OK thanks Indrin, it’s a pleasure for me to be here, especially since it’s the first time I will be giving a talk at the AAPM as a member of Carleton University instead of the National Research Council after many years. We’ve been asked to make disclosure statements and so there's the disclosure statement, I've been involved through the National Research Council with a variety of licenses of Monte Carlo to various companies. What is Monte Carlo transport? Monte Carlo is where we simulate the paths of many particles, we’re using random numbers and we make use of the known probability distributions from the physics of various interactions and then we can track these particles through geometry and keep track of various physical quantities. We learn about the average properties and the stochastic distribution of effects if we really want to,

although that’s not usually what we do in this kind of situation in Monte Carlo treatment planning. But here's an example of why we really need to make use of Monte Carlo I think. We’ve got a _10 MeV_____ photon incident from over on the right and it’s shown as this yellow photon on a slab of lead and this is the same photon looking down on the axis. Well I'm just going to concentrate on this view here (to right). The photon comes in and it creates an electron positron pair. As everyone knows, electrons are blue and it’s going off in this direction and as it slows down, it scatters and bounces around quite a bit and you can see it also in this direction from this perspective. At this point, it gives off a photon which is shown in yellow. It’s a bremsstrahlung photon and so there we have a fairly complicated set of processes and at the same time the positron follows a very

winding path. If you're up close to the front you can see there's a bremsstrahlung photon being given off here and at this point the positron annihilates in flight. You can tell that it’s annihilated in flight because the positron, the two 511 keV photons do not go off back to back. Then one of these 511 keV photons goes over here and undergoes a Compton scattering and gives energy to another electron. This is a very complicated process over all. It can't be solved analytically but we know all about it and in fact as you see here, this is a Monte Carlo simulation of exactly all of the events that are occurring. So that’s why we need Monte Carlo. Here's another picture which is just here because I think it’s interesting. We’ve got a ten MeV set of electrons coming in from the right again onto a water phantom, and you can see if you understand the physics of your

electron beams in the clinic, you can see that by this, could we have the lights back out? That just made the colors much harder to see from here. The fluence is increasing with depth and therefore this is where you're seeing the peak and the dose maximum in the depth dose curve. Now just to give an example of a very simple photon simulation, let's say we've got a very simple model in which we have only two possibilities; we can either have a Compton interaction, or a pair event. And we know the total cross-section is equal to the sum of the two cross sections (the inverse is the probability per centimeter), Now let's take two random numbers, we’ll call them R1 and R2, which are uniform between zero and one. I won’t say any more about this except to note incidentally that people have spent their entire careers creating random number generators for us which give us basically an infinite number of random numbers in a sequence and so we'll take that as a given. If we now ask the question, how far does the photon go before it interacts , I assert that if you take the number X being equal to minus the natural logarithm or R1 divided by the total cross-section and use that value as a distance, if you then sample a large number of these, these distances will be exponentially distributed between zero and infinity, they will have a mean of one over the cross-section to give you the mean free path. So we have now had a very simple example of how to sample the distances the photons have gone, making use of, first of all, the random number and second of all, our knowledge of the cross-section which in this case is just the given by the single number.

We've now got the fact that the photon is going to interact at a given position. The next question is, is it either going to be a pair or a Compton interaction and we make that choice by looking at our second random number and saying: if R2 is less than the Compton cross-section divided by the total cross-section, we'll declare that a Compton scattering event occurred, otherwise it's a pair production event. So this is a very simple Monte Carlo simulation, but you see we've used random numbers to make two different decisions and we've made use of our knowledge of the physics or cross-sections of it and we can put these things together. To do a full Monte Carlo simulation with electron transport, things get a lot more complicated in some ways, but in many ways it's just bookkeeping, we're using the same principles in order to sample from the known

physical distributions but to get the answers and to track the events. So how do we, actually use the simulation? Well, we score whatever data we want, in this case it could be the average distance to the interaction or how many of each type of interaction or the energy deposited by each type or interaction for example. We could also answer all of these analytically in this case, but in more complicated geometries there are more complicated cases taken into account all of the different interactions, the Monte Carlo technique becomes much more useful. So how do we use Monte Carlo in radiotherapy? Well, we're already using it almost on a daily basis, Monte Carlo calculations are the basis of much of what has been used in clinical dosimetry for years, namely the AAPM's dosimetry protocols. TG-51 and earlier TG-21 for accelerator or external beam

dosimetry, both make extensive use of Monte Carlo for the various factors that are used. The TG-43protocol for brachytherapy dosimetry is almost entirely based on a lot of information from Monte Carlo calculations about brachytherapy seeds and similarly TG-61 on x-ray dosimetry is just chock full of Monte Carlo results so in that sense, we're already using it daily in clinical radiotherapy. However we're here about Monte Carlo treatment planning and lets start by just reviewing some of the major codes that have been available, many for many years. Back in 1963, Martin Berger of the then National Bureau of Standards wrote a seminal paper about condensed history technique for doing electron Monte Carlo transport and along with Steve Seltzer, who's sitting in the second row, they developed a code called E-tran and it has very many derivatives that as it were

that have taken the ETRAN electron transport and patched it in so the Cyltran series and the Cyltran code which is part of the integrated tiger series and more recently and more extensively now, MCNP all make use of the basic ETRAN transport physics for electrons. A similar general purpose code started off as the EGS or electron gamma shower series of codes ,originally from Stanford Linear Accelerator Center where it was used for high energy physics applications and then in 1984 the EGS 4 system came out along with PRESTA and more recently, with all of the physics improvements that Iwan Kavrakov has made, EGSnrc was released in the year 2000. The code MCPT is Monte Carlo photon transport was , developed by Jeff Williamson and a larger number of his collaborators over the years. It's a photon only code but it's very, very useful and very,

very efficient for many brachytherapy calculations. More recently, the Penelope code has been developed which is a general purpose electron/photon code, but it's of the same genre as ETRAN and EGSnrcand then more recently, there's been a series of much faster Monte Carlo codes. The first one was VMC in 1995, it's morphed into XVMC for photons and more recently VMC++, this was worked my Mattaias Fippel and Iwan Kawrakow. There's also been other codes like DPM which Sempau et. al. have developed. In addition to this, in terms of general purpose codes, there's the BEAMnrc code, which is for modeling linear accelerators. So why is Monte Carlo important for treatment planning and Indrin's already said basically the real answer here is that in principle, Monte Carlo gives us the right answer. There are no significant

approximations. I'll just leave it at that, I think that that's one of the reason we're here. There have been a large number of benchmarks showing if you're careful and input enough information about the situation that you're looking at, you'll find that you get very good agreement between the experiment and the calculations. In fact, I worked in a standards lab for many years where people were some of the world's experts in measuring radiation and we actually got to a situation where if there was a difference between our calculations and their measurements, they were as uncertain of their measurements as we were about our calculations. So we've got it to that stage anyway. I'd like to show you a couple of examples now from practical situation where Joanna Cygler developed a very nice test of treatment planning systems in general where she

had an air cylinder in a basically solid water phantom with a series of hard bone discs behind it with the electron beam coming in from the top. And a few years ago, George Ding and Cygler et. al. published some results from this and you can see where they were comparing it to CADPLAN, which has a state of the art electron Monte Carlo, not no, sorry, electron pencil beam algorithm in it and you can see the measurements showed very sharp peaks between the gaps, when a profile was taken behind those discs whereas CADPLAN basically hardly sees the geometry in front of it at all. With the VMC++ implementation with by Nucletron and a paper published this year, Cygler et. al. showed that you could get very good agreement at different positions behind this, although you can see here that the spikes from the Monte Carlo don't quite agree with the maxima that

you see in the measured values . That was literally a problem with the size of the voxel. We're looking at things here which change over the r the distance of one millimeter by a very large amount and yet the smallest voxel they could look at was of the order of four millimeters so with smaller voxels, it was, it was demonstrated, but only over restricted regions, that they could get almost perfect agreement, so the agreement with the experiment, the measurements is very good. So there's been a lot of research done on Monte Carlo treatment planning, many people have published, many groups have published about extensive research systems capable of doing full Monte Carlo treatment planning and there's been any number of talks at this conference about it, the Medical College of Virginia, Fox Chase, Stanford, Seville – Spain, Tubingen in Germany, the University of Michigan, etc., etc. There are many others with active Monte Carlo research programs at this point. Now what are the major steps in Monte Carlo treatment planning? Well first we have to model and simulate the patient independent part of the radiation source above the jaws and above the multi-leaf collimator. You can create a phase space either through a model or an actual large file, you then have to transport the beam through patient dependent part of the collimator systems. So as you change the jaws or you add blocks or whatever of the multi-leaf collimator, which is dependent on the patient, you have to do the transport through that collimation associated with each patient. Then you have to import the CT patient data into the Monte Carlo code where you need both the densities and the materials in order to do the Monte Carlo calculations unlike previous calculations, which just use densities basically. Then you have to transport the beam in the patient and calculate the dose and then you might possibly smooth the data because of the statistical uncertainties. And I just have a couple of pictures here of accelerators but given the time, I'm gonning to just skip through them. Some of the issues related to Monte Carlo treatment planning are about statistics, we have to worry about the statistical fluctuations and possibly apply smoothing.There are issues about voxel sizes, there are issues about CPU time which can be very dependent on voxel sizes and beam energy and which code you're using and the point down at the bottom here is about the accuracy of the information you have about your accelerator. This is often a major issue especially if you're trying to do an ab-initio Monte Carlo calculation of the accelerator. But here is a simple example just of a 10MeV beam of electrons where we've either run ten million histories, tracked ten million separated electrons or just 100,000 histories and you can see that the blue line with the very large number of histories and nearly two hours worth of calculation time is very smooth, has uncertainties of the order of one percent and the 100K histories which only take 52 seconds have ten percent uncertainties and you can see that they clearly are giving the same answer but there's a lot of statistical variation which one has to worry about. This is just a two dimensional representation of those same things, you can see extremely noisy isodose lines in this case, but very smooth ones in the majority of cases when you're away from the flat regions anyway for the one hour calculation. So here's an example where we have 20MeV of electrons with very strong heterogeneities, there's an aluminum block with an air block beside it, this is taken from a paper from Iwan Kawrakow back in 2002 and you can see that there's quite a bit of statistical noise after two million tracks. If we now take that same data, the one advantage when you're looking at Monte Carlo is that you just run it for a large number of histories and you can get what you consider the right answer and so that's what's been done here and what Kawrakow did was apply a smoothing. He run two cases, one for 600 million histories and one for the two million histories that I just showed you a minute ago and then he smoothed it using an algorithm that he had developed and presented in this paper and you can see that it clearly works very well, the smoothing can be done very accurately. So how long does it take? Well,

Monte Carlo gives the entire distribution first of all, not just a few points and the time for N beams is the same as the time for one beam. But timing is a complex question and it depends on a large number of things. It depends on what statistical uncertainty you want and how it has been defined, depends on the voxel size, the field size, the beam energy, depends on whether it's a photon beam or an electron beam, and the accuracy you seek. So also, the speed of the CPU and the optimizer that you're using can be important. So we, Radhe Mohan and I, proposed a   a well specified benchmark for comparing the in-patient part of the calculation only, so this is only one part of the calculation. But all of the previous issues were controlled in this specified case. And the results show that the state of the art codes are very fast, even on few hundred megahertz or 500 megahertz

machines, you were looking at calculations of the order of a few minutes to get typical statistics of two percent for electron beams and even faster with smoothing. So here's an actual results from the clinically implemented Nucletron version of VMC++ and you can see that running on a, but if we, now look at this currently used machine say of 2.2 gigahertz, we're looking at a one minute calculation for the two percent statistics. Now this is a question about how fast can we really do it if we're   in a research lab and we want to test things out, so here's a calculation that again Ivan Kawrakow reported. He's done 250,000 tracks in five millimeter voxels which is a little on the big side for electron beams but is still often used today and he smoothed it after he had done the calculations. This entire calculation, and that's the dash line. The solid lines are for the small voxels

and for a very long run. You can see that in this calculation, the short calculation which took five seconds, that's for the smoothing and the entire Monte Carlo calculation on a two megahertz machine, you've got an almost perfect result. Not exactly perfect, but very close, certainly good enough for practical purposes. So in summary, treatment planning based on Monte Carlo is coming and very close and is already here for electron beams. In principle it provides all the accuracy available from the physics and the major issues still are from my point of view, proper implementation, getting accurate beam models and how to commission the treatment planning systems for individual machines is still being debated. There are a variety of different approaches and we all have to get used to working with statistical uncertainties which are frequently obvious, but the advantage

of them compared to systematic is that we no longer have to worry about the systematic uncertainties in the older methods which are not so obvious. So at this point I want to first of all, thank a variety of people. Iwan Kawrakow, who provided me with the data for several slides and Joanna Cygler also for providing slides and Paul Keall for providing the couple of slides that Indrin actually used, so thank you all very much.