Northern Illinois University

An Analysis of a Distributed GPU Implementation of Proton Computed Tomographic (pCT) Reconstruction

George Coutrakon, Kirk Duffin, Bela Erdelyi, Nicholas Karonis, Caesar Ordoñez, Michael Papka, Thomas Uram
Department of Computer Science

The Bragg Curve

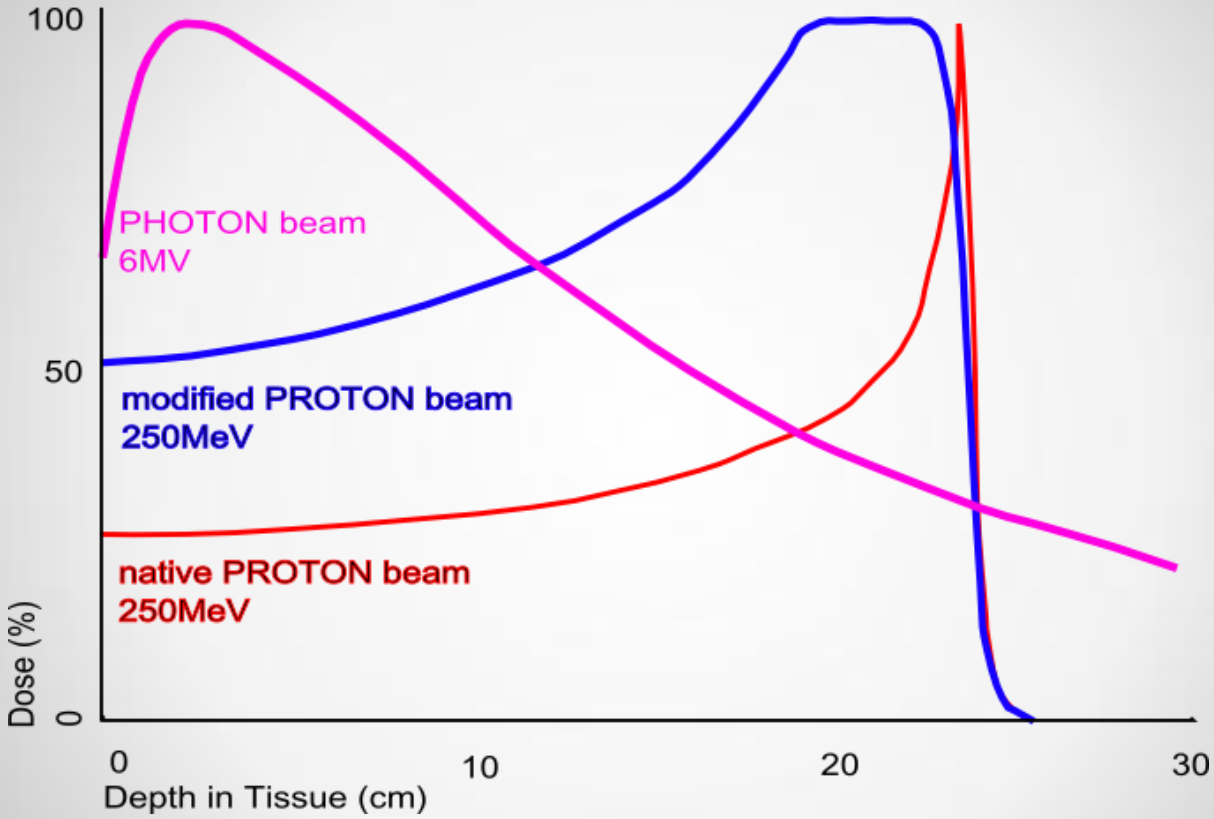

- Vertical axis depends on stopping power
- Relative stopping power – with respect to water

pCT: Proton Computed Tomography

- Imaging modality that uses protons as probe
- Direct measurement of proton relative stopping power (RSP)
- Images: 3D distribution of RSP
- Potentially more accurate than RSP obtained from X-ray CT (no need for conversion)
- Beneficial to proton therapy

Prototype pCT Detector

- Loma Linda University Medical Center
- Northern Illinois University
- University of California at Santa Cruz

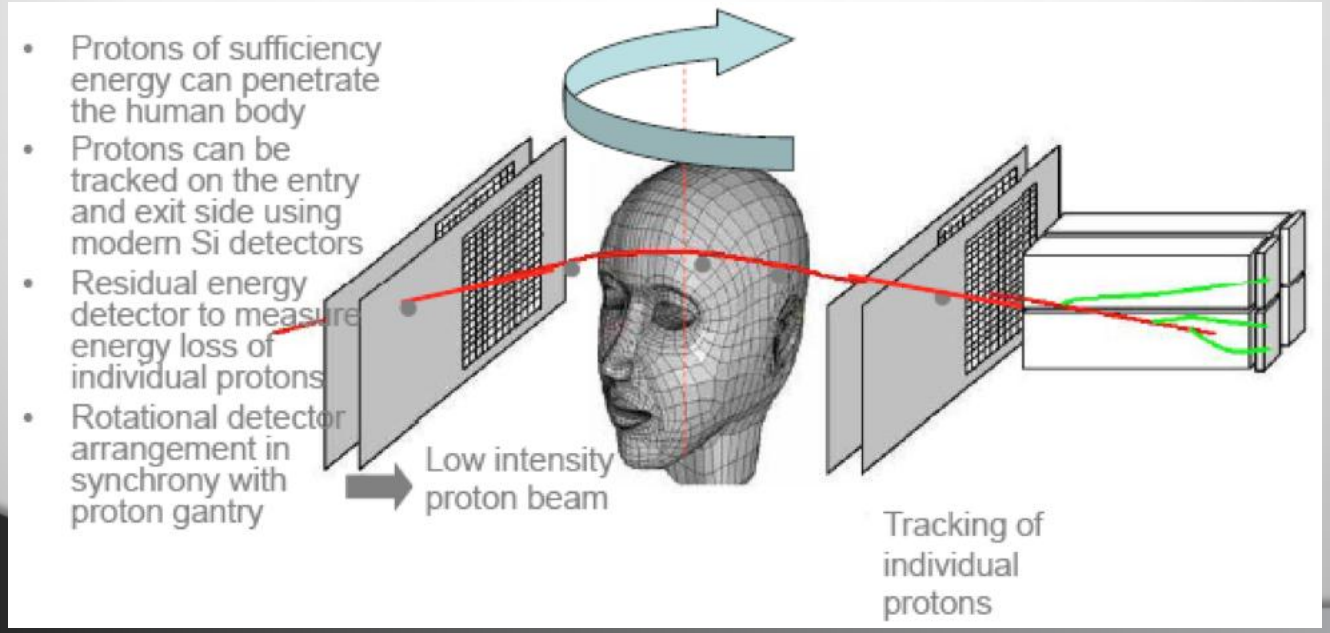

pCT: Challenges

- Large data sets
  - Estimate need 1 to 2 billion proton histories (events) to image objects the size of a human head - ~100GB input data
- Non-linear path of proton in material medium
  - Multiple Coulomb scattering (MCS)
  - Cannot use data reduction techniques such as those used in emission/transmission tomography (PET, SPECT, xCT)
  - Requires event-by-event processing
- Require lots of compute time
  - Almost 7 hours to reconstruct 131 million events on 1 CPU with 1 GPU (Penfold PhD thesis, 2010)

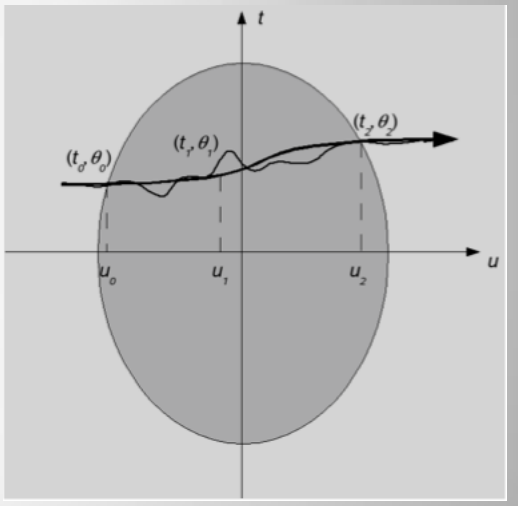

pCT: Solution

- Large linear system $Ax=b$
  - One proton per row (~10^9)
  - One voxel per column (~10^7)
- Naïve implementation – 160PB for A

pCT: Solution Simplification

- Memory Compression
  - 150 non-zero coefficients per row
  - 2.4TB for 2 billion events
- Path simplification
  - Most Likely Path (MLP)

pCT: Linear Solvers

- Block based iterative linear solvers
  - Block-iterative
    - Intra-block parallel, inter block sequential (e.g. DROP)
  - String averaging
    - Intra-block sequential, inter block parallel (e.g. CARP)

pCT Solution: Parallelize the Problem

- Computer cluster
  - Multiple compute nodes
  - CPU/GPU hybrid
  - Software technologies
    - MPI (Message Passing Interface)
    - CUDA (Compute Unified Device Architecture)
- Distribute data set to multiple nodes on cluster

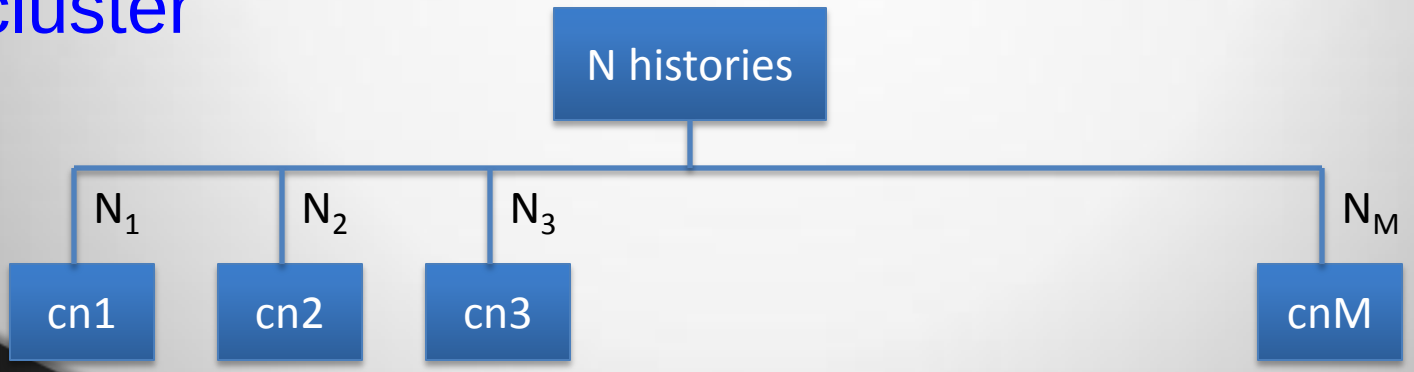

$$N_1 + N_2 + N_3 + \cdots + N_M = N$$

pCT Reconstruction Flowchart

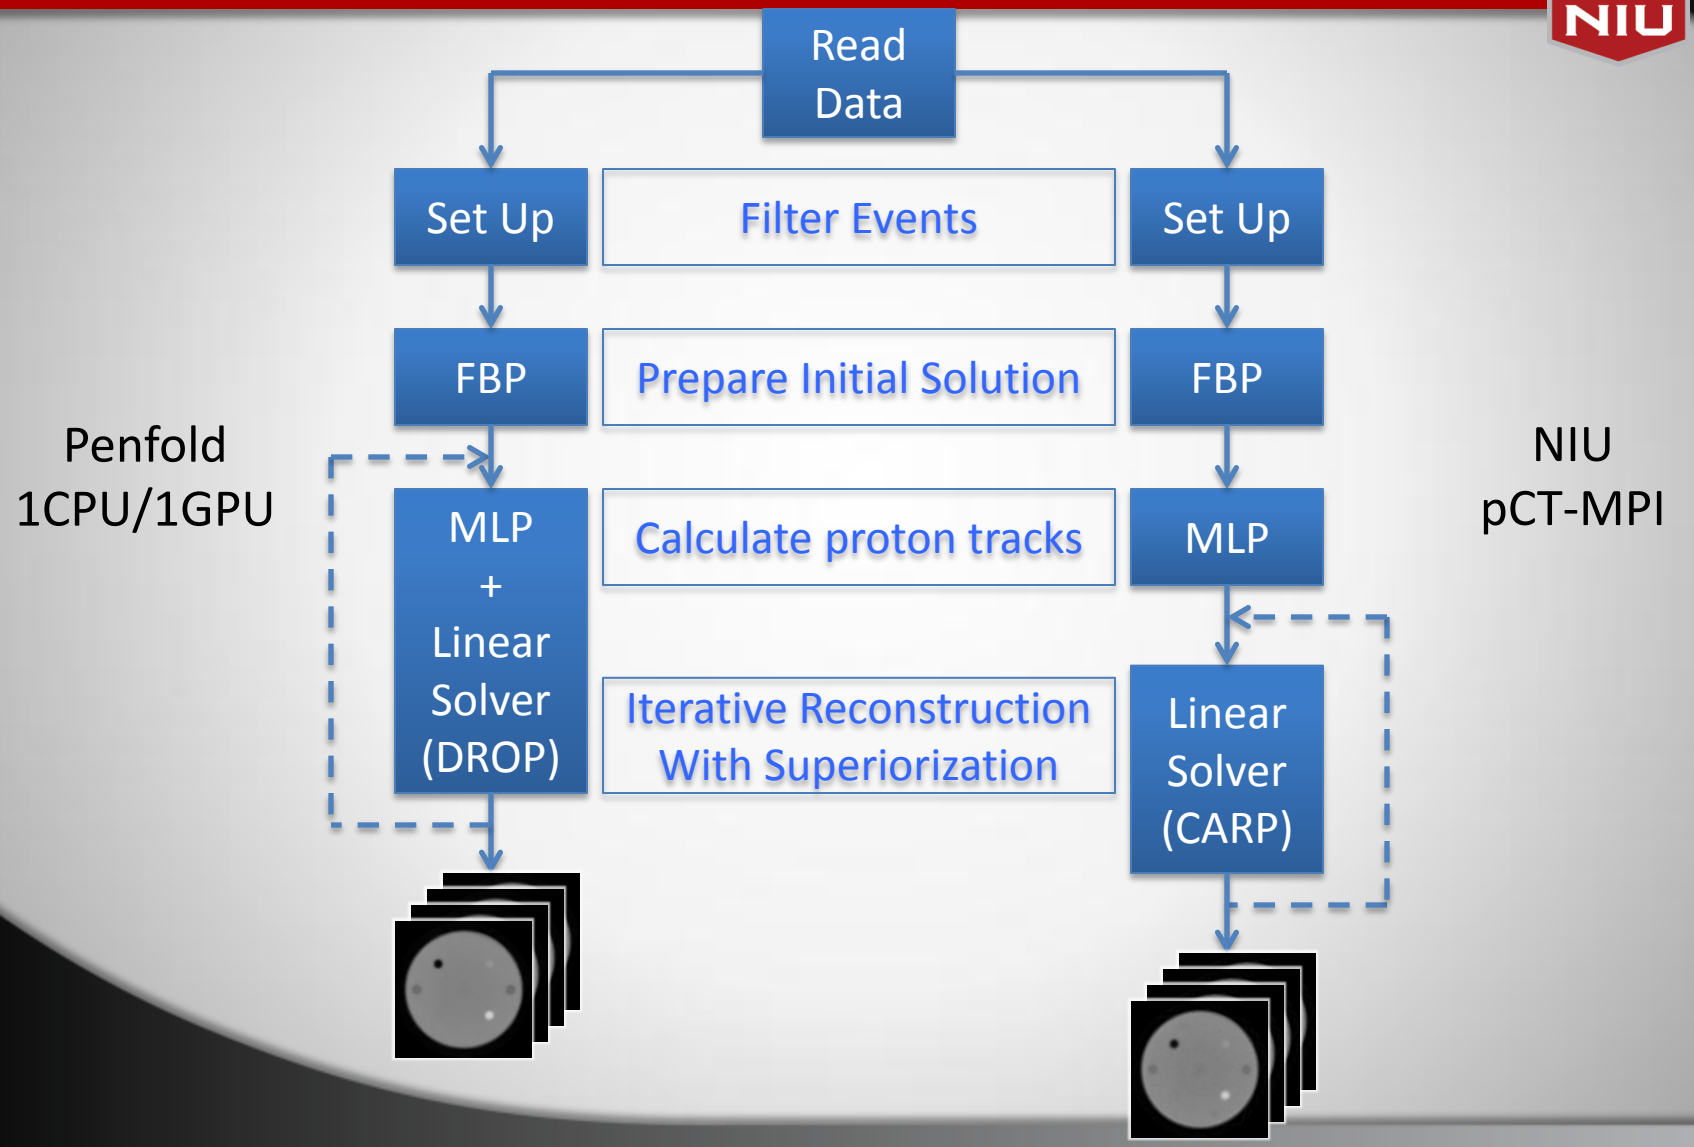

NIU Gaea HPC

- Power on: January 19, 2012
- 60 Compute Nodes
  - 72 GB RAM per node
  - 2 6-core CPUS per node
    - Xeon X5650 2.67GHz
  - 2 GPUs per node
    - NVIDIA m2070 (Tesla)
    - 6GB RAM
- 200TB storage array
- Infiniband network

Lucy Phantom for 3D Image Reconstruction

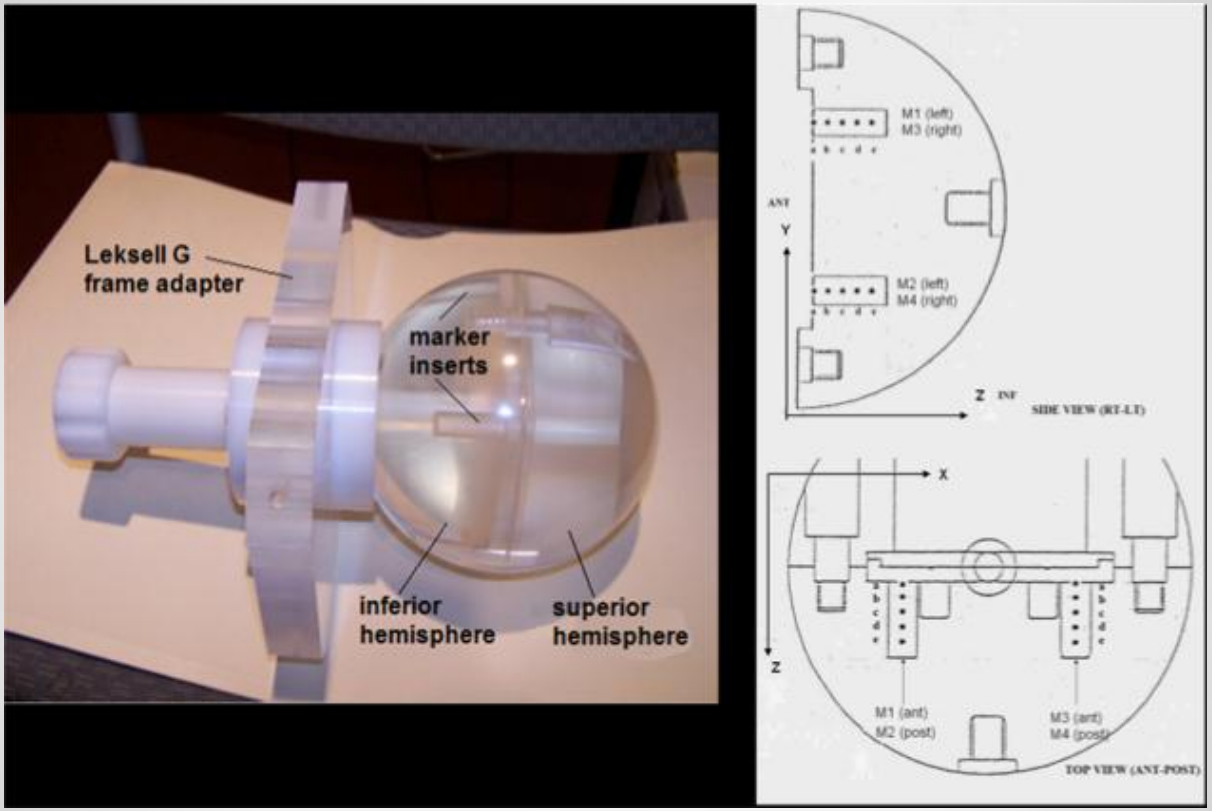

- 14cm-diameter polystyrene sphere
- 4 cylindrical inserts (air, lucite, polystyrene, “bone”)

Lucy Data Set

- Data acquired with prototype pCT detector at LLUMC (December 2010)
- 200-MeV protons
- 90 projection angles at 4-degree increments (2π-coverage)
- 131 Million histories
- Synthetic data sets generated
  - 1 billion histories: read Lucy data 8 times
  - 2 billion histories: read Lucy data 16 times
  - For timing purposes only
  - No image quality evaluation

pCT Reconstruction
131 Million Events

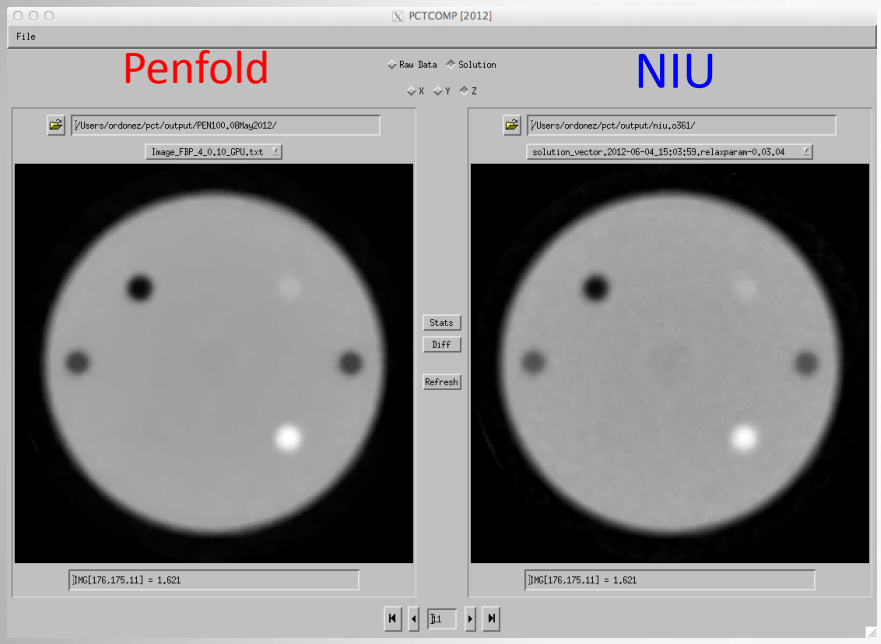

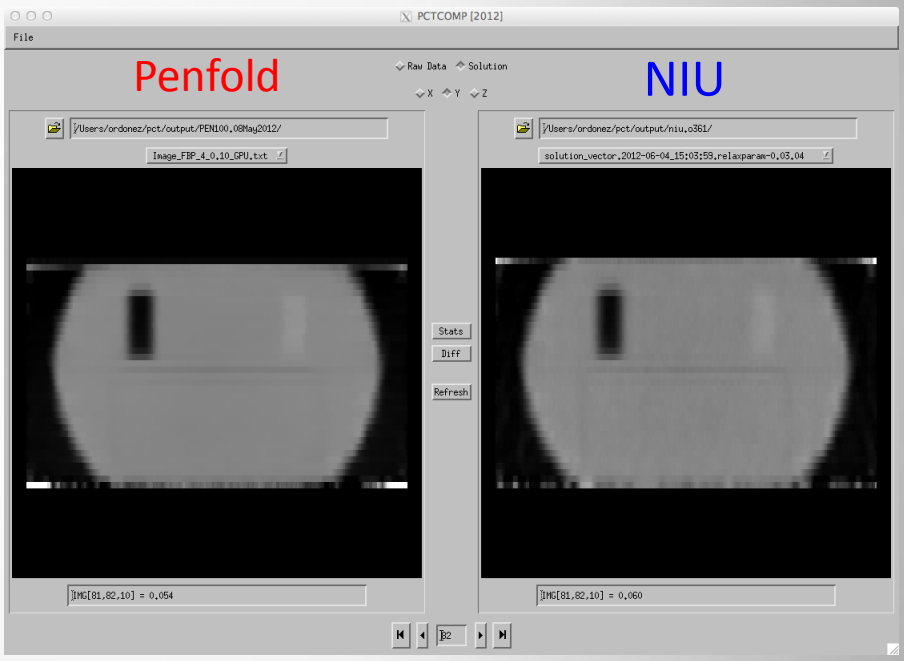

pCT Quantitative Analysis

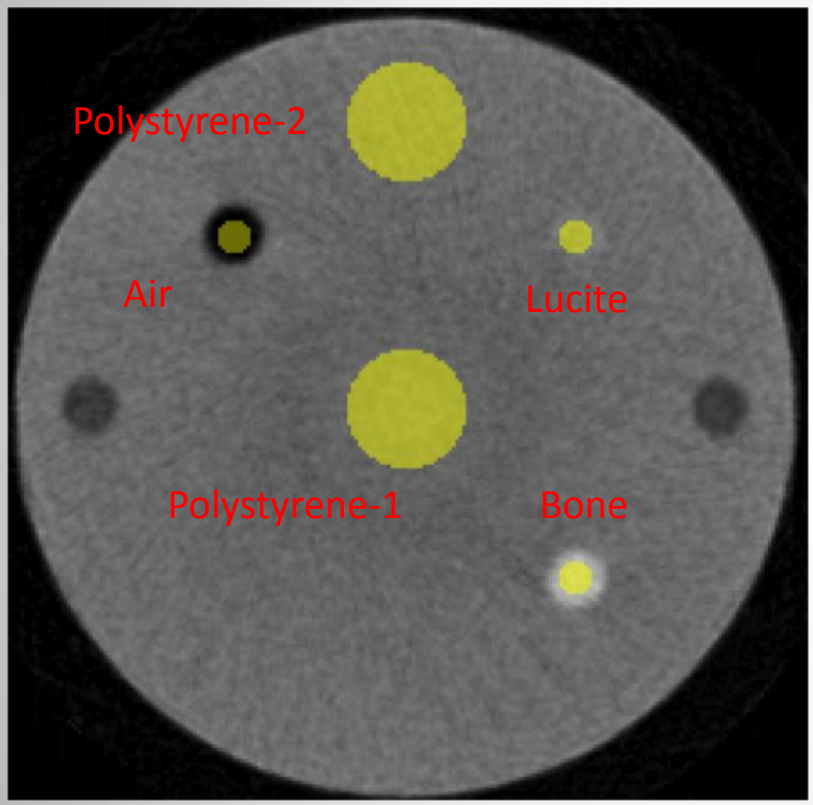

- Select 5 “Regions of Interest” (ROI) in Lucy Phantom (Sen and Duffin)
- ROIs actually volumes
- Each ROI has homogeneous density with known expected RSP (Schulte)

Material	RSP
Polystyrene	1.035
Bone	1.700
Lucite	1.200
Air	0.004

pCT Quantitative Analysis
131 Million Events

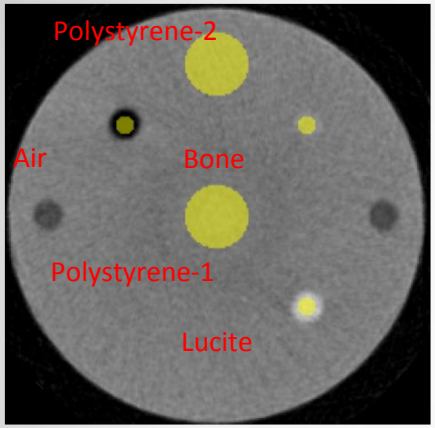

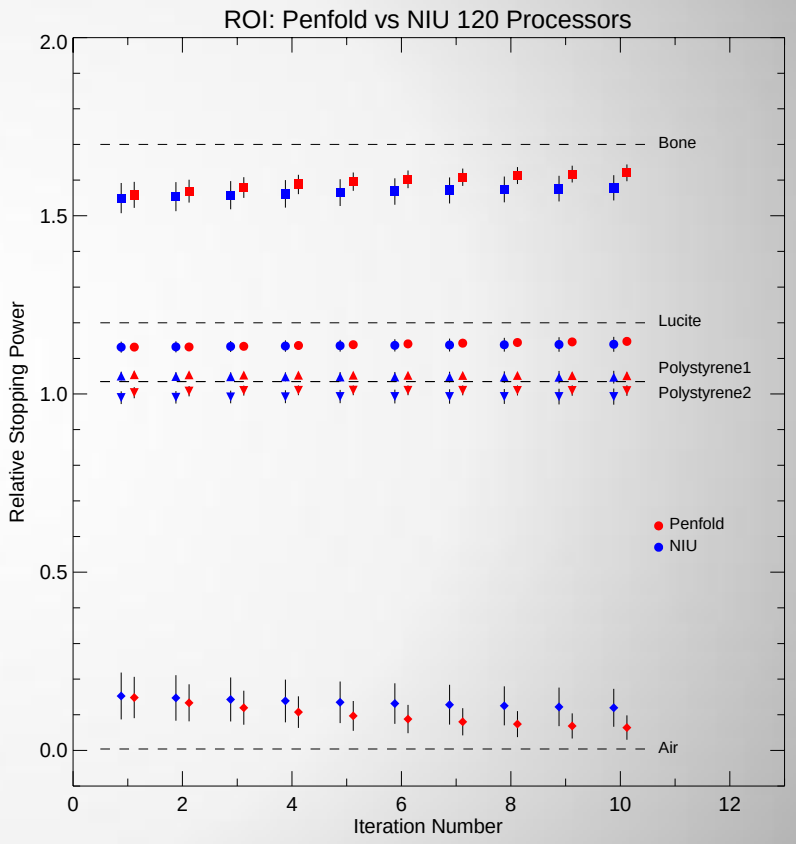

- NIU and Penfold RSPs agree
- Measured RSPs close to expected values
- NIU RSPs have greater variance
- Penfold compute time = 402 min
- NIU compute time = 53 sec

Processor Scaling
131 Million Events

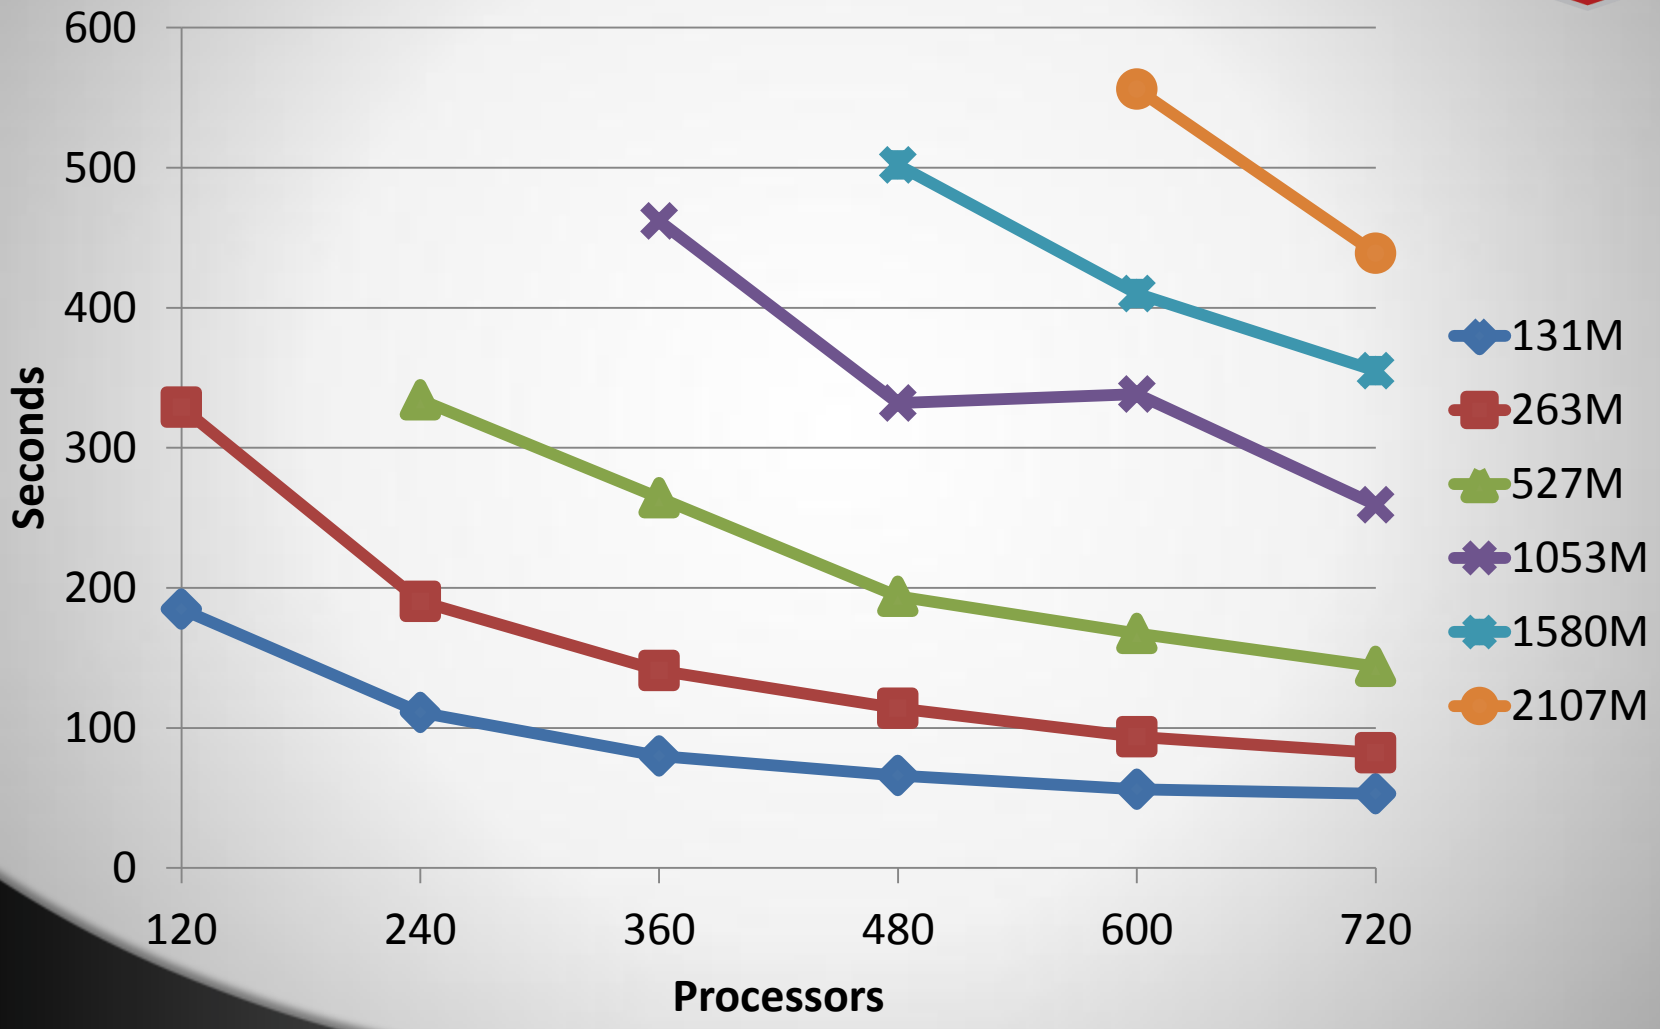

Processor Scaling
131 Million Events

Reconstruction time (sec)	Number of Processors (12 per node)					
	120	**240**	**360**	**480**	**600**	**720**
Read Data	1.006	0.949	1.048	1.160	1.213	1.380
Statistical Filter	12.805	13.302	12.712	12.618	13.088	13.796
Initial Solution	0.924	0.785	0.871	0.788	0.833	0.865
MLP	58.812	31.684	22.104	16.943	13.586	11.748
LinSol (10 Iters)*	111.752	63.318	42.689	33.549	27.105	24.174
Total Exec Time	184.875	111.000	80.000	66.000	56.160	53.000

- 68-92% of execution time spent in MLP + Linear Solver
- 46-60% of execution time spent in Linear Solver

Simultaneous Load Scaling

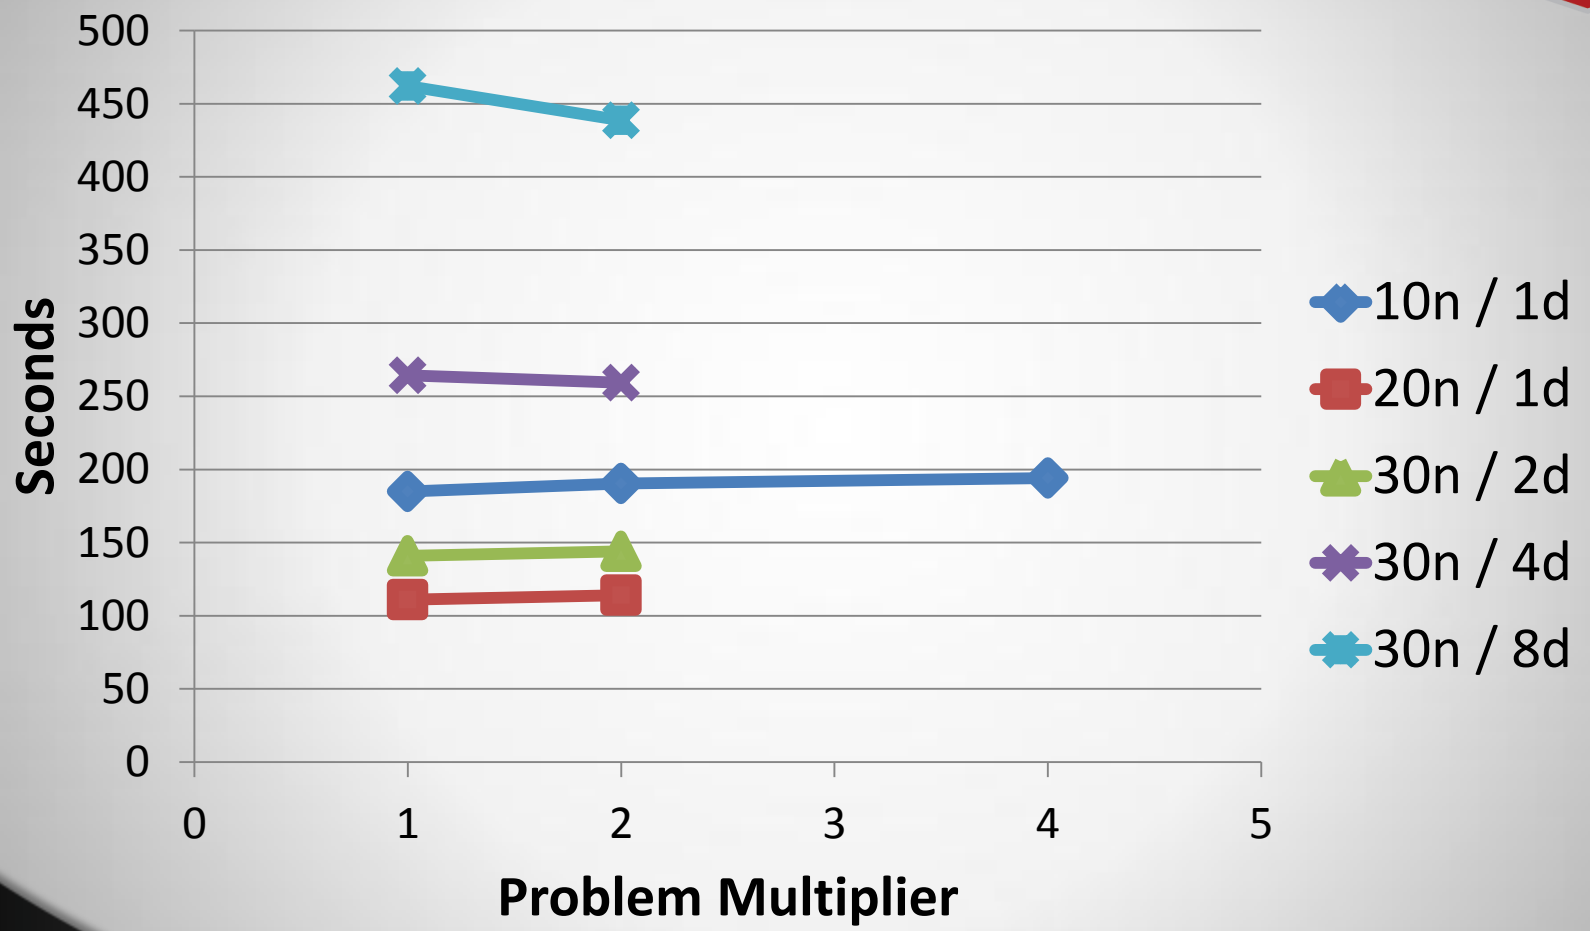

Data Scaling
720 Processors

Reconstruction time (sec)	Multiple of 131 Million Events					
	1	2	4	8	12	16
Read Data	1.380	1.671	2.827	3.734	5.452	6.488
Statistical Filter	13.796	12.490	13.078	13.357	14.421	14.526
Initial Solution	0.865	0.871	1.115	0.972	0.975	0.740
MLP	11.748	22.167	41.322	77.737	115.164	150.992
LinSol (10 Iters)*	24.174	44.566	85.170	162.810	217.239	265.512
Total Exec Time	53.000	82.247	144.00	66.000	354.983	438.778

- 67-95% of execution time spent in MLP + Linear Solver
- 46-60% of execution time spent in Linear Solver

Summary and Conclusions

- Multi-CPU/GPU speeds up pCT reconstruction

- Scalability
  - Scales linearly with number of processors
  - Scales linearly with problem size

- Promising Perormance

  With Penfold 1CPU/1GPU as “image standard”:
  - Image quality(NIU pCT-MPI) → Image quality(Penfold)
  - Time(NIU pCT-MPI) << Time(Penfold)

Future Work

- Don’t store MLP?
  - Calculate as needed
- Improve image quality
  - Other linear solvers (algorithm)
  - Relaxation parameter
- Path solution parameters
- More robust solution – no initial guess

Collaborators and Sponsor

- Yair Censor, University of Haifa
- George Coutrakon, NIU
- Kirk Duffin, NIU
- Bela Erdelyi, NIU
- Gabor Herman, City University of New York
- Ford Hurley, LLUMC
- Nicholas Karonis, NIU
- Caesar Ordoñez, NIU
- Eric Olson, ANL
- Mike Papka, ANL
- Scott Penfold, Royal Adelaide Hospital
- Reinhard Schulte, LLUMC
- Thomas Uram, ANL

US Department of Defense, Contract No. W81XWH-10-1-0170