Jessica Bell:
This is *ACM ByteCast*, a podcast series from the Association for Computing Machinery, the world's largest educational and scientific computing society. We talk to researchers, practitioners, and innovators, who are all 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, Jessica Bell.

Jessica Bell:
Welcome everybody, to another episode of the *ACM ByteCast*. Today, we have a really exciting guest. Vint Cerf is here to tell us about some of the history of the internet, his amazing career, and what he's thinking about for the future. So Vint, will you please introduce yourself to our audience?

Vint Cerf:
Thanks so much, Jessica. It's a real pleasure. My name is Vint Cerf. I'm the Vice President and Chief Internet Evangelist at Google since 2005, but my career goes pretty far back into the late 50s and early 60s.

Jessica Bell:
Awesome, great. So, yes, right off on that point, let's start at the beginning. I'd love to have you contextualize for our audience, especially for our younger members of the audience, what it was like to be in the computing field at the very beginning of your career, and talk about your path of how you got involved, and then how you moved through this computing world before there was all this stuff we take for granted, like all of the TCP IP stuff that you were so pivotal in, and the internet and things like that. So yeah, take us back and talk about that time.

Vint Cerf:
We need some kind of weird audio effects where we go back [crosstalk 00:01:41].

Jessica Bell:
Time machine.

Vint Cerf:
Let's go back to the late 1930s for just a second. Conrad [inaudible 00:01:48] in Germany is beginning to play around with computing based on switching systems, things that you would

have associated with the telephone network [crosstalk 00:01:58] switches and things like that, or vacuum tubes.

Vint Cerf:
When you get into the 1940s, World War II has hit, and there is a focused attention on computing, partly to do things like ballistic calculations, but most importantly, of course, code cracking. Everyone knows about Alan Turing and the cracking of the Enigma, a German encryption system [crosstalk 00:02:21] Bletchley Park. But I bring this up because John Von Neumann was the American engineer who worked with Turing and others and conceived sort of the basic structure of computing as we think of it today, a CPU, a bus, and a memory, and things move back and forth across the bus.

Vint Cerf:
That classic Von Neumann architecture emerges in the 1940s and shows up in commercial quantity with the UNIVAC machine in the early 1950s. So we're seeing two base machines and eventually, of course, the transistor gets invented in 1947 and eventually turns out to be a replacement for tubes, much more efficient, much smaller, so we started to see transistor-based machines coming out of IBM, for example.

Vint Cerf:
My first introduction to a tube based computer comes in 1958. I was all of 15 years old. My father got permission to take me to visit something called the Semi-Automated Ground Environment, which is a machine that was physically so big, it was made out of vacuum tubes, that you had to walk into the computer. It was literally [inaudible 00:03:29] building. You walked inside the computer to use it.

Vint Cerf:
The tubes are glowing red, it looked like Dr. Strangelove except Strangelove was four years in the future from the time that I was seeing this thing. [crosstalk 00:03:41] looking at 24 inch radar screens. The system, Semi-Automated Ground Environment, was taking radar information from the distant early warning radars in the northern part of Canada to detect Russian bombers coming over the pole. It was supposed to automatically detect an alert when that happened.

Vint Cerf:
At this point, I don't have access to anything. I'm just goggling. But years later, my best friend and I got permission to use computers at UCLA while we were still in high school, so we would

commute to UCLA and make use of a paper-tape based machine, a Bendix G-15 we would normally use for computer controlled milling, but we were programming it to do some interesting transcendental function calculations.

Vint Cerf:
You'd type up the program on a paper tape, feed it in, and it would run for a while and it would punch out a bunch of paper tape, and you'd put that into a flexowriter and print out whatever the answers were. We both got very excited about using those computers for that sort of thing. We're all of 16, 17 years old.

Vint Cerf:
Then I went to Stanford University as an undergraduate in the math department, but I took every computer science course that I could, and when I graduated in 65, after using Burroughs 5000s and 5500s,, which were very hated computers, I went to work for IBM as a systems engineer.

Vint Cerf:
I ran a timesharing system called Quick Tran. Now you have to understand that timesharing was invented in the early 1960s at MIT with John Murphy and others, so it was fairly new. And the fact that IBM had a commercial timesharing system running in 1965 was pretty amazing.

Vint Cerf:
Ran that for a couple of years and realized at the end of the two years that I didn't have the theoretical base that I really needed to pursue a career in computer science. [crosstalk 00:05:32] school at UCLA as a graduate student to learn, what's a compiler? How do operating system get designed? [inaudible 00:05:40] theory of computation. All of those things.

Vint Cerf:
But right in the middle of that process of working on my PhD, I got involved in a project from the US Defense Department called the ARPANET [inaudible 00:05:52] Advanced Research Project's Agency packet switch network that was exploring a way of hooking a wide range of different brands of computers together over a homogeneous packet switch net.

Vint Cerf:
At the time, this is the late 1960s, packet switching was heretical. If you were doing any kind of network switching, it was supposed to be circuit switching, which is the way the telephone system works. But that would have been really slow. We were hooking a dozen university

computer science departments together with their machines from all kinds of places, with digital equipment [inaudible 00:06:26] IBM and HP and so on. And so we, instead of having each machine dial one up when it needed to send something, which just took too long, we introduced this packet switching idea.

Vint Cerf:
This turned out to be just stunningly successful. It worked very, very well, and by 1971 or 72, networked electronic mail gets invented as one of the several applications that people [crosstalk 00:06:51] we could do remote log-in to a timesharing machine on the other side of the network. We could do file transfers. Then we could do electronic mail. And we could see emerging out of the electronic mail, the social aspect of that kind of communication.

Vint Cerf:
[crosstalk 00:07:09] the first one that I knew about was called SciFi Lovers, because [crosstalk 00:07:14] were all geeks and we're arguing over who's the best science fiction writer. The next one was Yum Yum, which was the Stanford University Restaurant Review. And so 50 years ago, we were already seeing sort of the roots of social networking emerging from email [crosstalk 00:07:33] there were a variety of things emerging in nascent form.

Vint Cerf:
That's the early 1970s. We do a public demonstration in October of 72 in Washington, DC of this [crosstalk 00:07:47]. Then I go Stanford University to join the faculty with a joint appointment between computer science and electrical engineering.

Vint Cerf:
In the beginning of spring of 1973, the guy that I had worked with on the ARPANET project, Robert Kahn, had gone from Bolt, Beranek, and Newman, which is the contractor that built the basic underlying packet switching network, Bob went to ARPA. And so he shows up in my office and he says, "ARPANET really worked well. We are thinking of using computers in command and control."

Vint Cerf:
But that means some of the computers have to be in ships at sea and in aircraft and mobile vehicles. But ARPANET was running on dedicated telephone circuits [crosstalk 00:08:31] connecting the packet switches together. So you can't tie the tanks together with wires, because they were old wires and they break and the airplanes never make it off the tarmac. So he had

already started working on a mobile packet radio system, which is what we use today, effectively.

Vint Cerf:
We carry our mobiles around, but back then, we're talking in the early 70s, this was amazing stuff. The radio was a cubic foot and cost 50,000 each. So you can imagine how much has happened over the course of the last [crosstalk 00:08:59].

Vint Cerf:
He had a packet radio system in the San Francisco Bay area and a packet satellite system over the Atlantic. Now we have a problem. How do we hook the packet radio, packet satellite, and ARPANETs together in order to make it look uniform? And that was the internet problem.

Vint Cerf:
Ironically, about a mile and a half from my laboratory at Stanford is Xerox Palo Alto Research Center, [crosstalk 00:09:28] Bob Metcalfe and David Bogs are experimenting with ethernet in the same timeline, May of '73, and that idea they got from the University of Hawaii, which had been running a program for a few years called ALOHAnet.

Vint Cerf:
It was called ALOHAnet because you just transmit whenever you want to and if there's a collision in the air, this is radio based system, if there's a collision and the data doesn't get to the central computer and you don't hear an acknowledgement, then you just retransmit. But instead of transmitting after a fixed delay, you randomize the delay so that you don't have another collision. Aloha is sort of hang loose, do whatever you want.

Vint Cerf:
Ethernet was a little more sophisticated because it could detect the collisions very quickly and then stop transmitting in order to make it more efficient.

Vint Cerf:
We had four different kinds of packet switch nets. Bob and I developed something we called TCP, Transmission Control Protocol. We then engage three different groups at Stanford University Bolt, Beranek, and Newman in Cambridge, Mass, and University College, London, in London, to build the first versions of the TCP protocols.

Vint Cerf:
We're staring going to do that implementation in 1975. We iterate through several instances of the protocol design. We split the internet protocol off from the TCP part in order to allow for real time, but unreliable communication so we can hit the radar traces, real-time voice and video, which I have to point out was part of our objective. So voice and video, that we're doing, well, we're doing voice right now, but [crosstalk 00:11:00] doing voice and video on a regular basis, we were planning for that in the 1970s. We were doing experiments with it in the early 1980s, but we just didn't have very much capacity to do it.

Jessica Bell:
I was reading an interview of yours and someone had asked you, "Did you have any idea what this would have become today?" And you were like, "Well, we designed this network to be really future proof." How do you go about breaking down a problem like that to be strong enough and flexible enough to accommodate a future that feels like it just exploded into this thing that we call the internet now. How did you break that problem down?

Vint Cerf:
Two things. First of all, we made some fundamental assumptions. The first one was [crosstalk 00:11:47] change any of the networks that were going to be part of the internet, because they had already been built. And second, we said, we don't want them to know that they're part of the internet. So we said, so they have to be interconnected with computers that interconnect the networks, we called them gateways, today we call them routers, those gateways had to know that they were part of the internet, even though the networks they were connected to didn't.

Vint Cerf:
The network addressing, the global addressing, of the internet, not known by the networks, but it was known by the gateways and the host computers that were talking to each other end-to-end. Second, end-to-end principle was important. Whatever you put into the net popped out the other side, no matter how many routers it went through or gateways it went through, just like when you throw a postcard into the post office, it may be carried in a variety of different ways, but it comes out intact at the other end, most of the time.

Vint Cerf:
[crosstalk 00:12:39] guarantee, it's a best effort system. And we said, we'll make the packet switching system the core of the system best efforts, but we won't make any guarantees. If you make guarantees, then you have to have an end-to-end process for detecting loss, retransmitting,

detecting duplicates. That's what TCP did. The IP layer and the adjacent user datagram protocol was sat on top of it, was a real time, unreliable service, didn't guarantee sequencing or anything else, but it was fast. And that's good, you want where the missile is now, you don't want to know where it was 10 minutes ago, so For real-time applications we needed that.

Vint Cerf:
But so the problems sort of dictated what the solutions looked like. Two very important principles that I think need to be understood in order to realize why this system has been so dramatically capable of scaling and of adopting and supporting new applications.

Vint Cerf:
The first one is that we didn't put a limit to the number of networks that could be connected. Although some addressing considerations that we had to deal with as the number of networks grew.

Vint Cerf:
Second, we said that the internet packets won't know technically how they're being carried. Just like the postcards don't know. That was important because since they don't know, they don't care, and when you add optical fiber, for example, which was not part of the original design, the internet protocol layer didn't care, didn't know, all it knew is that it just got dumped down into some network [crosstalk 00:14:08] the internet packets.

Vint Cerf:
The second thing, though, equally important, is that the packets don't know what they're carrying, just like a postcard doesn't know what you wrote on it. The consequence of that is that if you introduced a new application, the only place that needed to know what the bits meant that were in the packets, were at the edges of the net where the applications were, not in the core of the net.

Vint Cerf:
The net is actually application ignorant. You could have made the applications more efficient if the network knew about the details, but we didn't want that because we didn't know what the applications were going to be over time, and we didn't want the network to constrain the applications. And you can see from the origins of ARPANET at 50 kilobits a second in the backbone, to present day internet whose core backbones run at 400 gigabits a second, going [crosstalk 00:15:01] terabyte in the next year or two, the system has scaled by a factor of one to

10 millions [crosstalk 00:15:08] orders of magnitude. It's very rare to see an architecture that can do that.

Vint Cerf:
The number of applications, the number of protocols, now is in the hundreds, the arrival of the World Wide Web in December of 1991, a new layer of protocol was put on top of TCP IP [crosstalk 00:15:28] Tim Berners-Lee, and that opened up a whole new batch of potential applications. They were enhanced by Marc Andreessen and Eric Bina at the National Center for Supercomputer Applications around 1993 or so, when they said, why don't we make this a graphical interface instead of a tech interface, which is what Tim had produced.

Vint Cerf:
The graphical user interface of the Mosaic application, the Mosaic browser, was a stunning achievement because it transformed the network, it suddenly looked like a magazine with formatted text and imagery, eventually streaming audio and video. [crosstalk 00:16:08].

Vint Cerf:
Jim Clark, who was the founder of Silicon Graphics, takes one look at Mosaic and says, "Holy crap, that's a big deal." And he brings Andreessen and Bina to the West Coast and starts Netscape Communications in 1994, and in 1995 they go public. This is very unusual to go public after a year. And the stock goes through the roof. Suddenly everybody's throwing money at anything that looks like it might be part of the internet. That's the big [crosstalk 00:16:34] boom.

Vint Cerf:
Then in April of 2000, there's a big dot bust when a whole lot of those companies didn't have a business model [crosstalk 00:16:42] which they ran out of and the CEOs were scratching their heads saying, "What happened?" And the answer is, capital is finite, revenue is supposed to keep going, if you have a business model and some of them didn't.

Jessica Bell:
That doesn't feel familiar at all, does it?

Vint Cerf:
They just flat fell on their faces, but the Worldwide Web and the internet continued to expand dramatically. People kept throwing new content into the web. Not to get paid for it, but simply because they wanted to know that what they knew was useful for somebody else.

Vint Cerf:
Now we're awash in information and we can't find anything, which promotes the search engines, AltaVista originally from Digital Equipment Corporation, and then [crosstalk 00:17:25] and then Google, and Bing, and other [crosstalk 00:17:30].

Vint Cerf:
So, here we are. That's in the 1990s. And then along come the mobile phone, the smart phone. that has a long history, which we don't have time to talk through, but it started in 73 with [crosstalk 00:17:44] at Motorola. The same year, Bob Kahn and I are starting to work on the internet.

Vint Cerf:
Marty's working on mobile phones, we're working on internet, and my first son is born in 73 and he wants to know whether he's the brother of the internet. Everybody says, "Okay, so you and Bob are fathers of the internet, so who are the mothers?" Another long half-hour conversation.

Vint Cerf:
What's important though, is the milestone with the smartphone coming from Apple, Steve Jobs. The reason it's so important is that the two technologies, mobile telephony and internet, had been going in parallel for quite some time. Suddenly in the smartphone, they come together and they are mutually reinforcing. So the smart phone makes it possible for you to get access to everything on the internet wherever you are and can find a radio link. And of course the smartphone makes the internet more useful because you can get to it from anywhere you can find a link.

Vint Cerf:
The two are dramatically powerful. We see that today, as we see smartphones proliferating around the world [crosstalk 00:18:47] people experience the internet primarily through applications on the smartphone.

Jessica Bell:
I think that brings us now back to the present day and thinking about this extremely powerful network that has now been connected to us in so many different ways. Like you said, it's in our pocket, we can deal with it all the time. I'm curious to hear what you think the major challenges and problem spaces are around this network today. Do they feel very similar to the challenges

and problem spaces when you were starting out to think about this or do they feel new and different? Speak about what you think is our next big hill to climb [crosstalk 00:19:32].

Vint Cerf:
It's a very big hill. We, at the beginning, even though this was being done for the Defense Department, we were not focused heavily on security technology and cryptography [crosstalk 00:19:43] like that.

Vint Cerf:
I will say as a side observation that I was working with NSA in 1975 on a secure version of the net using classified technology to secure the packet switch system, but [crosstalk 00:19:55] commercial sector or in the university sector, I can't imagine relying on the graduate students to be disciplined about their use of cryptographic keys and other [crosstalk 00:20:07] do that.

Vint Cerf:
Security was not an afterthought, it's just that the technology of public key crypto wasn't available in the earliest periods when we were doing the design. They didn't become available until somewhat later in the mid to late 1970s and early 80s. But we retrofitted it in. That's why we have [inaudible 00:20:25] and we have TLS and we have IP sec and DNS sec, and all these other things, are retrofitable, so that's the good news.

Vint Cerf:
Security is still a big challenge. And I want to come back to that later in the conversation about why is security such a big problem.

Vint Cerf:
The second thing is information and the internet, and misinformation and disinformation, and the side effects of social networking, which have built in feedback loops, which lead to some fairly serious problems associated with people's behavior. And this is just a major problem that we are experiencing right now, because it's very hard for an algorithm to figure out that someone has spoken an untruth, for example, [crosstalk 00:21:17] represented something.

Vint Cerf:
We are now challenged by the social networking environments to figure out how do we protect users from the harmful side effects of social networking? Some of which are by accident, people spreading misinformation because they don't know any better, or worse they spread

misinformation and disinformation deliberately, for whatever is motivating them. It might be political, it might be a pecuniary.

Vint Cerf:
Of course, scientists will tell you that whatever you think is true now may not be true 10 years later when you discover your theory was wrong. We are challenged right now by the harmful hazards that show up in the net, including malware and distributed denial of service attacks and other [crosstalk 00:22:06] what am I thinking about? Identity theft and so on [crosstalk 00:22:11] bullying and so on.

Vint Cerf:
The reason that's such a big problem is that governments are looking at this saying, "Our citizens are at risk. We need to do something. Who do we blame? And who do we make responsible for fixing everything?" You see the companies that offer social networking being hammered on by members of Congress here in the US and Parliament elsewhere.

Vint Cerf:
You also are starting to see fragmentation of the internet where nation states are drawing boundaries around the network, claiming they have data sovereignty inside of their countries and they want to introduce rules that would frankly make the internet not work very well because now data transfer from country to country is no longer easily accomplished. There are rules that are in conflict as you cross from one international boundary to another.

Vint Cerf:
You start to hear a call for digital cooperation among countries in order to find a way for the countries to cooperate with each other and come to common agreements of how they will deal with abuse on the network, how they'll deal with law enforcement, how they might deal with extradition treaties and other kinds of things. How do you deal with digital evidence? How do we establish a chain of custody of digital content? How do we assure that digital evidence hasn't been tampered with? You can easily extrapolate to a wide range of problems that already exist in the physical world and have their counterparts in the online world. [crosstalk 00:23:42] struggling to figure out how to cope with those in a way that doesn't just essentially fragment the internet into a useless collection of islands.

Vint Cerf:

I would argue that we've already seen how powerful the internet can be in terms of enabling people to share information, discover information, to be educated, to do just-in-time learning like going to YouTube, saying, how do I cook Chinese eggplant?

Vint Cerf:
There is this huge upside and this very difficult downside, and that is the big arm wrestling match that we have to cope with. And I would say that also introduces one of the big problems. Why do we have so much trouble with safety and security and reliability? The answer is bugs in software. Now we are faced with the problem of teaching people who are interested in computer science, or just want to use it, that if they're going to build software that others are going to rely on, they have to take responsibility for and be accountable for [crosstalk 00:24:48] they make.

Vint Cerf:
We have to create incentives and we have to create technology that will help programmers discover stupid mistakes before they get out into the field, [crosstalk 00:24:59] won't be perfect either, so then we also have to build in mechanisms for updating software safely and securely. We have to know where the update come from. We have to know that it hasn't been altered on its journey from the source to the destination. [crosstalk 00:25:13] especially important with the internet of things, where you have boxes all over the place, full of software, able to communicate in addition to compute. And if there are bugs, we need to be able to fix them. And there are questions about how long will they be supported, what if this is a heating ventilation and air conditioning system with a lifetime of 30 years, will the IOT related aspect of it be supported by the manufacturer for that period of time? And if not, how do you get your hands on the source code to fix something after they say, we don't want to support it anymore? You turn to a third party.

Vint Cerf:
And I haven't even gotten into the intellectual property problem and the [crosstalk 00:25:54] digital content for hundreds of years, which is yet another huge area of concern.

Jessica Bell:
I'm curious, you talked briefly upon where the responsibility lies in creating these codes, is it the responsibility of the manufacturer, of the programmer, of the company, of the nation state? I'd be curious how you think this then ties back to our education of the kinds of people who are developing these technologies, writing code, dealing with these problem spaces. How do you

think these big challenges are affecting or shaping the way that we're thinking about computer science and computer science research now?

Vint Cerf:
Two ways. First of all, I think everyone should have the experience of writing a program, discovering how hard it is to write a program that doesn't have a bug, and how hard it is to find the bug to fix it. If that forces you into a certain mode of thinking, sometimes we call it computational thinking, it is critical thinking, it is break down the problem, find evidence, compare that with your theory. It's very scientific. And I think everyone should have that experience. Not necessarily because they're going to be programmers, but because it establishes a modus operandi which will serve you well in a wide range of disciplines, a wide range of jobs. So that's one thing.

Vint Cerf:
The second thing is better tools. The research community needs to help us build tools that will help us track down bugs or avoid making mistakes that are exploitable. So that's the big research thing.

Vint Cerf:
And the third thing is to figure out where and how accountability should be applied so that we don't lose the value of the enabling power of computers in our zeal to protect people from harm. We want to also provide them with the enabling power of computing to let them do things that a human being could not normally do.

Vint Cerf:
When you Google search, you're doing something no human being could do because the scale of the search is so big. When you do translation from among 100 different languages, very few people speak 100 languages, as far as I know, and even if the translations are not perfect, they enable you to do something that you would not otherwise be able to do, which is to gather some useful information, even if it's not precisely right, at least to get the gist of something.

Vint Cerf:
What about things like real time transcription so that people who are deaf can see what's being said, cochlear implants, which is yet another way of neural interfacing to electronics. The field of computing and electronics is an endless frontier, because you're limited only by what you can figure out how to program.

Vint Cerf:
From my point of view, this is a fantastic field to be in, but it does have some real challenges on the ethics side and on the technical side and on the business side, that will be a rich territory for students to contemplate as they try to figure out their place in the economy.

Jessica Bell:
Thinking about that, as we wrap up our time together, I always like to hear from guests, especially guests who've been as pivotal and involved in the creation of the internet, is what keeps them so excited? What are you just really pumped to hear about in the next five, 10, 15 years of computing? What keeps you here and continues to fuel your passion?

Vint Cerf:
Well, we haven't talked about artificial intelligence and machine learning, but it's going to be an incredibly powerful tool. Multilayer neural networks are doing things that we used to think were not possible. Moreover, they also make mistakes, and so figuring out why and how they make mistakes and being able to anticipate that and plan around it is a super important thing, because some of those mistakes could be fatal, literally. Self-driving cars being an obvious example of that. So that's one thing.

Vint Cerf:
The second thing, from the network point of view, which is my world primarily, we've already gone off planet. Starting in 1998, we began thinking about how to design and build an interplanetary internet, NASA, the Jet Propulsion Laboratory, and now the other space agencies, like ISA and [inaudible 00:30:17] and the Korean [crosstalk 00:30:18] have been working for the last 22 years to standardize a set of protocols that will work at interplanetary distances, unlike TCP IP.

Vint Cerf:
We now have, onboard the International Space Station, we have standardized interplanetary protocols. We have prototype software running on Mars right now in the [crosstalk 00:30:37] and in the rovers, and they will be available for the return to the moon in 2024.

Vint Cerf:
As we send out more scientific spacecraft, as they complete their scientific missions, we can repurpose them to be nodes of an interplanetary backbone. So for me, this is like the beginning

of a fantastic science fiction novel. I won't see the end of it, but I'm having a ball at the beginning.

Jessica Bell:
Well, Vint, I want to thank you so much for giving us your time. This was a wonderful conversation. I wish we had many more hours to talk about all these different rabbit holes that go down, but thank you so much for being with us today.

Vint Cerf:
It's a real pleasure, Jessica. Thanks for taking the time to chat and I look forward to another opportunity someday.

Jessica Bell:
Awesome. Thanks so much.

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