Scott Hanselman: This is ACM ByteCast, a podcast series from the Association for Computing Machinery, the world's largest education and scientific computing society. We talk to researchers, practitioners, and innovators who are at the intersection of computing research and practice. They share their experiences, the lessons they've learned, and their own visions for the future of computing. I'm your host today, Scott Hanselman.

Hi, I'm Scott Hanselman. This is another episode of Hanselminutes. Also, we're combining this with the ACM ByteCast. So you may be hearing this as a part of the ACM ByteCast podcast series or as a part of Hanselminutes, depending on where you're listening to it. I encourage you to listen to both podcasts because they're wonderful. We are talking today with Dr. Bob Metcalfe, graduated from MIT in 1969, and is the co-inventor of ethernet and founded 3Com. How are you, sir?

Bob Metcalfe: Great, how are you?

Scott Hanselman: I am very thrilled to be chatting with you, because we wouldn't be chatting with you, had you not done the things that you'd done. I don't know how we would be talking over the internet. Maybe we'd be using token ring at this point.

Bob Metcalfe: No, I don't. Yeah, but you're right. Were it not for ethernet, I wouldn't be here today.

Scott Hanselman: So I'm curious, we could do what people typically do at these points on the podcast, and we could start at the beginning, but I kind of like to work backwards. What are you working on a day-to-day basis right now? What are you doing now?

Bob Metcalfe: Well, it's complicated. I've taken a job to start my new career. My new career is computational engineer and I'm affiliated with MIT, so I'm a research affiliate at my alma mater, MIT, and my first engagement as a computational engineer is to model geothermal wells. I'm interested in how we can use geothermal energy instead of fossil fuels, and the reason it's complicated is that I've just barely started and I have a long way to go. Doing a new career is not easy. For example, here I am, I'm modeling geothermal wells, and you got to know thermodynamics to do that, and I don't know thermodynamics. So I have this big digression to go learn thermo.

Scott Hanselman: I'm at a point in my career where I think I'm thinking about retiring and I was on vacation recently for a week and I felt like that must be practice retirement. And then I went back to work and I realized that it's not about retirement, which is doing nothing. It's about retiring and then doing what you want to do. What made you want to do this? I mean, you don't have to work, but you thought geothermal was exciting because it could change the world and help us with our power issues.

Bob Metcalfe: I've tried retiring before, several times, and I see you'd tried it for a week and that's interesting because after a week you still think it's cool. It's two or three weeks where it begins to get boring, so you better extend your tests a little or you'll get the wrong impression. So having attempted a retirement before, I realize I like being on the steep part of the learning curve, it's just more fun. And what I like about the geothermal problem is that it's a very big thing, if successful, it will. You don't hear about geothermal that much, and in fact, it accounts for 0.4% of our energy, but there's a billion years of energy sitting several kilometers down and we just need to know how to get it out of the earth economically and life will change dramatically.

Scott Hanselman: I'm heading to South Africa for a month and they're having, of course, a big power issue there with what's called load shedding where they turn off the power grid on a neighborhood by neighborhood basis, and I'm taking some solar panels with me, pocket solar panels, and I've been testing these just here in Oregon, and I know it's obvious. I know that there are those who are excited about solar and alternative energies, but the idea that I took a big 40,000 milliamp hour battery and I stuck it in my front porch and I put out a little suitcase version of a solar panel and I came back later and I had free power and I'm like, now I can't stop walking around looking at all the wasted power that's hitting the earth from above. Why have we not blanketed all of Utah, all of Arizona, all of these places with solar panels, of course the minerals and whatnot, but now when you think about geothermal, and I did some research into this, you're right, it's just we're sitting on this giant battery of energy and it's just churning around in there, but we have to dig deep enough to get to it.

When you learned about geothermal, did that just kind of overwhelm you with like this potential is there and we have to get to it?

Bob Metcalfe: Well, I was dragged in kicking and screaming by a geothermal enthusiast named Jamie Beard, and she captivated me into the topic and the nature of the problem is interesting. And so you say the energy is right there, it's right beneath our feet. The trouble is the drilling's too expensive. We're all working on getting drilling to be cheap, and then suddenly William will unleash this enormous CO2-free, low-cost energy source.

Scott Hanselman: Jamie Beard, of course, is the founder and executive director of Project Innerspace, which is a nonprofit, focused on geothermal, and she's been kind of working with the geothermal entrepreneurship organization. Why? Is it just because it's too hard to use energy to dig deep? I mean, we've got several miles down, haven't we before? Is it just beneath our ability to drill or drill safely?

Bob Metcalfe: I have to get realistic. A mistake that a lot of energy hippies make is they think the energy just has to be clean. There's a lot of them in Portland and, "Oh, clean energy is all we need." Well, no, no, it has to be cheap. You won't get massive adoption until it's cheap, so it has to be cheap and clean, not merely clean. And so there's a thing called the Levelized Cost of Energy, the LCOE, and we're all

fighting to get it. It varies. It's 10 cents per kilowatt-hour or 20 cents, and we'd like to get it to one cent. My goal is one cent per kilowatt-hour, and that's viewed as very unrealistic, but I think when we achieve the drilling breakthrough, one cent LCOE is going to surprise us. It's going to come suddenly.

Scott Hanselman: Really, there's a lot of stuff that's been coming suddenly lately. Unprecedented stuff. It feels like we could use a win right now, and that would be a pretty wonderful win.

Bob Metcalfe: Yes, it would. Well, the connectivity comes suddenly, and the internet, for example, is 54 years old and signed up two thirds of the human race, and that's five billion people. You've heard that expression. We tend to overestimate what we can do in the short-term and underestimate. We're entering the long-term of a lot of things, and they seem suddenly surprising. AI is now surprising. The internet is now surprising. We have so much connectivity now, we don't know what to do with it all, and there's more of it coming. The next generation internet is on the verge. It's happening now, and so a lot of the pathologies of connectivity are attributable to the suddenness of how quickly we became connected. Suddenly we have all this bandwidth, but we don't know how to manage it yet. So we get these pathologies like hacking and pornography and advertising used to be viewed as a pathology of the internet until we figured out that advertising was going to pay for everything, and then spam was left and we don't talk much about spam anymore, and now we're onto new things like fake news and censorship and the modern pathologies of connectivity.

Scott Hanselman: The interesting thing about pathology, my wife is a nurse and she also had cancer, is that we get it not to zero, but we get it to an acceptable level. Spam exists, but it's largely in our spam folder and we now have it to a level that is manageable. I wonder if that's the thing. You introduce a new item, pathology increases, and then you push it down to an acceptable level and then it's up to society to decide what is acceptable, whether it be inappropriate behavior on social media or whether it be spam or whether it be fake news. We get it down to some percentage. It's like, "All right, we as a society agree that the cost of entry to the internet is this many banner ads or this many privacy."

Bob Metcalfe: In the 1970s, I worked for Citibank on one of maybe the first cash card and at gas stations and things like that, and I got to look at the business plan for my subsidiary of Citibank, and there was a line item for fraud, and that offended me because I was an engineer. It was my job to eliminate, make sure fraud was impossible. But the business people realize what you just said, which is you work down these pathologies to an acceptable level, then you move on. And so they had a line item for fraud. It was a single digit percent, and we tried to make it go down and down and down, but fraud could persisted.

Scott Hanselman: Is that a design philosophy? Because I'm call back to your idea that you want to have a Levelized Cost of Energy of 0.1, that's certainly aggressive, but your goals

isn't the goal to have an aggressive goal, an unwavering goal that then pushes you further than it would other people?

Bob Metcalfe: Yeah. What do we call it in Silicon Valley? I think it's called a stretch goal and it's good to have those.

Scott Hanselman: Well, when you worked on with David Boggs on the ethernet paper from 1976, was it true that you could move the bits around faster than the computers could handle them? So you had more bandwidth than they couldn't be sent or received, you had a bigger pipe than you needed initially?

Bob Metcalfe: Well, right after building the coaxial cable ethernet at 2.94 megabits per second, I went off with this other guy and we developed a fiber optic interconnect at 100 megabits per second, and it was completely useless because there was no computer that could take that bandwidth economically and put it into memory. Remember those memory was expensive then. And we were using the first DRAMs, the 1103s from Intel, and they were just too expensive to really load up. So memory bandwidth was a constraint in the design of these connectivity systems.

Scott Hanselman: And now it boggles my mind that my kids, my kids are 17 now, can take for granted our gigabit network at home here, and I actually have dual gigabit. I have both Ziply Fiber and Comcast running through a router and then 10 megabit throughout the house, which is ridiculous, because all we're doing is watching Netflix. It's so interesting that the Comcast of the world try to sell gigabit ethernet, gigabit speeds rather to people who are using maybe 20, maybe 30 megabits per second. That seems to be a bit of a scam. I wonder who needs that kind of bandwidth. Even myself, I like having it, but I don't really use it unless I'm moving a giant file around.

Bob Metcalfe: Calling it a scam as little over the top. I read the other day that 82% of the traffic on the internet is video, 82%, and the internet was never designed to carry video, but it managed to might evolve. 80% of it is just Netflix probably.

So Netflix is a kind of video. It's one way. But the big powerful new video is Zoom, is two-way, and then we've suddenly been introduced to two-way video. And so I expect that's going to be skyrocketing. And we were suddenly introduced to it by COVID. I remember there were all these professors, I was a professor, many of us were denying that we would ever teach a course on video because prevented us from reaching out to our students and developing the proper rapport. And then suddenly COVID came along and every single course at the University of Texas was taught over Zoom, so much for pressing the flesh.

Scott Hanselman: What did you teach classes over video? What were you doing during COVID? What were you up to?

Bob Metcalfe: I was a professor at the University of Texas.

Scott Hanselman: And you were doing Zoom classes?

Bob Metcalfe: Yeah, my last class was 18 freshmen over Zoom on startups. It was in the second semester of the year of their freshman year. They were all freshmen. And then it became obvious that not one of them had been to Austin, Texas. They were freshmen remotely. And to me, they were just little boxes on the screen. There was one little Indian guy. He kept showing up in the lower right-hand corner, one of my best students, and I was dying to meet him, this little guy over here in the right-hand corner, he was six foot four. When I finally got to meet him, a different presentation.

Scott Hanselman: I work at Microsoft in my day job, and I've been here now 15 years, but I have only ever worked from Portland, Oregon. So I was remote before it was fancy and I optimized my house for high bandwidth. I've always optimized my camera for high quality, and I'm always surprised even after all these years when I drive to Seattle, how tall these little boxes are when I meet them in person. They're not shaped at all like they are when I work with them all day on Zoom and Teams.

Bob Metcalfe: Well, an interesting but little-known fact is that I taught over video as a consulting professor at Stanford in the '70s. So we had the Stanford Television Network and local companies in Silicon Valley would subscribe and they would show up. So half my students were in the room and the other half were in the valley sitting in rooms watching TV, Stanford to do things ahead of everybody else. And there's a seventies video network.

Scott Hanselman: I think that as someone who really enjoys computing history, and I find myself when I am teaching, going back to history and reminding people that the thing that they think is new and fancy is not as new and fancy as they thought about. I assume you're familiar with Douglas Engelbart's, The Mother of All Demos from 1968.

Bob Metcalfe: Of course.

Scott Hanselman: And this is just an important thing for people to know about where when someone says, "Oh my goodness, Google Docs." And being able to type in a Microsoft Word document and see the cursor of the other person was being thought about by folks at the ACM and by folks at research centers back then, what was limiting us from doing all of those things? Was it memory, like you say, was it connectivity? Why did it take so long for it to then hockey stick?

Bob Metcalfe: Well, Engelbart's NLS, as I recall, and I was a user of it, ran on a PDP-10, and a PDP-10 was a 36-bit word and it running about a megahertz. So there's a limitation on Doug's dream. The reason Park exploded is we took a lot of Engelbart's ideas and moved them off the PDP-10 onto personal computers where they blossomed. By the way, is it that important about what happened first or who did it back then? Why is that so important?

Scott Hanselman: That's a good question. I don't think the who did it first or the light bulb was invented by two different people on different sides of the worlds. What I find to be interesting is that someone had the idea, it was a great idea, but there was a constraint, whether it be an economic constraint, a social constraint, and a constraint of the laws of physics that prevented it from happening like that. And then looking back with the benefit of hindsight, I find it overwhelming to look back and say, there was a time I remember very clearly when there were no iPads, and then I turn around and then there's an iPad, and now I have a pocket supercomputer that I can talk to and ask, "How tall is Brad Pitt?" Because that's basically all I use my phone for now is when I'm watching a show to figure who's that guy.

Bob Metcalfe: You have trouble remembering how tall Brad is?

Scott Hanselman: Well, so I was reading an article in Ars Technica a couple of days ago, where an individual bought the very last Encyclopedia Britannica from World Book, because World Book is the only place that's making encyclopedias, and they wanted to feel the 127 volumes of this giant book. But there's also a sense of maybe who are we if we're not offline? And I think about that, what does it mean when my kids maybe don't know who they are disconnected? So I try to take them out into the woods and out into a disconnected environment so that they might remember who they are outside the context of computers. That is overwhelming. And I only have 30 years of historical context. I can only assume it's either overwhelming for you or you just choose to not think about it.

Bob Metcalfe: Forgive the slight digression, but you just mentioned that World Book owned Britannica.

Scott Hanselman: I think I mixed my two things there.

Bob Metcalfe: But they used to be Arch competitors. Did one of them buy the other one?

Scott Hanselman: I think I was using it the way we use Kleenex to refer to tissue. So it is in fact World Book that makes the boxes. I don't think Britannica makes physical books anymore, but World Book does.

Bob Metcalfe: Okay, so I have three or four of those old encyclopedias, including Britannica prior to 1957. So there's no mention of Sputnik in it, and the Vietnam War had not occurred yet, and so on. But going back to the other question, it comes up a lot is who did what? So when you try to tell the story of ethernet, there's a guy from Boston, Texas, no, at other city, San Antonio, Texas, who was at Data Point and they introduced Archnet, 1970 something, and that was before ethernet and better than ethernet. And it was just a horrible twist. It was the old Betamax thing happened and ethernet won and Archnet. So that's the people giving importance to what happened first and who spread it for it. And so apparently that is important to people. It may be a source of learning that is maybe if you grovel over the events that happened and what came first and

what came second. You can learn something about innovation that would be a reason to do it.

Scott Hanselman: ACM ByteCast is available on Apple Podcast, Google Podcast, Podbean, Spotify, Stitcher, and TuneIn. If you're enjoying this episode, please do subscribe and leave us a review on your favorite platform.

So that's a good attitude, that's a very positive, healthy attitude that you have. What does it feel like when you got a touring award? But I noticed in the interviews you acknowledged immediately like you're receiving this award, but there's so many people that have been involved in the evolution and the growth of this, but then they choose to say, "Here's Bob's award." And then you said, "Well, there's hundreds of people involved in the ever-growing standard of ethernet." And I thought that was very magnanimous of you to do that.

Bob Metcalfe: Well, it's a feature of ethernet. And its design is that it's part of a, you mentioned the open source community is there's a different community of which ethernet is a part and a feature of the ethernet brand is its susceptibility to new technologies. It's like built in backward. Compatibility and interoperability are two big ethernet values, and they allow ethernet to persist. So I'm asked, what's going to obsolete ethernet? And the answer is tricky. Every time a new technology comes along that threatens to obsolete ethernet, they have a meeting and decide to call it ethernet. It will persist forever, but it'll just be changing underneath.

Scott Hanselman: That's a great point. I think about it in the context of what's going to be the thing that replaces books. eBooks didn't replace books because there's still books, and the IEEE has over 1000 standards and they're building on top of the land man group of standards. This is ethernet and that is the legacy is the building on top of the standard rather than trying to make yet another standard, and still things work. I'm sure that you and I are talking over a million different switches on a million different brands between my network and your network and all the things in between. And it all just works because we all agreed that it should just work.

Bob Metcalfe: And not just one ethernet, but a dozen of them. Your traffic is reaching me and mine and is reaching you

Scott Hanselman: Even from the 20-foot coax that goes from my switch into the garage, then quickly switches into fiber and then into the ether, presumably, and off into where you are on the other side of the country. I am impressed that I still have a sense of awe over this, that every once in a while on a random Tuesday, I'll be on a video call and you're very clear. You're in high definition. I go, "Oh my God. I'm talking the speed of light to the gentleman over here in the other side of the country." How amazing is that? And my kids are utterly unimpressed. I wonder if am I supposed to also be unimpressed or is it okay to have that sense of awe? Do you have any sense of awe?

Bob Metcalfe: Kids are a little older than yours. They're in their 30s. Yeah. In fact, my son had time, my daughter had time to work for eight years for Facebook.

Scott Hanselman: Oh, wow.

Bob Metcalfe: Son, coincidentally, got to work for eight years for Apple computers. So both of them are so immersed in what you and I find to be awesome. They don't find it so awesome. It's just like air. It's just everywhere for them.

Scott Hanselman: That's so wonderful. Yeah, there was a young person at work who was having an issue with their GIT repository, and it was a carriage return issue, and I think they were in their 20. Yeah, exactly. So you just facepalmed. And I said, "Oh, well, that's a carriage return line feed." You know about that? And I said, "There's a very funny history about that." And they said, "What's a carriage and where is it returning to?" I was like-

Bob Metcalfe: [inaudible 00:23:15] in a line. What is that?

Scott Hanselman: So I didn't want to turn into professor mode. So I always ask permission before turning into professor mode, and I said, "Would you like to know?" And they said, "All right. All right. I got 20 minutes. Let's hear it." And we talked about asking, we talked about the carriage on a typewriter and how it returned, and we went on and on. It's 2023 and you're on a Mac or a PC and you're upset about a carriage return. And it's because of decisions that were made in the 30s and 40s about how to encode text, and what a wonderful thing that was. And I just loved that kind of history. And they just really didn't care. They just tolerated the whole lecture.

Bob Metcalfe: Since you raised character return line feed. So 1970, we were developing a internet protocol called Telnet, which allowed to login to a remote host. And some of the hosts had the convention that they only use carriage return. They didn't use line feed, and others had the convention. They insisted on carriage return and line feed. And so the displays were all screwed up as you went from the 10 x PDP-10s or the Multics, which was a typewriter that only had line fed. I forget how we resolved all that. I think maybe we just switched to ASCII and accepted line feed.

Scott Hanselman: Well, I remember doing when I was in networking class in college 30 years ago, having conversations about different networking cards using big-endian and little-endian, and whether or not it would be the CPU that would do the flipping or whether that was now built into the hardware, should the network card be responsible for worrying about big-endian and little-endian, does ethernet care versus the computer cared, and thinking about those bits flipping on the way off of the network and into the CPU and back. And it's in my brain and I know that stuff, but I don't think anyone's thought about that in many, many, many years.

Bob Metcalfe: Many years. Big ending. And I still don't understand the difference.

Well, I built a piece of hardware for the ARPANET, and it was the seventh one built a very low number so no one really knew what to do. So I made mine hugely functional so I could had a mode, big-endian and a little-endian, and eight bit 16 bit, 36 was a very popular word size then. So I had a register where you could specify the word size and you was full duplex. It could send and receive at the same time, and the card was three feet square to do all that functionality. And then so part of the progress in making them cheaper was eliminating all that generality and finally settling on bytes, instead of 36 bit words is as the feature. Oh, well.

Scott Hanselman: Yeah. It's funny, even now I've see an article from IBM, which is endianness guidance for open source projects, where they have a whole conversation about thinking about Indianness because you're coding usually for little-endian processors. That's basically the way things are working now. But people still have big-endian support guidance. So every once in a while you'll hit a bug in something, like a couple of weeks ago, I was dealing with some home networking stuff and I had to think about jumbo frames, and I hadn't thought about jumbo frames in many years. It's always an abstraction until it leaks, isn't it?

Bob Metcalfe: Yes, jumbo. We still have jumbos.

Scott Hanselman: I guess most are-

Bob Metcalfe: That's the wrong way to say it. The jumbos are the new thing. And-

Scott Hanselman: Well, jumbo frames could do 9,000 bytes apparently in a frame.

Bob Metcalfe: Yeah, I don't remember the numbers. But I do remember setting the maximum packet size.

Scott Hanselman: Welcome to your new computer and max out the packet size, because you need to do that, because the defaults are wrong.

Bob Metcalfe: You know what the consideration was?

Scott Hanselman: What's that?

Bob Metcalfe: I don't know. It may surprise you.

Scott Hanselman: Okay.

Bob Metcalfe: The Alto personal computer had a disk and we imagined that disk, you'd want to put data on it and you'd want it to go really fast as fast as you could. And that would be the dominant method of transfer would be just to disk over the network. Well, there were sectors. The disk had sectors, and so the maximum packet size was set to be large enough to contain a sector plus the metadata,

the overhead associated with that sector. And that's how long, so we made the packet size as small as possible, but still capable of carrying a sector at a time.

Scott Hanselman: And that Alto had a two and a half megabyte, one platter disk cartridge.

Bob Metcalfe: That you put in the pizza oven.

Scott Hanselman: That's so fun. I know I've mentioned this a little bit before, but I am finding as I am now in getting into my 50s that knowing this stuff is overwhelming and I wonder if there's what to do with understanding the full stack or whether I just let it go and accept that these abstractions don't matter anymore. At least for people who aren't hardware designers. I'm sure that for networking engineers, they always will matter. But as a user, I keep wanting to say, when I learn about ChatGPT or AI, I want to walk the stack all the way down and I want to understand what's happening.

Bob Metcalfe: So how far down in the stack do you go in your automobile?

Scott Hanselman: I love that you did that because I use that analogy all the time because I'm making my kids drive stick shift, but they want to just take an Uber. That's even higher level than an automatic. They just want their body to be teleported. And Uber is the closest thing we have to teleportation, but I want them to build a kit car. Let's make a car from scratch in the garage. But then you start getting ridiculous like, well, you have to smell your own iron. Are we going to become blacksmiths? How far down do you go?

Bob Metcalfe: And where are you going to get the oil from?

Scott Hanselman: I'm going to have to drill into the geothermal at this point.

Bob Metcalfe: But then you'll need a hard metal to go into the granite, so you'll probably need to do some metallurgy.

Scott Hanselman: And now I'm at the Renaissance Fair trying to figure out why it's so hard to make a internal combustion engine.

Bob Metcalfe: Yes. So anyway, you need to apply that same thinking to the computer stack, but apparently you and I are stuck on that. There's something special about that stack and we want to know about that stack all the way. And then there are these programmers, they're called full stack programmers. Is that what it's called? Full-

Scott Hanselman: Yeah, full stack engineers.

Bob Metcalfe: My son's a full stack. He does both. Yeah, he does JavaScript on the server, but also on the client. And that makes him full stack engineer.

Scott Hanselman: JavaScript on the client and the server.

Bob Metcalfe: I've got to ask him about big-endian and little-endians, and what-

Scott Hanselman: Should do that because I'm over here-

Bob Metcalfe: [inaudible 00:30:16] that, I mean, it's possible to live a rich full life without running into that concept.

Scott Hanselman: And that right there, I love that you said that, it is possible to live a rich full life and never know that stuff, because my dad made me learn how to change my own oil and change my own tire, and then he gave me a AAA card, and AAA has never failed me and they've always changed my tires. But there was a joke that a friend of mine told me that he says, "There's not a lot of software engineers on the Walking Dead. So when the zombie apocalypse comes and the zombies eat us all, it's always the preppers that have the ability to chop wood and build stuff and blacksmiths, and there's no one there to reboot the router because that skill doesn't matter anymore." So I think my dad tried to prep me for the zombies, teaching me to chop wood and carry water when I am just obsessed with the computer going all the way down to the bottom.

Bob Metcalfe: I bet you, Portland has that in common with Maine, which is there's a lot of people up here who do five different jobs.

Scott Hanselman: That is very true. That is very true. There's even digital preppers, I learned out about this recently, where you can get a raspberry pie and a couple of terabyte disk and go do what's called the internet in a box. So you download the Wikipedia, Stack Overflow, you download several terabytes, you put it on a disk, then you fire up the raspberry pie and you have a little intranet that presumably you would run when the zombies eat us all. And you'd have a little local network and you could look up, you could Google for stuff except on your little raspberry pie, and that'd be a way to digitally prep.

Bob Metcalfe: Okay. I better get working.

Scott Hanselman: You should get on that. So as we get towards the end of our chat here, and thank you for rambling with me and playing. I do appreciate it. Do you feel that you're going to see the breakthrough that you want in geothermal in your lifetime, in the next 30 years?

Bob Metcalfe: Oh, yes. Jamie Beard and I worked together. We founded GEO at the University of Texas and were funded by the Department of Energy to do startups that do geothermal. And our goal was this decade, commercial scale, this decade, and that implied something that was complicated. In order to get scale, Jamie figured, and I came to agree, we're going to need the oil and gas industry to help. They have scale. The trouble is the hippies running geothermal hate the oil and gas industry, and that's counterproductive. So one of the goals of GEO, as

supported by the Department of Energy was to encourage startups in geothermal and then encourage the partnership between those startups and the incumbents, the big oil companies, and get them all working together. And then going back to your question, and we had a hope of achieving scale in this decade, not 2050 or one of those silly years that they keep naming. No, in the 20s. In the 20s, we should-

Scott Hanselman: In the 20s. And you think this is going to be an overnight success in 40 years where it's not in the zeitgeist right now. My parents aren't thinking about geothermal, but then the stories will start happening and it's going to pop like that. We've seen AI pop.

Bob Metcalfe: Well, we saw the internet pop in 1995, when the worldwide web came public through Netscape. I guess that was '93 or four. It popped. Suddenly there was a reason to have it all, and it popped. So I've been around for a few personal computers popped at one point, the internet popped at one point, and AI has popped four or five times. Marvin Minsky was my undergraduate thesis advisor. I knew about AI back in the '60s, and it popped and then it popped again, and it's popping now, and it might break through this time. My theory of the past pops was that the AI ran out of data, but now we have the internet, so the AI really doesn't have to run out of data. Maybe there's a lot enough data to feed its growth. Incidentally, here's a fact. Microsoft, your employer, has placed an order at Nvidia to buy every processor they make.

Scott Hanselman: Yeah, it's insatiable. They want all the GPUs.

Bob Metcalfe: And they want it for ChatGPT. And so AI is cycle starved. I also claim it's connectivity starved. We have 10 to the 11th neurons with on average 10 to the fourth synapsis per neuron. So that's 10 to the 15th, whereas GPT four has a trillion. It sounds big, a trillion parameters, a trillion synapses, but a trillions only 10 to the 12th. So there's a 10 to the three difference between the connectivity of the human brain and of the CPTs. So AI is connectivity starved, it's compute starved, it's probably memory starved too. I don't know. Or that's-

Scott Hanselman: The InfiniBand networking standard that they use to talk between all of the things when they're training these models. You just can't be too fast. They need more, and then more and then more. They need more GPUs and they need more speed. But InfiniBand is 400 gigabits or something like that.

Bob Metcalfe: Well, Ethernet's 800 now.

Scott Hanselman: Yeah.

Bob Metcalfe: I met with Jensen Huang, who's an awesome guy, and I got him to agree that the power of his products is not in the Gs, it's in the PUs. No, not in the PUs, it's in the Gs.

Scott Hanselman: Yeah.

Bob Metcalfe: Processing units are beside the point. It's the connectivity among them that has made the difference.

Scott Hanselman: Indeed.

Bob Metcalfe: And we got to the subject of InfiniBand, and he was a little embarrassed to be talking to me because ethernet has been a constant enemy of InfiniBand, and he seems to think that his own company, Nvidia, has chosen InfiniBand over ethernet, and I think he's about to get surprised. [inaudible 00:36:32].

Scott Hanselman: Well, thank you so much for spending as much time with me as you did, sir. I very much enjoyed this conversation and your time.

Bob Metcalfe: Oh, a great honor. Thanks for inviting me.

Scott Hanselman: We've been chatting with Dr. Bob Metcalfe, the recipient of the ACM Touring Award, recognized for the invention, standardization, and commercialization of ethernet. He's now working on making geothermal happen, and I'm looking forward to hearing about that soon. This has been another episode of Hanselminutes in partnership with the ACM ByteCast, and we'll see you again next week.

ACM ByteCast is a production of the Association for Computing Machinery's Practitioner Board. To learn more about ACM and its activities, visit acm.org. For more information about this and other episodes, please do visit our website at learning.acm.org/bytecast. That's B-Y-T-E-C-A-S-T. Learning.acm.org/bytecast.