Intel PSG (Altera)
Enabling the SKA Community

Lance Brown

Sr. Strategic & Technical Marketing Mgr.

lbrown@altera.com, 719-291-7280

Agenda

- Intel Programmable Solutions Group (Altera)
- PSG's COTS Strategy for PowerMX
- High Bandwidth Memory (HBM2), Lower Power
- Case Study – CNN – FPGA vs GPU – Power Density
- Stratix 10 – Current Status, HLS, OpenCL
- NRAO Efforts

Altera == Intel Programmable Solutions Group

Intel Completes Altera Acquisition December 28, 2015

STRATEGIC RATIONALE

- **Accelerated FPGA innovation** from combined R&D scale
- **Improved FPGA power and performance** via early access and greater optimization of process node advancements
- **New, breakthrough Data Center and IoT products** harnessing combined FPGA + CPU expertise

POST-MERGER STRUCTURE

- Altera operates as a new Intel business unit called **Programmable Solutions Group (PSG)** with intact and dedicated sales and support
- **Dan McNamara** appointed Corporate VP and General Manager leading PSG, reports directly to Intel CEO

Intel Programmable Solutions Group (PSG) (Altera)

- Altera GM Promoted to Intel VP to run PSG
- Intel is adding resources to PSG
- On 14nm, 10nm & 7nm roadmap with larger Intel
- Enhancing High Performance Computing teams for OpenCL, OpenMP and Virtualization
- Access to Intel Labs Research Projects – Huge
- Will Continue ARM based System-on-Chip Arria and Stratix Product Lines

Proposed PowerMX COTS Model
NRC + CEI + Intel PSG (Altera)

Moving PowerMX to Broader Industries

Proposed PowerMX COTS Business Model

Approaching Top COTS Providers

PowerMX CEI COTS Products

- A10 PowerMX Module with HMC (Existing)
- Backplane PowerMX Module (Existing)
- S10 PowerMX Module with HMC (New)
- S10 SiP PowerMX Module (New)

RFQs

Quotes

Orders

Shipments

Support

SKA Consortium Members
Australia
Canada
China
India
Italy
New Zealand
South Africa
Sweden
The Netherlands
United Kingdom

Identifying & Working with Standards Bodies for Broader Adoption

High Bandwidth Memory

10X Improvement, Lower Power

Need for Memory Bandwidth Is Critical

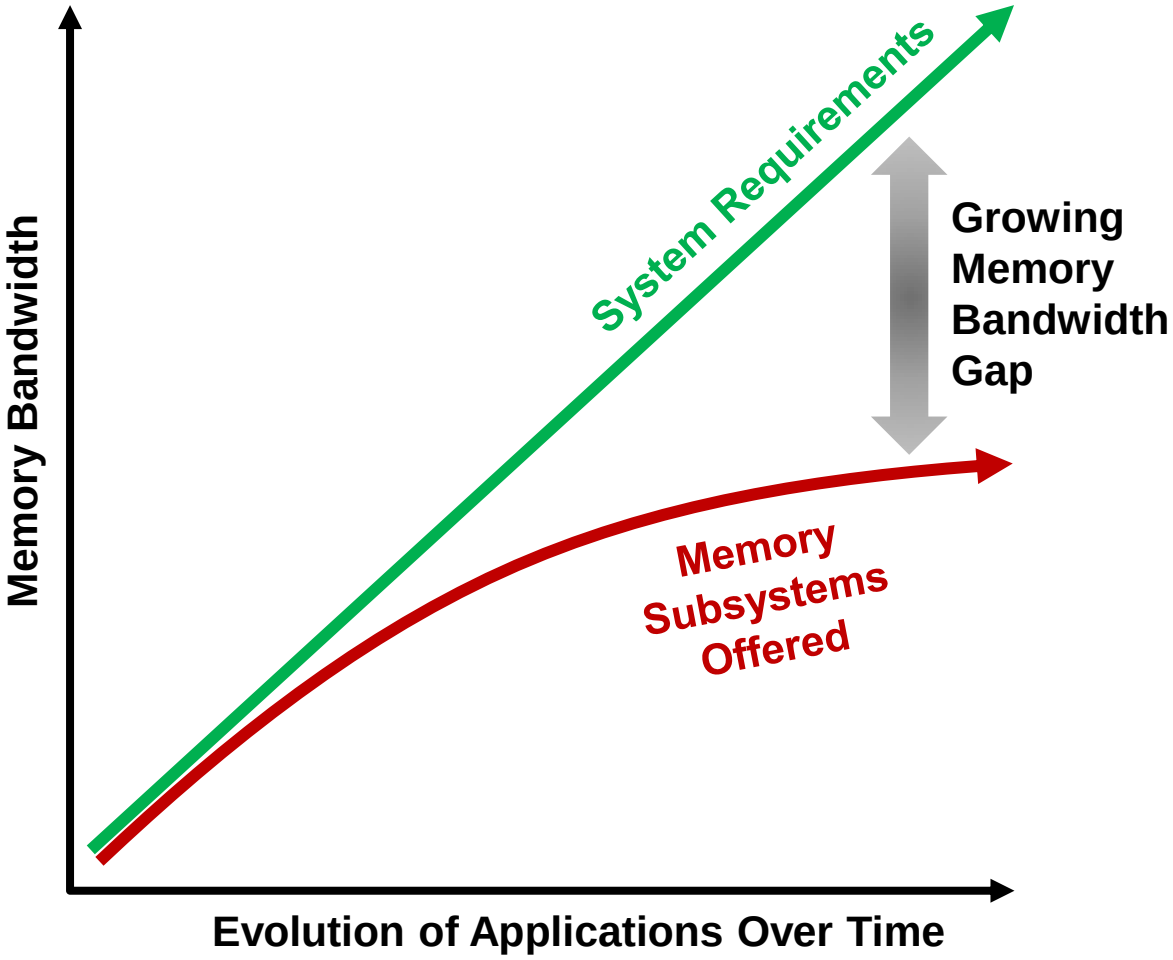

Key Challenges to Meeting Memory Bandwidth

1. End of the DDR roadmap
2. Memory bandwidth is IO limited
3. Flat system level power budgets
4. Limits to monolithic memory integration

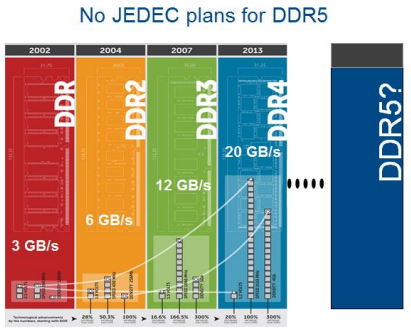

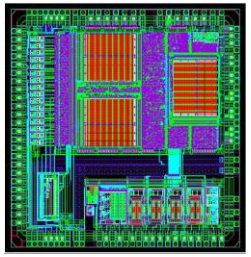

Innovation Needed to Meet High End Memory Bandwidth Requirements

“Far” Memory with Discrete DRAM

Discrete

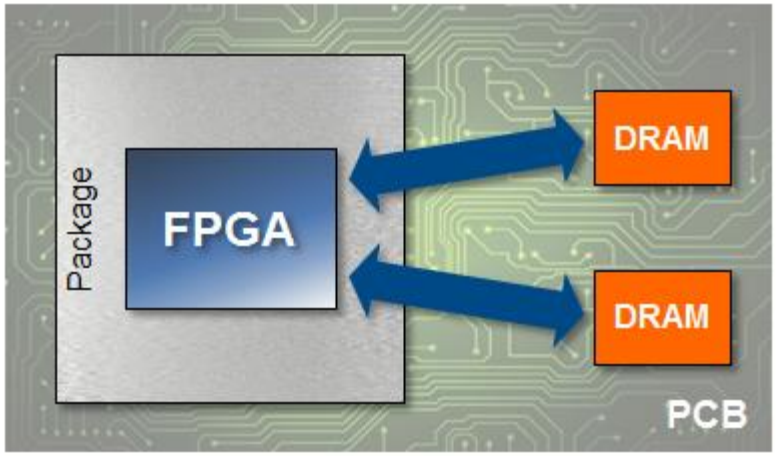

- ✘ Lower bandwidth
- ✘ Higher power
- ✘ Largest footprint

Cannot meet requirements of next-generation applications

“Near” Memory with DRAM SiP

System-In Package

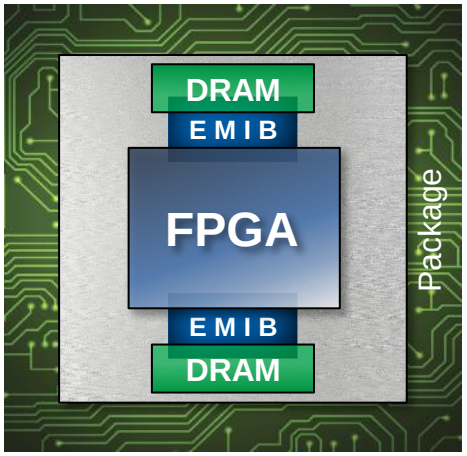

- ✓ Highest bandwidth
- ✓ Lowest power
- ✓ Smallest footprint

Meets the memory bandwidth needs of next-generation applications

Adoption of SiP Extends Beyond FPGAs

AMD

SK hynix

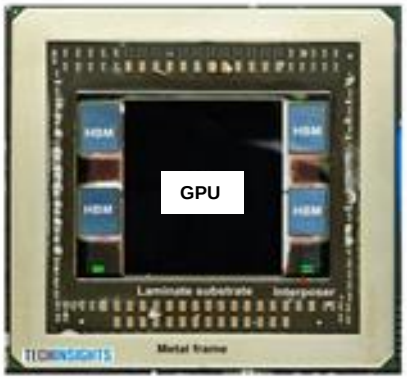

AMD Radeon

- GPU +Memory
- AMD Fiji GPU
- 4 SK Hynix HBM1

intel

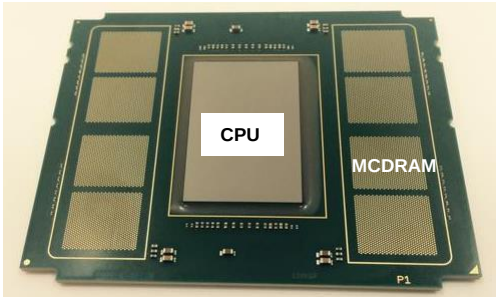

Intel Xeon

- CPU + Memory
- MCDRAM
- 16GB RAM

NVIDIA

SK hynix

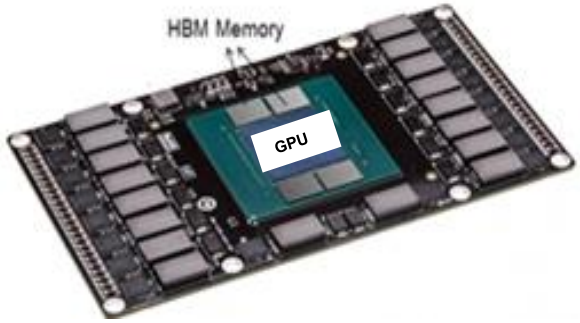

NVIDIA Volta

- GPU +Memory
- 4 HBM1
- Planned for 2016

ALTERA
now part of Intel

Stratix® 10 – Industry's Only FPGA-based DRAM SiP

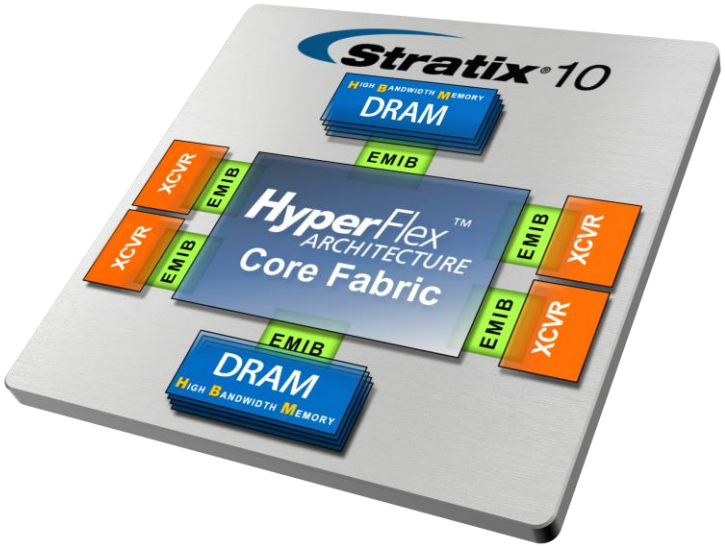

- **10X bandwidth versus discrete DRAM**
  - 256 GBytes/second per DRAM
  - Not possible with conventional solutions
- **Multiple system-level advantages**
  - Lower system power
  - Smaller form factor
  - Ease of use
- **Intel EMIB technology**

Solves the Memory Bandwidth Challenge

Radar Data Processor Application

Required Bandwidth: 400+ GB/s

FPGA + DDR4-2666

5 FPGAs

(FPGA: 42.5 mmX42.5 mm , 680 IO, 48 Transceiver)

20 DDR4 DIMMs

21 GB/s BW / DDR4 DIMM

FPGA + Hybrid Memory Cube 30G VSR

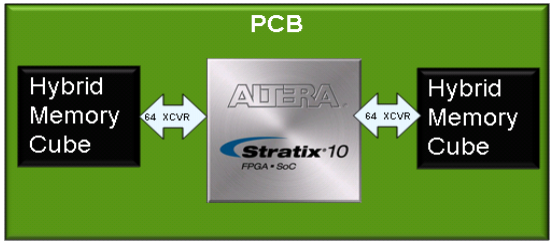

1 FPGA

(FPGA: 52.5 mmX52.5 mm , 480 IO, 144 Transceiver)

2 HMC Memories

(HMC: 31 mmx31 mm)

320 GB/s BW /HMC

Stratix 10 DRAM SiP

1 DRAM SiP

(DRAM SiP: 52.5 mm x 52.5mm)

Memory: 0 (Included)

512-1024 GB/s BW

Drawings are NOT drawn to scale

High Performance Computing PCIe Acceleration Card Application

Required Bandwidth: 1000 GB/s

FPGA + Hybrid Memory Cube 30G VSR

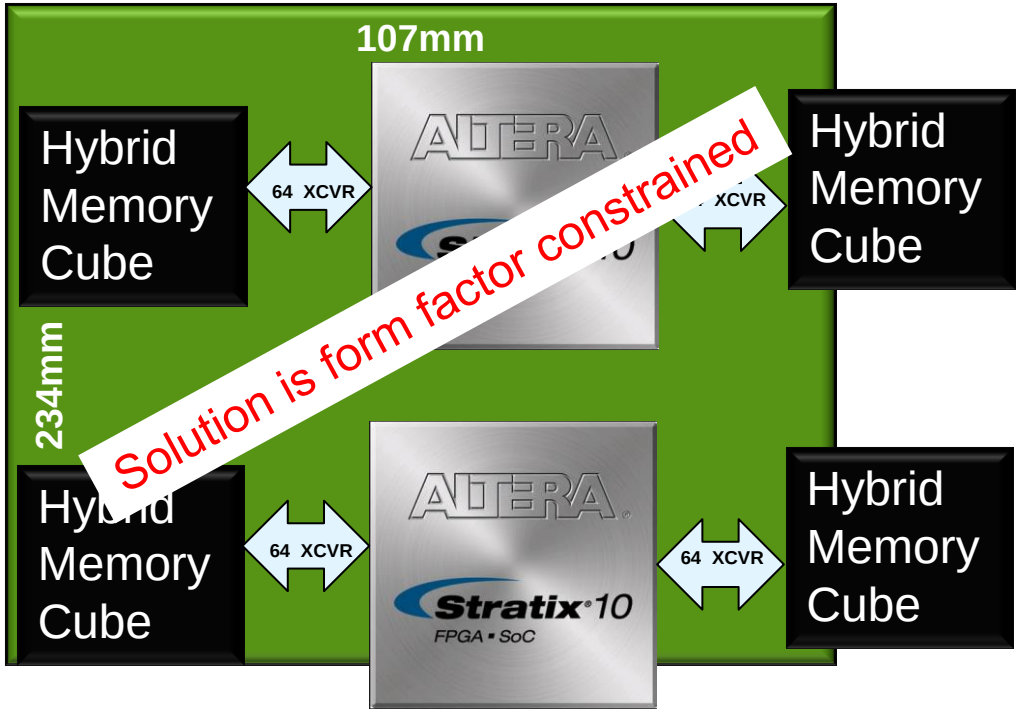

2 FPGAs, 4 HMC Memories

(FPGA: 52.5 mmX52.5 mm , 480 IO, 144 Transceiver
HMC: 31 mm X 31 mm)

320 GB/s BW /HMC

Stratix 10 DRAM SiP

1 DRAM SiP, 512 – 1024 GB/s

(DRAM SiP: 52.5 mm X 52.5 mm)

Memory: 0 (Included)

Drawings are NOT drawn to scale

ALTERA now part of Intel

Ultra HD 8K Viewer (8KP120) Application

Required Bandwidth: 431.52 GB/s

FPGA + Hybrid Memory Cube 30G VSR

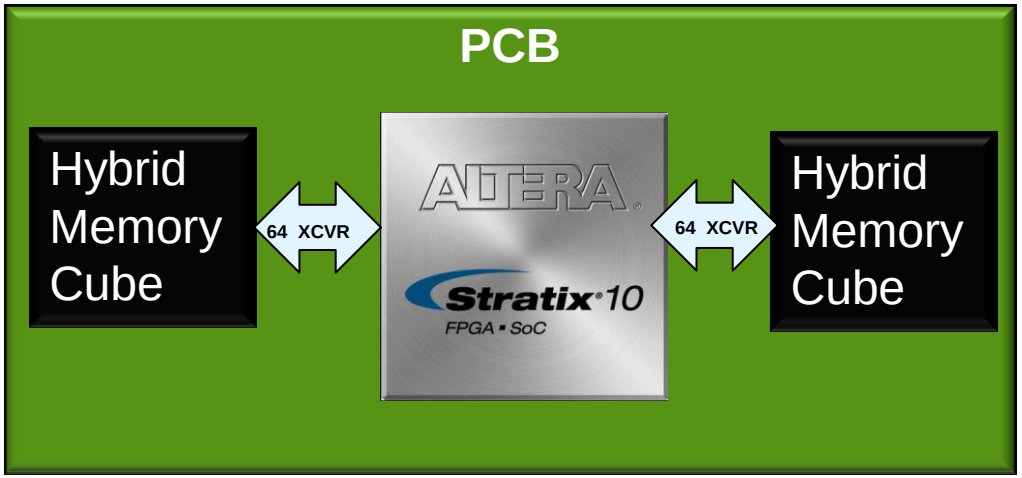

1 FPGA

(FPGA: 52.5 mmX52.5 mm , 480 IO, 144 Transceiver)

2 HMC Memories

(HMC: 31 mmx31 mm)

320 GB/s BW / HMC

Stratix 10 DRAM SiP

1 DRAM SiP

(DRAM SiP: 52.5 mm x 52.5 mm)

Memory: 0 (Included)

512-1024 GB/s BW

Drawings are NOT drawn to scale

Ultra HD 8K High End Camera Application

Required Bandwidth: 172 GB/s

- Form factor critical application

NHK 8K camera
in 10cm housing

Hybrid Memory Cube 30G VSR

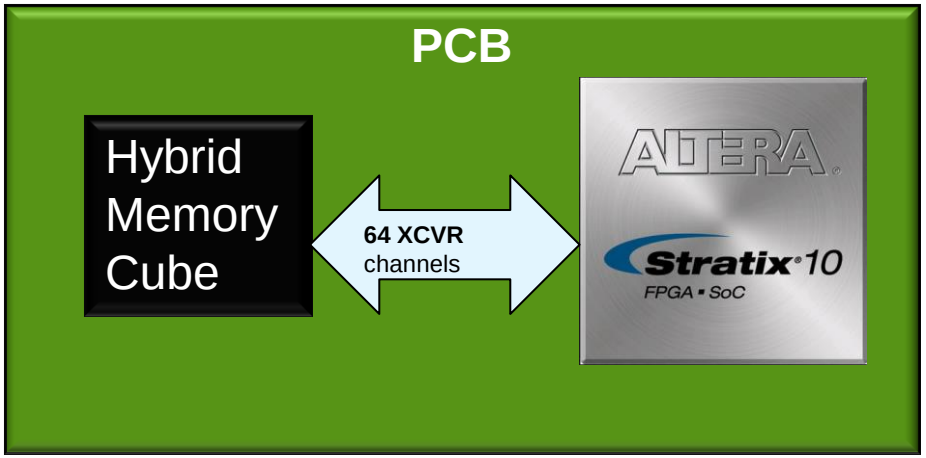

1 FPGA

(FPGA: 52.5 mmX52.5 mm , 480 IO, 144 Transceiver)

1 HMC Memory

(HMC: 31 mmx31 mm)

320 GB/s BW

Stratix 10 DRAM SiP

1 DRAM SiP

(DRAM SiP: 52.5 mm x 52.5 mm)

Memory: 0 (Included)

512-1024 GB/s BW

~50%+ Board space savings

Drawings are NOT drawn to scale

Summary – Stratix® 10 DRAM SiP

- Solves the memory bandwidth challenge
- Provides 10X memory bandwidth versus discrete solutions
- Enabled by innovative 3D SiP, EMIB and next-generation Stratix 10 FPGA technologies
- Enables key applications including wireline, broadcast, military, HPC, test, and more

CNN Case Study
FPGA vs GPU

Performance & Power

Deep Learning:
Convolutional Neural Network (CNN)

- **Convolutional Neural Network**
  - Feed forward artificial neural network
  - Inspired by biological processes
- **Applications**
  - Classification & Detection
  - Image recognition/tagging
  - Big Data pattern detection
  - Datastreaming Analytics
  - Targeted ads
  - Fraud & face detection
  - Gaming

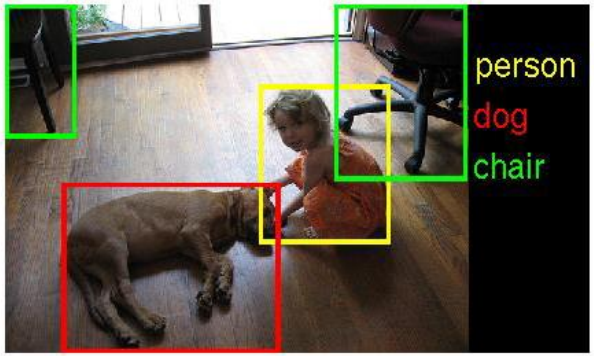

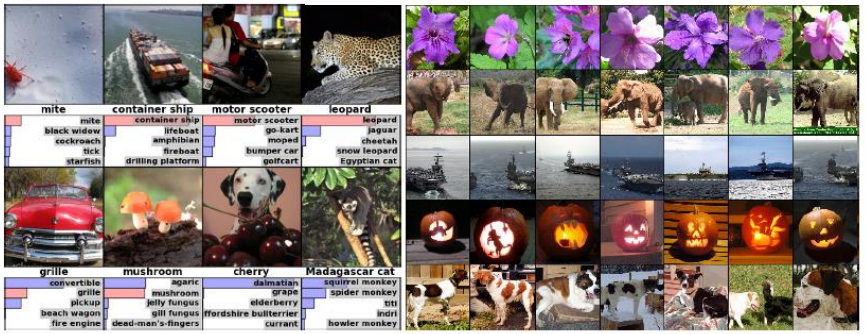

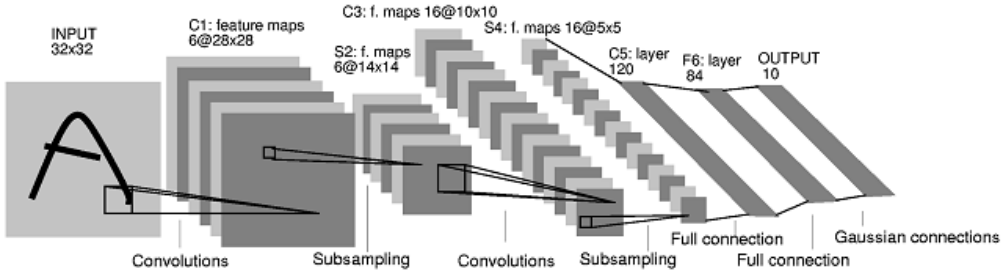

ALTERA
now part of Intel

ImageNet

- ImageNet is a yearly competition held since 2010 – Large Scale Visual Recognition Challenge.
  - 1.2 million images for training, 50,000 for validation, 100,000 for testing.
  - 1000 different image classes.

(a) Siberian husky

(b) Eskimo dog

  - 2012 Winner: **AlexNet**, top-5 error rate of 15.3%, 5 convolution layers
  - 2014 Winner: **GoogleNet**, top-5 error rate of 6.67%, 22 layers in total
  - 2015 Winner: **Microsoft** with 4.94% (Baidu was 4.8%, but was disqualified)
  - Top-5 Error Rate: How often is the correct answer not in the top-5 results?
    - Trained human result: 5.1% Top-5 Error Rate, at *1 minute per image*

AlexNet Competitive Analysis – Classification

System	Throughput	Est. Power	Throughput / Watt
Arria 10-1150	600 img/s	~60W	10 img/s/W
2 x Arria 10-1150	**1200 img/s**	~90W	**13.3 img/s/W**
Caffe on NVIDIA TitanX with CUDA	1000 img/s	~250W	4 img/s/W

- FPGA competitive with GPUs on raw AlexNet throughput
  - Dominates in throughput per watt, for similar node GPU
- Expect similar ratios for Stratix 10 vs. NVIDIA 14nm Pascal

Stratix 10 Updates

Schedule

- Near the front of the line with direct Intel mfg support
- Next Update on Feb 17th
- S10 Common Element (S10CE) is meant to help with early S10 work
- Early Power Estimator (EPE) updates frequently
- Certify Partners in all GEOs on HyperFlex – 2 Week Course

S10CE - Preliminary

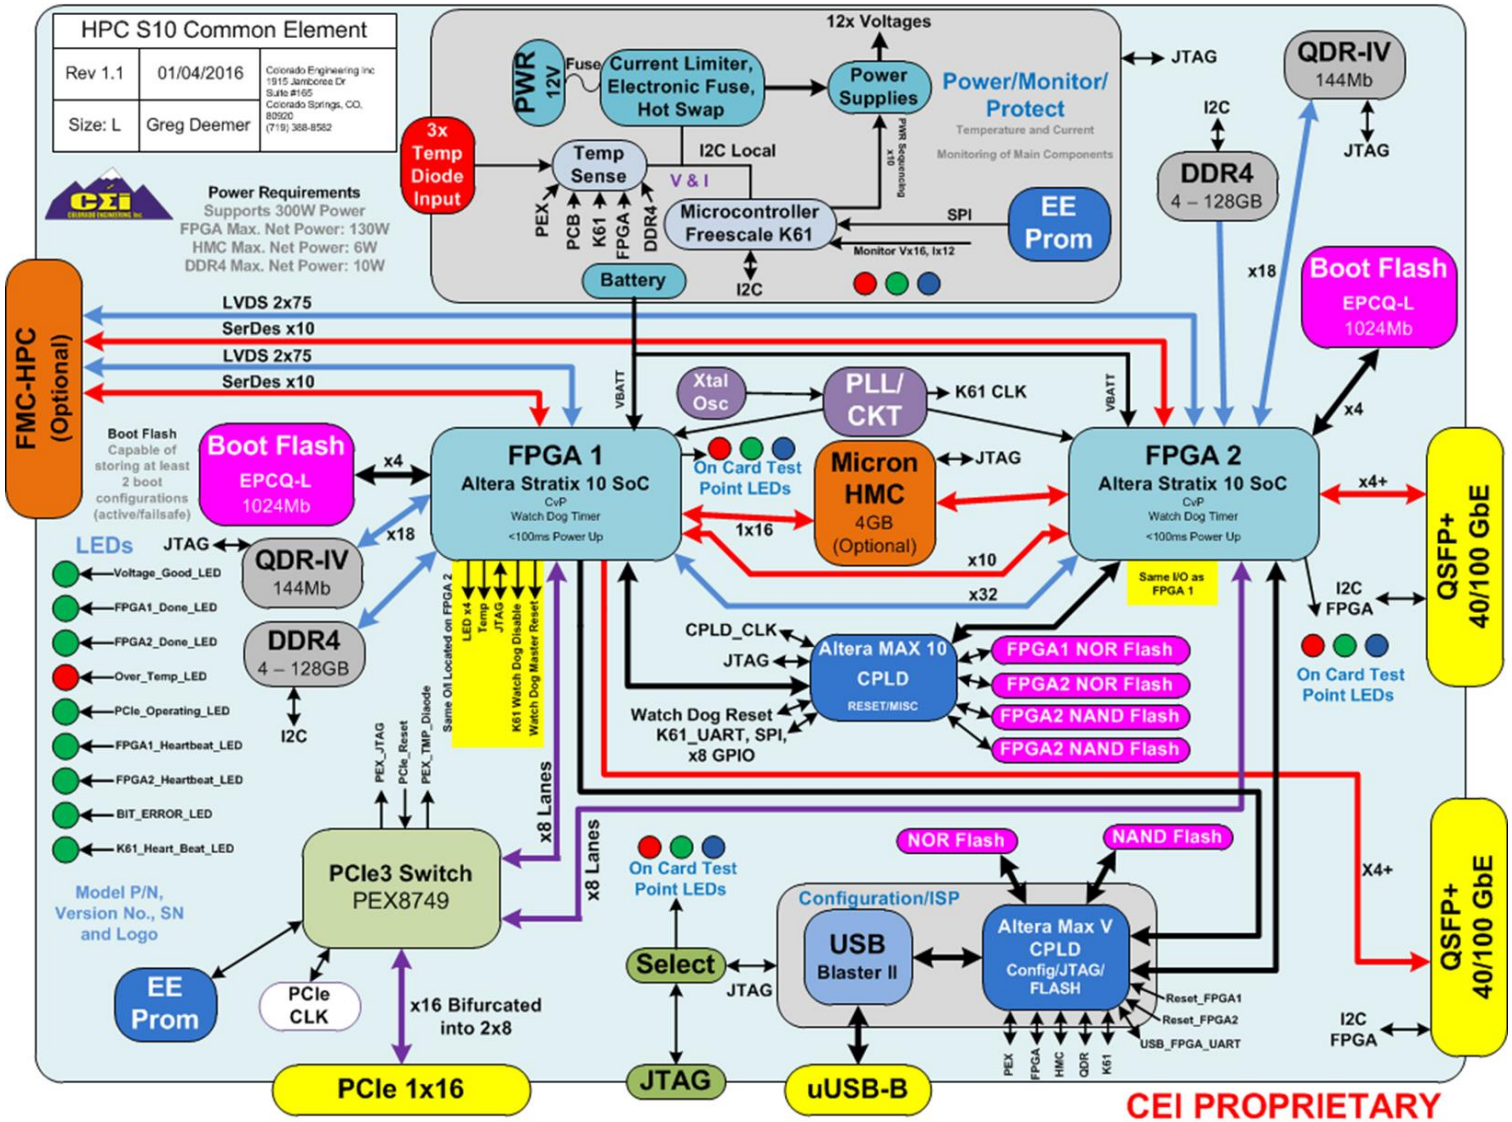

Up to 70% Reduction in Power Consumption

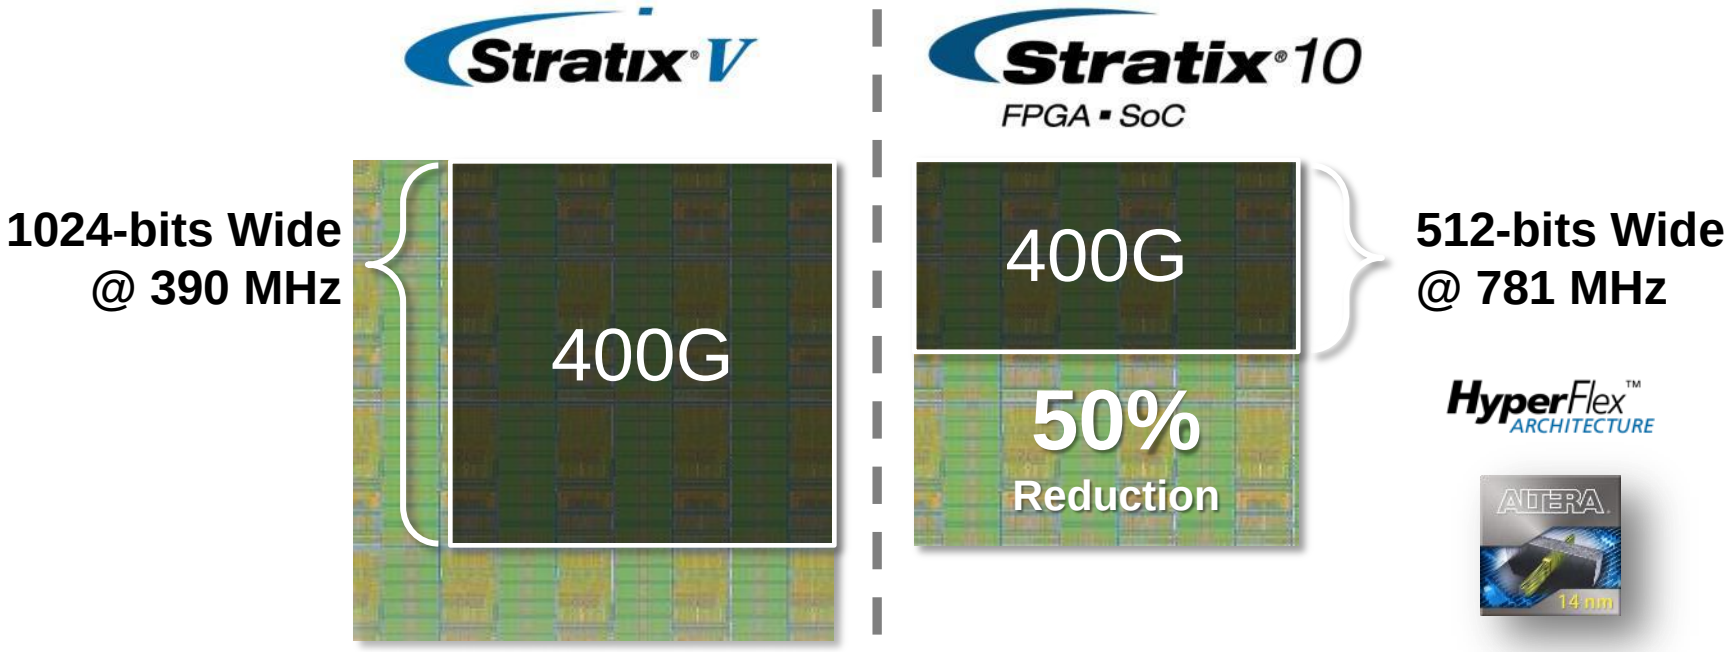

Customer Designs	Power Savings Stratix V vs. Stratix 10
Wired Network	**40%** Less
Wireless Network	**52%** Less
Data Center Server	**63%** Less
Enterprise Storage	**70%** Less

Enables Higher FPGA Adoption in the Data Center

120W vs. 44W

Stratix 10 FPGA ▪ SoC

HyperFlex™ ARCHITECTURE

- **Five** Stratix V FPGAs
- PCIe Gen2 x8
- DDR3 x72 @ 800 MHz
- FPGA performance – 250 MHz

- **One** Stratix 10 FPGA
- PCIe Gen3 x8
- DDR3 x144 @ 1.2 GHz
- FPGA performance – 500 MHz

HyperFlex Delivers

- 2X core performance of FPGA
- 63% power reduction

OpenCL vs a++ Compiler Summary

A++ Compiler
for HLS

Targets CPU, GPU and FPGAs	→	Targets FPGAs
Target user is HW or SW	→	Target user is HW
Implements FPGA in software development flow	→	Implements FPGA in traditional FPGA development flow
Performance is determined by resources allocated	→	Performance is defined and amount of resource to achieve is reported
Builds the entire FPGA system	→	Builds an IP block
Host Required	→	Host is optional

Altera Efforts for NRAO: Gridding + Image Deconvolution Updates

Project Potential

- Gridding and Image Deconvolutions are the current **bottlenecks** in post processing
  - Single Image Data: 100GB-400TB (Double Precision FP)
  - 100 CPUs take 1 day to process data. 1 CPU takes 10 days.
  - Final Image Resolution: ~1k*1k – 15k*15k pixels
  - All processing done offline in batch
- Potentially scalable solutions using Altera FPGA for algorithm acceleration in conjunction to NRAO SW toolkit CASA
  - Higher data processing needs are a reality

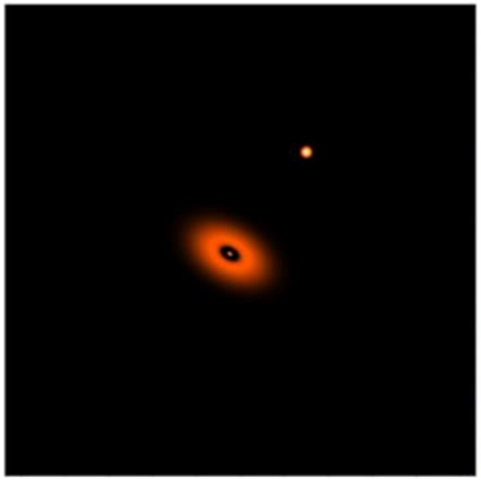

Technical Overview

How do we analyze the data?

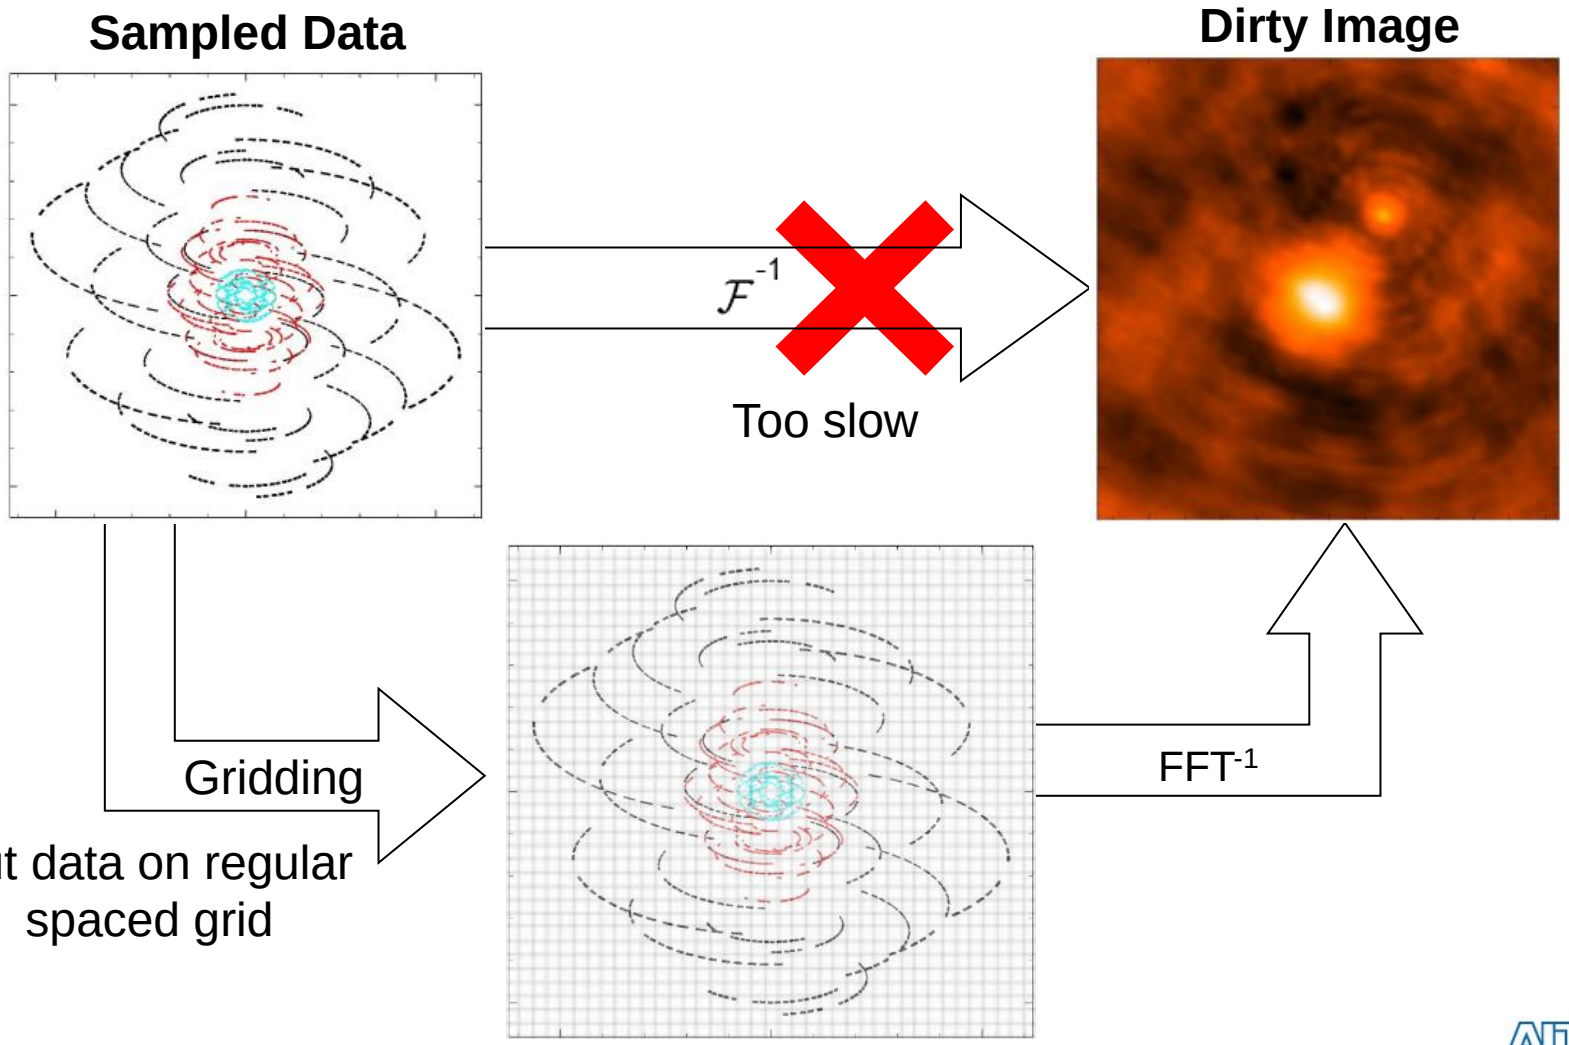

Technical Overview

How can we make the data look better?

Dirty Beam * **True Sky Image** = **Dirty Image**

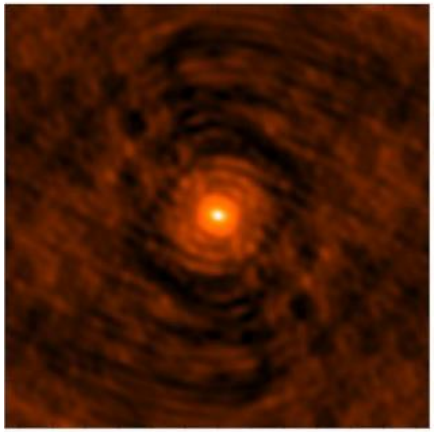

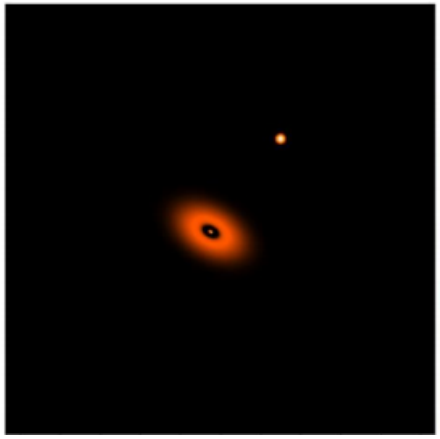

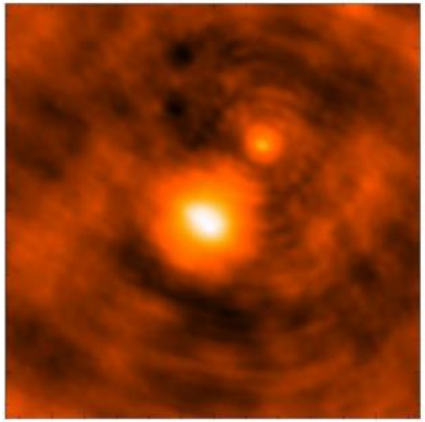

$\mathcal{F}\uparrow$

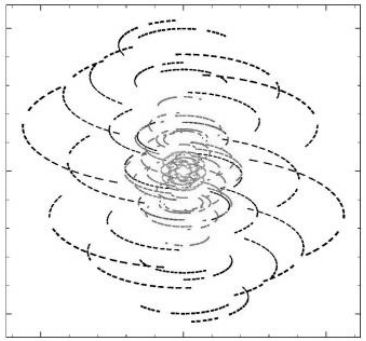

(u,v Coordinates)

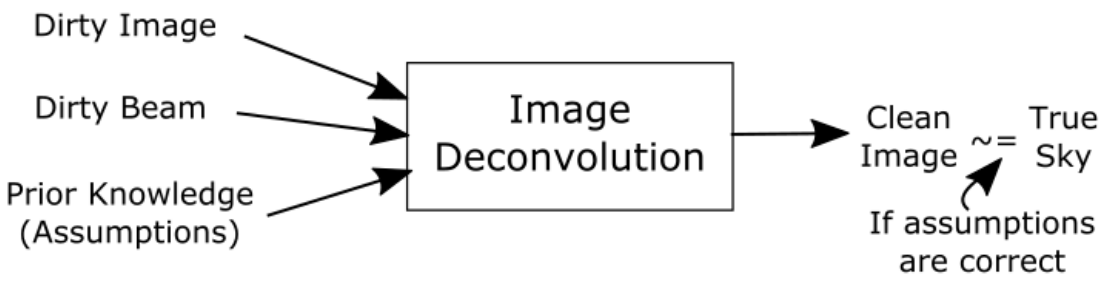

Project Scope

- Implementations of Gridding algorithms using OpenCL should provide 10x-20x performance boost
  - Altera working with partner ImpulseAccelerated for benchmarking source code
  - Need to test on ReFLEX Arria10 Dev kit
  - Testing planned on Arria10 or Stratix10 (due to Floating Point DSPs)
- Deploy and Test implementation within NRAO's development environment
- Develop and implement parallel algorithm for Image Deconvolution

Multithreaded gridding algorithms tested

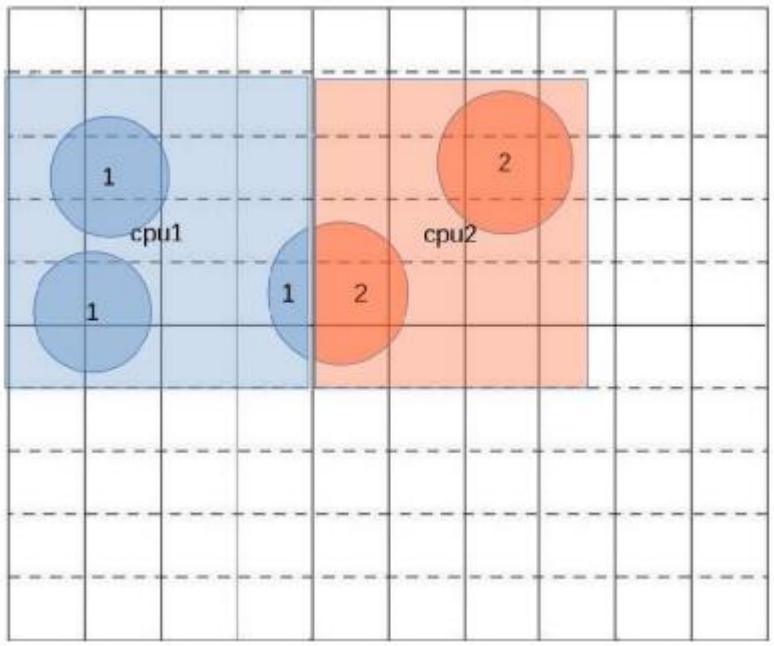

Kumar's

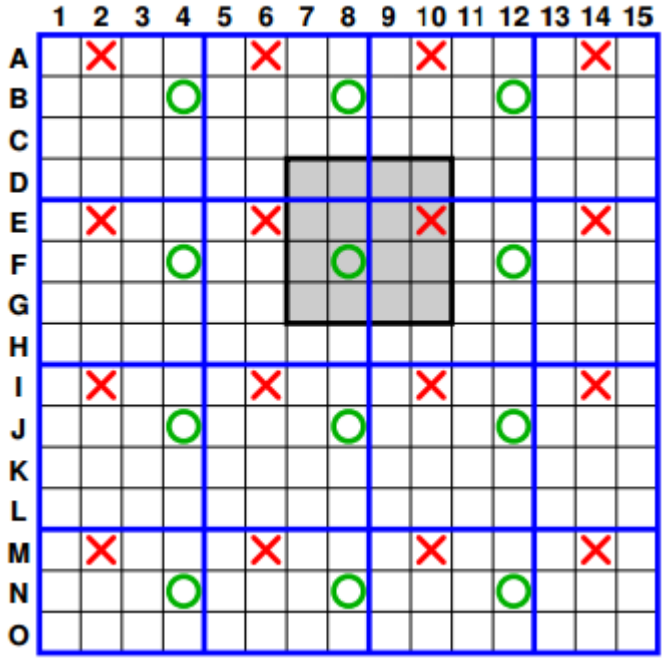

Romein's

Implementations/Progress

- Single-Threaded
  - Ran on S5, integrated with CASA. Slow- no parallelization
- Multi-Threaded (Kumar Golap - NRAO)
  - Load balancing – soln. partitioning, W-only projection
  - Ported from FORTRAN to C & integrated with CASA
- Multi-Threaded (John Romein - ASTRON)
  - Sorts input data on host to increase locality
  - Trying to figure out if IDG is based off Romein's alg
- IDG (Bram Veenboer – ASTRON)
  - Standalone (not sure if practical solution)
  - OpenCL implementation exists (from ImpulseAccelerated)
  - Still need results on A10 board
- All
  - Not double-precision floating point numbers

Next Steps/Learnings

- Need to target multithreaded solutions
- Figure out which implementation is best-suited for FPGAs
- Simplify CASA integration by creating a more standalone application – this will allow faster development
- Implement the best (or more) implementations on AOCL using standalone app
  - Requires data re-organization if using Romein/IDG
- Profile timing and resources used
- Impulseaccelerated benchmarking projections:

Platform	Runtime (s)	Clock	M Vis/Sec	Watts	K Vis/Watt	
Stratix V - 5SGSMD8K	12.3	202 MHz	2.2	49	46	include
Intel I7-4970K @ 4.0GHz	15.3	4.0 GHz	1.8	120	15	no dat with In
GTX980 GPU	0.45	1.2 GHz	61.2	165	371	include
Arria 10	2.05	400 MHz	13.4	50	269	estima
Stratix 10	0.41	600 MHz	67.2	70	960	estima

Thank You