

Good morning. Thank you very much for staying until Thursday, so, that's pretty good, and I see still quite a number of people here, and for getting up early enough. You know, usually I find it difficult to catch the first session, but today I did. So, it's very nice. Today's topic is Monte Carlo applications in radiation therapy. I will cover a few aspects of the Monte Carlo applications. The first one that I will start with is the rationale for doing Monte Carlo simulations at all, and then I will talk about some details like, why we need Monte Carlo in radiation therapy dose calculations. Then, I will talk about a few applications starting with accelerator simulations and radiation detector simulations, treatment planning dose calculations and finally, I will show some examples of beam delivery and dosimetry verification. I'll start with... the definition of Monte Carlo. Of course, first we have to mention the originators, Von Neumann and Ulam, in what they quoted as

in 1949, they proposed that almost, you can derive a solution to any problems if you can precisely define a probability distribution. By doing random sampling or, you know, simulations, you can find a solution. Not necessarily exactly, but you can find something close enough. So, that sounds pretty good. Of course the name of Monte Carlo was quite a thing because we could've called it with quite different names, but we end up calling it Monte Carlo. Of course, you know the city of Monte Carlo, where you can really enjoy your time especially on a Saturday or Sunday or maybe on the last day of a conference, but here, we all staying here talking about radiation therapy. So, obviously the casino there (in Monte Carlo), is very famous and in 1997, after the World Congress (on Medical Physics), we actually had a symposium on Monte Carlo and we went, all went there and became the real losers, you know, because we know most time you lose

when you go to casinos. However, we can demonstrate indeed that, if you can construct a problem and use random sampling, you can get some kind of solutions. So this work is done by Steve Jiang when he was at Stanford. He shows that we can calculate PI using this way, not necessary using many other ways. In this way, say if you go to the casinos or any English pubs, you usually see that, you know, if you play darts, you can actually, throw them on here randomly, and if you do, especially after a couple of pints of beer, you can certainly throw them randomly. Or most time, you know, if someone is so good, he'll always get in here (the center), then it's biased, it's not so good. So, if you think, the ratio of the number of times the dart will fall into the circle to the number of times falling into the whole square is proportional to the ratio of the area of the circle to the area of the square. And of course you know that's PI over four (from the equation

here). So in order to calculate PI, we can simply calculate, count the number of times the dart will fall inside the circle and the number of times the dart will fall on the square, and times four, you have PI. So, believe or not, we can try that. So... start drinking and we will get some numbers. So if you use computers to get a random number and to see whether or not, the dart will fall into a circle, then you can see the number goes up and down like this, with the number of times you sample. By the time you get to about two thousand, four thousand to ten thousand times, you can see the number gets really, really, really close. So, that's pretty good because even... almost seven years ago, with a very slow computer, we get that solution very quickly. If you look at this, by the time you get

to hundred thousand times, that's 3.1431 and by the time you get to ten million times, and a hundred million times this number is almost the same as the one, the same

accuracy as the very clever Chinese mathematician, who derived the value to this accuracy about fifteen hundred years ago. So, that's Steve Jiang who wrote a little program in about 30 seconds, on a very slow computer. You see, you can do it very well. Obviously, that's quite good. However, in radiation therapy or radiation physics, we have to do a lot more than that. What we do is, we try to simulate the interactions or the transport of radiation particles, in medium or in a patient. What we really need is a number of random numbers, probably you wanted to have a set of random numbers, which you can repeat, you know exactly what they are, but if I don't tell you, probably you say "Uhhh indeed, they are random." And if you can ensure that you can use them to randomly sample the interactions or events from the probability distributions, which govern the physical processes. Of course, we are not here just to simulate the random events. That sounds

good for a student, but for us, what we really want is a quick answer, you know, a quick answer that can be used by the physicians, by the dosimetrists or by ourselves. So what we really want is the mean quantities. The quantities that represent what we really use routinely. That's the energy deposited or the dose or the fluence and so on, but if you simulate an enough number of histories, you'll find the value gets... really, really close to what you want. And that's what we are really doing here. So that already raises a question whether we want to simulate randomly, at all. If you know how to derive a precise answer using an analytical equation, just do it. Don't use Monte Carlo, it's going to take too long to get there. However, sometime we found the boundary conditions are so complex; you may not be able to construct something so easily. Then Monte Carlo is nice. Actually, the Monte Carlo method is a method for the lazy guys. They construct a

very simple geometry and it let you do simulations and then you get the answers, and many people even spend years, still can't get an answer, right? So in that sense, it's very nice. So let's see what we can do if we want to simulate photons interacting in a cubic of water. See, in this case, I assume we know the energy of the incident photons, now first, we need to find out where the photons're going to incident and how far the photons are going to go before we reached the point where you have an interaction. Of course, in this case let's assume the photons gonna have a Compton effect. That means the photons gonna to be scattered and the recoil electrons will be generated. Of course, you have to use random numbers, go to the probability distribution to find the angle for the scattered photon and the angle and, the energy of the recoil electron. Of course, now you start following the particles, possibly start with electrons, if they've lower

energies, you can follow them first until the electrons stop, then we get rid of these particles or we can follow the photon. In this case the photon gets out of the geometry very quickly. You lose interest, why should I follow it, just terminate it after the photon disappears. Assuming there's no further effect from that photon escaped from the geometry. But if you come back here, you can see we will follow this electron. An electron can generate another electron. So, that's a delta ray, let's call it. Of course,

nobody knows which one is the real delta ray simply because we find out which one has a lower energy we call it delta ray. Alright we follow the delta ray and probably the delta ray can generate another thing like a bremsstrahlung photon. Of course, then we follow it until it gets out. So, by now, everybody already is either stopped or gets out of the geometry of interest and we say, okay, done, this is a simulation, for one particle, one incident particle. We call the whole thing now here, a history. So now we know when we talk about histories, meaning we start with how many particles. You can start with an electron or you can start with a photon. Doesn't matter, you can start with a positron if you like, for photon or electron transport in general, but that's the way we do it. Of course, now we can do more. In this case, we see what happens with pair production. You have photons to go in to generate your electrons and positrons, and also we can have a photo-electron effect, where you have a photo electron generated. By

simulating many of them, now you obviously can estimate the dose deposited by the photons or the dose distribution from a photon beam like this. Of course, if you have a spectrum or some other things changed, then the dose distribution will be different, i.e., the depth dose curves, the profiles and so on, so forth... However, this is all very easy if you are a Ph.D. student listening to someone really good like Dave in the audience then, after three lectures you can build a program to do that. Not necessarily the electron part of it, but the photon part of it is quite easy. You can write a little program and you run it and you can simulate the tracks of the photons. Things become quite complex, however, if you have a geometry like a linac, it's not so easy. You can spend years just to construct the geometry in order to simulate it accurately and if the geometry is flexible enough and you can change it easily so that it can simulate linacs from different vendors,

even makes some changes. So, the geometry like this can be quite difficult and here we show you the tracks of photons and electrons and so on... I think here we have different colors, but I can hardly see it from here. But, anyway you can see, the geometry is quite complex. So by the time we started doing Monte Carlo we already have the general packages, which you can use to simulate the transport of photons and electrons, and in radiation therapy dosimetry most of the problems can be tackled by these general packages. However, the user code or the program used to run for a specific problem can be different. The more you work on it yourself, you probably make it more efficient for your specific problem. But if your problems are very general or can be very different every time, it's better to use a code, done by other people, which is general and easy to use and probably can tackle different problems. Usually, the result of that is, it takes

maybe a little longer to run because it is general, you know. If it's too general that means maybe it's a little slow, but it doesn't matter. If so you can use it differently. If we go back and look at all these applications of Monte Carlo in radiation therapy or radiation physics altogether, we suddenly realize, almost everything we used including the numbers you saw from the previous speaker, most of them are from Monte Carlo. It's just amazing how we can do it. Even a lot of the numbers we used in the standards in dosimetry, everything is from Monte Carlo. Well, let's look. In the early days you can easily simulate say a few thousand histories and derive some quantities like fluences or spectra, etc. So in the early 70's and so on you see quite a few guys did that and they

used that to calculate stopping powers, stopping power for electron beams and stopping power for photon beams. Especially when you have a photon beam, it goes into the

medium and you have the interactions, giving out an electron spectrum, which is quite complex and the stopping power calculation for that problem can be quite difficult. You can't, really calculate it by hand. Of course, further down the road we see people did calculations for backscatter factors, head scatter factors, phantom scatter factors, primary to scatter ratios, etc. All of these problems can be done once and for all. You do it and you use it in the future. Dosimetry response was difficult. I think starting from the 70's, we know the papers, by Nath et al and also later on by Roger et al and further on, you know, I joined in and we did a lot of calculations on dosimeters. Again, this kind of things may take billions of histories to get a reasonable answer, but you do it once and then you use it all the time. It's very nice. It's very nice to do it and publish it and people will use it for five or ten years. However, for treatment

planning calculations, things become difficult because you almost have to do it every time for a specific case. We use the data from treatment head simulations. Again, these are already better than this (i.e., patient dose calculation), because probably you can simulate for one machine, then you can use it for quite a while, for a lot of patients. And if you have another machine you can simulate it again. But this one (i.e., patient dose calculation) is really difficult. You really have to do it for every patient. So speed really becomes an issue. It's not like something that you can run for three months then come back from your vacation and look at the results. This one people start talking about the number of coffees. How many coffees? How many cups of coffees you drink before you get the results. Of course further down the road, we started looking at dose calculations and also the end results, whether or not we can deliver the dose accurately when

you have to simulate the geometry of the patient-specific beam-modifiers and also the device after the patient, and you can see, you can utilize that to look at the fluence to reconstruct the dose. So it's very interesting to see all these applications. Today, I'll only focus on a few things here. Starting from accelerator and detector simulation and a little bit on treatment planning dose calculation and over to here on beam delivery and dosimetry verification. So now let's look at accelerator simulations. There are two papers, the first one talks about the electron linac and the second one is also electron linac, but it mainly focuses on photon beams. So we can go there and look. There're a lot of papers listed in there and a lot of the results here also were from there. So let's start from a linear accelerator head. You can see the geometry starting from the electron gun, down into the wave guide, to the bending magnet and the target and the scattering foils or

the flattening filters in here and collimators and MLCs. Obviously for a Monte Carloist you can try it starting from here (i.e., the electron gun). But probably you don't have to do it, because if you look at the electrons coming out of the vacuum exit window the distributions maybe, can be simplified, and because you can get reasonable results at the scoring plane. So a lot of people would do that, assuming some very simple geometry and distribution like a pencil beam or a small beam with a small diameter or they all come in at the same angle and so on, so forth or we can use a distribution with a little bit

of angular spread. And probably in the future, we can go back and simulate it for accuracy in the exit vacuum window. It would be very interesting. So far, we haven't done that yet, if we look at the linacs that we simulated, a lot of the linacs have been simulated. If you look at them in the 95 BEAM paper there're a lot of them and later on people

simulated even more linacs, kV units and others, and so on. So you can see it's quite interesting. This is just a schematic diagram, showing the geometry of different linacs. They can be quite different. The good thing about Monte Carlo is, not only we can calculate the fluence, the spectra and also the dose distributions, we can also look at the characteristics very closely. In this case, what we're showing you is a scatter plot of the particles of the photons, the last interaction sites of these photons before they reached the scoring plane somewhere down here. So if you look back and that's the Y-axis, you can see that's the direction wrtx-jaw movement. If you look at that, you can see if you have electron coming down here. We need electrons to interact with the target to generate bremsstrahlung photons and starting from there, some of the photons may just... pshuuu, one step they reach the scoring plane and see some of them interacting here. The last

interaction site is here. And some of them will be in the primary collimator or some of them will interact in the flattening filter and then the chamber and so on and the jaws. You can see from here, you can easily reconstruct source distributions. Oh yeah, really, if you call them sub-sources. You can see particles are actually coming from here, not necessarily you look at where their parents are from, but you can just look at that and help you understand the commissioning process. People talk about the extra focal source, and sub-source distributions, head scatters and so on. It exactly explains how many particles coming from where, which will help you understand and come up with a model. It's very nice, I think. Of course we can also look at the details. Here, you see, this is an electron linac, in this case. We see the applicator here, I can easily tell. Now we can look at the energy spectra from the electron beams and the photon spectra in the

electron beam and the fluence of the electrons, planner fluence of the electrons and also the bremsstrahlung photon fluence from the electron beam. In any case, see here, I can actually draw all these components in the electron beam. Some of them are from the collimator; some of them are from the applicator, some are from the lowest scraper of the applicator. By looking at this, you can easily tell what's going on in there and which component is more important. So I even remember one of the ABR questions on Monte Carlo. I didn't give that question. It actually was asking, which component is more important in the Monte Carlo simulation of an accelerator head? Of course the correct answer will be difficult. No! The question was which one was the least significant in the simulation among a number of things like the scattering foil ...the ion chamber, the protection window and so on. Probably, you can find out which one is the least

significant. This is a very interesting figure showing you the correlation of the electrons reaching the scoring plane, and we plot the angles as the Y-axis and the position as the X-axis. You can see all these flowers or leaves that represent particles coming from the bottom scrapers. The highest is from the bottom scrapper and the second one from the

second scraper, you know, the second scraper from the bottom up and you can see these patterns. And all these patterns are different for different linacs. It's very interesting if after you have done all these simulations and look at them, right away I can tell that's a Therac machine or a Varian machine. So the details are quite significant, although dosimetrically, you really don't need them to calculate depth dose curves. However, maybe one day we'll find out in some applications where you really need to know the angular distribution and so on. For example, if you want to use a telescope to look at

the angular distribution of particles, especially when you calculate the head scatter factors that may become quite interesting and important. You can use it to validate your measurements using a telescope. Now, let's look at the depth dose curves. Depth dose curves for a 25 x 25 field I remember, yeah, until quite recently, the early two thousand, we had publications still talking about it's very difficult to get the 25 x 25 fields with electron beams. And here you see some results from Michael Lee and in his thesis you can see that you get quite good results. We get even better results once we started doing more calculations. This is up to how much time you want to spend on it. You can still see some differences for some beams. This one is for the smallest applicator on the Varian linac, you can see the agreement is very good. If you look at profiles, here is a 6 MeV beam with different applicators. You can see it shows quite nice

agreement, although you may see it here, it's up to about 2%. The point is we do like to get to 1%, and yes you can if you keep trying, you can get a better result. You may have to play with the parameters, like the beam energy distribution, the angular distribution and the spatial distribution of the incident electron beam, hitting the exit window of the accelerator. So if you do that, you can get much better agreement. At some point I will stop, because I don't want to waste my time because your chamber measurement probably has uncertainty, which is comparable. Of course, you see that's a 20 MeV beam. Look at that! You see that the variation in here is actually caused by the staircase shaped scattering foils, so it can be very interesting. If your beam is wrong a little bit, or you move the location of the scattered foil, it may give some problems. So obviously we'll pay attention to that in order to get all these nice results. So in summary, Monte

Carlo simulations have been commissioned for different linear accelerators for Monte Carlo radiotherapy dose calculations. In principle, if you know enough, and in general, we know enough, we can do very accurate calculations. Sometimes you get something wrong, you know, even the vendor may give you wrong information and you just keep simulating. I still remember at Stanford, one of the graduates was simulating something but couldn't get it right. Until Art Boyer came and he said, you know, I have a geometry for the flattening filter from MD Anderson that uses tungsten rather than copper. Then we tried it. Right away the agreement became 1%, you know, originally it was 10% different around the edges. So if you don't get the right geometry it can be very difficult. However, you can't guarantee you know everything, especially when the engineers come in and they'll tune the machine or maybe replace a part and then you may not

get the same thing. You know the simulation that I just mentioned actually it's because the flattening filter was changed. So sometime you do have to play with the material or

geometry a little bit in order to get it right. Some people may say is that right? Are you still a real Monte Carloist? Because now you start playing with parameters, almost like changing the rules, you know, in a casino in order to win. But to some extent, if you still keep everything else random, I look at the results, I am still right. I think if the end result is okay, in this case, I think I would accept some playing of the parameters. Of course, the most important thing is, you have to make sure your final simulations are okay, you know, and that means you have to have measurement based simulations. And then the next thing is to establish measurement based source models. As I said, if you know the distribution, if you know the physics behind it, you can establish a source

model that can be commissioned using measurements, then I think, you can achieve what you want. And you don't really have to do the simulation for every linac. You just do it in the same way. By the way, I always believed the commission of conventional dose calculation algorithms such as the correction-based and also the model-based algorithms like the convolution or pencil beam or super-position convolution. They all can get dose distributions right in homogenous phantoms. If you think about it, that means they can simulate the source and so on okay. So as a Monte Carloist, probably I should pay more attention to the inhomogenities and so on, or maybe to the details of the beam defining components, rather than on the source distribution. Right? Because if they can get it right that means probably we can get it right too and we don't have to do too much about it. Now let's look at detector simulation. Alan Nahum has a paper

reviewing the perturbation factors and so on, you can look at it, and it quoted a lot of simulations and so on, and also measurements. That's a very good paper; you can go there and look. The reason we want to simulate the detector response is very simple because we all wanted to build a detector that is independent of energy and angle, and the geometry and material and everything. So if you have one detector, you'll go to the standards lab to get one calibration and go home, measure everything, the KV beams, the photon beams, the electrons, protons and so on. But we don't have such a thing. The thing is you always have something, which is not perfect, but it's pretty good. And you can look at the calibration curve as a function of angle or energy. You can see, indeed, that it varies a little. So how you can correct for it, if you don't measure your beam exactly at the same beam quality or angle or exact set up as you did it in the calibration. So you

need to have some kind of correction factors. You need to investigate the perturbation effect. And one way to do it is to use Monte Carlo, because Monte Carlo can help and you don't have to do that experiment. Most people believe their experiment and I do too. But the thing is sometime you can't even do it (the experiment) and then the better way to do it is, of course, to do Monte Carlo simulations. And also Monte Carlo calculations can help you understand the physics behind it, if they agree with the experiments. So let's look at here. If we assume the detector is an ion chamber, we know that... even the TG51 (protocol) is still based on the principles of, you know, a Bragg-Gray cavity, although no equations showing that anymore. But if you think about the Bragg-Gray relations, they still have these parameters in it. Then the dose to water can be

derived from the dose to the air in the air cavity, multiplied by the stopping power ratio for water to air. Of course, you know, that's not always true. We need something else. And that's what we usually call the spices, you know, the perturbation factors to make any dishes taste reasonable, otherwise just pure, raw ingredients, usually don't taste so good. Well, we need that perturbation factor in some cases, even the TG51 (protocol) still need you to do this, to do that. The temperature -, the electrometer - and you know, the ion recombination correction, and together with lots of other modification factors. Some of the perturbation factors can be calculated by Monte Carlo because you may find it difficult to derive them experimentally. Of course, that's very interesting. Now, we don't look at ion chambers. Let's look at one interesting problem we investigated quite a number of years ago. I still remember this when I did my Ph.D. and I

submitted an abstract to a conference in Japan and Dave Rogers looked at it (he was a referee actually). He look at it and say, Oh, that's an interesting work. He wanted me to work with him on Fricke dosimeter perturbation factors and we did the simulations. Because the Fricke dosimeter uses Fricke solution, which is very close to water. However, you can't really just use that in water, you have to have a container for it. It doesn't matter what container you use, that container may affect the response of the Fricke solution. Right? If it is say something close to water, that's fine, but if it's something like glass, in this case, the glass material may generate more electrons and so on, or causes more attenuation. So the dose may be wrong if you measure it using a glass container. So what we try to do is, to measure that. Of course, it will be very difficult. How can you measure that? Can you measure it without a wall? Well that would be difficult. Or you

measure it with different containers. That's a good idea. You can measure it with a container made of plastics and with a glass container or with something else. But how do you know which one is the absolute correct one? So one way to do it, I say okay, why not you measure the ratios, although I don't know the absolute answers to each container, but I know the ratios and then I can do a Monte Carlo calculation. So, this figure shows the correction factors calculated by Monte Carlo using different containers and measurements using different containers. I know the physicist there at NRCC, Ken Shortt, was very confident about his Fricke measurement. He said my Fricke dosimetry here has roughly .3, .4% accuracy. I can guarantee you. Other people can't get that accuracy, but I can. The precision is something like that. He just made sure it is super clean. Everything is very clean and the results are more reproducible. If you compare the ratios

of the correction factors, because one is... I remember this is for glass, and the other is for polyethylene, I think, if I remember it right. The ratios, here you can see, agreed to within .3%. That's the first time with a very precise measurement to confirm that the measurement and Monte Carlo calculations agreed to within .3%. The paper was published in 1994, I think, and a few years later they did the measurements again, and they came back, I think, with the same numbers. So, I think they still believe these numbers are true. So, indeed you can see, Monte Carlo calculations can really help you

understand and derive the correction factors. Now we know, we can go back and say okay if the Monte Carlo calculations agreed with measurements, I probably can use the Monte Carlo results, i.e, the individual correction factors (for different materials). Now probably, I can trust the corrections from Monte Carlo alone, i.e., the

perturbation factors for the glass container. So, that's what we did. So in summary, the Monte Carlo detector simulations can be very good. For this case, the Fricke wall effect we know that changes up to 2% so that in some standards labs, they changed the calculations in there to give you the right dose by up to 2%, that includes the NPL in the UK, PTB in Germany, and also the NRCC. Also, we found a very good agreement, that's the, you know, first hand proof that Monte Carlo's working and you can do it very precisely to .3%, if you think about it. We have used variance reduction techniques to make it right and to make it efficient and also we found, that's mainly because of electron transport and, you know of course, photon transport as well, but relatively speaking it's very difficult to simulate High Z materials. If you think about glass it has

High Z materials in it, and if you can get it right the electron transport must be right in there. Now, let's look at ion chambers and here is for low KV x-ray beams. We use a detector recommended by TG61 and we recommend that you can, actually measure dose for KV beams in water, using the formulism we recommended. But you need to have a waterproofed chamber and if you don't what happens? You have to have something else. You'll need a sleeve to protect the chamber. However if that sleeve material is not the same as water, say PMAA, or polystyrene, then the attenuation can be different or so is scatter. So you have to calculate that effect. We did that by doing a calculation with and without the waterproofing sleeve and we did measurements with a chamber, which is already waterproofed and we put a cap in there, it's like a sleeve in there, and with and without it, we can look at the difference. Look at here. The correction factors, are again,

about 2% but you can see they agree. Measurements and Monte Carlo agreed to within .2, .3% again. But it, indeed proved that if you do very precise simulations, meaning you know the geometry very well and you do measurements very accurately, in this case, of course, we have to do experiments very accurately. We know for very low energy KV beams depth dose curves drop very quickly. You have to know the precise location of your chambers. In that case, the ion chamber measurements were mainly done by Jan Seuntjens and I was mainly responsible for the calculations. We worked together at NRCC at that time. So, we would perform measurements on Saturdays and Sundays; we went there to use a telescope to find the exact chamber location and did the measurement. It was quite difficult and the results turned out to be very good. And we can see that... all these materials have different correction factors. These results were

included in the TG61 report, agreement was .2% and that's very good agreement...anyways, I think it's quite impressive. And this one, of course is a good test of the photon transport. At very low energy, actually, it's only photons that matter. Electrons don't go very far, so you can almost, you know, forget about their transport. Treatment planning dose calculations. Now we know that, commercial systems begin to

have Monte Carlo dose calculation algorithms. However, if you look at early days, we have the results from Achterberg et al presented at the ESTRO meeting in 1999. We can see for Pinnacle (with a pencil beam algorithm), I'm not saying that this is Pinnacle, because almost all of the commercial system now uses the pencil beam algorithm. It's a very good one. The new pencil beam redefinition algorithm is very accurate, but again, it's not really implemented in the systems. So are many of the systems, you're looking at the differences

here. I'm not trying to say that it's the case you do it clinically. Nobody's gonna treat like this, but you can see the difference here. The only thing you want to see is significant differences between Monte Carlo using BEAM and dose distributions from Pinnacle with the 3D pencil beam. And that's another case showing for lung treatment. In this case it's about three or four beams, I don't remember and that's from Todd Pawlicki. He calculated dose distributions. In here, because you have low-density materials next to the tumor target, you can see, probably, not able to see very clearly, but if you trust what I said, here in general, you see the 95% is different, quite different from the 95% shown by Monte Carlo. The point is, actually recently, we have done about 30 or 50 cases, for lung treatments. We just bring back, pick up the patient's geometry and redo the calculation. We found out, actually, that the dose changed in the target, but overall is

small. It's only about 3 - 5%, with or without inhomogeneity correction. The biggest difference is you begin to see cold spots in some of the regions near the edges. So maybe it's interesting to look at the recurrences. But it's very difficult. Most of the patients, who we'd looked at, we didn't document where, exactly the recurrences were. It's very difficult to find out the correlation if any. Of course, we also looked at lung pneumonitis and that again, that's a difficult thing. Usually it's defined inside the treatment volume. However, you can't really see... or correlate it with a hot spot, if you do Monte Carlo. So, again it becomes a little difficult. Here I thank Indrin Chetty for the slides. He actually showed you the dose calculations in here from Monte Carlo and also the dose calculation from the original treatment planning dose calculation, I think is based on

equivalent pathlength correction and then the recurrence is here. I don't think he meant the recurrence is due to the miscalculation, but somehow, if we can go back and look at all these cases or do further dose calculations for such cases, we can probably predict whether or not Monte Carlo plays a role in terms of predicting whether we can achieve better local control. In this case, probably, I think it's mainly due to the definition of the tumor or treatment volume, rather than the dose calculation. If you look at here, of course a big chunk of the recurred tumor is in the treatment volume. Obviously, we have to say something's wrong with the dose or maybe something's wrong with radiation therapy altogether, but anyway, that becomes really interesting. Of course we can also look at head and neck cases. In this case you can see the Corvus distribution, of course again, it doesn't matter if it's Corvus or not. You can use other

RTPsystems, you know, maybe CMS, maybe Pinnacle. Now Pinnacle uses the superposition-convolution and maybe the final dose calculate is more accurate. Almost all of the inverse planning systems now, use more accurate dose calculations. For Corvus, the NOMOS now have the PEREGRINE Monte Carlo to back it up. These are from early plans. Now you can see the distributions are a little bit different. So I am not even so sure this is due to the dose calculation and maybe due to beam set up or due to some other things because I know one small bug in their programming. When they switched the beam orders, the leaf sequence files were not switched properly. So, that only happens when you have many beams, with multiple couch angles for the same gantry angle. Then it will cause a little problem. I know that. I was informed later (by NOMOS). I went back to look at them. We had two cases that had that problem.

Because one of the physicists was very intelligent he was trying 28 beams for a head and neck case and in each gantry angle he had two or three different couch angles. Now do you see the small problem like this? If you look at the details you'll see differences in the target now. So to some extent it's important. The other day I already showed in the analysis of dose modeling for IMRT, I showed that after we calculate many of the IMRT cases we see more fluctuations in the mean dose, the minimum dose, the maximum dose in the target for head and neck, and in the thoracic regions, less for the prostate treatment. However, the point is that we can hardly predict which case is going to be problematic. What we found was, you can calculate to 50 cases, and you see no problems, but the next one may have a huge problem like this. This is the same when we do the IMRT QA and I shouldn't talk about this because I am here to talk about Monte

Carlo. We have treated over 1500 IMRT cases at Fox Chase. We had only three cases from our treatment planning systems. It doesn't matter how we measured, using film, ion chamber, at different points and so on. It just doesn't agree, the difference was more than 9%! What we have to do was re-plan and we used the new plan that met our criteria to treat the patient. So sometime we still have problems. We probably don't know exactly what has caused the difference, but it happens. So, in my opinion, even we know superposition convolution is good and even we know, that the simple correction based calculations are good enough for prostate, I think I would prefer to use a more accurate dose calculation, which can provide more assurance when I do the dose calculations. Now let's look at the beam delivery and dose verification for IMRT. So that's a certainly a hot topic, you know. And we can look at the reasons why we want

to simulate that. It's because we know the effect of leaf leakage and scatter. They can be significant, especially when you have a lot of leaves and a lot of, a lot of segments. And we know the leaf size can have different effects, so can the leaf shape and collimator jaw positions, because that would change your, you know, to use the conventional terms, the head-scatter and so on, although in Monte Carlo, you don't care about head-scatters. We do everything, all that you want, through random sampling of particle transport, interaction and everything. We also know the leaf sequence effect, dynamic or step-and-shoot or with synchronization. You know synchronization would change things ... so it's quite interesting. Here it's just a schematic diagram showing you the geometry of a

Varian Linac and then you have the jaws in here and also MLC down here. So the size of the jaws, will affect the extended source size and the shape of

the MLC would change the dose, giving you a tongue-and-groove effect, for example. A graduate student, Ajay Kapur, did that simulation in 1997 and he graduated in 1999. He simulated very detailed geometries of the Varian MLC and now the component module is included in the EGS4 BEAM package, the NRCC BEAM package. You can see the geometry of tongue, groove, and everything simulated very accurately and to celebrate his graduation, he simulated this (a wineglass). Of course, I think, he deserved that, a drink is nice. But if you look at here indeed, you can see the tongue and groove. In fact, everything is quite interesting, quite remarkable. It means, if you do detailed simulations, you can get very nice results and agree with measurements. Recently I think Jeff Seibers did more in, you know, the MCV group. They did more calculations for the MLC geometries and so on. And here, another student Michael Lee did

simulations for the NOMOS mimic collimator; you can see here that the film distribution and dose distributions. We will be able to look at this if we do it in detail. They agreed with Monte Carlo, very, very nicely, very impressive. Of course, he also did measurements in a cubic phantom, put a film in there and trying to see dose distributions at different locations. You can see this agrees very well. If you can simulate everything very nicely and do your film measurements very, very, accurately, you can get really, very nice agreement, oh about one or two millimeters maximum difference, two millimeters. You can see even from the very low dose regions the results are very, very good. Of course, you can also simulate one, just one segment, or one field. Now, you do a film measurement like this, and you compare with the original intensity map used by the inverse planning system. Of course, all of the inverse planning systems would do a final

correction, if not a final dose calculation. They'll add the leaf leakage or other effects from the leaf sequence. However, Jun Deng did the simulation showing that, although you'll see all these tongue and groove effects like in this case, we have to simulate them using Monte Carlo and the program with which we can do ray-tracing and find out exactly the effect of the tongue and groove effect, but, if you have more beams, say more than four beams, the overall dose distribution doesn't change at all. It is just, generally speaking, a little shift of the average dose roughly about one to two percent. So probably you don't have to correct for it unless you have very special applications like the way you treat the breast. You have only two beams, tongue groove there may change the results although, you know, due to set up or patient breathing the effect, also is, kinda smoothed out. But in this case you can certainly see. I would not spend more time on the

exact details of the leaves in this case to simulate IMRT. Because you can verify using other experiments: just design a phantom, put your inhomogeneity in here, assume your target is in here, that's your critical structure, you can put your ion chamber anywhere you want. And you can compare your calculations and measurements. And you can see indeed, you get good agreement, doesn't matter whether you have included tongue and groove effect, or used synchronization or not. So my view towards synchronization is,

you know, probably it's not necessary, because in fact, that could be very small. It's a small shift. It's like a one or two percent shift (in the DVH curves). Of course, also, I'd like to thank Jeff Seibers. He gave me this slide, showing that you can simulate, if indeed you know the details of the geometry of the EPID. Now you can simulate the response of the photons going through the phantom or going through

the patient, or use that to validate the beam delivery. Now he simulated a detailed geometry in here, you can see that you can hardly tell which one is the measured and which one is Monte Carlo. He can certainly fool you. Of course, you know he told me here, the little thing in here (the little dot)? That's an artifact, so you can tell from that, you know, that's measured, although he wrote it here. So that's indeed, very, very interesting. You can do it very accurately and very precisely. The only thing I think would be difficult is, although he truly showed his capabilities in this case of simulating tongue and groove effect and so on, but you can see here he has to assume the leaves have the same gaps, you know, that's not necessarily true. So you can find some of the tongue and groove effects, are not exactly the same and so on, but that's true. So

that put another question, as whether or not we should simulate all these to that extent, because to some extent you are assuming something, which is not necessarily physically true anyway. But it's very important I think, if you know the transmission map, like this, and if you know your patient geometry on the day if you use a CT-on-rails or a cone-beam CT, you can use this one to reconstruct the beams and you can re-calculate dose for that patient, for that day, for that particular treatment, etc. That would be very interesting, right? We'll close the loop in this case. That's what he showed here for a breast treatment. You can see here, measured and calculated, they look almost the same, although we can see probably the shades are slightly different in here and he showed that, you can check the setup with the chest wall here, and the breast tissue, etc., in here. If you look at the geometry in here, it's quite interesting. You can do a lot with it. We are

almost there. We are almost there, starting from head simulation to patient calculation, to the detector behind it and so on to close the loop and do a very accurate dose calculation or dose prediction. We don't have to even, you know, irradiate a patient. We can just do this, and figure out what we may get. Of course, assuming the MLC will behave properly, you can do a weekly MLC QA to make sure the MLC is doing the right thing. And the next thing is following the moving target, the CyberKnife. You all heard that, you've, you saw that... that's a 6 MV X-band LINAC on a robotic arm. It looks very interesting and you can see, you have the orthogonal panels to receive x-rays from here. So you can locate your target if you have seeds or bony structures. Now you relocate the target and you can move your robotic head to point at the right location, they call it a cruise missile technology. So you can track the tumor, it doesn't matter if

that tumor goes to Iraq or come back here, we can find it... Usually we say, you almost can, we cannot of course, allow the patient to go pee, and then we have a cup of coffee, so but... that's not the case. The thing is with Monte Carlo we can simulate again. Because even you can simulate quite accurately the dose calculations are so long, and

still how to track the movement of the tumors, especially, when the robot start tracking? I'm not here to promote the CyberKnife. I want to see what we can do with the calculations. So these calculations were done at Stanford by Tom Guerrero. He was a resident, but he was very happy to do a physics project. I encourage our residents now at Fox Chase to do physic projects. So we have more abstracts at the AAPM rather than at ASTRO, so that's nice. So, what you see here is actually motion measured by the optical markers on the surface of the patient and also at different time

intervals...I'm sorry... at different time intervals we do the x-rays. So then the x-rays will tell you exactly where the seeds are and optical markers on the surface monitoring the motion of the chest wall and so on. You see the locations. Of course, these are the absolute measurement because the x-rays tell you exactly where the seeds are with orthogonal x-rays. So, what we look at here? These are the observations of the target motion using the surface markers. And that's the residual errors when you use the robot and based on the periodical corrections from the x-rays. You say okay, I know you moved and I'll correct for it, now you move again, I'll follow you, using the optical marker, but periodically I change that, because I realized, actually the target in the lung is gradually moving. We call it a very, slow, slow shift. It's probably because when the patient was set up, when patient was a little stressed, so you holding it like that, gradually when

the doctors are gone or someone is more relaxed, and the music's up and he started having more diaphragm motion. So then you can see the target location is changing gradually, over a period of about, I think in this case, about 10, 15 minutes. So that means if you just use very simple technique to gate based on chest wall location, you may get it wrong, because at some point it can't even deliver a beam anymore. So it's very important to periodically, look at a location of the target in the lung and you find out whether or not you can correct for it. So these are the residual errors we measured, using the x-rays. So that means even if the robot can follow the target, you still periodically find the mismatch of the beam and the target location in here. So, if we are certain that the positions are wrong like this, we can evaluate the problem. If you don't use tracking, what would be the problem? Or use this, you know, the residual error to see what the

CyberKnife can achieve. You know, you don't have to use a CyberKnife, you can use a MLC, to move the leaves to do the same now. We have a proposal to do that. So, here what you see is the dose distribution, the first plan is based on a static patient, so usually we hope the patient will be static. Never mind, they're probably not alive, if that's the case, but that's what we do and we look at the orange line here, we covered the whole target, except for a little bit, you know, maybe a little out. In this case, he planned without any margins. Just look at the CTV and then just, you know, plan it and see what happens. Find the right isodose line, give a dose, say how much I want and he'll look at it. Now, assume you have the motion in this case. So there's no following, if you recalculated those using Monte Carlo, just following with different locations, so he did a lot of calculations, about, I forgot... say over a hundred locations, you don't move the patient, you

move the beams. So it's quite easy, in this case with the CyberKnife, you can do that and with Monte Carlo, you can do that too. And you see, you miss a lot of the tumors, in this case. But if you deliver with motion, you know, motion tracking, and using the residual error from your ray imaging, you can see we recover most of it. Most part of it, you see somewhere near here, we missed by about one pixel. Here, because we used 128 by 128 voxels in the XY direction. That means each voxel is about 3 mm. So that means, if we add a 3 mm margin to the original CTV, this treatment would be almost perfect, if you use the CyberKnife. So that's the indication, from what we found... how much we need to put a margin, rather than 5mm, rather than 8 mm, we can use 3 mm. You can achieve something nicely. So here's the DVH. The original deviation is in blue. The green DVH is when the patient is moving and then the red is the one with the

CyberKnife's tracking. But because you have no margins, you missed a little bit. By adding a 3 mm margin to it, and you see, you can actually recover that. That line will go back to here. Now you have nice coverage. You still get hot spot. You shift the curve to here, you still have the hot spot. The hot spot is due to the averaging effect of your beams over the target and the target's moving a little, due to residual error. So you can reduce the hot spot but you can certainly remove that cold spot by adding margins. So overall the results are very interesting. So I'd like to summarize my talk by saying Monte Carlo is becoming a practical tool for radiotherapy dose calculations and treatment verification. Either use your own homegrown software or you can use commercial software. However, if you have your own homegrown software, you can do a lot more. You can, say, simulate new hardware, new software, or you can check different treatment

techniques. With the commercial systems, you still probably can't do all sort of combinations of brachy or photon beams, or electron beams or using different things. For example, to simulate, say the CyberKnife, you probably can never do the Monte Carlo calculation using PEREGRINE to simulate the CyberKnife, you see what I'm saying? If you have your own, a home grown system, you can do that. Another thing is... Monte Carlo calculations have found a lot of applications, in brachytherapy for example, I didn't talk about it at all, in proton therapy, a lot of calculations, although a lot of people believe that you don't need Monte Carlo for protons, you know, they don't scatter so much. But if you look at, study beam scanning with detailed geometry, especially for lung. You have inhomogeneity in there. Maybe Monte Carlo may find a role in there because the calculation is becoming fast. And radiology? How to simulate PET and all these kind of things

together and also outcome studies. A lot of people use Monte Carlo to predict, you know, look at the utility parameters and so on with the patients and predict outcome. That's another thing that becomes very interesting. I can't really finish off my talk, without mentioning the people, who taught me how to do Monte Carlo, you can see the names here, and some of the colleagues I worked very closely and you know, did the calculations together, and also students in here, and also people who provided the slides for me because I actually send out an e-mail only a week ago and said you know, if you have some slides I can share... we can look at the most recent results and so on... I thank

you for providing the slides. And finally I'd like to tell you that Fox Chase is not that far, you know. It's only a five-hour drive from here. It's equally green! So if you have time after the conference come to Fox Chase and have a visit. Thank you very much.