Catalyzing Computing Podcast Episode 4 - What is Thermodynamic Computing Part 2

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, which took place January 2019 in Honolulu, Hawaii. In this episode, I interview workshop organizing committee member Natesh Ganesh, a PhD student at the University of Massachusetts, Amherst. He was awarded the best paper award at IEEE ICRC '17 for the paper A Thermodynamic Treatment of Intelligence Systems. I also speak with workshop participant Gavin Crooks, formerly a senior scientist at Rigetti Quantum Computing, who developed algorithms for near-term quantum computers.

This is the second episode on the Thermodynamic Computing workshop. For more about the background and inspiration for this workshop, listen to my interview with workshop proposers Tom Conte and Todd Hylton in *What is Thermodynamic Computing Part 1*. Enjoy.

Interview [00:01:01]

Khari: So you're listening to the CCC podcast coming from Honolulu, Hawaii, where we just concluded the Thermodynamic Computing workshop. Here with one of the participants and organizers, Natesh Ganesh.

Natesh: Hi. Thank you for having me on the podcast.

Khari: Yeah, for sure. So Natesh, can you give us a little bit about your background and how you got involved with the workshop?

Natesh: I am currently a PhD student at the University of Massachusetts, Amherst. I have been working on non-equilibrium thermodynamics, information theory, Landauer limits to computing and I happened to run into Todd Hylton, one of the other organizers, about a year and a half ago talking about the intersection of these topics. We kept talking and we felt like we had some ideas that were common where we thought there was a lot more that we could do from a thermodynamic perspective when it comes to computing. From there, him and I worked out a lot of ideas together and I have been presenting on some of these topics, especially the thermodynamic conditions as it relates to inference and learning, at other conferences. Him and Tom were talking about putting together a workshop and they reached out to me, because I do have a certain sense of background in a lot of these intersecting topics. Yeah, that's how I got here.

Khari: Can you talk a little bit about your research using thermodynamics for inference? How does that work?

Natesh: There's been a lot of work, it's hard to just cite a few authors, but we do have a few people here…people like Susanne Still, David Wolpert, and Jim Crutchfield. All three of them were here. They had done work where they were interested in essentially the lower limits of dissipation when it comes to learning – when a system is learning or when a system's performing inference. Susanne has done a lot of work, especially when it comes to predictive inference.

So I was interested in the same topic, and I had taken some of their work and tried to apply it to neural networks for example. I was very interested in what are the lower limits to computation with respect to neural networks. I was also curious about adopting some of the recent results that I've obtained in these areas of non-equilibrium thermodynamics and tried to explain some of the general thermodynamic conditions and how they map on to specific learning algorithms.

Khari: Ok

Natesh: I was trying to then ask the question, if…instead of building systems that satisfy certain algorithms could we build systems that satisfy certain thermodynamic conditions, and if we did that, then because these conditions and algorithms are equal, will we get the algorithms as a byproduct?

Khari: Right.

Natesh: That kind of gelled well with what Tom and Todd were also looking for. So I guess that's like more of my research.

Khari: Ok, so the workshop was kind of divided up into breakout groups, three main areas, physical systems, model systems and theory. You were one of the leaders of the theory group. Can you talk about what kind of conclusions came out of those breakouts?

Natesh: So it's very interesting because…to put all the background in it, even though I do a lot of work in non-equilibrium thermodynamics, I still feel like I'm more on the engineering side of things than a lot of these people. So I was trying to understand from their perspective. It was great to have Suzanne there as a co-leader, and let me divided into a couple of parts that we kind of realized. There is this huge history of people working on something called thermodynamics of computation, and for a lot of people, it was that area. They thought that, you know, thermodynamic computing was the same thermodynamics of computation.

Khari: Can you explain the difference between those two things?

Natesh: Thermodynamics of computation: we have a lot of these existing computing technologies and a lot of them do not take into account thermodynamic inefficiencies. There are a lot of thermodynamic costs that can be optimized in these existing technologies. For people in theory group felt like for the longest time, maybe we didn't have the tools to do it, now, more and more, we have the tools to apply these new ideas and remove those inefficiencies across the stack in existing technologies, and that in itself should provide a lot of energy benefits and that should reflect in one aspect of the final report.

But there are others, which is where I stood: I was more interested in asking the question, say, given a particular problem statement, some kind of problem that we want to solve, can we determine a set of thermodynamic conditions that is equal under that problem statement, and then, starting all the way from the material level, put together a system that solved the problem by satisfying those conditions? So they…essentially it could be thought of a little bit as like top down versus bottom up, and obviously there's going to be overlap in certain aspects of it.

So, there were people who were interested, in the theory a group of, you know, working on those aspects of the theory and improving the tools to do the top down part and improve the thermodynamics of computation. There are others who were particularly interested in extending the theory that we have to improve the bottom up part. For example, Michael DeWeese, who I was only happy to meet, talked about how a particular system – for example, Gavin Crooks, David Sivak, other great people in this area – have figured out, say, given a particular initial state and final state that figured out how to like…as an external controller can tweak the knob to go from A to B, the path to go from A to B, so that you say minimize dissipation or something.

Khari: Ok

Natesh: Some kind of optimal path. Michael DeWeese was very interested in saying, you take that framework, and then can you extend it and say you have to go from A to B, you don't want to have any external person tweaking the knobs, but you set it up, set the whole system up in such a way that will go from A to B while optimizing something, and that is kind of what we wanted to do from that bottom up perspective.

So those are the kind of areas that people noted that needed to be solved, and those have like wide implications in a lot of other cases outside of computing, but I think it gels really well with what this group wanted to do particularly. There was also detailed aspects of like wanting to improve, for example, Susanne focused on asking the question if we really pushed ourselves to the thermodynamic limits do new rules for information processing kind of come out of it? And, you know, some of these rules might already be existing rules, which is great, then we have done our job in those areas, but there might be other rules that we didn't expect, and gives us more ideas on how to compute.

Khari: Ok. Do you have any examples of what kind of rules those would be?

Natesh: Ummm…I mean, not exactly. If I was venturing a guess or something like that, I would think that if I had a system that was…and this is kind of based on what Susanne has been working on and some of the work I've been working on…I would think that if you had some kind of like plastic system, self organizing system, that was, say, maintaining homeostasis. If it was maintaining homeostasis in a very energy efficient manner, I would think that the only way for to do it is to kind of learn its environment and predict it.

Khari: Right.

Natesh: So it tells me that pushing it to the thermodynamic limit of efficiency in certain kind of systems produces the learning inference aspect of the behavior. So that is one such rule. I'm sure that there are…instead of just focusing on the minimizing dissipation part, non-equilibrium thermodynamics is much more richer and you can characterize a lot of different behavior. So I wouldn't be surprised that we can extend it and come up with all kinds of new learning rules.

Khari: Ok, so I guess outside of your specific breakout group, what do you think was sort of the big conclusion to come out of the workshop that you think is gonna be highlighted, sort of, in the report?

Natesh: I think what was really important was that they pointed out that for something like this to succeed going forward, we have to identify, say, a class of problems. That we have to be willing to show that these kind of ideas will be better than existing technologies, and I think that was important for me to hear. Also I think that part and also talking about at the end of the day you can build all of these technologies but you have to have a way to properly interface with it. You need ways to let people, different people be able to access it and use it and do experiments and test it. So those kind of fundamental issues that need to be solved, having that clarified was I think was very useful.

Khari: Yeah, I guess that's big because even if you have an abstract idea, logistically, you still have to have a system that anyone, well not anyone, but someone trained could use and be successful with it.

Natesh: Right

Khari: So building off that, what do you think the impact of thermodynamic computing will be in the future? Like if you had to imagine a sort of moonshot, 30 years from now, what kind of application would a thermodynamic computer be doing?

Natesh: So I'm a little biased here. I think like I said, this is where it helped having different people because Tom talked about randomized algorithms, which I thought was very interesting. I didn't particularly think about it when I was working on these ideas, but I think that definitely should be a big part, but the area that I'm specifically interested in, and I think these kind of ideas will really help is the area of artificial intelligence.

I really do think that these kind of like thermodynamic constraints on biological systems is what produce specific kind of like learning and influence behavior in biological systems. So…especially in this current state of artificial intelligence and machine learning, we want…these areas are really big now and we want to build intelligent systems and we want them to be energy efficient, and if it happens to be the case that pushing limits of thermodynamic efficiency gives us the learning as a byproduct, then we got both of...best of both worlds there.

In that sense that is the area that I'm most excited about. That is area that's my personal research. Maybe I maybe I am really biased, but that that is the area that I think that we will really make a lot of headway into. For example, areas like neuromorphing computing are going to provide a huge energy boost, right? And these ideas of thermodynamic computing are like extensions going past that and asking the question, can we maximize these efficiencies everywhere? Or maybe start from the efficiencies and go the other way around and get the learning. So I would see…and if you talked to Todd, maybe he might have also mentioned this…a lot of these areas, which need real-time learning; you know in a dynamic environment something like a chip that you can stick into a robot or self-driving cars, you need them to be both intelligent and efficient.

Khari: Right.

Natesh: And I think like…that's not even the start of it. There are all these IoT devices as well. So, those are the areas that I think that thermodynamic computing can really provide a lot of benefit in going forward. And I'm sure as we work our way through this we'll find out more and more specific areas where these ideas might provide a lot better or more efficient solution than existing classical computers.

Khari: Yeah, that's definitely a big thing. Kind of off-topic but speaking of AI; I know we were talking at dinner about self-driving cars and the criteria you think is needed for that to be successful. Could you talk a little bit about that?

Natesh: [Laughter] OK…

Khari: I don't know if that ties into thermodynamic computing at all.

Natesh: I don't know if it specifically ties into thermodynamic computing yet. I mean we're still very early to tie it into…like you said in an ideal situation 30 years from now I can say maybe a thermodynamic chip is what you need to like stick into a self-driving car and I'll do the job for you.

But self-driving cars in itself…how do I put it? I think how much, how far we've achieved is like great, but this is more of an anecdote where I was talking about how I was recently in India. And people who drive there are like acrobats. They're like amazing. They make vehicles do all kinds of things to navigate with each other, and that, I think is like amazing. And I kind of asked the question, if we can…to me, at least, I was partly joking, partly true…solving self-driving cars meant solving self-driving cars in India.

Khari: Right.

Natesh: It felt like if you can solve it in the hardest environment, then, you know, you should be able to solve it everywhere else. That kind of like…maybe I'm thinking too far ahead, but that kind of like…that seemed like a weird criterion for me.

The other part I guess I was curious about – and I'm sure if there are any self-driving car people who are working on these things already and can let me know that'd be great – it seemed to me that self-driving cars kind of like learn how to drive b=y kind of learning the rules of the road, which is what it seems to me, but it feels like humans learn how to drive or drive by knowing how to break the rules kind of, but choosing not to do it, in terms of risk-reward.

So it seems that solving the latter problem is much harder than solving…even the earlier problem is extremely hard, but the latter problem is much harder than that, so it seems to me that maybe that's the way we need to do it. And I think that there is starting to be work in that direction. So will thermodynamic computing solve it? No idea. [Laughter] But, you know, it's an exciting time to even be talking about this. So I think that's great.

Khari: Yeah. I mean, yeah, seems kind of unrelated, but it was an interesting point. Maybe there's some way thermodynamics can overlap with that in terms of how you assess risk, like within the system, right?

Natesh: Yeah, I think that that would be interesting. I think the part that thermodynamic computing can help there, at least to me, initially is purely that if we could build much more efficient inference chips just that is going to be huge. Because there is no way we can put in something with say…a set of computing chips that take up about a kilowatt of power or something on a car and have it run a self-driving car. I don't know how that is viable.

Khari: Right.

Natesh: But if thermodynamic computing allows us to create an inference chip much more efficient than anything else out there, then you know, that is a start. But then again, like…I feel like to get into this real notion of at what level…how much more thermodynamic computing chips can help the A.I. area needs a lot of work, because…I feel like I’m going to go into a long rabbit hole here a little bit. But, I would be interested in understanding if we to really build some kind of like thermodynamic inference chip, we'll have to extract the thermodynamic principles of intelligence in existing systems. Like what we're really trying to do is kind of build…extract principles of intelligence from the brain and kind of achieve it in an ideal scenario. So if you're trying to extract the thermodynamic principles of intelligence from existing biological systems and then implement them in some kind of artificial system, it might be the case that the complexity of human intelligence does not exist at the neural network level.

There might be some intelligence at the neural network level, but the complexity might be more existing in another either upper or lower level, and ideally we would want thermodynamic conditions to extract that. So if we build thermodynamic computing based on these newer conditions of intelligence that did not necessarily map onto a neural network…say in the brain it was neural network plus glial cells plus all kinds of stuff, right? Then suddenly it might be the case that purely from an intelligence performance perspective thermodynamic computing might give you a better chip.

But I think like at this point, I simply do not know, which is why we need to figure those questions out. Because from the neuroscience community there’s continuous work coming on how we have all these newer parts that contribute so much to our own intelligence, and there is work from the A.I. and machine learning community to put in these new ideas into existing neural networks. That's what we want to do as well. So hopefully it will work out.

Khari: Yeah, we'll see where it goes. So just logistically, how do you think the workshop was organized and how easy do you think that process was from other people that might be interested in holding a workshop or something like that?

Natesh: I think was great. I mean, thank you so much for all your help setting it all up. I also like the way we split it into initially. We had these people who were linked in similar areas to work on work together, and then we kind of randomized the groups. That was great because I think all these groups having different coming from different areas of expertise, working together in a single group was really good for people to come in and say, “oh, these ideas have already been solved” or “this is an issue you really need to care about.” So I really loved that aspect of it.

I've not been to a workshop where we like randomized groups like that. It was also good in a workshop sense because…maybe more people have been to workshops…. I mean, I'm still a grad student, so I of like navigating all these areas. [Laughter]

You know, a lot of the people here were heroes of mine, so that was great for me. Maybe it I was like, “oh, this is my chance to see all my heroes in person.” But I think the part that I really liked was it was more focused on groups and discussions rather than everyone coming in and just like giving what they're working on. So that was great.

And finally, the last part of the writing. Because we were working towards starting off the report on the last day. I think that was really useful. Again, I haven't done workshops where we were going to write a report on the last day to start it off. So because we had…there was an end goal saying on the last day, we need to have a set of ideas and something on paper that we could translate down to a beginning of a report.

As someone who is writing his dissertation, I know it's very easy to put it off like, oh, I'll write it later, I'll work on something else now. And it's even harder when there's like 40 people going in 40 different directions, getting everyone to talk about what to write and how to write it, it's hard. And it was good that we started off here because I think what we have now, the skeleton of a report is in pretty good shape to write, kind of like mold it into something very useful going forward.

Khari: Yeah, there's definitely a lot of content that came out of this. So it should be relatively easy, I guess in the grand scheme of things, for the organizing committee to kind of turn that draft into final.

Natesh: I mean I'm glad to hear that. So, I completely enjoyed my experience. I think this was a great workshop. I like the format, and I was telling Tom and Todd about maybe these are the kind of formats we need to have in other workshops going forward, because I felt like I was exhausted at the end of every day, but it was a good kind of exhaustion. I was happy that I really was forced to rethink everything through the day, and then let's start all over again and go through the process. That was great.

Khari: Yeah. CCC workshops, for anyone interested in attending one or proposing one, definitely can be demanding because really working all day, but also very heavy on breakouts. Not a lot of time where people just talk at you. So there will be information linked somewhere on the page where this is posted if you want to learn more. Any final thoughts, Natesh?

Natesh: I realize that there's a lot of work to be done moving forward, but you know, if you're a researcher you want to be at the beginning of these things…like you want to have a lot of work. I'm finishing up grad school, so I'm going to balance just focusing on my own dissertation while also like working on these reports. But I'm very excited.

I think going forward, we're going to see more kinds of these workshops, hopefully, and I think we have sparked interest in different people here to think about these ideas and be willing to introduce them to a wider community who are doing the same. So, yeah, let's see. Maybe I'll be back here in five years or something, like thermodynamic computing is great. [Laughter]

Khari: Yeah, all right. Well, thank you for being here. Go enjoy Honolulu. [Laughter]

Natesh: Yeah, finally.

Transition to interview with Gavin Crooks [20:32]

Khari: You're listening to the CCC podcast here with Gavin Crooks right after the Thermodynamic Computing workshop. Gavin, how you doing today?

Gavin: I'm great. Little tired.

Khari: Yep. Thanks for being here. So you tell me a little bit about your background? What interested you in coming to this workshop?

Gavin: I've worked on non-equilibrium thermodynamics for a long while, trying to understand the fundamental limits to thermodynamics, particularly when it comes to microscopic systems and information processing. So this workshop is very relevant to those issues.

Khari: OK, and you're a chemist by training, is that correct?

Gavin: My PhD is in chemistry, yes, although I tend to sort of cut across various areas.

Khari: How did you get more into computing?

Gavin: Uhh…into computing? I guess computers were always my original passion. Chemistry was kind of a side project.

Khari: This sort of a return to home.

Gavin: Yeah, computers are fun! [Laughter]

Khari: All right. So, can you kind of summarize what the breakout groups that you participated in at the workshop talked about? Any conclusions that came out of those?

Gavin: That we…you know, we discussed what it really means to have a thermodynamic computer. Part of the process is trying to define that more clearly, sort of two different ways of really thinking about that. One is a computer that's operating at the limits of thermodynamics and another is to think about a computer that is in some sense using thermodynamic principles to run an algorithm; or to…using thermodynamic principles to construct the computer rather than having to design it in some sense.

Khari: Do you think either of those areas is more likely or more compelling, at least personally?

Gavin: Likely. It depends on the timescale you're talking about. I think trying to build computers operating close to the fundamental limits in their processing is something that we can do now, or beginning to be able to do now, on a prove of principle level. How far away we are from being able to do that on a practical level is unknown, and that's why we, we're getting together and talking through these issues to try to see where the technology might be leading us. If we can sort of get ahead of that curve and develop the fundamentals that we're going to need to understand computers operating at that level.

In principle, we can have computers that use tiny fraction of the energy to do the same computation compared to computers today, but it's also very possible that these hypothetical computers, they're not going to be, necessarily, sort of general purpose machines in the same way that the current laptop is. It will be a specialized device, and part of the issue facing is what are the sort of killer apps. What would actually use these computers to…what problems would we actually attack with it?

The other side of it is sort of the more self assembly…can we instead of trying to have to design a thermodynamic computer…conceptually, a computer at this scale could have a huge number of components in the same way we have huge number of neurons in our head. And the neurons in our head are non pre-programmed exactly which neuron connects to which neuron. The overall architecture of the brain is obviously programmed at some level that all genes, but the individual connections are not predetermined. So you know, can you build computers using that more sort of self-assembling kind of paradigm?

And you know, that principle is a radical thing. In practice, it is really hard to see how to do that.

Khari: Right. So if you had to project like your wildest dream idea for a killer app using a thermodynamic computer in 30, 40 years, what would it be?

Gavin: It would be a laptop-sized device that could perform exascale calculations. Exascale is right at the limits of what current generation supercomputers can do.

Instead of requiring a warehouse of computer hardware and an entire power plant to power the thing, you'd be able to do the same kind of calculations on a laptop-sized device using a few tens of watts, about the energy expenditure of our own brain. And you might say,” well, why on earth would I need to do that much calculation on my laptop for?” The answer is I don't know. This has always been the case with computers that it's been hard to predict what, given very large computational resources, what people actually put those resources to.

So I don't know what we'll do exascale calculations on our laptops for, but I'm sure there will be amazing applications if you could do that. That kind of energy budget, a few tens of watts for an exascale calculation is sort of the reasonable limit that we can hope to achieve if all the physics and device design actually comes together.

Khari: Yeah, I guess that's a powerful thing. It's not too dissimilar to the fact that most cell phones are as powerful as a huge computer that would have taken up a room in the 40s or 50s.

Gavin: Oh yeah. I mean there’s always the joke about the IBM guy saying, “well, 50 computers that should be our worldwide market. What would anybody need more computers than that for?” Or Bill Gates allegedly saying “you need 60040 K”. Yeah, we just keep finding new interesting things to do with all that computer resources.

Khari: So, I guess, do you have any research that you want to plug outside of what we've talked about here today?

Gavin: Research to plug?

Khari: Or just interesting projects, you want to let people know about.

Gavin: Oh, well, I mean my current research, apart from thinking about the thermodynamics of computation, also thinking a lot these days about quantum computers. Quantum computers, it's the same inspiration that we're running up against the limits of what the conventional computer technology can do, and in terms of how fast you can go and how much energy it needs. Quantum computers is novel way around that limitation by taking advantage, again, of physics and trying to just find a whole new paradigm of how we do the calculations.

And it all connects together. So there’s aspects of thermodynamics that crop up in quantum computers also.

Khari: Did you meet any people at this workshop that you hadn't met before that you might establish a potential collaboration with in the future?

Gavin: Collaboration, I don't know.

Khari: I know it's been one hour since the workshop ended.

Gavin: It has been one hour. [Laughter] but no, no, no, I mean, one of the nice things about workshops and conferences and what have you is you actually meet people who you may know professionally from their papers…you don't know necessarily personally.

So I have met Seth Lloyd, for example. I know his work very well. He knows my work. I met him once a long time ago when I was a grad student, which he won't remember. So, no, it's good to actually meet him in person. There's a lot of overlap in the things we're thinking about.

Khari: Well, that's all the questions I have. So thank you for being here and I hope you enjoyed the workshop.

Gavin: Okay. Thanks.

Outro [30:45]

Khari: That's it for this episode of Catalyzing Computing. To learn more about the Thermodynamic Computing workshop, visit the workshop web site under the Visioning Activities tab at cra.org/ccc. I'll be back soon with more interviews from members of the computing community, including University of Maryland iSchool Dean, Keith Marzullo and more visioning workshop recaps. If you've enjoyed the podcast like, subscribe, and rate us five stars. Until next time…peace.