Introduction to Parallel Programming in OpenMp

Dr. Yogish Sabharwal
Department of Computer Science & Engineering
Indian Institute of Technology, Delhi

Lecture - 01
Introduction to Parallel Computing Architectures and Programming Models

This modular going to be on Introduction to Parallel Programming in OpenMp; and my name is Yogesh Sabharwal.

(Refer Slide Time: 00:34)

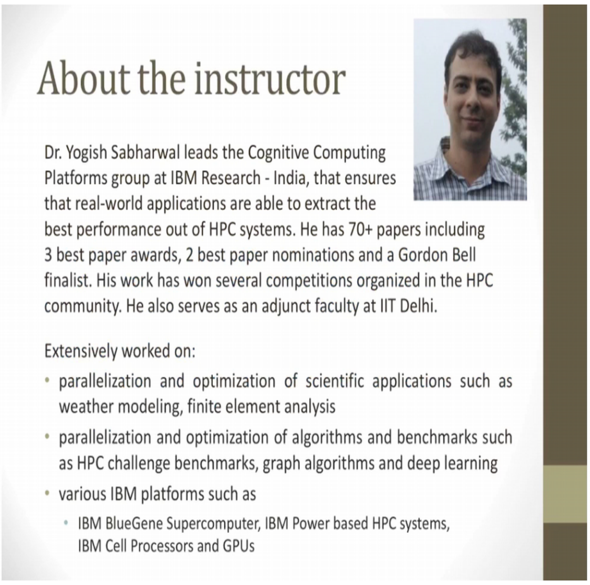

I did the Cognitive Computing Platforms Group at IBM research India; we basically look at how to optimise scientific applications and algorithms on massively parallel systems. So, I have done extensive work on you know parallelization and optimisation of scientific applications such as weather modelling codes, finite element analysis and so on. And also we have looked optimising algorithms and benchmarks such as the HPC challenge benchmarks various graph algorithms and also you know deep learning frameworks.

So, we have expertise in various IBM platforms such as the IBM, blue gene supercomputer the power based HPC systems and cell processors and also a lot of parallel programming on GPIs.

(Refer Slide Time: 01:21)

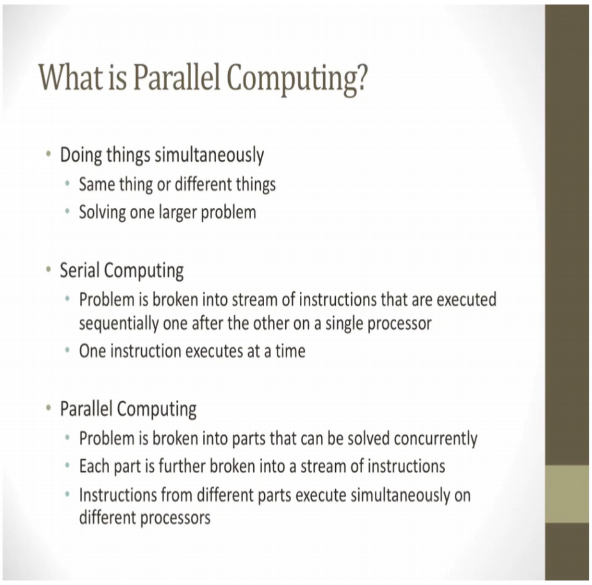

So, what is parallel computing right? So, parallel computing means doing things simultaneously. So, what are we doing simultaneously, we might be doing the same thing or something different set of things simultaneously right for example when you are driving a car right you might be driving at the same time you might be talking on the phone not a great idea, but it might be happening right and at the same time your breathing as well right see all these things are happening in parallel they happening simultaneously right. So, that is one form of parallel computing.

So, in this case what you are doing is different right you are doing three different things at the same point and time. Now in different situations you could be doing the same thing right when you are exercising maybe you are putting your hands up and so on right. So, you are performing the same action in parallel.

So, when we do things simultaneously we may be doing the same thing or different things, but what is the important in parallel computing, we are doing we are solving one large problem I mean for instance at this point of time you are breathing and I might be you know writing something on a piece of paper, but that we do not generally call that parallel computing, because we are completely doing to separate thing which are not related to each other at all right. So, in parallel computing generally we referred to the scenario, where we are trying to solve one large problem and while solving that problem we are doing things simultaneously.

So, in serial computing basically the problem is broken down into a stream of instructions, and then these instructions are executed sequentially one after the other on a single processor right. And the important point is that only one instruction is executed at a time at a given point and time right. And now when you move to parallel computing what changes? So, in parallel computing the problem is broken down into multiple parts that can be solved concurrently simultaneously, right.

So, you have to bring your problem in to different parts, and each part now is further divided into a stream of instructions which can be executed on a processor right one by one. So, the important part over here what differentiates it from serial computing is that your executing different instructions, of the different parts simultaneously, right.

(Refer Slide Time: 03:35)

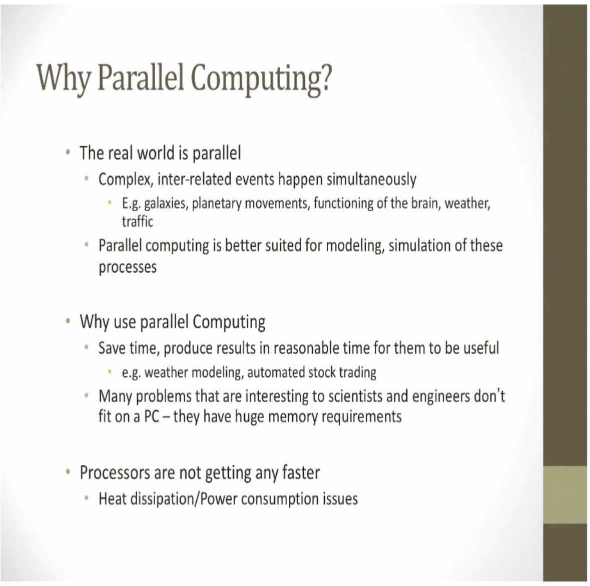

If you look at galaxies, planetary motions, how the brain functions, how the weather modelling is done, how the traffic moves right all of the real world phenomena are parallel by nature right. So, these are complex interrelated events that are happening simultaneously. So, it makes sense if you want to model them, then it makes sense to do the simulation in parallel. And why do we want to use parallel computing because first of all it saves time the producer results in reasonable amount of time. So, that they are useful to us right for instance suppose that we are doing whether modelling. So, today when we try to do weather modelling right at a good resolution let say it 10 cross 10 kilometre resolution for a region like India right then the weather simulation itself can take you know if you run it on a single processor its can take days to run right.

So, what is the point of that if you are predicting tomorrow's weather, and it takes you of a three days to predict that whether it is of no use to you right. So, you want to be you want your results to be useful and if you cannot do them on a serial computer then you have to basically use parallel computing, but the other issue is that there are many problems that are interesting to scientists and engineers when you doing the simulation you basically want to do it at a very high resolution right.

So, if you wanted it were very high resolution then there is a lot of data involved, and this data generally does not fit in the memory of a single computer of a progra pc. So, what do you do? Either you can store it on the secondary storage on the disk and did you keep on fetching it, and your algorithm has to somehow done in parts right or the other option is that use a parallel computer. In parallel computing you have multiple nodes you have memory associated with every node and therefore, the total memory that is available to you as much larger.

So, you can load the entire data that you have on the combined memory that the systems are. The other issue is that the processes are not getting any faster, the clock frequency is not improving anymore and the primary reason is that there are heat dissipation issues and power consumption issues. So, what does that mean? That means, that you have to resort the parallel computing right you are not going to get more speed out of the same processor. So, you have to use multiple processors where do we use parallel computing.

(Refer Slide Time: 05:49)

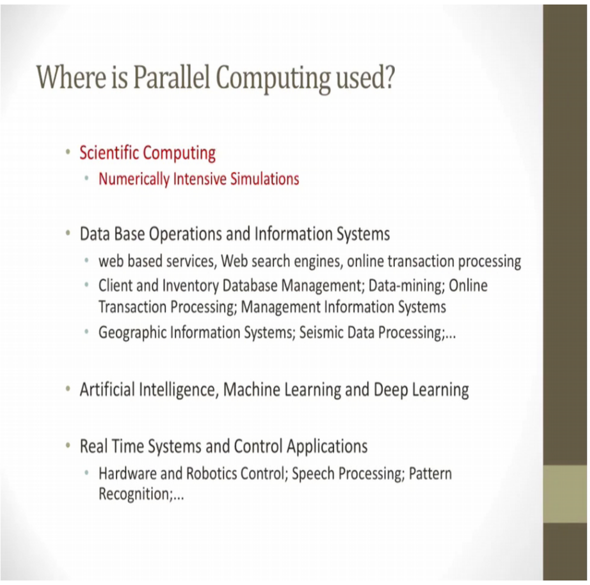

Student: (Refer Time: 05:52).

So, scientific computing is one huge area where parallel computers are used, this is primarily to solve numerical intensive simulations. So, I will I will come back to this, this is one of the most important areas I will I will cover this in more detail, but other than this parallel computing is huge heavily in database operations and information systems rights. So, for instance a web based services they made scalable. So, there are so many requests coming into web service these days they have to be load balance, they have to be sent to different servers the competitions have to be done and parallel, web search engines, online transaction processing right just imagine the number of requests that credit cards or these of the companies that are doing transaction right.

So, all of them have to be done in parallel right then there are other systems like client inventory database management systems, data mining online transaction processing mis system and so on. Then another area is artificial intelligence machine learning and deep learning. So, again these areas are making heavy use of parallel computing, if you look at something like neural networks right. So, there are lots of nodes which are getting imports they are doing some competitions and then producing some new output which is then consumed by other node. So, it is a graph layered graph and lots of nodes are doing competition than parallel. So, all of this you know is very I mean able to parallel comp computing and then other various like real time systems and control applications where this is used heavily.

(Refer Slide Time: 07:19)

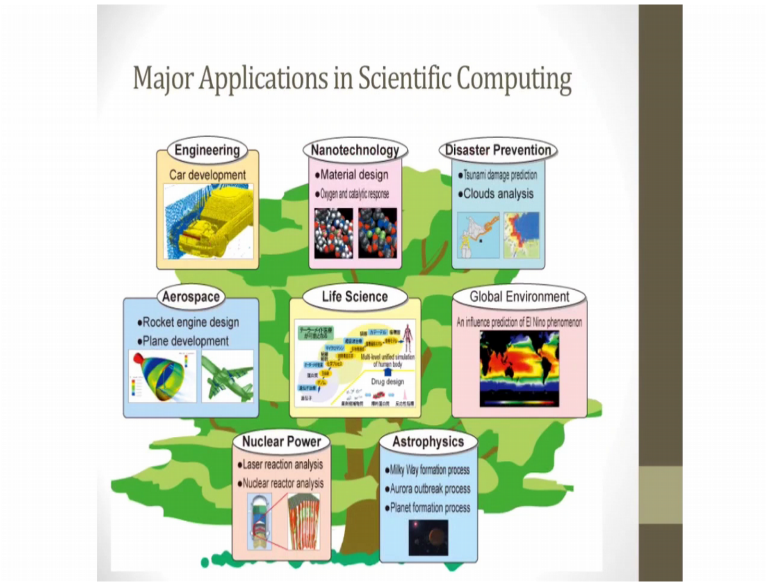

So, coming back to scientific computing right so, that is one of the focus areas where high performance computing systems are massively parallel systems or majorly used.

So, for instance in engineering rate in car development. So, as you know right these car companies they are required to pass certain safety tests. So, for instance they have to ensure that if the car crashes into wall at a certain speed, with the dummy settings inside the car what is the impact of the dummy they have they have to show that you know it is a limited or no impact. So, what does that involve that actually involves actually crashing the car into a wall, but that is costly right. If let us say you are trying to build that car you are not going to car crash that car every if every second day into the wall just to see that whether the dummy is getting hurt or not right.

So, again this is actually done via simulations. So, you use parallel computing again because the simulations are very computer intensive, and it can take very long say you would not use parallel computing and you actually do the simulations on the computers first. And once you are satisfied or you can reasonably sure that you know it is know that is when you actually take a natural car and go cash it into the wall. Similarly in aerospace it is very very important to understand; what is the effect of the wind or turbulence on the wings of an aircraft right all of that is done via simulations.

In nanotechnology for material design for instance if you are trying to design some blue you want to figure out, what is the right mix of chemicals that you have to use, what are the properties that it will have a lot of this is done by a material design using parallel computing, disaster prevention and global environment. So, things like weather modelling, tsunami prediction, cloud analysis right all of these are a very computer intensive simulations right involving lots of physical and chemical processes that need to be simulated. Then in life sciences for drug design and so on it is used heavily for nuclear reactor analysis.

So, you want to simulate the nuclear reactions right what will happen? In astrophysics to figure out the formation processes for the Milky Way or the planet formation process and so on right you figure out that these are the laws that should apply, but you want to simulate to figure out that that is actually what happens right. So, these are some of the predominant areas where parallel computing is used in scientific computing.

(Refer Slide Time: 09:39)

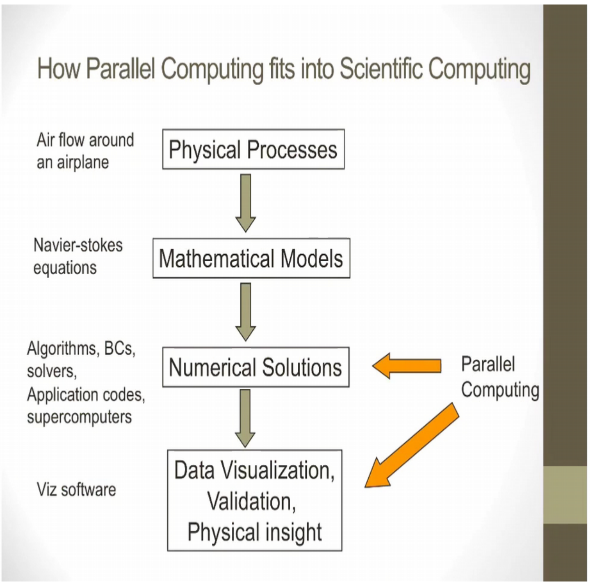

So, how do you actually go about using parallel computing in scientific computing? So, basically you start with some physical processes like for instance you want to study the air flow around an airplane right and then you build some mathematical models to do the simulation right. So, for instance new model it in various numerical solutions and then you basically solve that right.

So, solving it involves basic various numerical solutions. So, that involves lots of algorithms and you know using lots of solvers, and that is very very computer intensive right lots of numerical equations have to be solved. Scientists and engineers are becoming more ambitious. So, they want to you know do it at a very high resolution with lots of grid points. So, that makes it very very computer intensive and finally, once the output is produced you want to visualise that data and after you valid date at you figure out the physical insights, you may want to go back and modify your models and so on right. So, the numerical solution and visualisation is the part that heavily used parallel computation right.

(Refer Slide Time: 10:47)

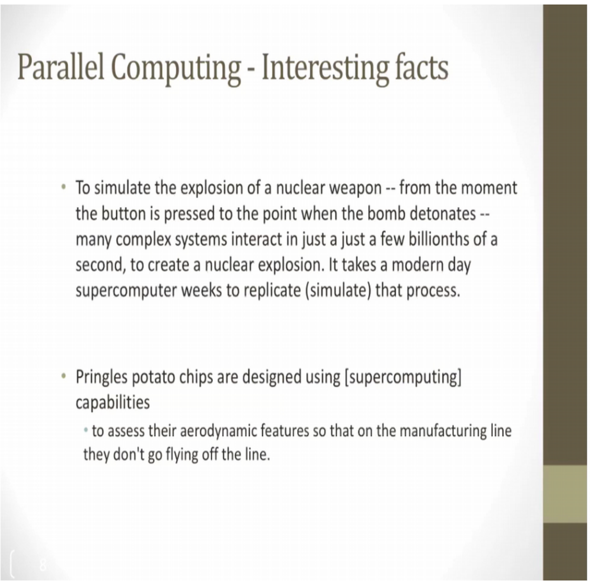

So, here are some interesting facts about parallel computing. So, people actually study the simulate the explosion of a nuclear weapon. And from the moment that the button is pressed to the point that you know the bomb detonates. Actually in real time that happens in few billions of a second, but if you want to simulate that process even supercomputers today take weeks to replicate that process; it is that computer intensive.

And another fun fact Pringle potato chips right, they have to actually use supercomputing capabilities, parallel computing to do simulation, what rate to have the assembly line move, because if they move it too fast and the chips are going to fly of the assembly line. So, they have to do all those studies and they do that using supercomputing capabilities.