Now, when we talk about quality control, we're referring to those procedures required to ensure that the distribution of radiation emitted from a patient is accurately reflected in the measured raw data. Now, there's several aspects or facets to quality control. There's quality assurance, where we want to verify that the scanner is operating properly, and we want to identify any problems with the scanner prior to scanning patients. Quality control can also refer to those corrections, or the acquisitions of those corrections, that are required to compensate for known imperfections in the measured data. Before we get to quality control aspects, I just want to review a bit of PET physics. I think we all know this, but in PET we've got different types of coincident events. We've got a

true event, where we have an annihilation. The gamma ray is traveling at opposite directions and detected in detectors and a line of response is drawn in between those two detectors. Each individual annihilation photon is sometimes referred to as a single. We can have a scatter event where one of the annihilation photons scatters, is detected, and still has energy within the energy window of the scanner. In that case, now, we're going to draw our line of response in between these two detectors. You can see that it does not pass through the point of annihilation. It's uncorrelated with that, that location, so scatter is going to add a background and reduce the resolution of the scanner. Unique to PET imaging is the idea of randoms, where we can have multiple annihilations occurring

simultaneously. So we can have the annihilation photon from, or photon from one annihilation being detected in coincidence with an annihilation photon from another annihilation. Again, line of response drawn between these two detectors and you can see that this line of response does not pass through either of the points of interaction or annihilation. So, in PET imaging we are requiring prompts, that's the data that we acquire, is composed of the true, plus the scattered, plus the randoms, and scatter and randoms are the noise, the true is the signal, that's the good guy. So for PET imaging then, data acquisition, we have one detector that's in coincidence with another detector at the opposite side of the gantry and, in fact, it'll be in coincidence with all those detectors on the opposite

side of the gantry whose lines of response subtend the transaxial field of you. So this detector is acquiring a fan beam acquisition. Similarly there'll be another detector at the same time acquiring its own fan beam acquisition, and so we have to organize these, this data into sinograms, so what we do is we sort parallel lines of response. So all those lines of response that are parallel to each other are mapped into the sinogram in one row. Similarly, if we have another group of lines of response at a different angle, all parallel to each other, they'll map into another row in the sinogram, and any one given detector will map into the sinogram space as a line, with either a positive or negative slope. So for this detector, for the blue one at the bottom here, I'm just drawing it arbitrarily as that.

It's important to keep in mind that for a sinogram, the pixel in the sinogram is related to one line of response, for a line of response in the, in the gantry, and associated with that line of response are two crystals, one at each end of that line of response. In PET imaging we use block detectors, which is a 2D array of crystals attached to four photo multiplier tubes, so the block can be just a big chunk in which we've got cut lines of different length or we can have individual discrete crystals that have had their transmission and reflectance properties altered and kind of all grouped together into this 2D array, but the idea is the same in that we're creating light guides in the block and these light guides define the pattern of light distribution in the block at the, at the kinda, where the photomultiplier tubes

are located.So this determines the light sharing or the light output to the photomultiplier tubes. So, if we have a simulation event occurring here, because of this cut line the light will not spread out as it moves towards the photomultiplier tube, whereas if we have a simulation event in the center we have maximum spread of the light as it moves towards the photomultiplier tubes.Then at the back end of this where we've looking at it from the bottom where we've got the array of crystals and the four photomultiplier tubes. For a scintillation event, the photomultiplier tubes as you know will detect a certain signal. We can compute an x and a z value or coordinate location for that signal or for that scintillation event just like in a gamma camera with anger logic. So each crystal

then will produce a unique combination of signals to the photomultiplier tubes, and then we can compare these signals, x and z, to preset values in a two dimensional look up table, what's called a crystal map. So in this way we're mapping the x and z values for that one scintillation event to a single crystal, and we'll talk a bit more about that in a minute. So, we've got a whole bunch of blocks and crystals which are grouped into blocks which are then grouped into either modules on a GE system or what are called buckets on a CTI system, but they're not all going to be operating the same, try as we may. So we have to do some corrections, one of which is a normalization correction. The purpose of the normalization correction is to correct for the variations in efficiency in the lines of

response in each slice of the sinogram. So you can think of it as a uniformity correction for SPECT imaging. For 2D imaging we can do a direct measurement because there's, you know, a manageable amount of lines of response and this is a direct measurement using a low activity source, either a rod source or a uniform phantom that we acquired for a long period of time because we want good statistics in the normalization correction. You know, so you have to run this thing for the order of twelve to twenty hours. For 3D imaging, there's an order of magnitude, more lines of response so we can't really do a direct measurement in any, you know, good amount of time, so it's an indirect measurement where we use a uniform phantom, and this normalization correction is acquired quarterly or

after system maintenance, and you can see normalization from our scanner using a rod source, so you can see the structure in it, you can see no, it's not uniform, but that's the point that this correction will even out all the variations in efficiency. So, here I show a sinogram of a uniform cylinder before normalization shown here, and again you can see that structure from the different crystals and then after normalization we remove that structure. So, back to quality assurance now. I think for quality assurance test requirements we need a well defined regimen of measurements. These measurements should be quick and easy to conduct because, you know, we've got a technologist that comes in the morning. At our site we, the first patient's, is injected at six in the morning so the

technologist is arriving at five in the morning. We don't want to make it too onerous for the technologist to do a bunch of QC before we start scanning, so, quick and easy to conduct. The measurement should be sensitive to the mode to failure of the scanner. I'll show an example of that in a minute, but if we want to look at a block and to see if it's going to fail or not, we better have a test that will be sensitive to the fact that the block could fail. Also, preferably, we would want quantitative measurements, not qualitative measurements. And, again, we've got a technologist that's in early in the morning. I don't want that technologist making a subjective call as to whether or not the scanner is good or bad. Some technologists I could trust, and some I probably wouldn't trust. So, I prefer a

quantitative measure where we can say yes or no the scanner's good or bad. Now one source of data for quality assurance are blank scans. On a stand alone PET system, a blank scan is used with the transmission images to create attenuation correction factors. And, I said, it's, it's acquired daily so it's a good source of quality assurance data. Newer PET/CT systems really don't have rod sources any more for transmission imaging, but on GE systems, for example, we still have a rod source for normalization so we could still use the rod source to acquire a blank scan. CTI systems don't have rod sources but you can do an emission scan of a uniform cylinder. The point is you're putting some activity in the scanner daily and looking at the performance of the blocks. So, here's a good blank scan.

You know, it looks like the normalization. If we have a bad block, shown here, it'll have a black streak in the sinogram. If we have some type of a miscoded event where we're not positioning the location of that simulation properly, it will, we could have one area shown here where we might have increased counts. So we can acquire blank scans or, or something similar daily, and I'm a strong advocate of that we need quantitative analysis of blank scans for the proper QA of a PET scanner because this quantitative analysis of blank scans will allow us to validate the system calibration. I'll show some examples of this in a minute, but the engineer does some tuning on the system, the calibration of the system, we want to make sure that the system's ready to go. Quantitative analysis of blank scans

will provide information on the crystal, the block, module or bucket efficiencies with respect to a system average, so we can kind of see how the individual blocks are behaving and which ones are detecting less counts in the system average. We can also monitor the stability of the system. You know, over time the system might drift. You know, we want to watch that. Quantitative analysis of blank scans will also ensure the validity of the normalization. Remember I mentioned that the normalization is acquired quarterly, so who's to say after one month of, you know, since we've acquired the normalization scan that that normalization scan will still properly correct for the efficiencies in the lines response. The system might have drifted, and we might not have a proper normalization scan.

Well, if we do quantitative analysis of blank scans, we can watch the system's stability and be confident that the normalization correction will do its job. When we first got our PET scanner, there really wasn't any software to do quantitative analysis of blank scans, so we wrote our own. We have a GE system. CTI at the time did have software to quantitate blank scans. GE now does have software to, to do stuff but we still use our own, our program that we wrote. It's got two, operates in two modes. One is just a daily QA where we calculate a global and a sliced normalized mean square error by using a reference scan, and this reference scan is defined as the blank scan that was acquired immediately after normalization. So, we calibrate the system, do a normalization correction, and

then immediately do a blank scan, which we save as a reference scan, and any subsequent daily blank scans will be compared to this reference scan. And we can plot the global normalized mean square error as a function of date here, plotted or listed in a table, we can show the normalized mean square error for each slice, and we also listed here we alert the user if there's any slice that has a normalized mean square error which is too large. So, for example, a slice twelve here, we're getting a block that's starting to fail, and it lights up, so the technologist can kind of watch this and you can kind of see it here too, it's a point that's higher than it should be. The system can also, this program also operates in, in an advanced QA mode. The service engineers use this, where we can compute

the bias and efficiency for each crystal block or module, and in those crystals blocks and modules with a bias greater than some preset limit are listed in a table. So here we're showing the bias for the individual modules and here we've listed the modules that are greater than a certain percent, the blocks that are greater than a certain percent. This is a 3D rotatable bar graph where we show the bias in, for one block, the crystals in one block, so we can see what crystals have less counts or more counts than the system average, and we can also compare the module bias with the reference scan so we can see what modules have drifted over time. So for quality assurance, then, I think the first step is to check to see whether or not the engineer calibrated the scanner properly. So, this might be where you've just

received the scanner, before you do acceptance testing, you know, is the scanner ready to operate or perhaps you have a scanner, a block fails, the engineer comes in and replaces the block but did he calibrate the system properly, and what he's going to be doing, what the engineer is going to be doing, is acquiring detector and electronics characterizations, and these are the corrections that are applied to the data on the fly as the data is being acquired. We can do singles mode detector calibration, where we're just looking at single photons we don't care about coincidence, so we'll acquire crystal maps, do some photomultiplier two gain adjustments, on some systems we might have to do an energy map, and its also involves coincidence time and calibration. So, the crystal map, then, the purpose

of the crystal map, is to map the position of the detected event to a specific crystal. So, that's that 2D look up table I showed earlier, and this crystal map is obtained quarterly or after any detector maintenance. So here I'm showing the count distribution of the block, so this is just an image of the block that's been irradiated with the uniform flux of activity, and the system has algorithm on it which will calculate the crystal map so it just kind of segments out the individual peaks which represent the individual crystals. So, for example, any pixel in this region right here will be assigned to this one crystal. Now, there is some false positioning of these events because, you know, the tails from this, this peak will bleed or kind of fuse into the adjoining or adjacent crystals, but it works quite

well. So, we can use the, we have to do QA on the system and validate or verify the crystal map. For example, here's one block in our system where the crystals aren't all in a perfect grid, there's separation right here from the centerline. The algorithm really can't track the fact that this crystal is shown so far down, so it draws a line that kind of bisects that peak. So, if we look at a blank scan, then, this pixel here or this crystal here is this kind of thin black line because there's less counts in this crystal than there should be. On the adjacent sinogram, you see this crystal here as a miscoded event, it's, it's higher than it should be, and if we look at the three dimensional map of the crystal biases for that block, you can see the fourth crystal down on this row is, one, two, three, four, right here, it's

showing decreased bias, it's detecting 60% less counts than it should. Then the next crystal down is this one here, which is detecting about 40% less counts than it should. So, this is a good tool for the QA of crystal maps. Another calibration that the engineer does is the photomultiplier tube gain adjustment, the purpose of which is to balance the gain characteristics of the photomultiplier tubes in a block. So, this will compensate for photomultiplier tube gain drift with temperature or with age. Vendors usually recommend that it be acquired weekly. We acquire ours daily or, again, after any detector maintenance. So, this is kind of a graph of the bias for one block in our scanner over time where we do a update gains in the morning and then we'll do a blank scan, and then we'll look at the performance of

the individual blocks. So, for that one, for one week here we had the update gains running and then we do, like I said, then we do the blank scan and you can see that the bias was quite stable. It was about minus 2%. Then there was one week where we didn't do update gains, but we still did blank scans in the morning and then we would do blank scans in the afternoon after we were done scanning patients. You can see that there was a large drop in the bias, and then over the course of the day that the system would drift, and then we turned on the update gains again, and did our blank scans in the morning, and everything was stable, and just to convince myself that things weren't drifting in the afternoon we did another week with update gains, blank scans in morning and afternoon, and you

can see that it was quite stable. Now this was the worst block in our system. We did have some blocks that didn't, did drift like this, but it was, it's, you know, it's important to know which blocks are drifting and, and by how much. Coincidence timing calibration is important. This correction adjusts for timing delays so events from all blocks are time stamped equivalently. This is acquired weekly or again after detector maintenance. Remember, in PET imaging the name of the game is coincidence detection so we want to make sure that a signal that originates here and a signal that originates at the same time here make it to the coincidence comparison circuitry at the same time. You don't want any timing delays, so this calibration does measure those timing delays and make sure that

events that are emitted or propagated from here and here at the same time reach the coincidence comparison circuitry at the same time. Here's the, a blank scan that was created with an incorrect coincidence timing calibration. Now this is a real crude example. We have two GE advanced PET scanners, so I just took the coincidence timing calibration from one scanner, put it on the other, and acquired a blank scan. But, you can see that it's not streaks for individual detectors like it was with a, with a bad block. Here it could be individual pixels, and it's just, remember each pixel is a line of response that defines two crystals so if the timing's off for those two crystals, you could have decreased counts. So, that's kind of the quality assurance of the scanner, you know, and tuning and the characterizations. What about daily quality assurance and, the question we're asking fundamentally is, can the scanner be used today. Potential problems for the PET scanner: system stability or system drifts, we've kind of talked about that a bit already, detector module, photomultiplier tube or preamp failures, loose cables or connectors, sometimes we have inoperable gantry motors, we sometimes have problems with our source loader, and I mentioned in passing our daily QA regimen is, we have an automatic gain adjustment, so this kicks off at four in the morning and then it's done by the time the technologist comes in who then does a blank scan and we quantitate the blank scan. It's important to note that the, this, this QA will detect but will not prevent all these problems. So, I've already

showed you this where we can use our, our daily QA to assess the stability of the system and the effectiveness of the update gains. I mentioned in passing earlier that any QA test must be sensitive to the modes of failure of the scanner. When we first developed that QA program, we had the technologist monitor the global normalized mean square error. Remember, we're comparing a reference blank scan to a daily blank scan, computing global and slice normalized mean square errors. So here for one block I'm plotting the, the global normalized mean square error as a function of time in green, the slice normalized mean square error in magenta, and then on this axis on the right is the block bias. So we had one block in the scanner that was quite stable, operating at a minus 1%, but then all of a

sudden it started to fail, and then it really started to fail in one week. You can see that the global normalized mean square error was insensitive to any changes in that block, whereas the slice normalized mean square error and the slice in which that block resided was, was sensitive and it did start to increase as the block started to fail. The, we did some calibrations, tried to bump up the, you know, tried to fix it. It worked for about a day and then it started to slide again. We tried a bit more, and then it started to slide again, so we basically just, at this point, replaced it. So, the new block actually has a bias that's quite high, 5%, but it is stable. So, again, QA must be sensitive to the modes of failure of the scanner, so we have our technologist look at the slice normalized mean square

error. We can also use the blank scan to, as an estimate of system sensitivity, so, you know, just the, the acquisition rate, the trues for the, for our rod sources, so you can see when we've replaced a rod source, but we also just kind of monitor it for any changes. You know, here's one deviation here where we had a problem with a block and it picked it up as a decrease in sensitivity by few percent. Now you don't need any fancy software to figure out that you've got a loose cable. For whatever reason, we get this once in a while where you'll, the scanner was okay one day; you come in the next day and a cable is loose and we've got an entire module not just, you know, a block but an entire module that's been disconnected from the system, so we get a big black streak here and here and there's, you

know, there's some signal being transmitted but we still have, you know, an adjacent sinogram decreased signal because of this loose cable, and then we just reconnect the cable and things are better. Sometimes if we don't know whether to scan or not we'll, we've got Germanium phantoms which we will image for a short period of time just to see whether or not there's artifacts, so this is uniform, a uniform Germanium phantom reconstructed images and you can see streaks that are emanating through the or passing through the phantom up to the position of that module in the detector, and then after we've tightened the cable it's much better. So QA schedule, then, for detector and electronics characterization, we do the crystal map quarterly, photomultiplier tube gains we do ours daily or you can,

depending on your scanner, you can do it weekly, coincidence timing we do weekly, systems corrections, normalization we do quarterly. I didn't really mention this. This is scanner calibration. This is the correction that converts counts per pixel into activity per concentration. We'll do that quarterly, or if we're doing a lot of quantitative work for research protocols we'll do that as needed on a, on a weekly basis. Switching gears a bit now to NEMA and acceptance testing, specifically NEMA NU2-2001 performance measure of positron emission tomographs. We'll use NEMA testing for, or for acceptance testing, and again the question we're asking is can we use our new scanner. It's been installed and looks calibrated. The engineers have done a good job, but can

we start scanning patients with it. We'll also use NEMA for any annual QA where we're just asking is the scanner still performing within specifications and as expected. The phantoms for NEMA 2001, there are three phantoms, the main one is a scatter phantom which is a 203 by 700 millimeter phantom with activity in the line source. So, it's 700 millimeters long, which is quite long, so it usually comes in different sections, in sections, three or four sections. I like this one from CIRS because it has an overlap in between the sections so there's no chance that the gamma rays or photons from the line source can pass straight through and not be attenuated by the phantom. Also, it only comes in three sections so you don't have a seam right in the middle of the field of view, axial field of view. There's a

sensitivity phantom which is, again, the 700 millimeter line source with aluminum sleeves that can go around it, and there's an IEC image quality phantom as well, and we'll talk about these but this phantom's got spheres of prescribed dimensions and a lung insert. So NEMA 2001 measures the performance of PET scanners under conditions that attempt to represent whole body studies, so the phantom is of greater length than the old NEMA standard, NEMA 1994. So because of, we've got a phantom of greater length, we've got out of field activity, which is important for three dimensional imaging. NEMA 2001 also standardized the oblique lines of response manipulation that's required for 3D acquisition. So, if we're doing any quantification of sinograms, of 3D sinograms, we have to use

single slice re-bending to collapse those 3D sinograms into 2D sinograms. And, a good paper to read on kind of the rationale behind NEMA 2001 was written by W. Witherspoon and others published in <u>The Journal of Nuclear Medicine</u> in 2002. So the performance measures for NEMA 2001, spatial resolution in the transaxial and axial directions, sensitivity, if there's a scatter fraction, count losses, a count rate correction, and then this image quality, and we'll just talk about some of these. For spatial resolution of a system represents its ability to distinguish between two points of radioactivity. So, for this test we're going to use F18 one millimeter point sources in air in six locations. So, we'll have it in three locations in two axial positions. So, we've got it, these are the x and y coordinates in the scanner for one axial position, so it's 0 and 1. We don't put it right up the center, we kind of move it up one centimeter, 0 and 10 and then10 and 0 centimeters, and we'll position this point source right in the center of the axial field of view and then we'll repeat all these three measurements one quarter of the axial field of view from the center of the axial field of view, if you know what I mean. We reconstruct using filter back projection and a ramp filter, and it's important to have the image pixel less than one third the expected full at half max. So once we have our reconstructed point then, we draw a profile through that point. The profile width is approximately two times the full at half max. So, we're not doing like a nice single pixel profile, we're going to have a big thick

one that goes, you know, centered on the maximum, and the idea behind this is that we want to reduce any variability in the measure. We don't want the measure to depend on where that point source is located within that voxel. So, we won't get the best measure of resolution but the most repeatable, and we report the full with half max and tenth max and the radial, tangential and axial directions. Transaxial resolution, then, for full with half max at 1 cm. and then at 10 cm. for a few systems. I just grabbed some product data sheets and listed them. For the Phillips Allegro we've got about 5.5 and then a 5.6 mm resolution. The GE DST has 6.3, 6.0, I forget the, millimeter pics, pic or crystals, so it's resolution 6.2 at approximately the center and drops off a bit to 6.7. The CTI High Rez has 4 mm

crystals. It operates in two different modes so you can get the good resolution of 4 mm with the purchase of optional software for reconstruction, but 4.6 mm at 1 cm, 5.8, 10 cm off axis, and then kind of the lower resolution mode of 6.5 and 7.5 mm. The sensitivity of a scanner represents its ability to detect annihilation radiation. So, this is the rate of true coincidence accounts per unit activity in the absence of attenuating media. So the sensitivity is expressed as counts per second per kBq, using SI units, but sometimes you might see it as a thousand counts per second per MBq, but, and the rationale behind this test is, you know, you need material around the source to ensure the annihilation of the positrons but then, unfortunately, this material will also attenuate the annihilation photons, so you

want to measure sensitivity without attenuation, but you kind of need attenuation. And the technique is based on a paper by Bailey, published in the European Journal in Nuclear Medicine in 1991, Here's the sensitivity phantom so we do successive measurements with this 700 mm line source with only about 2 MBq roles of F18 surrounded by nested, known observers. So, here's a graph where I've shown the acquisitions with the different sleeves, you know, with all the sleeves, and then when you take off a sleeve, take a measure, remove one of the sleeves, take another measure, you know, do this five times. You do a least squares fit to this, to these points, to calculate R_o the count rate with 0 sleeve thickness and you also are fitting for the slope, which is the attenuation coefficient

and then the system's sensitivity is just this R_o divided by the activity in the line source. It's the activity in the entire line source. And, sensitivity is measured at radial locations of 0 and 10 cm, and we report the system's sensitivity, the total system's sensitivity that I just showed you, and you can also report a slice sensitivity profile where for here I'm just showing the slice sensitivity as, or the sensitivity as a function of slice, for our kind of older GE Advance scanner operating in 3D. And some numbers for the newer systems in counts per second per kBq for the GE DST can operate in 2D and 3D. 2D we've got about 2 counts per second per kBq, in 3D 9.3. They use all rings are coincidence with each other in 3D. The Phillips Allegro is about 4.4 and the CT High Rez with Pico

Electronics is 4.5 counts per second per kBq. The scatter fraction now. The scatter fraction is a measure of the relative system's sensitivity to scatter, so the scatter fraction is equal to the amount of scatter, divided by the scatter plus the trues, so there's no randoms involved. So, again, we're going to use this long 700 mm source with a line source located, you know, running through it and it's the, scatter fraction's measured with low activity to avoid any random coincidences, dead time, and pulse pile up. So low activity, and NEMA stipulates that we need a randoms to trues ratio of less that 1%. And the scatter fraction for, again for different scanners, I also show here the lower level discriminator which is the, you know the energy level, of the lower window or the lower energy of the, of the

energy window. On the GE DST it's 19%, you know, makes sense because we've got a collimator in there, the 3D's 44% with a 375 keV lower level discriminator, Phillips Allegro's 40% with its GO, GSO detectors and a lower level discriminator of 410 keV, CT High Rez operates at 36% and it's got a 425 keV LLD. Count re-performance. The measurement of count performance gives an indication of scanner performance as a function of activity. So, again, we're going to use this long polyethylene cylinder, and we need a high initial activity of F18 to really saturate the system. In 3D typically we, you need around 800 MBq, for 2D you need about 5 GBqin 3 ml which is pretty hard to do so usually we don't, you know if you've got a GE system you just measure the count rate

performance in 3D and not really worry about the 2D, and we're going to acquire the data until the randoms and dead times losses are negligible so this is going to run between 14 and 18 hours, and actually the tail end of this when the randoms and dead time losses are small is the scatter fraction measure, so we combine the scatter fraction with the count rate performance into one acquisition. And what's computed out of this or calculated is the noise equivalent count rate. The noise equivalent count rate is a figure of merit relating the scanner performance to image signal to noise ratio after we've corrected the sinograms for randoms and scatter. It assumes, you know, ideal and correct randoms and scatter corrections. This is the formula for the noise equivalent count rate where we've got the rate

for the trues squared divided by the rate for the trues plus the rate for the scatter plus the randoms multiplied by this K value. K can either be equal to 1, if we're doing calculated randoms. So, if we don't measure the randoms independently but we just calculate it, K is equal to 1. If we do a delayed event window in subtraction, K is equal to 2. But, for NEMA, they stipulate that K is equal to 1, so we're doing a calculated randoms. And, what's reported is the peak noise equivalent calibrate at, and the effective activity concentration at the peak. So, just to show you 2 graphs, again on our older GE Advance, the magenta is the noise equivalent count rate where we use the NEMA, where we do a calculated randoms correction, the green is the noise equivalent count rate where we've

measured the randoms and, so again, we're going to report the peak noise equivalent count rate and the activity at which that occurs. Now I should mention this effective activity concentration. So, we've got activity in this line source that we divide the activity in the line source by the volume of that entire polyethylene phantom which is 22 thousand milliliters or cc's, and you can just see from this to that the noise equivalent count rate does change, that the peak and the value at that peak does change if we do a 1R versus a, you know, a measured correction for randoms. So, count performance then, again, is the peak noise equivalent count rate and the activity at which that occurs for the CT High Rez, with the PICO electronics, it operates at about 85 thousand counts per second and an activity at that peak of 21 kBq per mil. The GE DST's 84 and 49 in 2D so we

need a lot of activity to saturate the system in 2D. In 3D, 63 thousand counts per second for the peak at 12 kBq per mil, and the Phillips Allegro is 30 and 9.] And, finally, we'll talk about image quality measurement, which is standardized imaging situation that simulates a clinical whole body imaging condition. So, we've got a phantom that consists of a torso phantom with hot and cold spheres and a warm background and a long insert, so we kind of assemble this thing all together and then we also take that long scatter phantom, fill it with some, fill the lines first with some activity, and abut that to this image quality phantom just to, again to represent activity outside the field of view. So in this phantom we have hot spheres of diameter 10, 13, 17, and 22 millimeters internal diameter and the

two largest spheres of 28 and 37 millimeters are cold, they're just filled with water. There's a lung insert and the activity in the hot spheres is 4 and 8 times that of background and the activity in the background is 5.3 kBq per mil. So this represents, what this, where they get this number, this represents a 370 MBq injection into a 70 kilogram person assuming there's no excretion and there's no time for decay, there's no allowance for any decay or uptake time. And again as I mentioned we're going to have 3 millicuries in that line source for that scatter phantom abutted to this image quality phantom. And, we're going to simulate an acquisition of 100 centimeters in 60 minutes. So, given this, you know, we're not going to actually scan 100 centimeters we're just going to pretend that that's what

we want to scan and we're just going to acquire 1 acquisition centered on that image quality phantom, so the time, then, of the acquisition is 60 minutes divided by 100 centimeters multiplied by Delta Z where Delta Z is the axial movement of the table, how far that table moves, so it takes into account any slice overlap. If we didn't have slice overlap between our bed positions and then Delta Z would just be the axial extent of the field of view and you know, for most systems with, say, around a 15 centimeter field of view, axial field of view. This, this is, this comes out to 7 minutes. Now this T acquisition time, this was the, you know, when NEMA first met in 2001, this was the time that was required. Seven minutes was, is required for both transmission and emission imaging, but with

PET/CT now, you know, CT is done in a few seconds, most of this time we can do for emission imaging. So it seems, you know, it is rather high for a PET/CT scanners to image for 7 minutes because that's what, we don't do that clinically, but still that's the NEMA specifications so people follow that. I also like to do this test with, you know, what we're doing clinically, just to get some kind of a comparison. Repeat the acquisition three times for each kind of background, and then reconstruct using the clinical protocol. Now the analysis is quite involved and they're not going to really describe it too much, but what is reported is image contrasts and signals to noise ratios for the hot and cold lesions, we have to put regions of interest of the size equal to the internal diameters of these, these

lesions. We also have to put in 12 background regions of each size of region for each lesion kind of concentric, and we put these background regions also in one slice above or, or 10 and 20 centimeters above it, or 10 and 20 millimeters above and below this central slice. We also look at the residual error in the lung and we report the variability in the background. It's also good to do a visual inspection for artifacts because the, kinda there's a variability in background won't pick up any, you know, really won't quantitate any artifacts. And the manufacturers really don't specify right now, you know, what they get for this 'cause they kind of say it, it's a clinical, you know, reconstruct with a clinical protocol but they don't really give any definitions so I encourage you to look at some papers in the

Journal of Nuclear Medicine, for example, or the European Journal where people may have published some of their results and compare how close you are, but the, the vendors right now aren't, aren't including in this measure in their product data sheets. One last point to consider is NEMA and those scanners with LSO, lutetium oxyorthosilicate, because LSO is inherently radioactive, about 2.6% of LSO is lutetium-176 which is a beta emitter, so we do have background radiation just from the scintillators themselves, and this background radiation gives rise mainly to randoms, but there can also be some true coincident events, so the implications for NEMA, then, is that we cannot obtain a randoms to trues ratio of 1%, for example if we're doing the scatter fraction. And a paper by Watson this

year in the Journal of Nuclear Medicine kinda shows some changes that he proposes to the NEMA specifications to deal with LSO scanners. For example, for the count rate measure and the scatter fraction, he recommends that we acquire delayed events to measure, you know we're going to measure the randoms instead of do it calculated. For sensitivity we acquire a blank scan, I mean a blank scan in the sense there's no activity in the scanner we just kind of scan for a while, to estimate this trues background from the radioactive lutetium. So, just in summary then I did a quick review of some PET physics, I discussed a need for QA regimen and described, you know, how the QA regimen can pick up problems with the maps and the, the photomultiplier tube changes, and we talked about the

corrections and the calibrations that are required for PET imaging, and then we discussed the performance measures and, one more thing, I would like to apologize for not having a handout ready, I just got off vacation and the dog ate my homework and everything, but hopefully within a week I will have a handout that will be posted on the AAPM website for this thing so, anyways, thank you very much.