High Performance Processing with MONARCH – A Case Study in CT Reconstruction

Kenneth Prager
Raytheon Company
2000 E. El Segundo Blvd
El Segundo, CA 90245
keprager@raytheon.com

David Rohler
Multi-Dimensional Imaging
31300 Bainbridge Rd.
Solon, OH 44139
rohler@mdivac.com

Pat Marek
Raytheon Company
2000 E. El Segundo Blvd
El Segundo, CA 90245
psmarek@raytheon.com

Arthur Dartt
Multi-Dimensional Imaging
31300 Bainbridge Rd.
Solon, OH 44139
dartt@mdivac.com

Lloyd Lewins
Raytheon Company
2000 E. El Segundo Blvd
El Segundo, CA 90245
llewins@raytheon.com

High Performance Embedded Computing (HPEC) Workshop

23−25 September 2008

Problem Statement

- Develop an embedded system for mobile, real-time reconstruction of cone beam CT imagery
- System must provide 1 – 2 TFLOP/S of processing throughput while only consuming around a kilowatt of power
- Spiral cone beam reconstruction requirements:
  - 1.125 TFLOP/S of throughput
  - 100 GB/S Read/modify/write bulk memory access

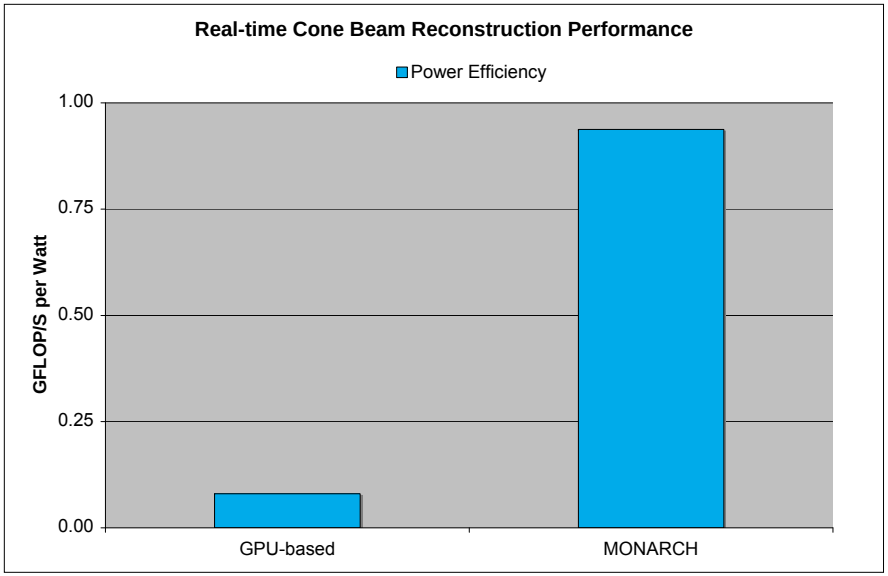

Compact VAC Overview

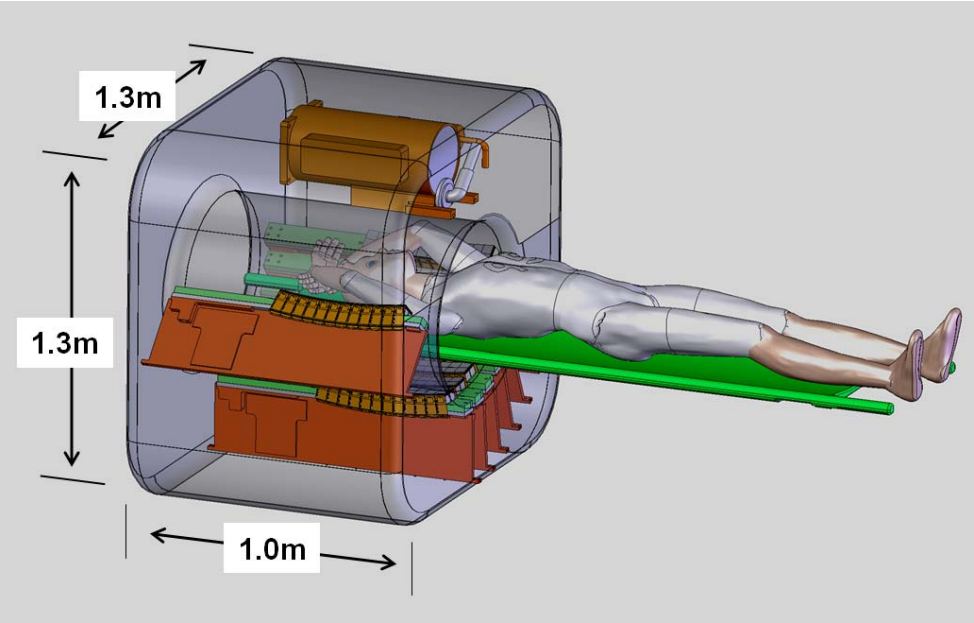

- Compact VAC (Volume AngioCT) is an advanced Computed Tomography (CT) system proposed for development by DARPA to allow combat casualty assessment in forward battlefield positions
- Size, weight and power enable mobility and portability suitable for use in multiple vehicles and enclosures for forward and rear military applications

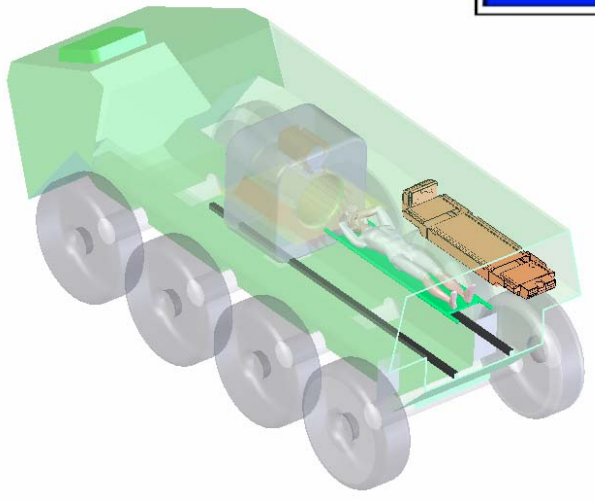

Example transport vehicle - Stryker

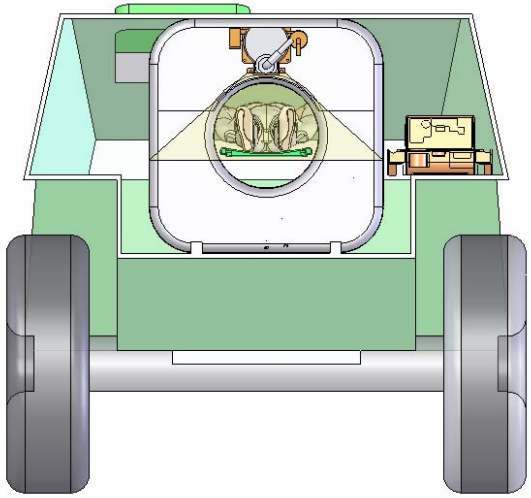

VAC Algorithm Overview

- CT cone-beam reconstruction is a computational algorithm that transforms a sequence of raw views into volume images.
- Volume images are typically organized as a “stack” of cross-sectional 2D slice images.
- Reconstruction consists of two components:
  - cone-beam filter process
    - Each view is independently filtered
    - No storage required
    - Low latency
  - cone-beam backprojection process
    - each reconstruction slice has contributions from a large range of views
    - each view impacts a large number of slices.
    - The range of slices and/or views is too large to fit into cache or local memory
    - Algorithms must be designed so that slice data and/or view data constantly flow in and out of local memory

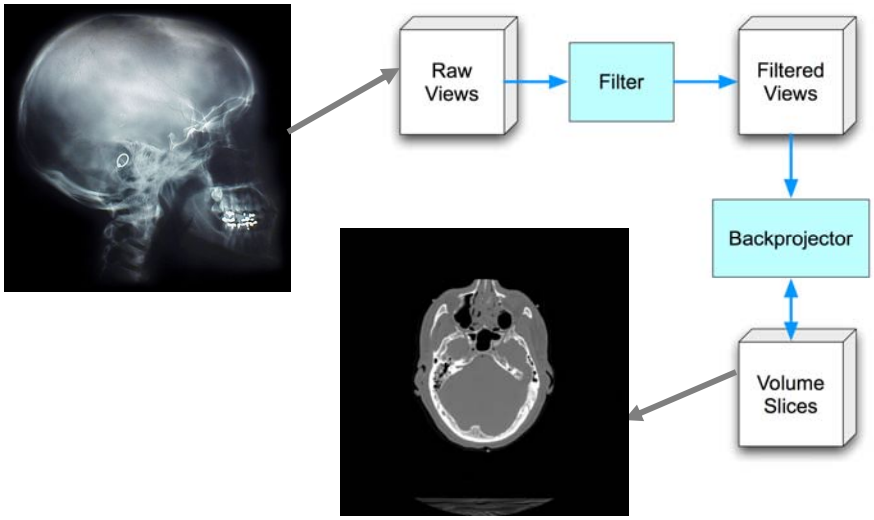

- Input: Raw Views
  - 384 x 896 (768 x 896 after filtering)
  - 5270 total views
  - 6 second collection time (1.1 mS per view)
- Output: Volume Slices
  - 512 x 512
  - 1800 total slices
  - 6 – 60 second reconstruction time (3.3 – 33 mS per slice)

Backprojector: Mapping Views to Slices

QuickTime™ and a
Apple Intermediate Codec decompressor
are needed to see this picture.

- Approximately 550 view images are needed to construct a slice image.
- A new slice is started every 3 views.
- Additionally, each view image is used by approximately 220 slice images.
- Therefore, all 220 slices are held in memory for concurrent processing.
- As a slice is completed, its resources are freed, and a new slice is started in its place.

Slice Number	View Numbers Needed for Reconstruction
1	1-550
2	4-553
3	7-556
□□□	□□□

- Basic concept…

```
for each new view
   for all active slices
      read slice portion
      for each active pixel in slice
         interpolate view data pixel
         update slice pixel
      end
      write slice portion
   end
end
```

Backprojector: Mapping Views to Slices

QuickTime™ and a
Apple Intermediate Codec decompressor
are needed to see this picture.

- Approximately 550 view images are needed to construct a slice image.
- A new slice is started every 3 views.
- Additionally, each view image is used by approximately 220 slice images.
- Therefore, all 220 slices are held in memory for concurrent processing.
- As a slice is completed, its resources are freed, and a new slice is started in its place.

Slice Number	View Numbers Needed for Reconstruction
1	1-550
2	4-553
3	7-556
□□□	□□□

- Basic concept…

```
for each new view
   for all active slices
      read slice portion
      for each active pixel in slice
         interpolate view data pixel
         update slice pixel
      end
      write slice portion
   end
end
```

Backprojector: Mapping Views to Slices

QuickTime™ and a
Apple Intermediate Codec decompressor
are needed to see this picture.

- Approximately 550 view images are needed to construct a slice image.
- A new slice is started every 3 views.
- Additionally, each view image is used by approximately 220 slice images.
- Therefore, all 220 slices are held in memory for concurrent processing.
- As a slice is completed, its resources are freed, and a new slice is started in its place.

Slice Number	View Numbers Needed for Reconstruction
1	1-550
2	4-553
3	7-556
□□□	□□□

- Basic concept…

```
for each new view
   for all active slices
      read slice portion
      for each active pixel in slice
         interpolate view data pixel
         update slice pixel
      end
      write slice portion
   end
end
```

What is MONARCH?

MOrphable Networked micro-ARCHitecture

- **Single, programmable chip type which can replace custom ASICs**
  - ***Saves $20 – 30M development cost and 18 – 36 months of development schedule per ASIC***
- **Power efficiency equivalent to 10 Pentiums[1]**
- **Power and performance similar to custom ASICs > 3 – 6 GFLOPS/W**
- **Programmable and Scaleable**
- **Balanced I/O & Processing uniquely suited to Signal Processing needs**

System-on-a-Chip: Standalone solution for embedded applications

Enabler for ultra small UAVs, hand-helds, munitions, etc.

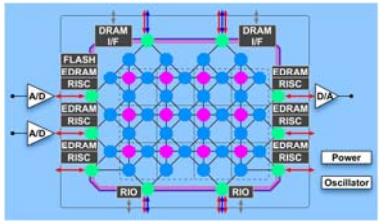

Tiled Array: TFLOP performance in a self-contained network

Enabler for embedded advanced, adaptive signal processing

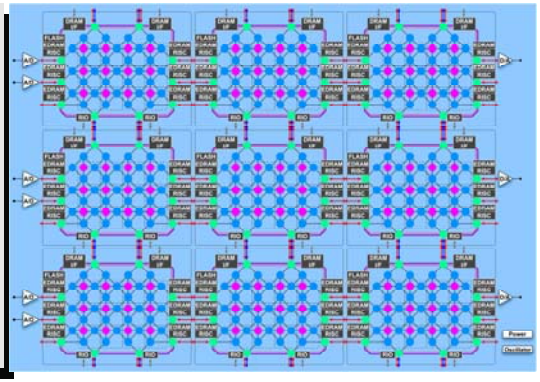

1. Based on FFT benchmark comparing 333 MHz MONARCH to 2.3 GHz Intel x86 quad-core Harpertown

MONARCH Chip Overview

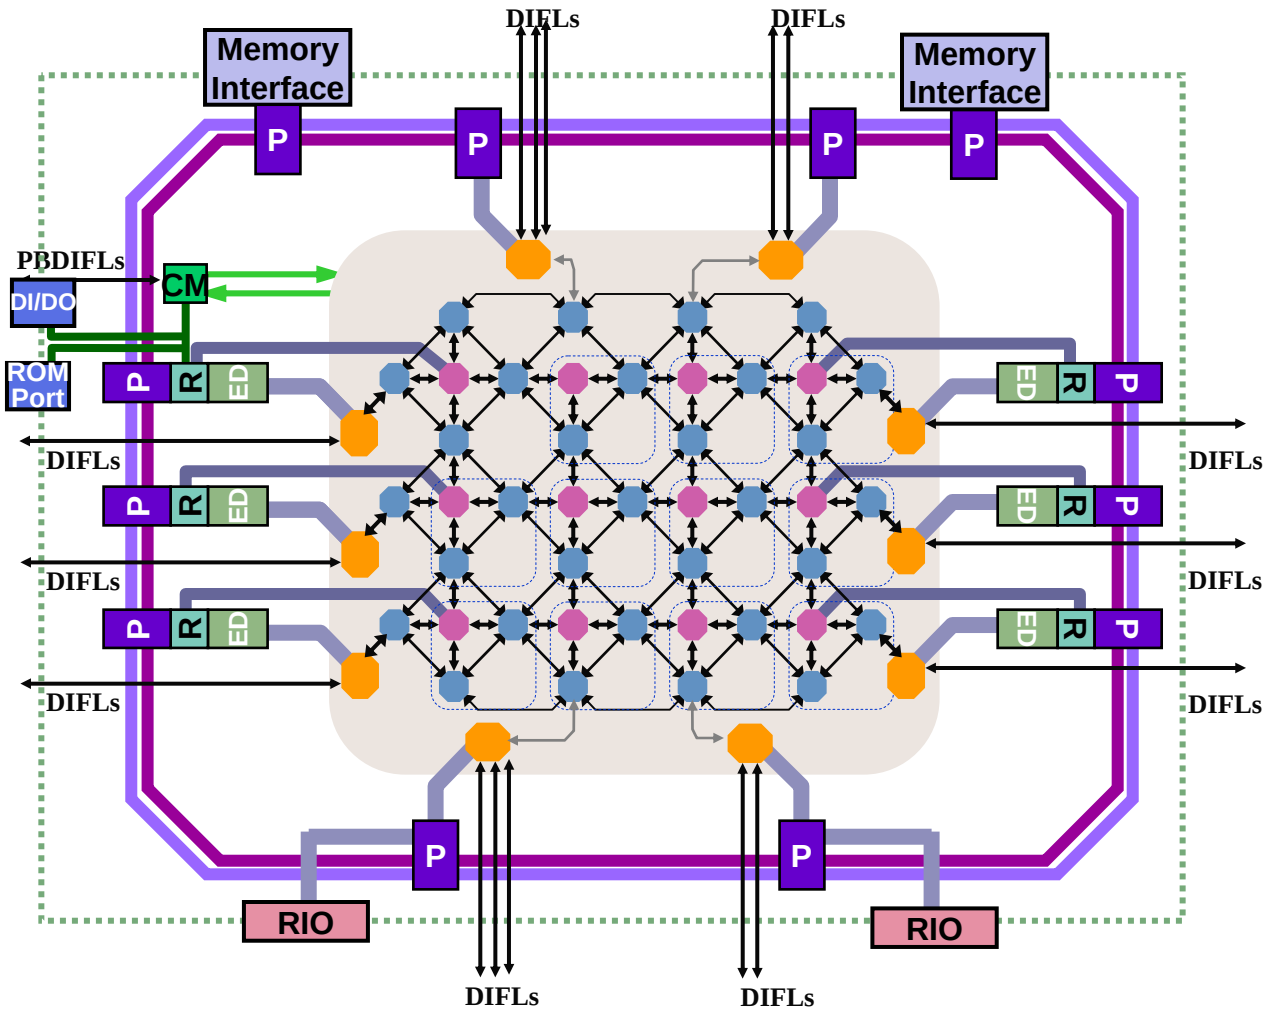

- 6 RISC Processors
- 12 MBytes on-chip DRAM
- 2 DDR2 External Memory Interfaces (8 GB/s BW)
- Flash Port (32 MB)
- 2 Serial RapidIO Ports (1.25 GB/s each)
- 16 IFL ports (2.6 GB/s each)
- On-chip Ring 40 GB/s
- Reconfigurable Array: FPCA (64 GFLOPS)

MONARCH Building Blocks

- Module Specs
  - 256 GFLOP/s throughput
  - 48 Mbytes EDRAM
  - 8 GBytes DDR2-DRAM
  - 381 Gbits/sec I/O rate
    - IFL & SRIO combined
  - 80 – 140 Watts
  - VME/VPX form factor

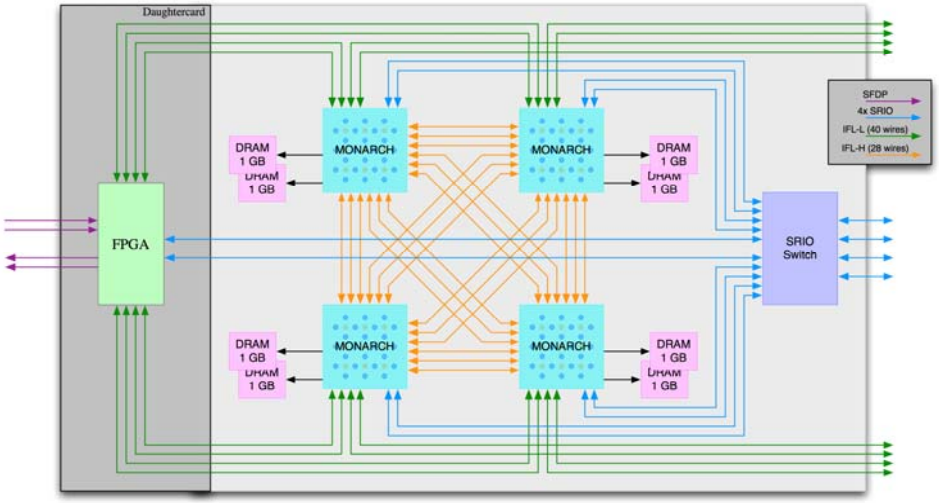

- Chassis Specs
  - 2.3 TFLOP/s throughput
  - 432 Mbytes EDRAM
  - 72 GBytes DDR2-DRAM
  - Flexible I/O options
  - 720 – 1260 Watts
  - VME/VPX form factor
    - Conduction-cooled

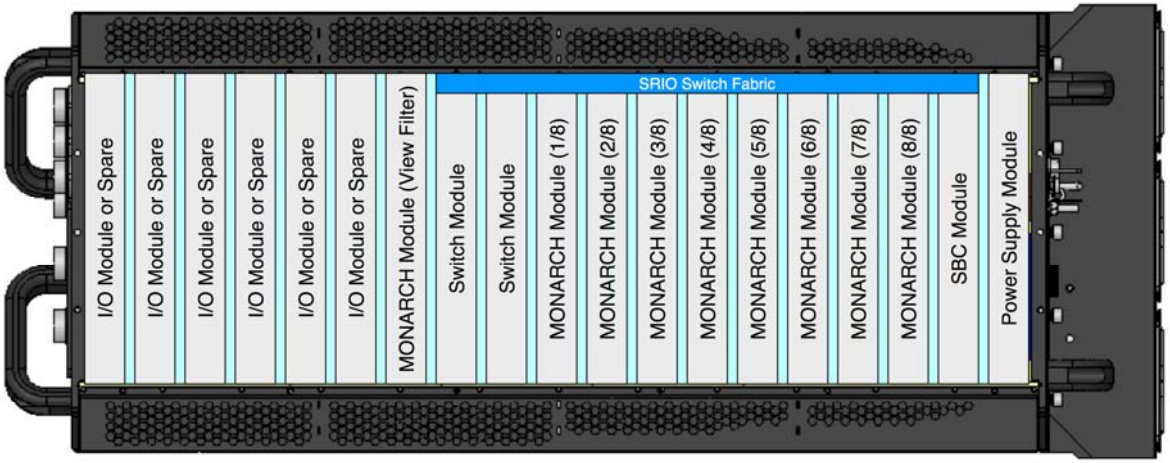

Mapping Checklist

- Examine the system I/O requirements
- Examine the system DDR2-DRAM requirements
  - Capacity and bandwidth
- Examine the system EDRAM requirements
  - Capacity and bandwidth
- FPCA Processing operations
  - Estimate the ops/channel for each function (FIR, FFT)
  - Calculate the total ops per unit time…

$$\frac{\text{ops/sample} \cdot \#\text{samples}}{\text{Sample Time}}$$

  - Calculate the number of MALU elements needed to compute the required ops, given the clock speed
- FPCA Memory operations
  - Estimate the memory for each function (FIR, FFT)
  - Determine the number of data transfers per data clock for each memory type
- System layout
  - Calculate the number of chips required to process all inputs given the number of math and memory clusters available for the required operations
  - Calculate the number of DDR2 and EDRAM elements required
  - Layout system based on stressing requirements

Processing Key Requirements

- View Filter – Receive 5270, 384 x 896 view images in 6 seconds; store, process, and distribute them over 12 seconds
  - Input data rate: (5270 x 384 x 896 x 2) / 6 = 605 MB/S
  - Input storage: (5270 x 384 x 896 x 2) = 3.6 GB
  - Output data rate: (5270 x 768 x 896 x 2) / 12 = 605 MB/S

- Backprojection – Receive 5270, 768 x 896 views in 12 seconds; Read 220 512 x 512 slices per view, process, and write back to DRAM, over 12 seconds
  - Total slice data rate: (220 x 5270 x 512 x 512 x 4) / 12 = 102 GB/S
  - Number DRAM ports needed: 102 GB/S / 3.56 GB/S = 29 ⇒ 29 chips
  - Slice storage per DRAM port: (220 x 512 x 512 x 4) = 231 MB
  - FPCA ports per DRAM port: 3.56 GB/S / 1.33 GB/S = 3
  - View data storage (EDRAM): (768 x 896 x 2) = 1.4 MB

Functional Block Diagram

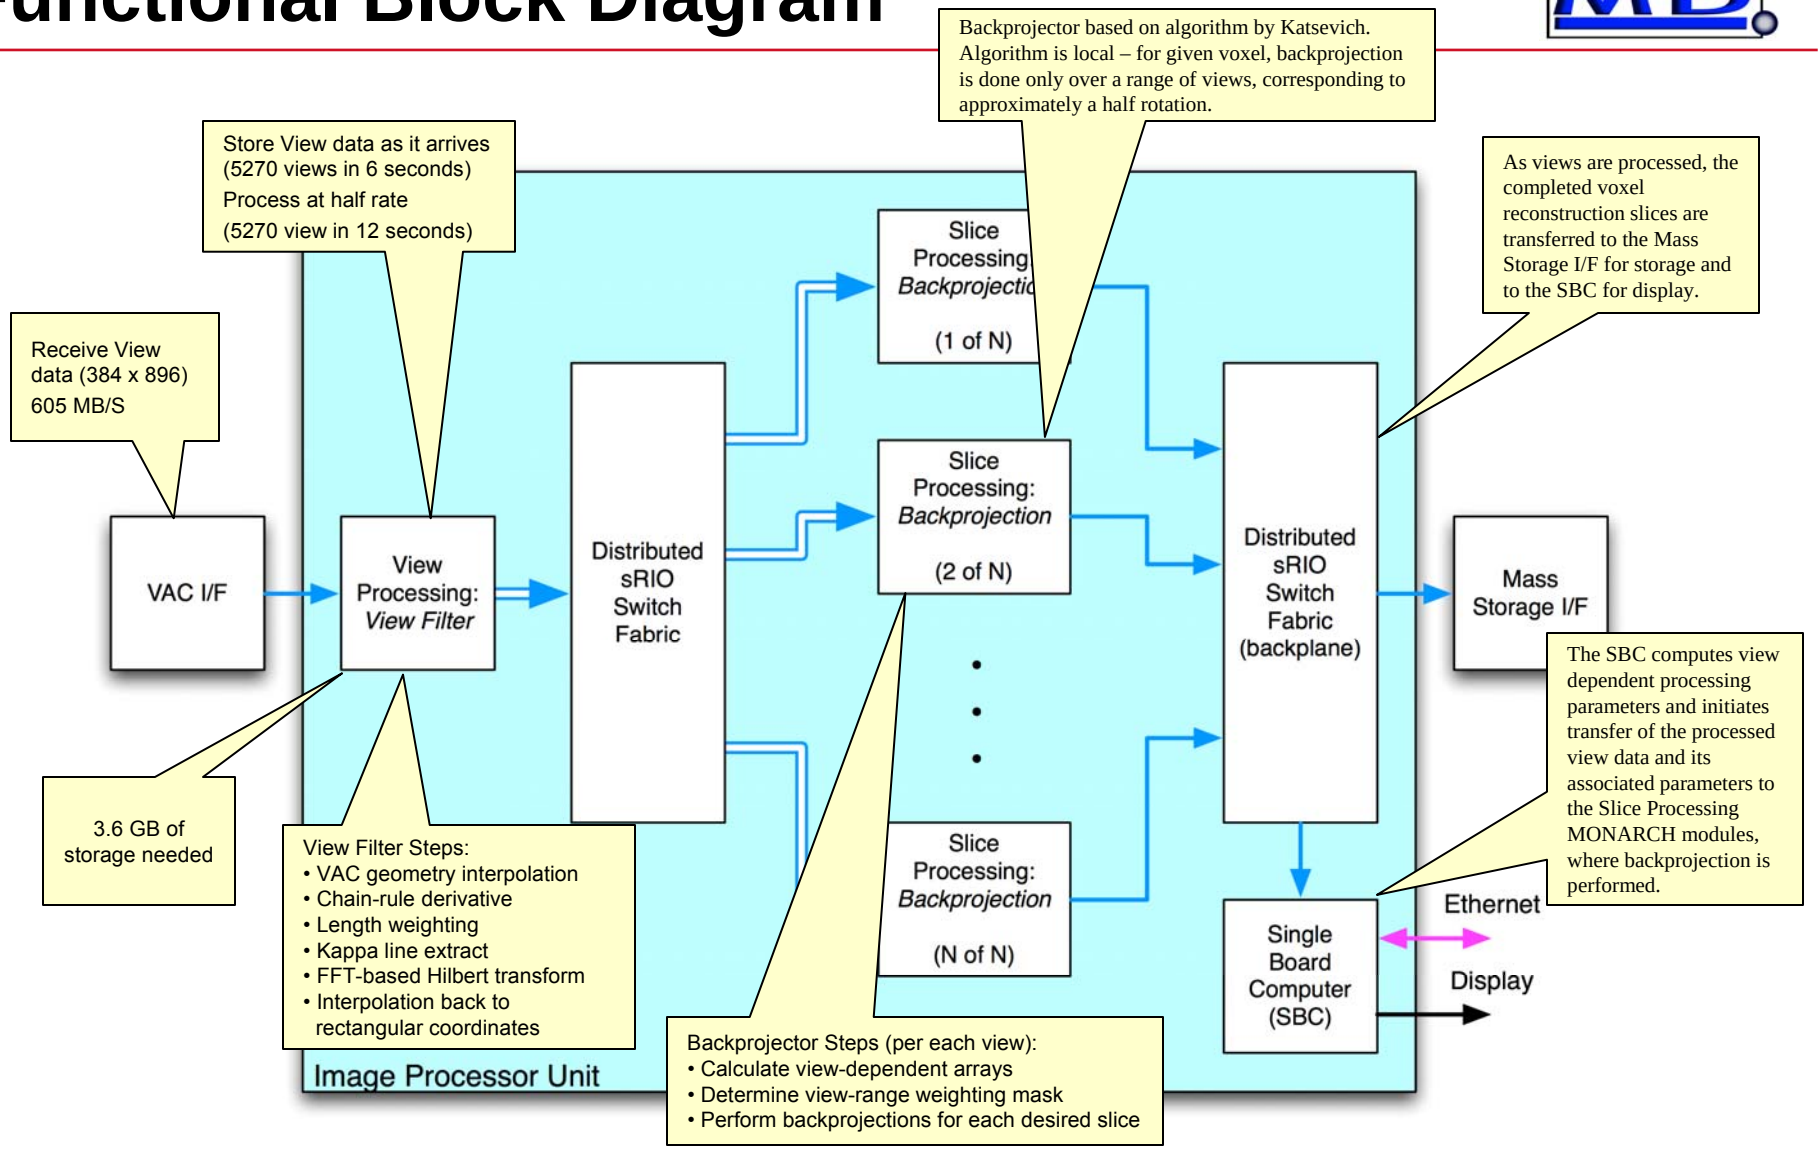

Final Mapping of Algorithm Chain to Signal Processing Hardware

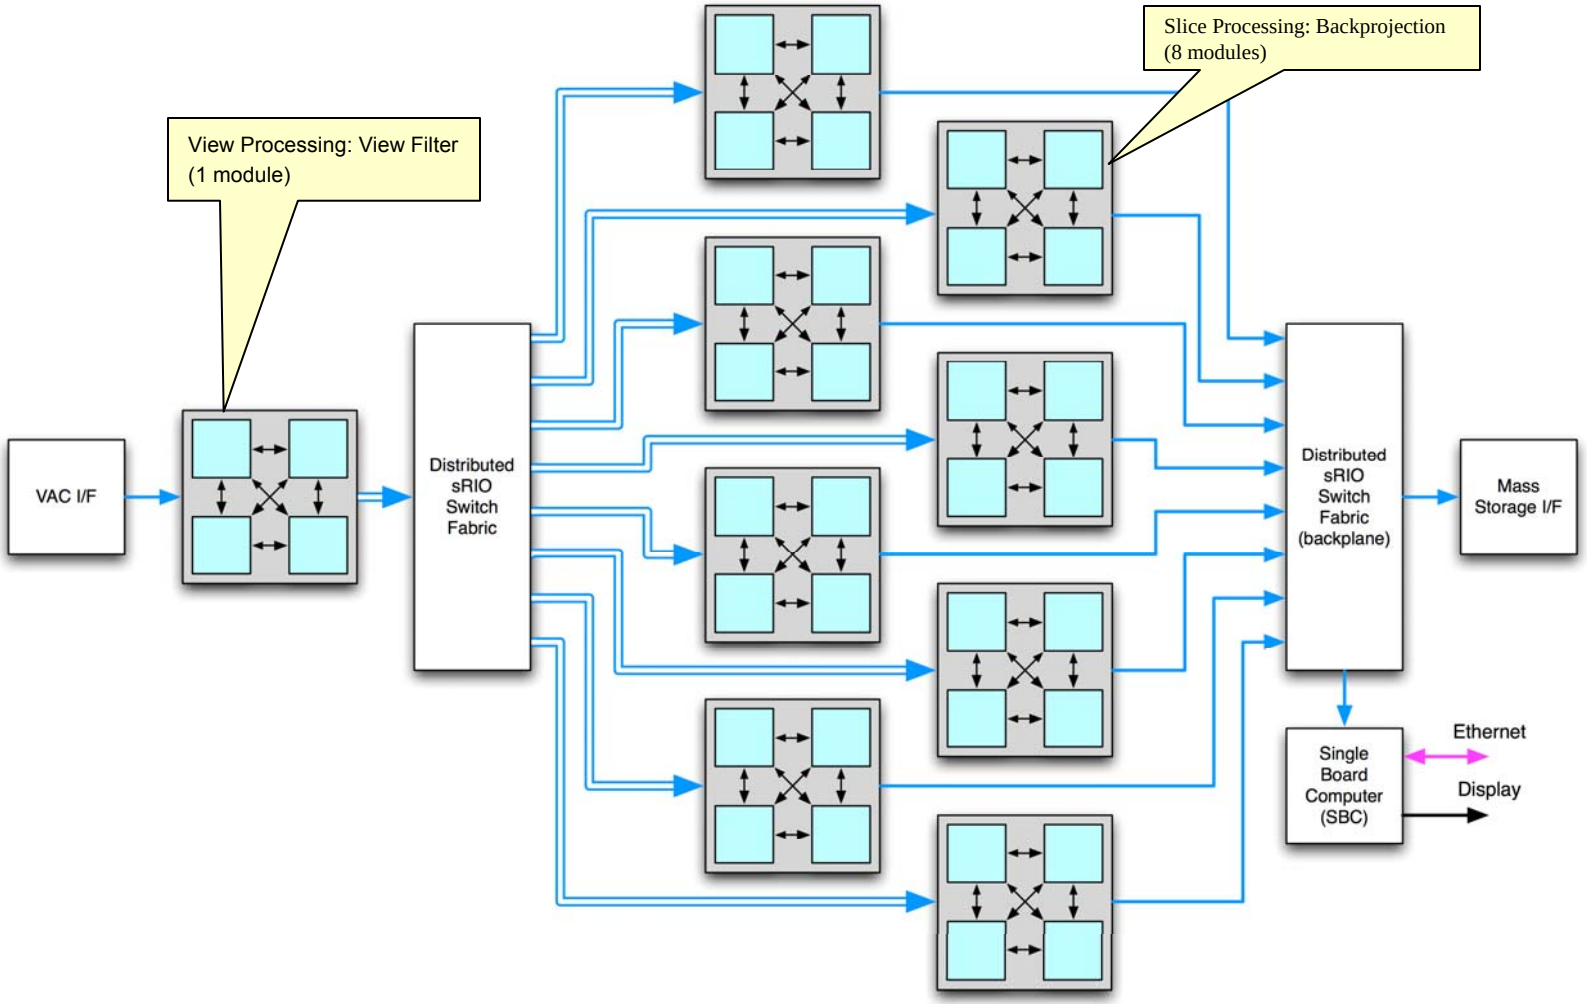

Backprojection Implementation Details

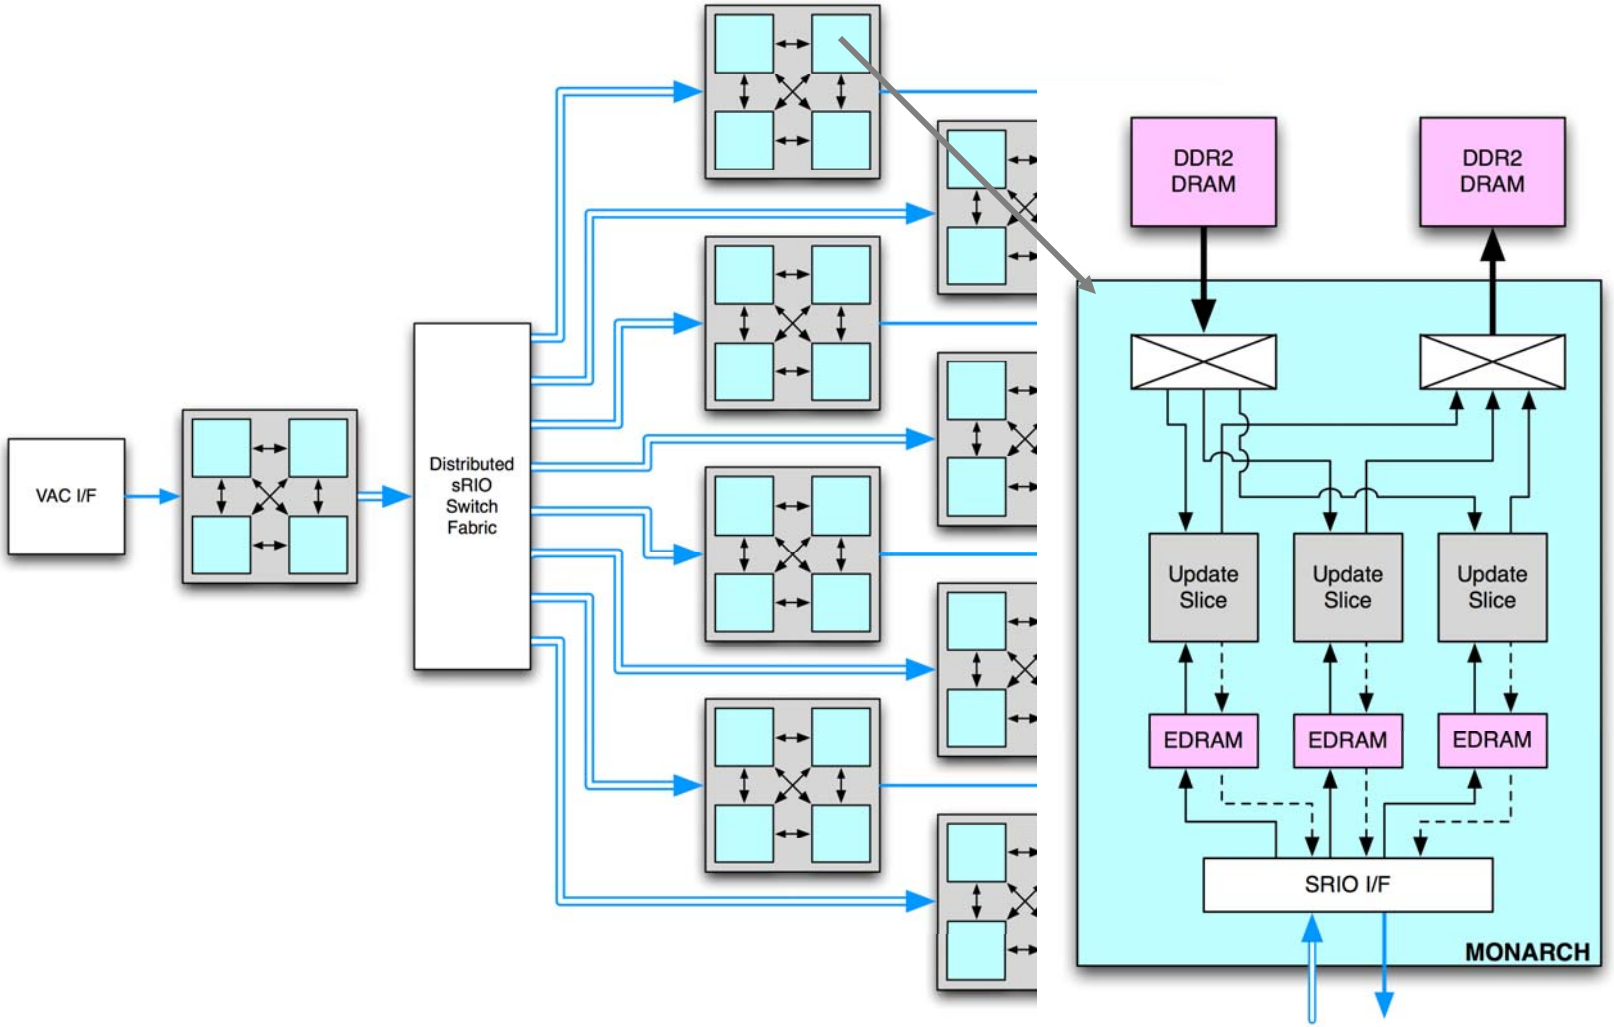

Comparison to GP-GPU

- Current model is to use GPU as a co-processor
  - Data is collected by main processor and then moved to GPU via PCI Express
  - PCI Express limits performance
  - Realistic performance is ~ 1.6 GB/S per port
- Total slice data rate of ~102 GB/S results in needing 64 GPUs
- Each GPU consumes ~200 Watts of power
- 64 x 200 W (plus an additional 800 Watts for support) = 13.6 KW
- Therefore, system efficiency is 0.08 GFLOP/S per Watt
- Compare this to ~ 1 GFLOP/S per Watt for MONARCH

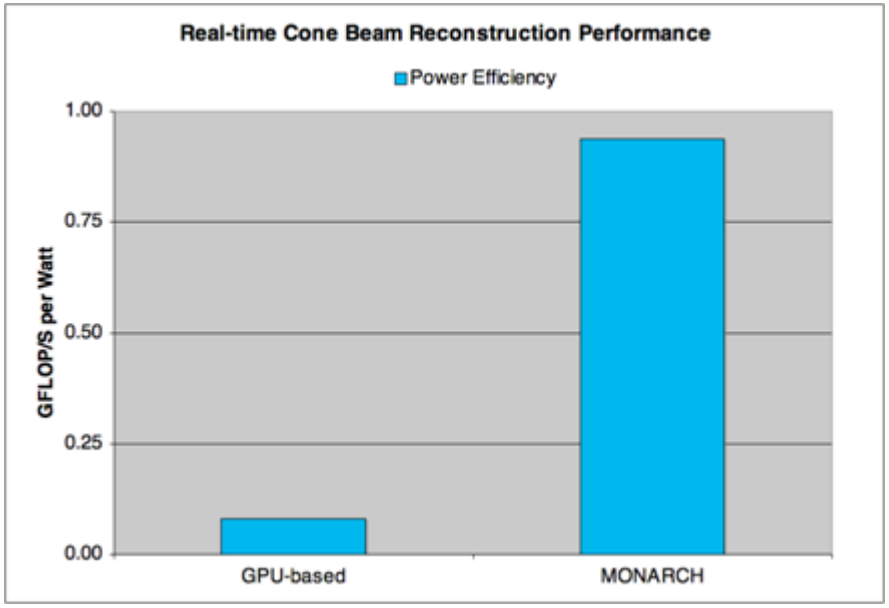

Conclusions

- To achieve the goal of mobile, real-time reconstruction of cone beam CT imagery, it is necessary to develop a computational platform that can supply 1 – 2 TFLOP/S while only consuming around a kilowatt of power.
- The MONARCH processor was developed under a DARPA contract with the goal of providing exceptional compute capacity and highly flexible data bandwidth capability coupled with state-of-the-art power efficiency and full programmability.
- We have presented:
  - an overview of a cone beam CT system
  - an overview of the reconstruction algorithm
  - the parameters for Compact VAC
  - a processing solution based on MONARCH
  - a demonstration showing that MONARCH is a natural fit to provide a mobile, low power solution for real-time cone beam CT reconstruction