

An Introduction to Medical Imaging Devices

Martin L. Purschke



About me

You heard me talk about DAQ design yesterday...

So what qualifies me to talk about Medical Imaging?

Medical Imaging

Medical Imaging has a lot in common with nuclear physics (and accelerator/beam physics)

Yes, there is the part where an actual M.D. comes in and wants to see something in particular, but up to that point, it's physics and engineering

Radiotracers, nuclei, detector technology, readout, analysis, All the natural habitats of physicists and engineers.

As a matter of fact, the by far most common diagnostics method using Te^{99} was developed at Brookhaven National Lab – you have accelerators, nuclear physicists, and the infrastructure

But also, I was deeply involved what I think is until today the coolest application of PET (Positron Emission Tomography) – the RatCAP

For the first time, the RatCAP allowed the imaging of the brain of an **awake** rat

I'll talk about this later some more, but this is my foray into Medical Imaging, and why I know all this stuff.



Medical Imaging Technologies

Without any claim to completeness (and some of what I'll cover)

SPECT (a little)

PET

MRI

X-Ray

CT

Multi-modality imaging

(there are many more, often variants of a common theme)

You will see that nuclear physics, DAQs, and data processing plays a prominent role.

There is a place for you in this field if you can do TDAQ, or are a “data engineer”

Medical Imaging Technologies grouped by other metrics

**Not using ionizing
radiation**

MRI

Using ionizing radiation

X-Rays

CT

SPECT

PET

It's ironic that the only technology here NOT using ionizing radiation, formerly known as Nuclear Magnetic Resonance Imaging (NMR), had to be renamed because people were freaking out over the word "nuclear" in the name

What does one want, and how does one choose a technology?

- **Sensitivity** - How well can you actually see what you are after?
- **Selectivity** - How well does your method distinguish between, say, benign and malicious tissue?
- **Contrast** - What is the dynamic range between the different features in your image?

- For example, if you want to see if a bone is broken, X-rays (and by extension, CT) gives you a high value on all three – the contrast between bone and tissue is large
- X-ray for a mammogram has a low value on all three

Each modality has strengths and weaknesses

- For example, MRI has excellent position resolution but is “blind” to metabolic processes
- PET can show the metabolism but the position resolution is poor
- Think of different modalities as looking at the same thing in different ways

- MRI by and large images the density of protons (think water in tissue)
- X-ray and CT image the electron density in tissue (that’s why bones show up so nicely)
- SPECT and PET can be tailored to show metabolic processes in tissue


Why PET is so cool

- PET can be tailored to show *metabolic* processes in tissue
- Remember, the *chemical* properties of a radioactive isotope are the same as the stable isotope – it will be used by the body indiscriminately
- So you take a molecule that is used in your body for some process – sugar, dopamine, what have you, swap out an atom for a radioactive one
- You get “hot spots” where lots of said molecules are used by the body
- For example, a tumor that grows uses a lot of energy (sugar), so radioactively tagged sugar will accumulate a lot of activity there
- Or tag molecules that the brain burns when it “thinks” – active areas show up as hot spots
- Frequently used radiotracers are ^{11}C and ^{18}F (all beta+ = positrons)

An Example - ^{18}F

- ^{18}F -FDG - Fludeoxyglucose is one of the commonly used radiotracers
- It is a sugar, so it accumulates where the body burns a lot of energy (“uptake”)
- Unusual “hot spots” can point to cancer

Whole-body PET scan using ^{18}F -FDG to show liver metastases of a colorectal tumor



^{11}C Raclopride – Brain activity

- Raclopride is a molecule that binds to dopamine receptors in the brain
- Can show which part of the brain does what
- Different activities stimulate different areas in the brain (singing, reading, listening, etc)





Image of a rat brain taken with the RatCAP using ^{11}C -tagged raclopride

(I actually took those data!)

SPECT (*Single Photon Emission Computed Tomography*)


Ring of Photon Detectors



- Radionuclide decays, emitting γ rays.
- Gamma lies on line defined by detector and collimator (known as a *line of response* or a *LOR*).
- Single Photons are detected
- Low efficiency (next slide)
- Can cherry-pick the photon-emitting isotope (^{125}I , ^{60}Co) (not that much of a menu though)


- I have never worked with SPECT, but I see most groups abandoning this
- Some Alzheimer's-related research still shows up

SPECT - collimation and low efficiency



- Without a collimator, there's no sense of "direction" of the photon – the photon could come from anywhere in the field of view
- Only the Collimator selects a direction
- Dramatic reduction in efficiency – most photons are lost in the collimator

SPECT Image Examples



Phantom

Collimator



Detector



Example of a phantom – the holes are filled with radiotracers (similar to what was used in these images)

S. Shokouhi, Vanderbilt U

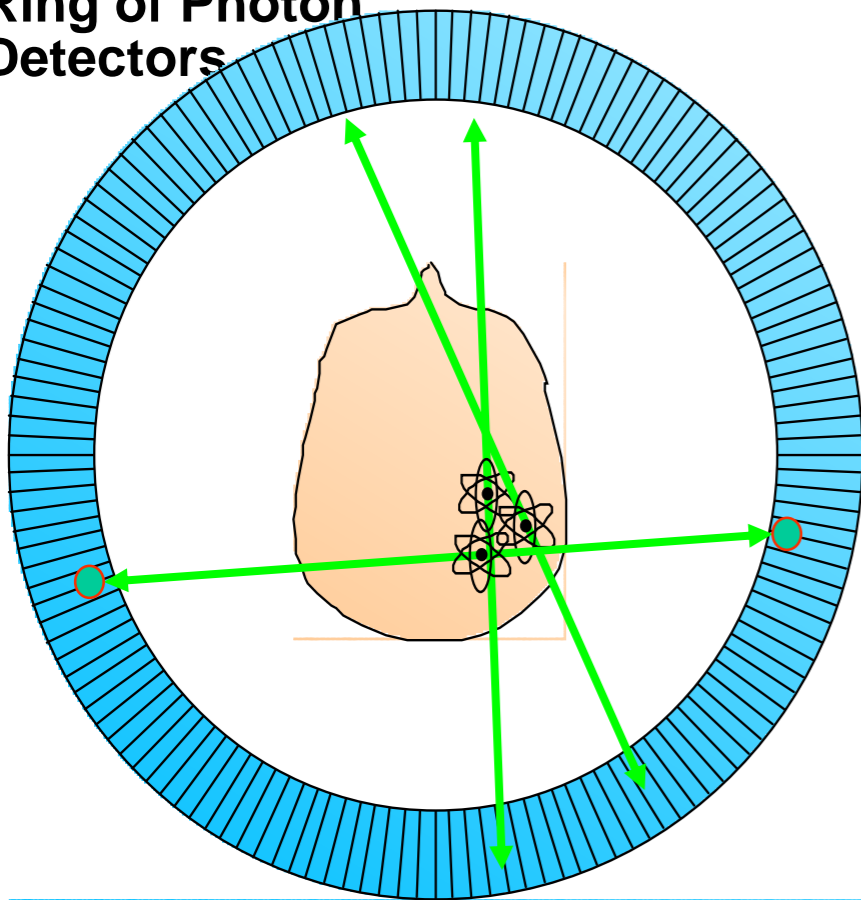
3rd-largest bores can no longer be resolved



Different collimator configurations



PET – Positron Emission Tomography

Ring of Photon Detectors



- Radionuclide decays, emitting β^+ .
- β^+ annihilates with e^- from tissue, forming back-to-back 511 keV photon pair.
- 511 keV photon pairs detected via time coincidence.
- Positron lies on line defined by detector pair (known as a *line of response* or a *LOR*).


- Back-to-back 511KeV photons are detected
- No collimator needed
- LOR defines the direction
- Best I can tell, all the action today is in PET, not SPECT




Phantom Reconstruction

The PET principle


- PET measures coincident photons from an electron-positron annihilation
- Each such coincidence establishes one “Line Of Response” (LOR)
- It is mostly intuitive for folks working with accelerators
- Commercial scanners often have sophisticated electronics or firmware for coincidence processing
- Or you record all photon counts with detector element info and a time stamp (“list mode”)
- When two photons are measured “simultaneously” within a certain time window, you consider them a coincident pair
- You can either record all the data and do the analysis offline, or the PET scanner has that ability built in and reports only the coincidences (commercial scanners usually do)



Offline Coincidence Processing



- You cannot know which ones are true and which ones are random coincidences
- But you can get an estimate of the random rate and their contribution by adding a delay to the window and destroying all true coincidences



- When all is done, the output from the coincidence processing is a ***sinogram***
- There are more corrections, such as scatter corrections, attenuation, etc...

Photon Detection

511KeV is a tiny amount of energy when it comes to detection

High-Energy physics tends to ignore anything below a few 100 MeV

The resolution of a detector is primarily determined by the “count” of what you measure – the famous \sqrt{N} error formula

511 KeV produce between 2000 and 20000 photons, the latter in the workhorse scintillator, NaI


NaI falls flat in basically every other metric – much too slow, it's very hygroscopic, not very dense

Many new scintillators on the market, some are *fantastically* expensive – sometimes hundreds of dollars per cubic *millimeter*

Readout through PMTs, APDs, SiPMs...

We are looking for fast signals so the coincidence detection works.

For that a few hundred picoseconds timing resolution is enough




Keep that number in mind for later...

A selection of scintillators

	NaI(Tl)	BaF₂	BGO	LSO	GSO	LYSO	LaBr₃	LFS	LuAP	LuI₃
Effective atomic no. (<i>Z</i>)	51	54	74	66	59	60	47	63	65	60
Linear attenuation coeff. (cm ⁻¹)	0.34	0.44	0.92	0.87	0.62	0.86	0.47	0.82	0.9	~0.56
Density (gm cm ⁻³)	3.67	4.89	7.13	7.4	6.7	7.1	5.3	7.3	8.34	5.6
Index of refraction	1.85	–	2.15	1.82	1.85	1.81	1.88	1.78	1.95	
Light yield (% NaI(Tl))	100	5	15	75	30	80	160	77	16	190
Peak wavelength (nm)	410	220	480	420	430	420	370	430	365	470
Decay constant (ns)	230	0.8	300	40	65	41	25	35	18	30
Hydroscopic	Yes	Slight	No	No	No	No	No	No	No	Yes

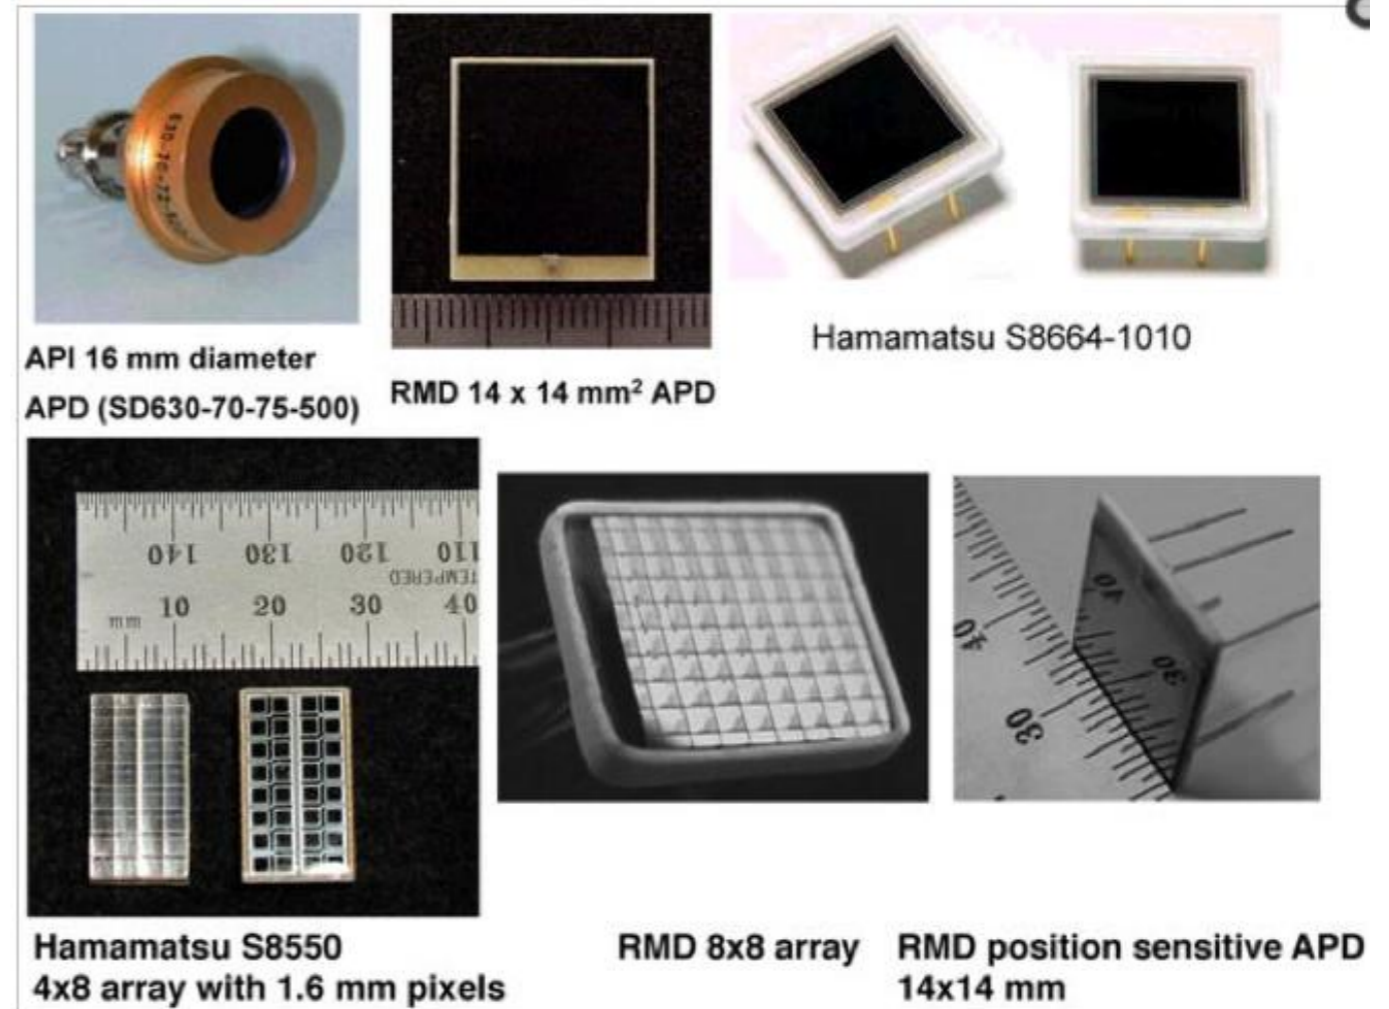
Readout

Photomultipliers



- Heat Dissipation – needs cooling
- Bulky
- Needs High Voltage (~ 2KV)
- Useless in a magnetic field
- Excellent linearity and noise

Avalanche Photon Detectors



API 16 mm diameter
APD (SD630-70-75-500)

RMD 14 x 14 mm² APD

Hamamatsu S8664-1010


Hamamatsu S8550
4x8 array with 1.6 mm pixels

RMD 8x8 array

RMD position sensitive APD
14x14 mm


- Small footprint
- Works in a magnetic field
- Modest supply voltage (~400V)
- Significant noise

A Real-Life Pet scanner – MiniPET (KTH, Sweden)



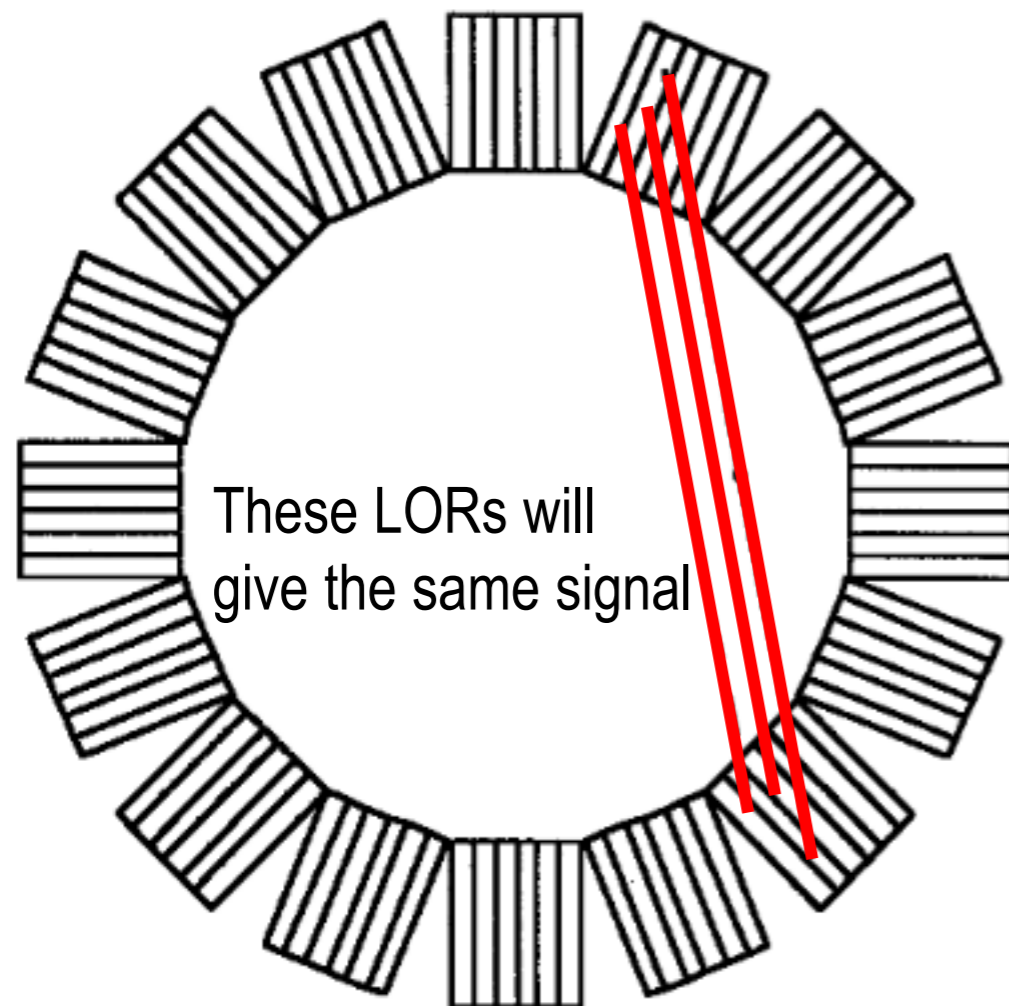
- 206mm opening
- Hamamatsu H9500 Position-sensitive PMT
- LYSO Crystals $1.27 \times 1.27 \times 12 \text{ mm}^3$
- 32×32 crystals per module, 12 modules
- FPGA-based data processing
- readout through Ethernet

Image of a rat heart




Depth of Interaction – DOI measurement

The Parallax Problem



The parallax effect limits the position resolution
Problem in particular in smaller scanners

- DOI can resolve the LORs better
- Usually requires some dual-ended readout for light-sharing
- More complex readout




Miyaoka, R.S.; Lewellen, T.K. Nuclear Science Symposium, 1997. IEEE , vol.2, no., pp.939-943 vol.2, 9-15 Nov 1997

Image Reconstruction - Sinogram


We are “looking” at the object from many angles instead of just one

This is what you do inadvertently with your eyes to see sharper





0 degrees

Image Reconstruction - Sinogram





0 degrees
22.5 degrees

Image Reconstruction - Sinogram





0 degrees
22.5 degrees
45.0 degrees

Image Reconstruction - Sinogram





0 degrees
22.5 degrees
45.0 degrees
67.5 degrees

Image Reconstruction - Sinogram



0 degrees
22.5 degrees
45.0 degrees
67.5 degrees
90.0 degrees

Image Reconstruction - Sinogram



0 degrees
22.5 degrees
45.0 degrees
67.5 degrees
90.0 degrees
112.5 degrees

Image Reconstruction - Sinogram





Image Reconstruction - Sinogram





Image Reconstruction – Sinogram and reconstruction

It's called a sinogram because any feature traces a sine wave




Image Reconstruction: Back Projection

Remember each projection only gives you a line
You fill the different projections in a histogram (each line)
Sources in the same spot in different projections accumulate





Image Reconstruction: *Filtered* Back Projection

A naïve back projection adds a lot of noise and artifacts to the image

The BP is the equivalent of a Fourier transformation

The filter clips “harmonics” that are beyond the Nyquist resolution and reduces noise



There a lot of different filters in use

A lot of papers written about that
You can get your PhD for finding a better one...

I'll get to another (and for folks like us, more interesting, I think) image reconstruction method in a minute... but first:

The RatCAP and Derivatives

- One wants to use PET to study the neurophysiological activity and behavior in laboratory animals
- Understand and treat illnesses in humans.
- However, animals needed to be anesthetized during PET imaging.
- Anesthesia can greatly depress the very brain functions and affect the neurochemistry that one is trying to study
- Cannot study animal behavior while under anesthesia
- The “Holy Grail” – study brain processes in the awake animal using PET

The Quest

To convert something like this...



... into something like this...




RatCAP: Rat Conscious Animal PET

A miniature, complete full-ring tomograph mounted to the head of an awake rat.

Compact, light weight (< 200 g), low power detector



Small field of view (38 mm dia. x 18 mm axial)

Attached to the head of the rat and supported by a tether which allows reasonable freedom of movement for the animal




RatCAP: Rat Conscious Animal PET

Ring containing 12 block detectors of 32 2x2 mm² x 5 mm deep LSO crystals with APDs and integrated readout electronics



4x8 array of LSO crystals (2x2x5 mm³) APD (Hamamastu S8550)



Actual RatCAP Ring

All the magic lies in the RatCAP ASIC, which takes care of all the signal processing. Reports channel # with timestamp


RatCAP Derivatives

Once you have a working ASIC, you try to re-use it for all kinds of other cool things...

This “MRI compatible” was a major piece of work – no ferromagnetic material anywhere



Original RatCAP V1



**RatCAP V2
MRI compatible**



**Breast Scanner
Prototype**



**Breast Scanner
UPenn Scanner**


The Technology went commercial




BNL licensed the technology to an
upstart company



RatCAP



PET Insert for MRI




Wrist Scanner



Breast PET-MRI


Other opportunistic applications – Plant Scanner

Here, the plant is fed ^{11}C -labeled CO_2 with a cuvette

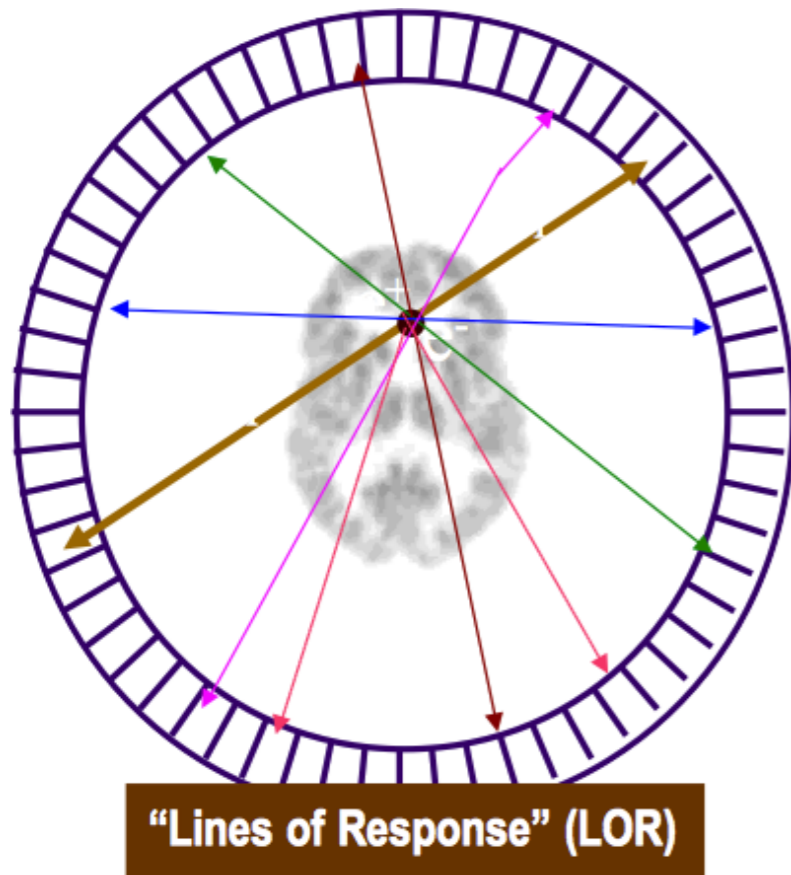


The large opening is needed mostly to allow the plant to fit into the scanner. Most of the FOV is just empty

**This supports research for better biofuels
You can track where the produced sugar is going in the plant**



Alternate Iterative MLEM-based Image Reconstruction



- One divides the Field of View (the active area of the scanner) in small “voxels”. Ours are 1mm^3 cubes
- The **System Matrix** A describes the probabilities that the decay at a given voxel position results in the photons ending up in a given pair of detector elements (“LOR”)
- Number of LORs go $\propto(N^2)$ for N detector elements
- FOV size = number of voxels goes $\propto(N^{3/2})$
- Matrix has dimensions number voxels x number of LORs
- Each matrix is specific for a particular geometry

This is what you measure
This is what you want


So: Activity Distribution $\times A =$ Detector Response

So all it takes is to invert the matrix, and we are set

Great! Activity Distribution $= A^{-1} \times$ Detector Response

Not so fast!

The System Matrix is HUGE



	Ratcap	BS proto	UPenn/BS
Voxels	15113	194427	1328360
LORs	72192	342784	22444800
Matrix elements	109 mill	66646 mill	29814774 mill

- No way to invert such huge matrices
- So we resort to an iterative process that “sculpts” an activity distribution until it matches the measured detector response
- This method beats Filtered Back-Projection hands-down
- But it is *really* computing intensive – we cannot yet do this for the full Breast Scanner

Computing Challenges

As small as the detectors look, they produce a lot of data

The largest UPenn / Breast Scanner can produce 450MB/s (1.5 TB/hour)

The computing challenge is two-fold:

Simulations

A quarter million CPU hours now and then (4 - 5 times a year) to obtain the "System Matrix"

Image Reconstruction

Iterative MLEM-based process that "models" an activity distribution which yields the measured detector response

Computing Challenges

Believe it or not, the ~250,000 CPU hours for a system matrix is actually the smaller problem

The simulation of any voxel is independent of all others

You can distribute the workload as much as you like over as many CPUs as you like

We have been using an online cluster (of my day-job at the PHENIX experiment at the Relativistic Heavy Ion Collider) opportunistically

We have joined the Open Science Grid (OSG) to disentangle us from the seasonal availability of our local cluster

We can get about 1000 CPUs / day and can get a new matrix in about 2 weeks

We do this only a few times a year

You are here: [TWiki](#) > [VirtualOrganizations Web](#) > [BNLPET](#) (23 Sep 2012, [MartinPurschke](#))

[↓ Introduction](#)

[↓ Overview](#)

[Our OSG Wiki page](#)

Positron Emission Tomography (PET) at BNL - Computations on OSG

Introduction

The PET group at the Brookhaven National Laboratory and Stony Brook University is interested in the generation of "system matrix", a simulated response model of the detector that translates into a matrix with a few billion non-zero elements. The computation is relatively straightforward but of massive-scale. For some detector systems the computations exceeding 50 CPU-years, above the capacity for dedicated and opportunistic local resources. This proof-of-principle phase aims at running some of these computations on OSG opportunistic resources.

Reconstruction Times

- The iterative, “activity-sculpting” process yields a usable image after about 100, a good one after 500...1000 iterations. Each iteration involves 2 matrix multiplications
- At one point, in 2004 or so, the small RatCAP reconstruction was perceived to be a big problem
- No more, we really have that optimized, 10-15 minutes per image.
- Caveat: it's not just one image. We do dynamic images, time-sliced, 2-15 minutes depending on the experiment. They get correlated with external events. Like a movie of the activity distribution.
- The plant scanner takes about 6-8 hours per image on a multi-core CPU. Multiple machines can work in parallel on independent images
- The Breast Scanner? We don't know, we cannot presently do it “our way”. Estimated 400 hours, but we didn't bother to make a system matrix yet


**If you think of clinical mammography applications, the patient expects to walk out of the facility with a preliminary result
No way this can be done – results come weeks later**

Switching gears...

I have been going on about MRI-compatible scanners now for a while...

Let's talk about MRI for a few minutes.




MRI in 5 minutes 😊



- Protons have a magnetic moment “built in” [not only protons, btw]
- Normally, that momentum vector can point whatever which way
- Apply a magnetic field, and those momenta align their directions along the field
- If you “push” them out of alignment, the a-momentum vectors start to precess
- They precess with a characteristic Larmor frequency $\omega = \gamma \times B_0$
- γ is the gyromagnetic ratio, 42.58MHz/T for a proton
- By setting the right field strength, you can dial in a particular Larmor frequency
- For example, with 1.5T, you get 65.16 MHz

MRI in 5 minutes 😊

- If you apply a radio wave with the exact Larmor frequency, you “tip” the momentum vectors in a coherent way
- As they relax, they emit RF at the same frequency
- If there are many protons -> strong signal; not so many → no or weak signal
- you can use this to measure the proton density




MRI is hard

Why didn't I think of that, and get the Nobel Price! Well...


- **First of all, if you have a totally homogeneous field, all protons have the same Larmor frequency – not very useful**
- **On top of your homogenous field, you apply a small gradient so the frequency becomes position-dependent – now we are getting somewhere**
- **You could send a number of pulses which “scan” the frequency range and measure the response – was done in the early days, way too slow, but it would work**
- **Rather, one sends a pulse with a narrow “white” frequency band, and measures the response for each frequency at the same time**
- **One shot gives you the proton density along the field gradient (one axis of a picture)**

Making Gradients


- 3 sets of coils let you make any gradient direction you want – loops for z, and “saddles” for x and y
- Modern scanners use “fingerprint” designs with high efficiency



- Remember, that's on top of a static high-strength field




But Wait! We are missing one dimension!



The points {A,B,C} and {D,E,F} radiate at different frequencies according to the field strength

However, A,B and C are indistinguishable! (and D,E,F, too)


Throw in a *phase-encoding* gradient:



By momentarily “speeding up the clock” for a region, you gain a phase shift for protons there

And you can assign a particular phase shift to a region

Phase and Frequency Encoding



Now you can get x by frequency and y by phase.
 N pixels in y form an equation system with N variables (so a 100x200 image takes 200 measurements)

You need N measurements with different phase encodings to solve this

Remember: by the time you measure, A and B have the same frequency again (only a phase difference)

2 pixels A and B take two "shots" S_0 and S_1

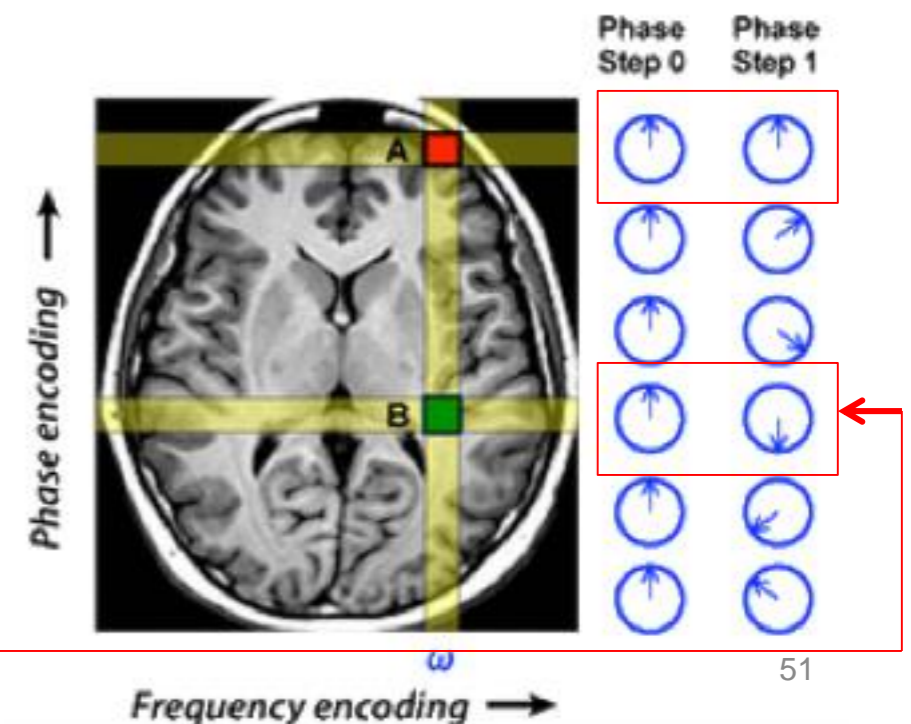
$$S_0 = A \sin(\omega t) + B \sin(\omega t) = (A + B) \sin(\omega t)$$

$$S_1 = A \sin(\omega t) - B \sin(\omega t) = (A - B) \sin(\omega t)$$

$$A = \frac{1}{2} (S_0 + S_1)$$

$$B = \frac{1}{2} (S_0 - S_1)$$


This is easy only for a 180° phase shift



Back to the MRI-compatible PET scanners


- **The trend in modern medical imaging goes to *multi-modal* imaging**
- **Combine two different modalities to combine their strengths**
- **PET-MRI – MRI gives you sub-millimeter resolution, PET gives you the metabolic information**

The 9.4T UPenn MRI for small animals




Our “UPenn” PET scanner mounted on the MRI tube

- Our UPenn PET scanner was designed to fit into the bore of the Varian 9.4T (as strong as they come)
- BTW – 9.4T \rightarrow 400MHz Larmor. You can buy them “by frequency” – 7T (300) etc



The Brutally Hostile Environment inside a MRI Magnet


- If you were to get too close to the bore with a steel screwdriver, you would not be able to hold on to it
- A number of early versions of boards got ripped apart because some ferromagnetic material was present in unexpected places (e.g. Ni in solder joints)
- **The RF pulses totally overwhelm the PET electronics.**
- **The plot shows counts registered in the PET scanner as a function of time – the spikes are the interference from the RF pulses**




**Still:
It did work!**

Simultaneous PET-MRI Rat Brain Images

MRI



PET



Overlay



This is sort of the “money shot” – one of the all-out-everything-worked scans of a rat

What you want is *simultaneous* PET and MRI. If you do it sequentially, it’s much harder to combine the images

Wrapping up with 3 quick items

Out of a stroke of good luck we were able to make an actual *clinical test on human patients* in Taiwan

Volunteering women who had already been diagnosed with breast cancer were scanned with our breast scanner prototype and MRI

The protocol and MRI magnet /gantry design didn't allow for simultaneous scans though

We had to resort to sequential PET-then-MRI scans

Our about 3cm deep FOV in the prototype allows to only image a thin “slice” of tissue

FDG scans – sugar uptake by the tumor

Fantastic images despite some compromises




**Breast Scanner
Prototype**

Human trials

Subject 3 –
8.67 mCi injection
Silicone Implant

MRI




PET

**PET-MRI
(sequential)**

Human trials

Subject 3 –
8.67 mCi injection
Silicone Implant



The volunteering women already knew they had breast cancer

In one patient, our scanner detected a 2nd tumor that had been missed (same breast)

Once / if the full-sized scanner can get FDA approval, new tools for breast cancer detection will be available

Time-Of-Flight PET – the Holy Grail of PET


Remember my “a few 100 ps” timing resolution to be able to find coincidences?

Well, there is another front. If we could use TOF to restrict the decay position along the LOR, the images would dramatically improve

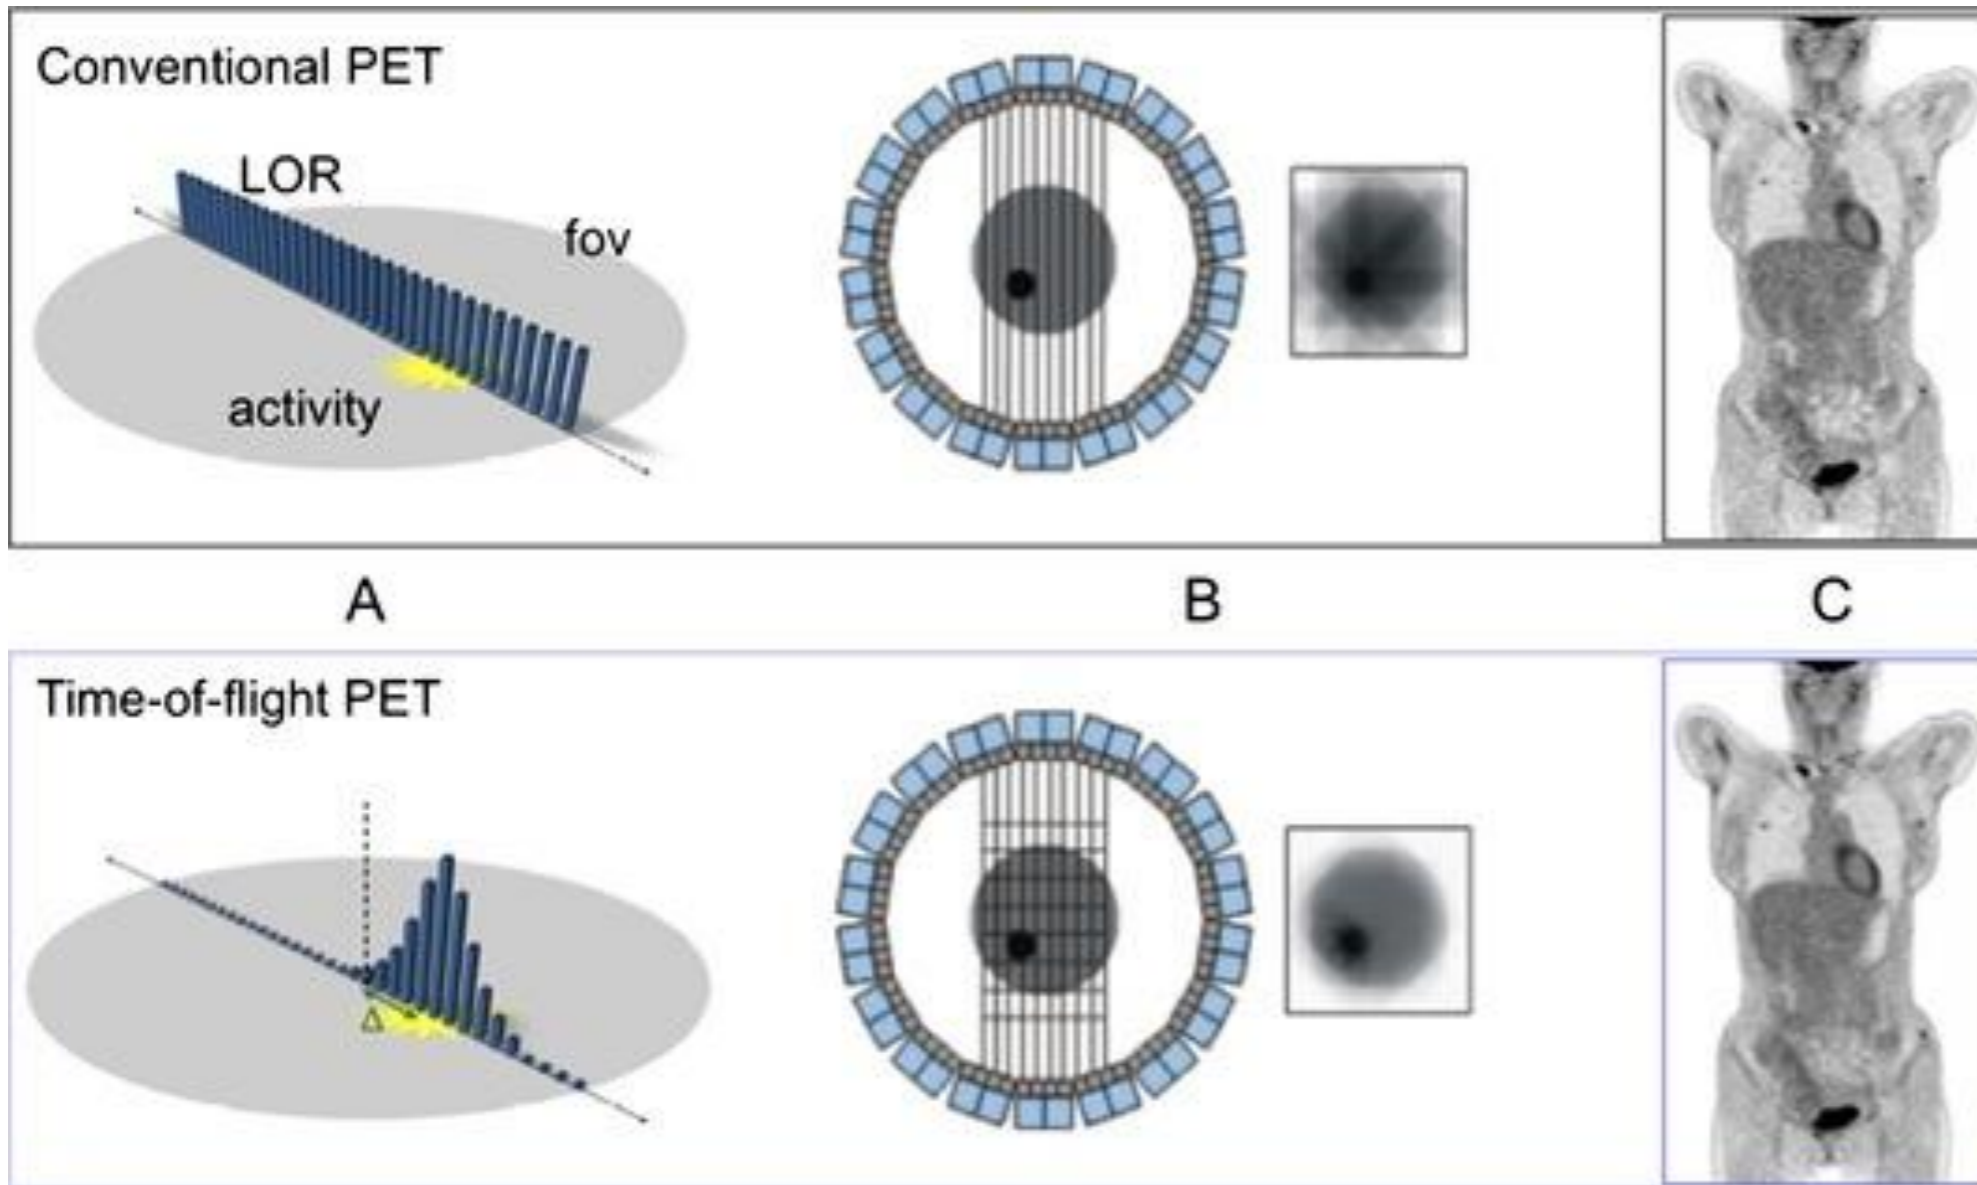
A human scanner is ~30cm

10ps timing resolution – 3mm spatial resolution

Sub-100ps timing is hard! Especially if you have just 511KeV to work with.



TOF-PET – the Holy Grail of PET



If you manage to restrict the position to just $\frac{1}{2}$ of the LOR, you gain a factor of 4 in image quality (or can use $\frac{1}{4}$ of the radiation dose)

Many groups are working on this

One promising technology are “Single Photon Avalanche Detectors” – google it...

In-Beam PET

Trend in cancer therapy: Proton beams (or even light-ion beams)

More complicated, expensive (hadron beams are much harder to make)

But:

Why?


Photons and electrons cause damage all along their path

Protons and ions cause damage at one narrow point

Energy
of tissue

Photon

Protons or ions




Using protons or ions, you can target a particular depth in the tissue

(read: where the tumor is located)

Hadron Beams are hard

This is the Paul Scherrer Institute's Proton Treatment Facility



Way beyond the capabilities
of a typical hospital!




<https://www.psi.ch/protontherapy/>

Brain Tumors



Here you can usually not afford to cause damage along the entire beam path

Proton beams ideally suited!

You “wiggle” the depth of the Bragg peak quickly by either adding more material in front, or changing the beam energy



Now you have a *very* powerful weapon!
And a new problem: “*Friendly Fire*”




If you miss your target area, that part of the brain (or the optical nerve, or whatever you hit) is dead


Throw in in-beam PET

The idea: Use the positrons produced by the beam to image the “target area”

More energy loss = more positrons



Creation of positron emitters via nuclear reaction



- Therapy Verification
- Mispositioning

- Correct for individual “body types” (tissue density)
- Organ movement correction

A brief message from people who helped 😊

Special thanks to

- David Schlyer, BNL, USA
- Paul Vaska, Univ. of Stony Brook, USA
- Réjean Fontaine, Univ. of Sherbrooke, Canada
- Steve Pickup, Univ. of Pennsylvania, USA
- Sepideh Shokouhi, Vanderbilt University, USA
- Fine Fiedler, Helmholtz-Zentrum Dresden-Rossendorf, Germany

THANK Y'ALL!