A conversation with CHCCS 2020 achievement award winner Dinesh Pai

Dinesh K. Pai
University of British Columbia

Abstract

The 2020 CHCCS Achievement Award from the Canadian Human-Computer Communications Society is presented to Prof. Dinesh Pai (UBC) for his numerous high-impact contributions to the field of computer graphics research. His diverse research addresses physics-based animation, multi-sensory displays including haptics and sound, and realistic digital human models. CHCCS invites a publication by the award winner to be included in the proceedings, and this year we continue the tradition of an interview format rather than a formal paper. This permits a casual discussion of the research areas, insights, and contributions of the award winner. What follows is an edited transcript of a conversation between Dinesh Pai and Doug James (Stanford CS Prof and former PhD student) that took place on 14 March, 2020, via Zoom.

The Interview

Doug: Hello Dinesh. Congratulations on being recognized as the 2020 CHCCS Achievement Award winner.

Dinesh: Thank you.

Doug: You have had an interesting career, so I guess a good question to ask first is "how did you get started in computer graphics and animation?"

Dinesh: Oh, well, I've probably had the one of the most circuitous routes to graphics and animation. I don't know if this is wise to say this... but, you know, I am interested in many things, and computer graphics is one of them. I'm really interested in how the world works. And, in particular, I'm interested in how humans work. And then, very specifically, I'm interested in how I work. In a mechanistic way. So that's why I've been interested in these kinds of topics. So, I didn't start off — and you didn't start off — in graphics; we started off doing something else.

In graduate school I was very much interested in robotics for this reason, and my PhD was hardcore robotics. Robotics is what I got tenure at UBC for — building walking robots [10], mobile robots that played soccer, and theoretical robotics. I was involved with the first RoboCup papers when I first came to UBC and started the Dynamo mobile robot project with Alan Mackworth [1].

But I was always interested in computer graphics and movement. Even my master's thesis on robotics was on computer graphics simulation. I built a simulator called INEFFABELLE [12] named after a story in a famous science fiction novel by Lem where we did robot simulation on an Evans & Sutherland Picture System 2, which was pretty state of the art at that time. I also wrote a controller for dynamic bipedal walking (in Lisp!) with 3D graphics.

By the way, if you want I have a little Canadian content story on the side there.

Doug: Sure, let's hear it.

Dinesh: So, I was this masters student at Cornell. My advisor (Prof. Ming C. Leu) said "you know, I have an account on this new setup called CADIF: Computer-Aided Design Instructional Facility [4] that a guy from General Motors just set up at Cornell." So I went and used his login to start writing code for graphics. And then I got hauled into John's office saying "Who are you? You're not allowed in here!" And that was a Canadian called John Dill, who at that time had just moved there from GM, and then later moved to SFU and was active in the graphics community in Canada. He did give me a login later. So this was my first encounter with Canadian computer graphics: to get into trouble with John Dill. [laughs]

Doug: [laughs] Nice start.

Dinesh: Yeah. So, I wrote an interactive 3D graphics simulator for robot programs using an Evans & Sutherland Picture System 2, but I was actually interested in just robotics. At that point, I was also interested in how the body works and how the brain controls its movement. So I went and took a neuroscience course with Avis Cohen thinking "Okay, now these people in neuroscience know how the body works, and robotics people will just get this information and implement it" but it turned out it wasn't so simple. I like to say that neuroscience is so hard that you should only try it if you have nothing to lose or nothing to gain. I keep going back to it though, because it's so fascinating.

So anyway, coming back to this, I was always doing simulations in robotics, and we were doing 3D computer graphics just as users but never published in graphics. I think the first transition was doing haptics with interactive 3D graphics. And then I think my first paper in graphics [5] was with you. So, we were doing real-time simulation and we thought we'd submit a paper to SIGGRAPH like our friends in Imager. It was probably the first animation paper submitted without a video because we didn't realize we had to. [laughter]

Doug: [laughs] It eventually had a video, but initially, we didn't have one. So I noticed that ArtDefo [5] is still your most cited paper and also mine. What is it about that paper that, you know, has had so much impact?

Dinesh: It's a good question. I think it was ahead of its time. First of all, it showed that you could do interactive 3D graphics with, in some ways, full-blown physics-based simulation, if you did the right kind of model reduction. That is, that you didn't have to hack things to get real time. You didn't have to start with an approximate physics solution, rather you could start with realistic physics and then see how to approximate it and precompute things. So I think that may be one thing.

But, I don't know, maybe it's the catchy title [laughs].

Doug: You always seemed like you had a unique perspective on problems in computer graphics and animation. As a student, I often attributed it to your non-CS background in mechanical engineering and robotics until I met more people in mechanical engineering and robotics and then I realized it was something else. Can you say something about the way that you select and think about research problems that has helped you to be successful?

Dinesh: I don't know. I think it is very much curiosity driven. I like to let the problems tell me: what is interesting about them? And what's the solution? I try to listen to the problem. And I don't say "I have this hammer and what am I going to use it for?" I think that's one. So that means that I'm willing to take the risk of learning new things, and to find a tool that's appropriate to solve the problem. I think that's part of it. I'm not a disciplinarily controlled person. Sometimes it may look like I'm working on too many different things, but in my mind, I'm working on the same problem, I'm just looking at it from different points of view. I'm just happy to use different tools, because they seem to be the right ones for a particular problem. For example, when it became clear that the biggest challenge in simulating physics is not the PDEs but the constraints, we developed Eulerian discretizations for solids that handle contraints beautifully [8, 9, 2]. Also, I should say, doing this is a form of risk taking. So it is uncomfortable and you need courage to follow your nose. But I've been very happy with how things turn out, after some period of uncertainty and fear [laughing].

Doug: One iconic project, when I think of you, is the Active Measurement (ACME) facility [11]. Recently there's been a lot of increased interest in 3D scanning to capture realistic, mostly visual models for use in virtual environments, VR, etc. Your work on robotic scanning of reality-based models with contact interaction behaviors for virtual environments [14] is now 20 years old and was, I believe, way ahead of its time—not just shape and shading, but you pioneered contact scanning of roughness, deformation, and sound for multi-modal display. Back then I don't think many people really appreciated what you were doing or why you were doing it. So, there needs to be a question here, but I'm not sure what the question is. What is the right question and what is your answer?

Dinesh: That's a good question. The answer is 42. [laughs]

No, I think it really comes down to the same thing, which is that I was interested in how things work, and, in particular, how they interact, and it was clear that, for a lot of things, what is interesting happens at the interface, like when you touch or interact with an object. So just measuring those properties with optical methods, which don't make contact or even see the contact, didn't make sense. So, again, I was happy to just jump in and make the right measurements, even for humans [7]. Maybe I went a bit overboard by building a full robotic system!

Doug: It was very ambitious for a pioneering effort.

Dinesh: Yeah, it was. It took a lot of work from a lot of people. But it was interesting. One key point was the idea that it isn't enough to just model the physics in the sense that you know the equations of motion. You also need to know the constitutive properties of the object to do the simulation, you need to know what the right boundary conditions are, the geometry, and things like that. So what we were doing then was to say, "Okay, we have to get the data on real objects. Not just equations of motion. We have to measure the material properties and responses." I think this kind of data-driven idea was a bit early at that time, because it wasn't easy to get data. Even today it isn't easy to get this type of contact information of objects. You can get a lot of YouTube videos, things like that, but even scanning people is challenging now.

So, yes, it was early to say that we're going to actually measure what things are like in reality, and use the physical model as a template with lots of parameters, and we can estimate these parameters from data. But this type of work now has a lot a lot of currency, with machine learning, big data, and things like that.

A second part was that we need to get lots of data, and so we needed ways to automate acquisition to make those things easy. I was thinking that we could provide a service where people would send us objects and we'd automate their measurements, like a factory in reverse, making virtual objects from real things. I think that the measurement automation was also a bit early, because robotics was also relatively undeveloped at that time. You didn't have self driving cars, drones that could go and scan huge areas of cities, and so forth. And the third one was contact, which is still, I think, out there as a challenge. And I should mention, by the way, after 20 years I'm back doing similar things now, except that now we are focused on measuring humans [13]. These days pretty much most of my time is focused on modeling and measuring humans.

Doug: What can you tell us about your current research activities?

Dinesh: So all the physics and simulation and neuroscience and other things [3, 6, 15] that I have looked at are converging on trying to build accurate physical models of the human body. Not just for visual effects. I want to create human models that have predictive value. A huge amount of what we produce as a civilization is related to the human body — you know, humans are self centered. We want to know things like, how does the clothing that we wear fit on our bodies, the ergonomics of our environment, how things feel when we're sitting on them, etc. To make those kinds of things, which have trillion dollar industries behind them, people are still using relatively old methods—you make it, you try it, it fails, repeat. And you fly goods and people across the planet to do it. So, my hope is that we are now at a stage where we can do what people did in aerospace, that is, you have realistic models of humans that you can simulate. You can simulate the heck out of something first, virtually, and see if it works, and then you make a few physical prototypes just to confirm that it's okay. So it's the same ethos, the thing that has been successful in the aerospace and automotive industries, that I'd like to bring to the human body so that you can make predictions. Predictions mean that you have to have real data, you have to understand stuff, not just tell a cool story.

For example, take the distribution of elasticity of the human body. First of all, there are individual variations. But even for average humans, when we looked at this, we couldn't find any literature that had measured, say, the elasticity of the human shoulder.

Doug: That's surprising.

Dinesh: Yeah. And even more surprising is even things that have had a lot of attention in the medical community for cancer diagnosis with elastography and stuff like that. The data is not usable for simulation of contact with the body, so we had to say, again, “Okay, we’re not just trying to publish a SIGGRAPH paper here. No, we need to roll up our sleeves and try to measure that.” So, this is my current research.

Doug: You’ve mentored a lot of graduate students who have gone on to successful careers in graphics. What is your secret?

Dinesh: I think the first secret is to get very lucky! I’ve been blessed with some incredible students. The second secret is point them in the right direction – they are smarter than me but not as experienced so I try to point them at good problems. Not for publishing the next SIGGRAPH paper, but problems that will be important in 5-10 years. To paraphrase Gretzky, to point them to where the puck is going to be. I am also willing to learn from my students, and they have taught me a lot. And I’ve been fortunate to have collaborators like Uri Ascher with whom I have co-supervised students. Finally, I like to pick students who aren’t just smart but great human beings. A lot of them continue to stay in touch with me and with each other.

Doug: Great, thank you. You’ve recently spun off some of your research into a company, Vital Mechanics. What has been your experience with R&D in industry versus the big R and little D that happens in academia?

Dinesh: That’s a very good question. Let me tell you, it’s been a very interesting ride! I am learning what is required to actually make research findings into something that is a real innovation that gets used. For me, one of the discoveries was that it is really a two-way street. You shouldn’t think of it as some waterfall model where you do research and throw results over the transom, and someone picks it up and it becomes useful. That never happens. So, what is really required is this feedback loop, between you, the company trying to produce a solution, and the people whose needs you are trying to address. There are some real needs for you which you have to say, “You know, it’s great that they want this, but it’s not what we are in business to do.” But there are lots of other things where, if you have your eyes open, and if you’re willing to let the problem tell you what the solution is, then I think there are tremendous research opportunities. And it’s actually much more satisfying because I feel like the research problems I’m working on are really real and not just some things that sound great in a grant proposal.

Doug: One of the areas that really inspired me was your early work on sound. Recently we’ve done a lot of work on sound synthesis. And students often asked me if it can be done in real time, or if real-time sound is possible. And I often laugh and think of you showing me the “Sonic Explorer” real-time sound demo on your SGI Indy back in the 90s from your work with Kees van den Doel [16], and so I think that you had quite an influence in that area. What do you think of all the work on sound, and the challenges of actually making sound synthesis practical?

Dinesh: A couple of things here. I think sound research had both big advantages over graphics and some big disadvantages. The big advantage was that everybody had really great high fidelity “display” devices which you didn’t really have in graphics, which was slower in coming. On the other hand, I think it was harder to get sound research to move to fully physics-based synthesis. Our early work was constrained by trying to do it in real time as part of contact rendering. But people are extraordinarily sensitive to subtle things about sound. I think your work has shown that yes, it’s much harder to get all these subtleties, but it is indeed possible and you just have to develop new algorithms. It’s more expensive, and maybe not real time now, but computing is going to cost zero in the future, asymptotically. So I think that it’s going to happen.

Doug: How did you first get interested in sound?

Dinesh: It was because one day I really understood that contact perception was a multi-sensory phenomenon. This is a true story of how it happened. I was doing a haptics project, and we’d gotten our very first haptic device from Vincent Hayward at McGill. We were trying to control it to make contact with a virtual object, with a cylinder. And haptic devices, as you know, have many features in common with loudspeakers in the sense that they’re meant to have very small inertia and high force bandwidth so that they can render sharp contacts. So they’re almost like loudspeakers.

Doug: Right.

Dinesh: So I had this master’s student at that time writing some code to try to touch the object... and the controller was going mildly unstable and it was vibrating. And when you went and touched the object you’d go “EEeeeenng” ... “DREEEEEE”. It sounded exactly like you are machining a cylinder on a lathe and it’s vibrating. The beauty was you could perceive it instantly. So I told him, “Hey, what you should do is turn this into a simulation of machining.” We made a simple 2D demo that I think, to this day, it’s one of the most immersive compelling demos even though it was so simple. You went and touched this turning cylinder and you felt the vibration in your hand, you could hear it, you could see sparks fly ... it was an amazing multi-sensory simulation with negligible latency between all the three modalities, perfectly synchronized. So that was amazing. After that I said, “Okay, we should really now look at how to synthesize the sounds directly instead of using haptic devices as...”

Doug: as a loudspeaker.

Dinesh: [laughs]

Doug: That’s perfect. So the first synchronized...

Dinesh: ... multi-sensory simulation of grinding, yeah [laughs].

Doug: That’s great. Well, I think that’s probably a good spot to end. Thanks so much for answering all of these questions.

Dinesh: My pleasure.

Doug: And congratulations again on the award.

Dinesh: Thanks Doug!

References

[1] R. Barman, S. Kingdon, A. K. Mackworth, D. K. Pai, M. Sahota, H. Wilkinson, and Y. Zhang. Dynamite: A testbed for multiple mobile robots. in IJCAI 93 workshop on Dynamically Interacting Robots, Chambery, France, August 1993.

[2] Y. Fan, J. Litven, D. I. W. Levin, and D. K. Pai. Eulerian-on-lagrangian simulation. ACM Trans. Graph., 32(3):22:1–22:9, July 2013.

[3] Y. Fan, J. Litven, and D. K. Pai. Active volumetric musculoskeletal systems. ACM Transactions on Graphics (Proc.

SIGGRAPH), 33(4):152:1–152:9, 2014.

[4] D. P. Greenberg, J. F. Abel, A. R. Ingraffea, W. McGuire, J. C. Dill, D. P. Loucks, P. N. French, and M. R. Taylor. Engineering: Cornell Quarterly, Vol. 16, No. 3 (Winter 1981-82): The Advance of Computer Graphics. 1981.

[5] D. L. James and D. K. Pai. ArtDefo: Accurate Real Time Deformable Objects. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99, page 65–72, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[6] D. M. Kaufman, S. Sueda, D. L. James, and D. K. Pai. Staggered projections for frictional contact in multibody systems. ACM Trans. Graph., 27:164:1–164:11, December 2008.

[7] P. G. Kry and D. K. Pai. Interaction capture and synthesis. ACM Transactions on Graphics, 25(3):872–880, July 2006.

[8] D. I. W. Levin, J. Litven, G. L. Jones, S. Sueda, and D. K. Pai. Eulerian solid simulation with contact. ACM Transactions on Graphics (Proc. SIGGRAPH), 30(4):36:1–36:10, 2011.

[9] D. Li, S. Sueda, D. R. Neog, and D. K. Pai. Thin skin elastodynamics. ACM Trans. Graph. (Proc. SIGGRAPH), 32(4):49:1–49:9, July 2013.

[10] D. K. Pai, R. A. Barman, and S. K. Ralph. Platonic beasts: Spherically symmetric multilimbed robots. Autonomous Robots, 2(4):191–201, December 1995.

[11] D. K. Pai, J. Lang, J. Lloyd, and R. J. Woodham. ACME, a telerobotic active measurement facility. In Experimental Robotics VI, pages 391–400. Springer, 2000.

[12] D. K. Pai and M. C. Leu. INEFFABELLE - An environment for interactive computer graphic simulation of robotic applications. Proceedings. 1986 IEEE International Conference on Robotics and Automation, 3:897–903, 1986.

[13] D. K. Pai, A. Rothwell, P. Wyder-Hodge, A. Wick, Y. Fan, E. Larionov, D. Harrison, D. R. Neog, and C. Shing. The human touch: measuring contact with real human soft tissues. ACM Transactions on Graphics (TOG), 37(4):58, 2018.

[14] D. K. Pai, K. van den Doel, D. L. James, J. Lang, J. E. Lloyd, J. L. Richmond, and S. H. Yau. Scanning physical interaction behavior of 3D objects. In Computer Graphics (ACM SIGGRAPH 2001 Conference Proceedings), pages 87–96, August 2001.

[15] S. Sueda, A. Kaufman, and D. K. Pai. Musculotendon simulation for hand animation. ACM Trans. Graph. (Proc. SIGGRAPH), 27(3):83:1–83:8, 2008.

[16] K. Van den Doel and D. K. Pai. The sounds of physical shapes. Presence, 7(4):382–395, 1998.