescope Observatory,Pune,India
dixitbela@gmail.com
Bela S. Dixit, Kaushal D. Buch, Rahul Aragade, Ajithkumar B., Jayaram Chengalur
Giant Metrewave Radio Telescope, NCRA-TIFR, Pune, India

The Expanded GMRT (eGMRT) project proposes to increase the field-of-view (FoV) of GMRT antennas using a 144-element Focal Plane Array (FPA) and 30-beam, 300 MHz bandwidth, 16384-spectral channel beamformer. Through architectural optimization and use of FPGA re-configurability, we have implemented a 32-element, 5-beam, 32 MHz bandwidth, 1024-channel prototype beamformer on a single ROACH-1 board using CASPER toolflow. The beamformer uses 64-element ADC board from CASPER. The correlator and beamformer are implemented on a time-shared basis as the correlation information is required only at distinct intervals. Nominally, the FPGA is programmed as a real-time beamformer and re-programmed as a correlator for gain/phase calibration. Compared to simultaneous implementation of both correlator and beamformer, this scheme reduces the resource utilization by about 50% and makes it possible to accommodate more number of antenna elements and beams on a single Virtex-5 FPGA. The correlator design is also optimized separately for row-wise computation of correlation matrix.

Beamformer weight calculation has been carried out using the maxSNR algorithm to achieve optimal SNR in presence of mutual coupling between the FPA elements.

The prototype beamformer is tested in aperture array mode using the ASTRON L-band FPA. Beamforming using the maxSNR method and the results from comparison between the beam signal-to-noise ratio (SNR) achieved through conventional phased-array beamforming and maxSNR beamforming will be presented.

Since the eGMRT configuration needs wider bandwidth (300MHz), we are working towards designing spectrometer and beamformer using RFSoC-based ZCU111 evaluation board and the initial results from these efforts would be described.

MeerKAT beamformer
Andrew Martens
SARAO
andrew@ska.ac.za
The design and inner workings of the FPGA-based MeerKAT B-Engine. This implements a beamformer producing 4 full time resolution beams per polarisation (for a total of 8 independent beams). It combines 856MHz bandwidth data (broken into up to 32k spectral channels) from up to 64 antennas. It is designed to be co-located with a CASPER X-Engine. The design is in Simulink and it is currently used on the SKARAB processing platform.

10:45 a.m.--11:00a.m. Pacific Time, 19:45--20:00 SA Time
Break

11:00a.m.--12:00p.m. Pacific Time, 20:00--21:00 SA Time
Discussion: Toolflow

Thursday May 20

7:30-8:00 a.m. Pacific Time, 16:30-17:00 SA Time
Invited talk: Chris Bochenek
Caltech
cbochenek@astro.caltech.edu
STARE2
Since their discovery in 2007, much effort has been put into determining the progenitors of fast radio bursts (FRBs). The durations and energetics of FRBs imply a compact, highly magnetized progenitor, making magnetars a popular progenitor candidate. In this talk, I will present the discovery of a fast radio burst with a fluence of 1.5 Mega-Jansky ms with the Survey for Transient

Radio Astronomical Emission 2 (STARE2) that is associated with the Galactic magnetar SGR J1935+2154. Based off this detection, the extragalactic population of FRBs, and the demographics of FRB hosts, FRBs can be explained by a population of extragalactic magnetars born through similar channels as the magnetars in the Milky Way.

8:00-9:00 a.m. Pacific Time, 17:00-18:00 SA Time

Session I

Session Chair: Kathryn Plant

James Smith
SARAO
jsmith@ska.ac.za

Recently we at SARAO have embarked on a course of developing a radio astronomy correlator using commodity GPUs rather than bespoke FPGA boards. This confers several advantages but also presents several challenges.
In this talk I will examine the concept of containerisation, present a short summary of its history and how it has gotten to where it is in the industry currently, including a considerations of advantages and challenges related to deploying applications in containers.
I will then present our findings for how this principle can be applied nicely in the domain of radio astronomy signal processing, both during development and deployment of real, working systems. I will then consider some of the challenges for deploying any high-performance computing application at scale, and highlight some of the challenges that we face in our specific use case. I conclude with a summary of the avenues of investigation which still remain, including the incorporation of some sort of clustered container orchestration into our correlator design.

The PANOSETI Transient Search Instrument

Dan Werthimer
University of California, Berkeley
DANW@SSL.BERKELEY.EDU

The PANOSETI experiment is an all-sky, all-the-time visible and near-IR search for nanosecond to millisecond time transients. The PANOSETI transient search will deploy domes at several sites; each site will have two domes separated by 1km. Each dome contains 80 telescopes and covers ~8000 square degrees, utilizing 80000 SiPM detectors, 1280 ASICs, and 320 FPGAs.

Data from the two observatory domes is cross-correlated data to distinguish between astrophysical transients and atmospheric phenomena (eg: cherenkov radiation). To implement cross correlation the PANOSETI instrumentation requires precision time stamping of events, accurate to a few nanoseconds. PANOSETI distributes time and frequency to each telescope over ethernet, using an open source technology called White Rabbit.

Currently, we have a prototype PANOSETI observatory at Lick Observatory; we plan to build a full scale system at Palomar observatory.

Jenkins automated testing in CBF for MeerKAT.

Ismail Jassiem
South African Radio Astronomy Observatory
ijassiem@ska.ac.za

This presentation discusses the use of Jenkins in the CBF team for automating the process of building instruments, performing integration/qualification testing, and generating qualification documentation.

Recent Updates from the Real-time RFI Mitigation System of uGMRT

Kaushal Buch
Giant Metrewave Radio Telescope, NCRA-TIFR, Pune, India
kdbuch@gmrt.ncra.tifr.res.in
Kaushal D. Buch, Mekhala V. Muley, Ruta Kale, Sanjay Kudale, Ajithkumar B.

Giant Metrewave Radio Telescope (GMRT), NCRA-TIFR, Pune, India

The Upgraded Giant Metrewave Radio Telescope (uGMRT) provides a near-seamless observing from 130 to 1450 MHz with a maximum receiver bandwidth of 400 MHz. The wider receiver bandwidth leads to an increased susceptibility to Radio Frequency Interference (RFI). The primary source of interference for uGMRT at lower frequencies (<800 MHz) is broadband RFI from gap discharge on high-voltage powerlines and associated distribution equipment. A robust threshold-based detection and filtering system has been designed using the CASPER toolflow and optimized for real-time implementation on ROACH board. The scheme works on Nyquist-sampled digital time-series per antenna and polarization and is used for mitigating broadband (time-domain impulsive) RFI from powerlines. It uses a variant of Median Absolute Deviation (MAD) called the Median-of-MAD (MoM) which is a hardware-optimized version to effectively mitigate longer bursts of RFI.

The talk would describe the filtering scheme in detail and various tests designed to qualify the system. These simultaneous comparison tests (unfiltered versus filtered) show a post-filtering improvement up to 10 dB in the signal-to-noise ratio (SNR). On astronomical testing front, continuum test observations on calibrator and extended radio sources show an improved image quality and noise RMS reduction by up to 30 percent. The incoherent array

beamformer mode show an improvement in the SNR of weak pulsars by a factor of 4. The real-time RFI mitigation system is a released feature for uGMRT observations on a shared risk basis.

A scheme for transferring the filtered data information (flags) through the system is underway. This talk would describe the ongoing efforts in this direction.

9:00 a.m.--9:15a.m. Pacific Time, 18:15--18:45 SA Time
Break

9:15a.m.--9:45a.m. Pacific Time, 18:00--18:45 SA Time

Invited talk: Bill Jenkins

Intel
bill.jenkins@intel.com
Intel PSG

9:45-10:45 a.m. Pacific Time, 18:45-19:45 SA Time

Parallel Discussions: RFI and Documentation

10:45 a.m.--11:00a.m. Pacific Time, 19:45--20:00 SA Time
Break

11:00 a.m.--11:30a.m. Pacific Time, 20:00--20:30 SA Time

More Discussion

11:00 a.m.--11:30a.m. Pacific Time, 20:00--20:30 SA Time

Conference Wrap Up