Catalyzing Computing Podcast Episode 3 - What is Thermodynamic Computing Part 1

Intro [00:00:10]

Hello. I'm your host Khari Douglas, and welcome to Catalyzing Computing, the official podcast of the Computing Community Consortium. The Computing Community Consortium, or CCC for short, is a programmatic committee of the Computing Research Association. The mission of the CCC is to catalyze the computing research community and enable the pursuit of innovative, high-impact research.

This episode of the Catalyzing Computing Podcast was recorded following the CCC's Thermodynamic Computing workshop. The workshop took place in Honolulu, Hawaii January 3rd to 5th, 2019. In this episode, I interview two of the lead organizers of the workshop, Tom Conte from Georgia Tech and Todd Hylton from University of California, San Diego. I also speak with workshop participant Christoph Teuscher, a professor at Portland State University. Enjoy.

Interview [00:01:01]

Khari: Hello. You're listening to the CCC podcast – coming from Honolulu, Hawaii, after the Thermodynamic Computing workshop. Here with one of the lead organizers, Tom Conte. Tom, how are you?

Tom: Good. How are you?

Khari: I'm good. So can you explain a little bit about your background and your involvement with the computer science community so far?

Tom: I am naturally an electrical engineer, but ended up in a computer science department somewhat by accident. So I'm joint appointed in computer science and electrical and computer engineer. My natural field is computer hardware or computer architecture. Over the last six, seven years, I've been looking at what we need to do to overcome the current constraints on silicon and CMOS (Complementary metal–oxide–

semiconductor). So I formed a committee inside IEEE (Institute of Electrical and Electronics Engineers), what we call an initiative inside IEEE, to study fundamentally different ways to compute. We call it Rebooting Computing. We've held four invitation only workshops. We've influenced federal policy, including the National Strategic Computing Initiative and DARPA's (Defense Advanced Research Projects Agency) Electronics Resurgence Initiative.

Khari: So that was kind of your initiative for proposing this workshop?

Tom: Yes, in general, when I was doing a lot of that activity, the incoming chair of the CCC, Mark Hill, contacted me and said he would like for the CCC to have some involvement in Rebooting Computing.

Khari That sounds great. So we're also here with one of the other organizers, Todd Hylton. Todd, can you say a little bit about your background and involvement with this workshop?

Todd: Sure. Well, I'm at UC San Diego. I'm a professor of practice there in electrical and computer engineering. I'm an executive director of a robotics institute. I spent most of my career in industry, but also in government. I was a DARPA program manager. I did a project in neuromorphic computing called SyNAPSE. It was actually I was there that I got interested in a lot of these ideas about thermodynamics and computing, because I was looking for the answer of how I get these technologies to like organize themselves.

Khari: So can you guys give me some background on sort of basic theory behind Thermodynamic Computing and how it would work?

Tom: So the basic theory is really the theory of open system of thermodynamics, which over the last 20 years has really undergone a revolution in our understanding. A lot of people think of thermodynamics as a very static field, but it's a field that we're still learning about. Given those advances, we've been able to hypothesize a new method of computing that is much closer to how natural physical processes work as opposed to the way we compute now, which is to really fight against those processes.

Khari: Anything else you want to add, Todd?

Todd: So I think I would mostly, you know, I would second all of those comments. I mean, we do have, I think, a new era of understanding in thermodynamics. We used to think of thermodynamics as mostly closed systems at equilibrium or close to that. But you know, what really matters in thermodynamics, what we really care about is not that. It's how…it's the non-equilibrium, open thermodynamics that pretty much everything is. Our computing systems are also thermodynamic. It's thermodynamics that drives the process of a computer down on the on the hardware level altogether, but it's still sort of in a box. We need to kind of unbox it so that it can sort of interact with the world, and so that we don't have to do so much work and spend so much energy doing it.

Khari: Right. So for people who might not know, can you summarize the difference between non-equilibrium and equilibrium, thermodynamics?

Todd: So in equilibrium thermodynamics…probably most people have a little bit of familiarity with the idea that entropy increases, that the world becomes more disordered over time. And so the basic observation of equilibrium thermodynamics is that if you take some collection of stuff – let's say, as a simple example that physicists always use, gas in a box. That if the box doesn't interact with anything else, that the gas will sort of maximally distribute itself, and in this sense it becomes kind of maximally disordered or maximally entropic. For example, if you were to take a box and somehow put all the gas on one side of the box, say by putting in a partition, and then you removed that partition, everybody knows the gas isn't going to stay on one side of the box. That it's going to go to this state called equilibrium.

So equilibrium is a very uninteresting state. [Laughter] It's essentially a state where nothing changes.

Khari: Okay.

Todd: Okay, and so although it's tractable, sort of, from a mathematical perspective…and so in some sense it's simpler than the things that go on in the real world – we'll call those things open thermodynamic systems. So open thermodynamic systems, that's a much harder thing to figure out. Open thermodynamic systems don't become maximally disordered. In fact, for some reason, which is what all the excitement is about now in non-equilibrium thermodynamics, for some reason stuff orders itself everywhere.

Khari: Interesting. So that's one of the questions we're trying to solve here.

Todd: Yeah, that's one of the questions we're trying to solve. And I think, you know, why should we do it right? If you look at where computing is today there's ideas of thermodynamics throughout computing. As I mentioned before, down at the hardware level you've got all this current and energy moving around, and transistors flipping and so forth, and that's what drives the process. But then if I go up the stack…let's talk about machine learning, for example.

So everybody thinks, and I agree that, if we're going to do stuff in the real world, like have cars that drive themselves and build cool robots we're going to have to learn a lot of stuff, we can't program it.

Khari: Right.

Todd: But if you look at ideas in machine learning, thermodynamics is everywhere. You've got like energy costs functions, you've got stochastics, you've got probability distributions, this stuff is all thermodynamic ultimately – in the natural world that's where those ideas come from.

So we go back to the hardware again. So now the devices have gotten so small, [laughing] that they're actually noisy, right? They actually are probabilistic. They're actually doing the things that we're trying to model in the software.

What if we didn't do it that way? What if we just were able to put the parts together? So that the stochastic and the probability distribution and the non-equilibrium formation of new structures that we do all the time in machine learning goes on, not only in the software, but also in the hardware in a vastly different and more efficient way. That's the vision for what we might do.

Khari: Right.

Tom: Let me just add a little to that. What we do today in a lot of ways is take that stochasticity, that randomness and fight against it, because we tried to build reliable switches at the gate level so that we can make the next level, the micro-architectural level, happy and they know how to reason about things, and the next level above that, and above that, and above that. And now we have very different models of how to do computation, than what that was designed to support, that naturally exploit randomness and stochasticity. And as Todd said, maybe we removed some of those levels, and in doing so we can get something that's fundamentally far more energy efficient.

Khari: Yeah, that's a pretty compelling vision for this workshop. So the workshop was mostly organized into breakout groups, sort of domain groups in terms of physical systems, model systems, theory, and then cross cutting. Can you guys talk about any interesting conclusions that came out of the breakout groups you were involved with?

Todd: Yeah, I'll tell you one. One that really struck me was that we were really struggling to find the right words to use to talk to each other, because we had a very diverse group of people. We had people who were hardcore computer architects, device people. We had people who were theorists, I mean at the highest levels of abstraction, world famous people in non-equilibrium thermodynamics and quantum mechanics, for example. We had people more from the biological and cognitive sciences who think about biological systems and brains and neuroscience and so forth. And so in talking to each other…it was interesting because even the simplest concepts of thermodynamics are pretty tricky, and not everyone feels very well versed or like they really understand thermodynamics.

I always tell this story that I took thermodynamics twice, once as an undergraduate and once as a graduate student. And I thought – in some sense, I'm a physicist by training and in many ways it should have been the easiest class I ever took, the equations are incredibly simple – I thought it was incredibly hard. It was really, really hard to get sort of that intuition about what you're talking about.

So, you know, these ideas like entropy you sort of finally, after a long time, get a little bit used to them. And even then, even among the experts in the world, they'll argue many fine points about what entropy is, and so that's where we are. That's also, I think, why there's just a fantastic opportunity, right. When you find that kind of overlapping interest, but disconnect in understanding, then I think you have the chance to do something really different. So Tom and I, and Susanne and Natesh and Erik DeBenedictis and Stan Williams and John Paul Strachan worked really hard to get a really diverse group of people here. So that's my first observation about sort of what I saw going on.

Khari: Anything else from you, Tom?

Tom: One thing that became interesting is we put people in these three natural groups: the people who are interested in realizing systems, the people who are interested in models of computation using thermodynamic systems, and then the theorists. And as we went on and we reported out there was a common core set of issues that we knew we needed to solve. In some ways I didn't expect it to be quite as unifying as it was. I expected there would be a couple of common nodes, but there were quite a few very common themes and needs among the very diverse groups.

Khari: So what are those common themes that kind of came out from the breakouts.

Todd: One is a real need for a measure of the complexity of the system so that we can understand how efficient it is to do a particular computation. That was, I think, an important need. Another one was a strong need for how to express a problem to solve.

People in the workshop had slightly different views of how to do that. We all knew it was possible or we wouldn't be here, but the need for, as Todd said, a common language; in this case, a common, if you will, computer language was a cross cutting theme among the different groups.

Khari: Anything else to add, Todd?

Todd: I think that from the theory side that there was enthusiasm for the idea that people would be really trying to take advantage of the stuff that they had done. I think more from the application side there was enthusiasm that there were people working on the theory. You could imagine the parts sort of fitting together.

Once you have the theory and some demonstrations then it becomes a whole lot simpler to sort of expand it, and teach it, and bring a whole bunch of people into it to make it all happen. It's going to take some time right. But I think, you know, we could get there. We're at the right spot in time to do it. That's what I also felt like.

People were really enthusiastic about it. It's one of the few workshops I've been to where I didn't see a whole lot of people doing their email. [Laughter]

Khari: Yeah, people did seem very engaged, especially ‘cause it was two and a half days, so we really…really put people on the grindstone.

So where do you see thermodynamic computing making an impact on the future? Like maybe you give me a realistic, or not realistic, but more pessimistic projection and then your sort of wildest dreams, thirty years from now.

Todd: So I have a hard time being pessimistic….

Tom: Me too, sorry. [Laughter]

Todd: …but I'll do my best. I think another theme out of the workshop was that even if we can't sort of figure out the magic dust that makes the machine organize itself, that uniting the thermodynamic themes that currently exist across the current levels of the computing stack would probably yield tremendous advantages in efficiency of computation. I think that's undeniably true, although it's not like we don't do these things now, we do, but I think there's opportunity, some theoretical guidance to maybe do it in a more rigorous way than we've done in the past.

So I think that's fantastic. If we can do that, we can probably extend the current computing paradigm for another few decades, get around some of these huge power problems we've got, and expand the role of computing everywhere. I mean, if we're going to have trillions of computers hooked up to the world, they’ve got to be energy efficient and we're getting there.

So that's my pessimistic outcome. [Laughter]

And so my optimistic outcome is that I'll be able to build something that essentially by looking at the world spontaneously creates a representation of it, so to speak – learns it. And it doesn't do it because I said here's the machine learning model that you've got to follow, and you've got to back propagate this way and that way and you've go this many layers. I mean I'll do some of that. But it's more like, you know, there are these experiments, I think, that people did on ferrets a while ago where they took the optic nerve and they cut it and they stuck it in to the auditory cortex, whatever you call that part of the brain, and, incredibly, the animal could still see.

Khari: Really?

Todd: Yeah. So that's what I mean, right. So imagine you had a computing system that was that flexible. Where you could feed it information, feed it inputs, and it would spontaneously create that sort of organization that you would need so that you could understand what you were doing. I mean we do we do that all the time. I walk into an environment I haven't seen before and pretty rapidly I can make sense of what's going on. So that's my dream that we'd be able to make technologies that could do that.

Khari: And you, Tom?

Tom: My enthusiastic view is a little closer in than Todd's, which is fine. In my view, we're at a place right now in computing where we've reached the limits of our current computing model.

We've, if you will, we've gone down a particular rat hole and we're at the bottom, and there might be a much better rat hole next door. So for me, the advantage of this thermodynamic way of computation is that it seems to have a lot of promise for being a fundamentally different way to compute.

Specifically, as I said I think a little earlier, the way we compute now goes against the grain of what is natural. We force processes to be reliable; we force processes to behave a certain way. If instead we went with the grain and we followed the way natural physical processes evolve, we can achieve certain classes of computation at a much, much…we're talking tens, hundreds, thousands of times more efficient levels of energy consumption than what we're doing today.

And as Todd said, we're at the end. We're really…we're at a point now where if we imagine computing everywhere, we can't power of that.

Khari: Right

Todd: Yeah, so if we're already using like five percent of the electrical energy in the United States to run our computers, you know, scaling up by a factor of 10 doesn't look like a very good idea.

Khari: Yeah [laughs], I guess that's bad both logistically and environmentally.

Todd: Yeah, pretty much every imaginable way is a terrible idea. [Laughter]

Khari: All right. So I guess just from sort of a more logistical standpoint, how did you feel the proposal approval process went? In terms of getting this workshop off the ground at the CCC and then getting it organized?

Tom: Well, I felt that the CCC is very easy to work with and they're interested in looking at very different things and trying new things and building communities around new ideas. And that was really a big advantage to us because we've had this, if you will, crazy set of ideas or perhaps I should say idealistic set of ideas.

Todd: I guess we're unconventional, optimistic people, and we think great things can be done. We think…we think it could be done, right? We can't do it by ourselves. There's not much that a few people can do.

So the opportunity to bring this group of people together was something that I had wondered for probably 10 years…how I would ever be able to do such a thing. So then Tom says well why don't we talk to the CCC? I think they'll I'll think they might be interested, and here we are.

Khari: Right, well that's great. So anyone else interested in learning more about this topic, what should they do?

Tom: I think they should read the report, but they should also feel free to reach out to us, express their interest. One thing that came forward over these last few days is that we've really built a community of fellow travelers, and if they want to join that community, there's enough room for everybody.

Todd: Yeah, quite a few people said to me, I didn't know that there were people who were actually thinking about this. I thought I was all-alone. I suspect there are probably a lot more people out there.

Khari: Yeah, so hopefully they listen to this. There will be links posted somewhere in the page that this is on, so follow those if you want to learn more. Thank you guys for being here. Any final thoughts?

Todd: No, I'd just like to thank the CCC, and you Khari for making this all this happen.

Tom: I'd like to second that and also thank all of the attendees for working so very hard. I was really impressed with how much, given the beautiful surroundings, how hard they really worked. [Laughter]

Todd: I didn't see much of Hawaii the last few days.

Khari: Keeping them in a windowless room and off the beach…

Tom: …was somewhat cruel. [Laughter]

Khari: Yeah.

Transition to interview with Christof Teuscher [00:20:00]

Khari: Hello, you're listening to the CCC podcast, coming live from Honolulu, Hawaii, where we've just concluded the Thermodynamic Computing workshop here with one of our participants, Christof Teuscher. Thanks for being here.

Christof: Thanks for having me.

Khari: So could you tell me a little bit about your background, your research interest?

Christof: Sure. My name is Christof Teuscher, I'm a professor at Portland State University. My training is in computer science, and I think for my entire career I've worked sort of across disciplines at the boundary between computer science, computer engineering, biology, chemistry, self organization, you name it, various fields, and I think that makes things really interesting for me, but also for my students.

Khari: Ok. So why did you decide to attend this workshop? What inspired the interest?

Christof: It sounded really exciting because it's something that I've been interested in for a while, and I'm always interested in these workshops because you meet people that think along the same lines, but also folks that may be thinking very differently about certain things, and doing sort of that force fit…just an example is the terminology. Some people understand something very differently; for example, for what they mean by an algorithm. If you ask a computer scientist, they have a very specific idea what that is. But if you ask a biologist, for them, that may be something quite different. So just getting on the same page is challenging, but it's also a really good exercise so that you understand each other and can think in the same framework.

Khari: Okay. So can you sort of summarize, I guess, what the breakout groups you participated in discussed at the workshop and what kind of conclusions came out of those?

Christof: Yeah. So I was mostly in breakout groups that were related to the physical devices and substrates. So we were talking mostly about what kind of physical substrates are what we call thermodynamic computers and which substrates may not be. How can we get there? How can we make these substrates or devices more efficient? How can we engineer them? How can we gain orders of magnitude of power efficiency, perhaps speed, perhaps increase the density by another order of magnitude or more?

So I don't know if there is an outcome, except that I think, and I can speak for myself, but I think I can also speak for many of the others, there is a buzz, there is excitement. There is, I think, a feeling that something is really there, even though we kind of still struggle with some of the concepts.

What we actually mean by thermodynamical computer, I don't think we have a formal definition by any means. We have a number of working definitions at this point, which is totally fine. [Laughter] But I think everybody agrees that there is potentially something really big beyond that concept, something that we need to explore, and something that has a lot of potential for bringing electronics and computers to the next level.

Khari: So I guess what kind of potential outcomes do you see for the future of thermodynamic computing? What would a future system be able to do that we can't do now?

Christof: Yeah, that's a bit of a loaded question, I guess. I would say more the question is how would a system, a future system, should do things differently with respect to energy consumption? I think that is the main goal of thermodynamic computation. Do things with less energy, have higher densities of devices in the same chips so that the chip doesn't burn, but there may be functionalities that we haven't thought of. There may be problems that we can solve that were intractable by traditional CMOS devices that we have in these days. So it's possible, but it's probably more unlikely.

Khari: Right.

Christof: Or maybe at least far further out. So we shall see, but I think society will benefit from any computer that is better, faster, has more transistors, has more capabilities, and certainly there's excitement about AI, machine learning, and what not. So if we can do these things faster, quicker, more energy efficient, instead of your cell phone that lasts a day, maybe it lasts a week, right. So these are all good things that will have certainly an impact on society and folks.

Khari: Yeah, they would certainly be beneficial. Did you meet anyone at this workshop that you feel like maybe you could collaborate with in the future?

Christof: Oh, absolutely. That is, I think, another really good outcome that you can potentially collaborate, put in grant proposals, write papers, and whatnot with people that you didn't know before. Yeah.

Khari: That's great. Do you have any of your own research that you kind of want to plug? Any thing interesting you've been working on lately?

Christof: Oh, yeah. We've been working on a lot of exciting things. Especially, I think what I realized is actually a thermodynamics computer to a large extent is a biomolecular system computing in DNA, because these reactions are governed by thermodynamics principles inherently and we harness them for doing computations. So it's sort of a natural thermodynamic computer that we already use, but still there is lots of things we can improve, we can do better, and extending those ways of thinking to larger sets of devices, larger set of physical substrates, I think will have huge benefits for the industry and for society.

Khari How does computing with DNA work for people who might not know?

Christof: You're basically harnessing the fact that DNA recombines naturally and you can code information into DNA strands and build logic gates out of these DNA strands that then recombine and then interpret the outcome of that recombination as a logic 0 or logic 1. So that's sort of it in a nutshell. There's a lot more to it and it's challenging in these days to build large-scale computers, and so that is something that needs to be addressed for sure in the future.

Khari: Great. Well, thanks for sitting down and taking the time out of your day to talk to me. Have a safe flight home.

Christof: Thanks for having me.

Outro [26:15]

Khari: That's it for this episode of the podcast. I hope you enjoyed it. I'll be back next week with more interviews from the Thermodynamic Computing workshop, including one with chemist Gavin Crooks and one with Natesh Ganesh, a member

of the organizing committee who's a grad student at the University of Massachusetts, Amherst. If you want to learn more about the workshop, visit the workshop web page under visioning activities at cra.org/ccc. Until next week….