Stanford eCorner

The Future of the Internet

Vinton G. Cerf, *Google*

March 30, 2009

Video URL: http://ecorner.stanford.edu/videos/2220/The-Future-of-the-Internet

Vinton G. Cerf, vice president and Chief Evangelist for Google, discusses the past, present, and future of the Internet. Cerf predicts that Asia's cultural influence will grow as the continent's Internet penetration rates reach European levels. He says that, while IPv6 will provide enough Internet addresses to last through his lifetime, the implementation of IPv6 creates difficulties for the Internet in terms of compatibility, security, and broadcasting. Cerf describes the trends and opportunities of the Internet in the 21st century: the transformation of information consumers into information producers; the rise of social networking; the emergence of new economic systems in online games; the development of user-generated advertising content via streaming IPTV; and the transformation of mobile phones into multi-purpose devices that provide geographically indexed information. In Cerf's view, the increasingly lower cost of storing and transporting bits fosters a new economics of digital information and the emergence of new Darwinian business models that challenge existing entities to "adapt or die." As a result, Cerf says the Internet is an unprecedented and unpredictable innovation engine because its infrastructure enables people to invent new applications simply by writing new software on the edge of the network without having to ask for permission.

Transcript

Let me begin by taking you back into time to the origins of the Internet. Before there was Internet, the Advanced Research Projects Agency in the Defense Department started a project that they called the ARPANET to try out this idea called packet switching, which is a very distinct way of switching data around than the traditional telephone system which uses what's called circuit switching. So, we built a four-node network. And I was lucky enough to be a graduate student at UCLA at the time I

wrote the software. They connected the Sigma 7 computer to the first packet switch of the ARPANET that was installed at UCLA in September of 1969. And by December we had four nodes up and running. The Sigma 7 is in a museum now. And some people think I should be there along with it. But I'm still here and having a good time. If we fast forward to a few years in 1977, at this point Bob Kahn and I have finished doing the original TCP design.

The detail specifications were written in December of 1974 on this campus with two other important graduate students, Carl Sunshine, who is now at Aerospace Corporation, and Yogen Dalal, who is one of the principals at Mayfield Associates, a venture capital firm here in the area. So, we finished writing that spec then we started implementing. By 1977 we were ready to do the first three-network tests. We had a packet radio network, which was an operation here in the Bay Area so you could drive around on the Bayshore Freeway with a big truck with some radios inside with radios that are up on the mountaintop. So, we were testing packet switching in mobile radio environments in the mid 1970s. And we had a satellite system over the Atlantic using Intelsat 4A that linked Europe and the United States with this 54-kilobit channel. Multiple ground stations competing for access to the satellite frequency so it was like an Ethernet in the sky. So, we had the ARPANET, which spanned across the United States, and actually a portion of it went into Europe. By this time, in 1977 I had already left the campus and gone to ARPA in order to run the Internet program for the Defense Department. And I asked the participants to please demonstrate that our new technology would actually link all three of these networks together.

This was an interesting test because the packet radio van was driving up and down the Bayshore Freeway, radiating packets which were running through a gateway, we didn't know we were supposed to call them routers and so we called them gateways, into the ARPANET. And then, all the way across the ARPANET through an internal satellite channel, through an internal satellite channel down to Norway and then by landline to University College London. Then, out of that campus through another gateway in Gunhilly Downs in the UK up through the Packet Satellite Net; that's this thing. And then, through another gateway back into the ARPANET again and then all the way down to USC and from Ancient Sciences Institute in Marina del Rey, California. Well, that's only 400 miles away. But the packet actually went over 100,000 miles because it went over two satellite hubs up and down, back and forth, across the Atlantic twice and across the United States. So, it was amazing that it actually worked. I remember, we were leaping up and down and saying, "It works! It works!" It's software. It's a miracle when a software works. It was very exciting.

I thought this was probably one of the most important demonstrations of the technology. You can make any pair of networks and reconnect by doing all kinds of hacks in-between. But getting three very, very different networks with different speeds, different error rates, different delays and everything else, to either work in a very transparent way with the same set of protocols was an important milestone. And of course, if you fast forward again, this is what the Internet looks like as of 10 years ago. It still looks like this, just bigger. The colors are actually different autonomous systems, different networks that are operated by different operators or different operating agencies. And what you're seeing is sort of the connectivity of all those different networks as they interact with each other across the globe. What's really important about this picture is not just that the Net got bigger and more connected and more colorful, which is true, but that it only works because of the willing collaboration of all of the operators of all the hundreds of thousands of networks that make up the Internet. I mean, you're part of that if you have a small network at home, even if it's just a WiFi system connected to a broadband modem. You're part of the Internet operating community.

And the fact that this works is primarily a tribute both to the technology but also to the motivations that people have to stay connected, to be connected. For business reasons, for personal reasons, for any of a number of different reasons. But it works because people choose to make it work. It doesn't work all by itself. And I found that to be quite surprising and also, of course, very satisfying. If we look at the statistics of the net, what we find are over 600 million visible servers on the Internet. That does not include laptops, desktops or things that are episodically connected like personal digital assistants. These are just the visible servers. Now, that number, 625 million, is low. We know it's low for the very simple reason there are a large number of enterprise networks on the system that hide most of their machines from the public view.

They use firewalls for that purpose and it's quite sensible that they do that. They want to protect those resources from the unwashed masses and the hackers and everybody else out there on the public Internet. So, this number is just the number of visible servers. I don't know what the total number of machines on the Net is because there's no place where they all have to be registered. The number of users on the Net is estimated to be almost 1.6 billion. And again, there is no one place where everybody has to register. So, these are approximations coming from Internet World Stats Corporation. The other observation to make, which is also quite important to the Internet today, is that there are 3.5 billion mobiles in use around the world. Not all of them are Internet-enabled. But many of them are, maybe as much as maybe 15%, possibly even as much as 20% of those mobiles are capable of interacting with the Internet.

And as time goes on, more and more of them will have that capability. So, for many people in the world, the first exposure that they'll have to this Internet environment will be through a mobile rather than through a desktop or a laptop. So, we need to be thinking about that. If you're in the business of delivering services through the Internet, you need to keep in mind that

mobiles are a challenging device. Here was a Blackberry with the screen size as big as a 1920 television set and a keyboard that's suitable for people that are 3 inches tall. So, it's a challenge to provide services through that kind of a device. Here's where the users are. And this should be somewhat shocking set of statistics for you. The Asian population is absolutely the largest number of users already on the Net. Ten years ago, it would have North America.

But as you see, we're number 3 on the hit parade. Our penetration is very high. We have the highest average penetration of any region in the world. But in absolute numbers we're quite small compared to India and China and Middle Asia and so on. But their penetration rate is only 17%. So, you can imagine how many users of Internet there will be if they get to the 74% penetration rate that we're at today. There are a lot of implications of these statistics. And Europe is almost 400 million people and about half penetrated. I've given up making any predictions about Europe, though, because they keep adding countries. So, I don't know what the definition of Europe is anymore.

What's important, though, is that the Asian populations are going to be a dominant influence on the content and use of the Internet. Their languages will be increasingly common on the net, whether it's Chinese or any of the 22 languages in India or Korean or Japanese and so on. And their styles and customs and interests will be reflected in the content of the Net and the products and services that they offer. It's their decisions that will influence what kinds of products and services show up on the Net. Of course, so will ours but this is a very significant population. Remember, about 55 or 56% of the world's population is recognizably in the region we call Asia. In the other parts of the world or as you see it, Africa has continued to be a fairly big challenge. There are almost over a billion people living in Africa and only about 5 or 6% have access to the Internet as we know it today. That datum I think is going to change materially as more mobiles become available in Africa. There are about 115 million mobiles available.

That's about 15% penetrated and I'm sure that number is going to go up. So, we'll see a significant transforming effect on the Internet as more mobiles become accessible or the Internet becomes accessible that way. This is just another picture of the usage of the network by region. And once again, you can see that Asia is outstripping everybody and it will stay that way. There is no question that that will always be the largest single group. And similarly, with regard to percentage penetration, the world average is about a little under 24%, which is pretty exciting, considering that the Internet didn't become accessible to the public until around 1992. Well, that's not quite right. I should say it more carefully. It became commercially available in 1989. It became visible to the public around 1994, when the first World Wide Web capabilities came out of Netscape Communications.

So, in some sense, the commercial publicly accessible Internet is only about 15 years old. So, to reach 24% of the world's population in that amount of time seems to be a pretty interesting statistic. Now, this is another important statistic. The only part of this graph that you should be worried about is the one that's going down. And that's not the Dow Jones Industrial Average, although that's going down, too. This is actually the amount of Internet address space that's available for allocation by ICANN, the Internet Corporation for Assigned Names and Numbers. Now, how could this be? We're actually running out of the Internet address space. When you type domain names like www.google.com, the Internet doesn't really know what that is, only it's a name. But you have to look it up in what's called the domain name system in order to translate that name into an IP address, an Internet address. And it's only with the Internet addresses that you can actually stir data around in the network like telephone numbers.

The problem is, when the Internet was designed and the standards were stabilized, somebody had to decide how much address space was needed. So, around 1977, when I was running the program for the Defense Department, after a year of argument among the engineers about how much address space the Internet should have, remember this is four years into an experiment that no one knows for sure is going to work, I want you to have the right context, after listening to a year of "Well, it should be variable length addresses," The programmers didn't like that. 128 bits of address space, everybody said, "That's crazy. That's 340 trillion trillion trillion address. Why do you need that much for an experiment?" So, finally I said, "Look, enough arguing. The 32 bits is fine. It's 4.3 billion addresses. That's enough to do an experiment. So, let's get on with it." So, we did. The problem is, the experiment never ended.

So, here we are. It's 2009, we're running out. By around 2010 or '11, we will have run out of address space. So, the solution in fact is something called IP version 6, which is now slowly making its way into the Internet. And I'm very proud to tell you Google has implemented IPv6 as well as IP V4. So, we're running in parallel and we urge everybody else to do the same thing. Three hundred forty times 10 to the 36th addresses, that's the number only the Congress could understand. I used to go around telling everybody, "Well, that's enough address space. So, every electron in the universe can have its own Web page if it wants." And then, I got an email from somebody at Caltech, "Dear Dr. Self, there's 10 to the 88th electrons in the universe.

And you're off by 50 orders of magnitude." So, I don't say that anymore. But it's vitally important that we actually introduce this larger address space because you want, I want, the Net to continue to grow. I mean, plainly with only 24% penetration, Google's chief Internet evangelist, has 75% of the world to convert. And I need more address space to do it. I mentioned mobility and mobiles before. Let me emphasize as hard as I can how critical these new devices are. They're clearly not telephones anymore. They're programmable devices. And they are used in just an amazing variety of ways for simple things

like texting and transmitting of data or using it to access the Internet or Web services like Google search and so on. Some people are using these things for making payments.

In fact, it's avery interesting phenomenon. Minutes in the mobile system have become a medium of exchange in places where people don't have bank accounts and don't have cash and don't have checks. But they have mobiles. So, they trade minutes back and forth as a means of payment because those minutes have value. I had mentioned innovative interface and like I say, it's a very limited screen and keyboard. The thing I would like, though, would be for a mobile to be able to detect available interface devices as it comes into a room. For example, this projector, if it had a radio in it suitably configured, my mobile could figure out that there is a projection device in the room and then use it in order to get away from the limited screen space. Or when you walk into a hotel room, it will be nice to notice that you have a nice, high-resolution flat-screen display and the mobile can make use of that. Or maybe a little Web TV keyboard is available and you could use the bluetooth interface on that to extend the keyboard of the mobile. At some point, the mobile becomes an aware device that uses anything else in the local area as a means of interaction.

So, I would find that very attractive. The other thing that we've noticed very interesting about mobile use is that because many of the mobiles know where they are through GPS or maybe through triangulation of the nearby bay stations, these devices induce people to ask questions about where they are or related to where they are. So, where's the nearest ATM machine or where is the nearest gas station or the nearest Thai restaurant or hospital or something. So, we've watched that phenomenon at Google. Many, many people ask questions that are geographically significant. And that simply exacerbates and emphasizes the value of geographically indexed information. So, if someone is producing information related to where things are and when things are happening there, either in the past or in the future or in the present, that information becomes very valuable because people are interested in it. And they are particularly motivated to ask those kinds of questions when they have their information window on their hip or in their purse. The other thing which is a clear phenomenon is that more and more devices are showing up on the Net. And I will admit to you that some of these I never imagined anyone would bother connecting to a Net.

In fact, we used to tell jokes in the early Internet development about Toasternet. Someday, your kitchen appliances would be on the Net. We all thought that was a joke. And then, somebody at an InterOp trade show actually put a toaster up on the Internet. And you could send a little, simple network management control packet to it to say how burned you wanted your toast. And so, we all thought that was really funny. And then, a few years later, somebody ran into my office saying, "Man, did you see the Internet-enabled picture frame?" And I thought that sounds about as useful as an electric fork. In fact, this is a very useful thing. For one thing, you don't have to log into it. You don't have to boot it up or anything.

It just runs. This one is made by a company called Ceiva. It plugs into a phone line and it has a built-in modem. And it calls up the Internet and goes through a website. Each picture frame has its own unique ID. And it logs in to that website and then downloads instructions. It says what images it should show or which ones it should remove from its local cache and how long it should show them before it switches to the next image. So, my family members and I have digital cameras. And we upload our imagery to the service website. And these little picture frames that we have scattered around in our homes automatically download pictures of the nieces and the nephews and whatever else is going on in the family.

So, you get up in the morning and you get some sense of what's happening to everybody. Now, those of you who care about security will recognize that if the website from which the picture frame is downloading images is compromised, then the grandparents may see pictures that they hope were not of the grandchildren. So, suddenly security becomes just as important at home as it is at work. There are telephones that look like telephones but they're really Voice over IP computers. There are refrigerators that are now sold that are connected to the Internet. I used to think, "What do you with an Internet-connected refrigerator?" One obvious thing is that they have touch-sensitive displays on the door. And this augments the standard American family communication system, which consists of paper and magnets on the front of the refrigerator going down. You can have a Web blog, email and Web pages and other kinds of enhanced family communication capability. But I like the guy in the middle who made an Internet-enabled surfboard. I guess he was sitting on the water, waiting for the next wave, thinking, "If I had a laptop, I could be surfing the Internet while I'm waiting." So, he built this thing and he put a WifI service back at the rescue shack on the beach and now he sells it as a product.

Pretty amazing. So, I'm predicting that there are going to be billions of these kinds of things on the Net. Sometimes, it goes by as Internet of Things. They will all be programmable devices. You'll be able to download new software into them to augment their functional capability. I don't have time to speculate about all the various things that you can do, although I've often been amused by the fact of what the Internet-enabled refrigerator could do. I mean, imagine if everything you put into the refrigerator had an RFID chip on it and you had an RFID detector inside the refrigerator, so when you put stuff in there, the refrigerator knows what it has inside. So, while you're up working and surfing the Net looking for recipes that it could make with what it knows it has inside. You come home and you see a nice list of things you have for dinner. It sounds like a pretty good deal to me.

You can extrapolate on that. You could be on vacation and you get an email that would be from your refrigerator. "I don't know how much milk is left but you put it in there three weeks ago and it's going to crawl out on its own." Or maybe you're shopping and your mobile goes off. It's an SMS, "Don't forget the marinara sauce. I have everything else I need for spaghetti dinner tonight." Unfortunately, the Japanese have done a really bad thing. They have invented an Internet-enabled bathroom scale. And so, when you step on this scale, it figures out which family member you are based on your weight. And it sends the information to the doctor to become part of the medical record, which on the surface is probably OK except that the refrigerator is on the same network. So, when you come home, you see diet recipes or maybe it just refuses to open. The one thing I do want to mention about mobiles again is that if you're like me, you have an entertainment system that consists of a whole lot of boxes and these little remote controllers, infrared controllers.

I fumble around, trying to figure out which controller goes with which box. And then, when I finally figure it out, that's the one with the dead battery. So, I suggest we just get rid of all that. All of the entertainment boxes should just be on the Net. And then, your mobile could control them, either directly, through the network in the house, or maybe even by way of a service company. So, you have a nice Web page and you can say, "Here, I want these movies and this music and everything else." It figures out how to tell these devices how to download and record or play back the entertainment that you're interested in. I did want to mention one other thing, though, and that's sensory network. That's becoming an increasingly common component of the Internet environment. I'm an example of that. I have a sensor system which picks up temperature and humidity and light levels in all the rooms of the house.

It's a commercially made product by a company here in this area called Orange Rock. It runs IP version 6 on a radio-based network in the house. It's a mesh so these little extensor devices are also forwarding packets around to each other. So, it maintains its connectivity over time. The reason that I installed it originally was to keep track over a year's time of how the heating and ventilation and air conditioning actually function. Each room in the house is either hot or cold. And rather than having anecdotal reports back to the HVAC, as an engineer I wanted real data to say, "Well, during the summer, what was the distribution of temperatures in the rooms or during the winter? Was it balanced or not?" Of course, one of the very important rooms in the house is the wine cellar. It needs to be maintained less than 57 degrees and preferably more than 60% humidity to keep the corks from drying out. So, that one has been carefully alarmed. So, if the temperature goes above 60 degrees, I get an SMS on my mobile.

This actually happened last year. I was at Argonne Natural Laboratory for 3 days. And at the beginning of my visit there, the mobile went off. It was the wine cellar calling, telling me that the temperature just exceeded 60 degrees. Unfortunately, I was away for three days and my wife was in Russia on a tour. And so, every five minutes for the next three days, my mobile said, "Your wine is warming up." By the time I finally got home, it was about 70 degrees in the wine cellar because the cooling system had cut out for some electrical reason. So, I called up the Orange Rock guys and I said, "Hey, do you make actuators as well as sensors?" And they said, "Yes." And I said, "Great. So, that's going to be the next weekend project. It's to install an actuator that I can remotely turn off and on." It's pretty clear that strong authentication is going to be needed for access control, right? Because I don't need the 15-year-old next door messing around with my wine cellar cooling system. That would not be a good thing.

Because the wine cellar is instrumented like the rest of the rooms in the house, I can actually detect when the lights go off and on because the light levels are being picked up every five minutes. So, I thought maybe that would be a good way to notice if somebody had gotten into the wine cellar when I wasn't there. But I realized that that didn't tell me whether any wine bottle had left the wine cellar without my permission. So, the next idea was to put an RFID chip on each wine bottle and an RFID detector in the wine cellar in order to tell whether any bottles have left. Setting aside the fact that RFID doesn't work very well through liquid, water in particular, you have to configure the room properly for this, somebody else pointed out that you could go into the room, drink the wine and leave the bottle. So, we have to fix that. And the next idea is to stick detectors in the cork so it can tell if there's any liquid left in the bottle. And while we're at it, we might as well put in detectors that can detect the kind of aromatic chemicals that make up the nice taste in the wine, esters and things like that, so we can figure out whether the bottle is ready to drink or not. You could also figure out whether by recording the history of each one of those bottles, if one of them had been overheated. That's the one you give to a friend who doesn't know the difference.

OK. So, I already talked about IPv6. Let me also tell you this year, 2009, is really a major evolutionary year for the Internet. Some significant core changes are taking place. The introduction of IP version 6 is part of it. Another one is to digitally sign the entries in the domain name system in order to at least resist some forms of fishing and farming where people have the address associated with the domain name incorrectly reported back to them because the domain name system itself has become compromised. And so, by digitally signing the information, you could check to see whether the data has maintained its integrity since the information was put in, the finding between the IP address and the domain name. In addition to the DNS SEC process, there's also a major ongoing effort, which we hope will come to closure soon to introduce non-Latin characters into the domain name system. Back into 2003, ICANN, together with the Internet Engineering Task Force, did begin introducing non-Latin characters into all domain names except the top-level domains. The top-level domains still consist of two-character codes

for country codes or multi-character generic slang, dot-com, dot-net, dot-org, dot-travel and so on.

Within this year, it's expected that ICANN will be permitting the registration of non-Latin characters into top levels of the domain name system. And so, I predict that that's going to introduce a significant influx of domain names that are written in scripts other than Latin. And that's important, remembering the statistics that I showed you at the beginning. Many, many of the people who represent large portions of the Internet population speak languages whose scripts are not representable in Latin characters. So, this is a very important activity. On the other hand, this is also really hard. The reason it's hard is that there are so many scripts that are used for writing languages. There's a system of encoding those scripts called unicode. And it incorporates 100,000 or more symbols in order to allow expression of all these different languages. And some of them look the same.

So, if you're thinking for a moment about Cyrillic and Greek and Latin, just a simple example there, many of the characters in those scripts used to write those three languages look exactly the same. So, the lower case "a" in Cyrillic looks very much like the lower case "a" in Latin. But they're coded differently. They are coded to be distinct. So, if we were to write PayPal, using a Cyrillic "a" instead of a Latin "a", the computer would think that was a different place to go than the PayPal written with Latin. And if you can't tell from looking at it, someone could put up a screen on the Web saying "Click here to go to PayPal," and you go to the wrong place. You'd be invited to type in your username and password. Whoever is running that site would quickly take that information and go off and drain your PayPal account. So, there are risk factors associated with introducing these new scripts. On the other hand, if we don't do it, the countries who cannot represent their languages today in the domain name system will go off and create their own rooms, their own domain name systems and now the Internet will fragment.

That's not good, either. So, the effort here is to put them together in a way that allows multiple scripts to be used in the safest way we can possibly achieve. So, now I want to talk a little bit about something else. This is a place where the anthropological views of this group would be very helpful. Tony Blair came to Silicon Valley in the summer of 2007. Some of you may remember that. He was out at Cisco Systems at their invitation and about 10 of us showed up for lunch that day. His ostensible reason for showing up was trying to figure out whether there was any way of recreating a Silicon Valley in England. I remember how this went. He went on about that and about education and everything.

And he got done with his speech and nobody said anything. And so, I thought, "Well, we can't have the Prime Minister sitting here after he finished speaking." We were all sitting there saying nothing. So, I said, "OK. Why don't we tell Prime Minister Blair what is it about Silicon Valley that somehow made us all successful or at least what similar experiences do we have?" And Steve Jobs puts up his hand and he says, "Well, one thing that we've all experienced here is we've all failed at one time or another." And he said, "But that didn't leave us with the mark of Cain on our foreheads." It was actually an insightful comment because in Europe historically if you fail in business, it's a very severe penalty. It's often very hard to get any backing from anyone once you fail in business in Europe. Perhaps that's changing now but certainly in the past it has been a serious problem. Here, it's considered a mark of experience. Of course, if you fail all the time, it's a different story. But one or two failures is not considered to be completely damaging; it's considered experience. We're also very fortunate here in the Silicon Valley to have such a strong source of educated workers, here in Berkeley, in San Jose and all the other colleges around the area.

And so, we have a continuous influx of educated people, skilled people, who can be put to work in the Silicon Valley enterprises. We also have very liquid markets, or at least we had liquid markets until recently. And now you're going to think that's beginning to pick up a little bit. But it's important because you can't have the level of innovation in industry without having a market that will allow public investment in these enterprises. You can't expect the venture capitalists to provide all the capital, although they do provide a significant part of it at the outset. They are willing to take risks in exchange for some kind of compensations. Venture capital was not common in Europe. In fact, most of the funding for European businesses used to come from banks which were notoriously reluctant to take risks. And so, creation of new business was and probably still is more uncommon than it certainly is here in Silicon Valley. We also have a talent pool that's very fluid in the sense that it can move from one business to another and it does.

In Europe, that's less so or at least it was. And certainly everybody seems to know everybody here. You either work for somebody or they work for you. Or you were partners together in some venture. And it has a kind of keiretsu-like feeling to it, where people know each other in the Valley. They are aware of each other's experiences, their strengths and their weaknesses. And finally, I felt compelled to tell a story of the "cargo cults". Some of you will remember that during World War II, especially in the Pacific Islands, where the Americans established bases, they employed the native population to help build the airstrips and buildings and other facilities that were needed during World War II. When the war was over, we left. And the level of the economies which had inflated significantly during World War II on many of those islands suddenly collapsed because there was no one there demanding work.

And so, some of those populations actually constructed out of palm leaves and crates and the like, things that looked like airplanes because they thought that maybe if they built these things that looked like airplanes that they would attract the

airplanes to come back. And they were called "cargo cults". I would use that term to describe countries that don't understand the dynamics and the environment that has to be in place to create a Silicon Valley and to sustain it. They build industrial parks. They put buildings together. They provide power and all this other stuff. And then they asked companies to come in there and populate these shells. The problem is, it's just like the cargo cult; it's an empty shell. And unless you have all the other desiderata to keep this economic engine going, it just doesn't work. And so, you can't just build the shell of an industrial base.

You have to have all the other pieces there working together. It's like an engine. You have to have all the parts tuned. It has to have gasoline and oil and all these other pieces. If you don't have all that stuff, it just doesn't run. And it turns out to be harder for some policymakers to fully appreciate that than I would have thought. And I don't suggest that Prime Minister Blair didn't understand that. But I think and I hope that it caused him to think a little bit about what the environment would have to be in order to create a Silicon-like Valley. Now, I know you're interested in organizational structures. And I'm pretty sure that I'm not telling you anything you don't already know.

But one thing which is very clear is that despite the current popular technologies like Twitter and instant messaging and the like, geographical dispersion is a problem. If you're not in the same time zone, it's really hard to work easily together. Email is not a substitute for sitting down over a meal or working at a whiteboard together. And I don't know about you, but video conferencing is still not the most wonderful means of having an interaction in my view. It's still a little structured. It's still a little artificial, to say nothing of the fact that the audio is often terrible for a lot of conferencing systems. So, I'm convinced that getting people working together in the same time zone is a very important tool. At Google, we're very dispersed. We have offices all over the world. But what we have learned in the last, say, three or four years is that putting small groups of people together in the same geographical location is a very powerful way of motivating.

And first of all, everybody sees what everybody else is doing. And if one person or two people are not performing, it's visible to the rest of the group, whereas if you're really physically dispersed, then sometimes it's hard to tell. I also think that sharing documents has helped a lot for us. I had one occasion where three of us were trying to write a paper together and we had a very short deadline. And so, we had all gotten to an audio conferencing system on the Net. And we all used shared Google Docs. And we could see each other making changes to the document and we could discuss it while we were writing. We actually got finished in a couple of hours. I never quite experienced that kind of remote working together. Of course, it only works when we're all up at the same time.

I mean, if you tried to do this you had to work serially, I don't think it would have worked as well. We all understand that organizations are getting flatter partly because of the ability for anyone in the organization to communicate with anyone else. And that I think has been helpful to the extent that the management is willing to listen to people's messages coming from literally any part of the organization. More people are able to work from home. I know I take advantage of that. But I have some theories about the communications that are popular today. Young people find email too slow. They consider that to be old hag. It's half 20th Century. I mean, it was invented in 1971, right? For them, it sounds like Civil War.

They can't tell the difference between 1865 and 1970. It all compresses, like the New York City view of the rest of the United States, if you have seen that New Yorker magazine. So, they fall in love with things like Twitter and texting and everything. I have the following prediction. Most of those younger people are working with folks or kids in school that are geographically nearby. So, real-time communication makes every bit of good sense. But as they get older and their friends begin to disperse, they're going to find that the ability to communicate that way is not very satisfactory. I'm not interested in reading my friends' Twitter at 3 o'clock in the morning. So, they're going to find that the time-differed communication methods are more attractive as their friendships expand geographically. That's a testable prediction so we can watch to see whether that actually happens.

The other thing that I'm concerned about is that more people are online. More people are participants in group communications, distribution lists, websites, blogs that have commentary on them. I'm watching more mob behavior taking place in blog kinds of environments where somebody puts up something outrageous, whatever it is, and then there is this raft of commentary, people becoming very angry and fulminating about it and everything else, without anyone doing much to validate the original assertion that showed up in the blog. Maybe it was just somebody's opinion. But I noticed that that kind of behavior, it feels to me like it's mob-like in its character, where everybody goes running off because somebody is running off in some direction or other. Those of you who are interested in anthropology, I'm thinking of you in particular, ought to say something about that. You could say whether or not we should anticipate this kind of behavior pattern in this online environment and what we can do about it. And then there's information overload. We're all sort of blizzarded with emails and Web pages and all these other online sources of information. And it's hard to figure out what to spend your time on.

Google tries a little bit to help here by trying to sort through all the billions of pages that it crawls through to help you figure out which might be the most relevant information that might be useful to you. But even there, at best I can say it's an approximation. I don't have any magic wand to wave here. But I can tell you that it's a serious issue for many of us. And finally,

I wanted to cover some of these material then give you a quick update on the interplanetary extension of the Internet. And that will be my final formal remarks. But I wanted to say a little bit about semantic link. This is Tim Berners-Lee. Since his original invention of the World Wide Web, he has always wanted to go beyond this syntactic character of the current World Wide Web into something that understands meaning better than it does today. I think he has been challenged to try to find a way and make this work.

But it dawned on me that there might be some mechanisms that he has invented that I didn't understand until recently. Everybody knows what a hyperlink is. When you make a Web page, you can put in a reference to somebody else's Web page. So, when you create the Web page and you put a link in it, you're basically saying that you're someone who's looking at that Web page. "I found something over here that I believe is relevant." So, it's just a pointer to some place else in the Net. Let's imagine for just a moment you have another kind of hyperlink that you can put into your Web page. Let me call it a semantic hyperlink. I don't know what Tim calls it but let me call it that for purposes of this discussion. Here's an example. You notice that the word "jaguar" is actually ambiguous because it could mean the cat, it could mean the car or it could mean an operating system from an Apple.

And so, because you noticed this ambiguity and you'd like other people and maybe other programs to notice this, you create a semantic hyperlink in your Web page. And it says, "Here's an example of 'jaguar' used as animal; point to this page. Here's an example of the word 'jaguar' in reference to a car and here's the pointer to the operating system." Now, imagine that you're the Google crawler. And you're crawling through the Web. And you encounter this semantic hyperlink. The program can say, "Oh, this is a warning about ambiguity of the term." So, now it incorporates that notion, this ambiguous notion of "jaguar", into its dictionary. So, when you or somebody types in the word "jaguar" in looking for jaguar, the system comes back and it says, "Well, did you mean the car, the operating system or the animal?" So, the crawling process begins to ingest semantic information as well as simply word content, strings. I don't know exactly how hard or easy it's going to be. But what Tim has done in my view is to create a mechanical way for each of us to contribute our understanding of semantics into the system and to have that understanding incorporated into the software that's crawling the Net and indexing it. The second thing I wanted to draw to your attention is that every day, when you use spreadsheets or word processing documents or presentation software or anything else, you're creating very, very complex files that are only interpretable by a piece of software.

And if all you have is the bits of those files and you don't have the software, the bits are meaningless. Now, the thing I worry about is that eventually we will accumulate a large pile of bits. But the software that interprets those bits may not survive. Now, there are some people who say that I'm simply crying "Wolf!" and I shouldn't worry about it and there will always be backward compatibility or it's trivial to keep the software. They're just more bits so just store them somewhere. I'm not very comfortable with this for a couple of reasons. First one, sometimes that piece of software that knows how to interpret the bits only runs in a certain operating system environment. So, maybe I have to save the operating system, too. What if I have to save the computer or an emulator of the computer to run the operating system that runs the application that knows how to interpret the bits? Isn't there intellectual property hiding in some of those various things? How can I make sure that the rest of the world will have access to intellectual property that belonged to a company that decided, "I'm not going to support this software anymore because it's not economic to do so"? So, my personal belief is that we do have a challenge ahead of us that we're poring digital bits into the Internet and we haven't figured it out, in my view anyway, how to make sure those bits have meaning 100 years, 500 years or 1000 years from now. I call that problem the "BitRock Problem".

And the sort of the snapshot example is the year 3000. You've just done a Google search and you've turned up a 1997 PowerPoint file. And maybe you're running Windows 3000. The question is, "Does it know how to interpret the 1997 PowerPoint file?" And if the answer is no, we have a problem. And it couldn't be not just 1000 years from now, not just 500 years from now. But it might even be two years from now. And so, if you're already experiencing the things I am, where certain JPEG files are no longer interpretable by current versions of photographic management programs, you already know what I'm talking about. OK. So, final update here is on the interplanetary extension of the Internet. This is not a Google project.

This is a project that I started with the Jet Propulsion Laboratory 10 years ago, 11 years ago, in 1998. Basically, we were thinking about exploring Mars because we had been there a number of times. And in fact, we built in 1964 a thing called the Deep Space Network that is used to communicate into deep space with satellites that are orbiting the planets or flying around the asteroids and things like that. Or some of them have landed on the surface of Mars like the rovers that are currently still in operation after five years, which is amazing because they were originally designed to last 90 days or at least that was the mission time. So, they're still in operation. Or the Phoenix lander that landed on the surface of Mars in the northern part of the polar regions in May of last year. So, we said, "Well, you know, we're going to have more and more of these spacecrafts out there. Some of them will have increasingly long lifetimes because we're returning now to the use of nuclear power sources which could mean missions could last decades." And so, there will be more things to service. And eventually, there will be enough things out there that it won't be easy to use manually scheduled communication systems, which is the way the Deep Space Network works today. So, we wanted to develop a kind of Internet in space so that we could support more complex missions with multiple devices, either in orbit or in the surface of the planets scattered all across the solar system.

We started out thinking we could use the GCPIP protocols for that and we were wrong. It turned out that there's a problem. The distance between the planets is literally astronomical. In time, it takes 20 minutes for a radio signal at the speed of light to go from Earth to Mars when we're at the farthest part in our respective orbits. And even when we're the closest together, that's 35 million miles, it still takes about 3-1/2 minutes to get from one planet to the other one way. So, the flow control methods of GCP don't work. And then, there's this other problem called celestial motion. The planets have the habit of rotating. And so, if you're trying to talk to something on the surface of the planet, it just rotates out of the way and you can't talk to it until it comes back around again. And so, that's very disruptive.

So, we ended up having to design a new set of protocols that operate better than this environment. And we call it delay and disruption tolerant networking or DTN for short. What's exciting about this particular effort is that it has been tested now terrestrially in several different environments. We got the Defense Advance Research Project Agency to test the DTN protocols in tactical military environments where disruption is common. People are jamming the radios or you're hiding in radio shadow to keep from being shot at. And we also tested it with the reindeer herders in Northern Sweden. We did two things with them. We put a WifI system in one of the villages. And then, we put the DTN protocols on laptops in the backs of all-terrain vehicles. And as they wandered into the village and wandered back out again, they were acting like little data mules.

They would drop off data and pick it up using the DTN protocols. And then, last summer they hang radio systems programmed with the delay and disruption-tolerant networking around the necks of the reindeer themselves. And they tracked the reindeer to see where they went during the summertime. So, we're now very confident that we have a functional set of protocols that survive in delay and disruption. Just as last October, we actually did deep space testing. Several years ago, there was a spacecraft called Deep Impact. And it ran with a comet and basically launched a probe in front of the moving comet. The comet ran into the probe and blew up. The idea was to get some of the internal material of the comet because that was the same material out of which the solar system was formed. And so, the scientists were interested in seeing that primordial material.

The project actually didn't work as well as they hope because a lot of the comet turned out to be water. And so, when the probe hit this thing and blew up, a big cloud of steam came out and it obscured the signal that you wanted to get from the spectrometer. So, they didn't get quite as much information as they wanted. The spacecraft is still in orbit around the Sun. And so, last October it was on its way back to Earth. It's going to meet another comet in 2010 or maybe early 2011. So, NASA gave us permission to upload our software to this spacecraft that was in orbit around the Sun. And so, 81 light seconds away, we launched our packets back and forth for about a month, testing the protocols in a work bay. Then, this summer we're going to put the protocols up on the international space station and test it again. And then, in the fall we're going to re-upload whatever the latest protocols are on the same spacecraft that ran into the comet before it has been renamed EPOXI.

And we'll have a three-node Internet running by the end of the year for test purposes on the space station in the comet chaser and on the surface of Earth. So, what we're hoping frankly is that by the end of 2009 we will space-qualify the new protocols. And they can be used by all of the space-faring nations in the world. We're already briefing the Consultative Committee on Space Data Systems about the standards that we proposed for deep space communication. If they adopt these standards, it means that every mission that gets launched will be compatible from the protocol point of view with every other mission. That means that we can re-purpose any of the mission equipment once it has completed its primary job if it still has computers and radios and computing capability and memory. And they can be re-purposed for communication purposes. So, we will end up building over time an interplanetary background. So, that's the plan. And it feels to me as if 2009 is the first step towards a serious interplanetary network.

So, let me stop there and let me thank you for letting me chew up so much of your time. And I'm very happy to spend time engaged in communication. So, thank you for that and we'll go into the Q&A if you like. I have to warn you that I'm hearing-impaired and I have an ear infection. So, I'm going to be very careful to use my hearing aid in the infected ear just a little bit. OK, I'm now functional. What is it? I'm operating within normal parameters. And I suppose they have a microphone turned on around. But do you need the microphone to catch the questions? Do they have to talk into my lapel? You've got one there? OK. All right.

So, I saw Terry with his hand up. So, thank you, Terry. Everybody else Terry Winnegret, right? I haven't been following up and I'm curious. There has been a long-term tension in Internet government between the sort of engineering-based organizations that grew out of the US community and the International Governance Organization from the UN and all that. Where is that moving as the population on the Internet shifts? That's a really good question. There was a period of time starting around 1998, when something called The World Summit on the Information Society was introduced by the International Telecommunications Union at the suggestion of the Tunisian government. And the question was, "What's an information society? What it should look like technologically and from the regulatory point of view and socially and everything else?" The immediate debate at that point was, well, how would you create an information society? And everybody looked at the Internet and said, "Well, that must be an example of the foundation of an information society. So, remember, this is a government

conference. So, immediately, the question was, "Who's in charge of the Internet?" We said, "Well, nobody actually. It's just a big collaborative thing." And they said, "Well, we don't believe that.

Nothing of this magnitude could possibly not be controlled by someone. We know better and we're governments. We control things." So, then they looked around and they saw this Internet Corporation for Assigned Names and Numbers. I was chairman of the board at the time that this was all happening. And they said, "Well, they're responsible for domain name allocation and Internet address assignment. So, they must be in charge of the Internet." So, then we became the target of this thing. And nobody could decide what it meant to govern the Internet. So, they had to create a working group on Internet governance. And then, they argued over, well, what does it mean to govern the Internet. That went on for awhile.

The result of all of that, which took place over a period of about four or five years, was the creation of something called the Internet Governance Forum, which was not a decision-making activity. It was simply a place where people could bring issues together and have discussions. And it was intentionally created as a multi-stakeholder activity, very much like ICANN itself, which is also very multi-stakeholder. That was an interesting breakthrough by itself. I realize I'm so long way of getting to where you want to get to. The interesting thing about the Internet Governance Forum is that governance actually allowed others from simple society, from the engineering community to operators and so on to participate on an equal footing in the discussion. Now, probably the reason they allowed this is that no one was making any decisions in that organization. It was simply a discussion. And there were some voices representing countries like Russia, China, Syria, South Africa, Brazil, who weren't satisfied with having a discussion. They wanted to make decisions about this thing called the Internet.

And so, they have been lobbying for some other parallel organization or something that could make decisions on a global basis about this thing. I'm not comfortable with that. But I will say that I can't argue entirely against some kinds of multilateral agreements. Let me try to say why. If you think a little bit about the risk factors that are showing up on the Net, if you watch 60 Minutes on Sunday, you know that, as overblown as I think the 60 Minutes was, they still have a point. There is risk out there. There are people who are on the Net who are deliberately trying to harm other people or at least who find ways of defrauding people of money or stealing their names and passports, creating false credentials and things like that. So, there are harms and crimes that take place in the Internet. We shouldn't be surprised at this. People commit fraud through the telephone or through the postal service, face to face.

The Spanish prisoner fraud was invented in the 18th Century. And it has reemerged as Nigerian letters coming across the Internet about the $25 million that you are supposed to help move from Nigeria to the United States. So, this kind of fraudulent behavior is harmful. And the thing that's hard about dealing with it is that the party committing the crime may be in country A and the victim may be in country B. And so, dealing with it means some kind of bilateral agreement of some sort and even an agreement about what constitutes harm and what constitutes illegal action. So, I believe that it is probably worth the effort for multilateral discussions about what is commonly thought to be harmful to society to be agreed. The problem, of course, is that not everything that each society thinks is harmful is agreed to be harmful by everybody else. So, there is going to be some lowest-level agreement on some of these things. Child pornography is another example of something that most countries consider completely unacceptable and therefore would cooperate in apprehending people who are abusing children that way. I think that one could speculate about a Law of the Net, something like the Law of the Sea, if you think of the sea as a common resource.

There are lots of things that people don't like about the way the Law of the Sea came out. It took 20 years to get there. But the fact is, we follow, most countries anyway are following a lot of the provisions in order to deal with events that occur in the open oceans, who has access to fishing rights, what do you do about salvage, all those sorts of things are largely agreed on the multilateral basis. So, I can see a common kind of set of agreements related to the Internet arising. And I would consider for the protection of our respective societies, if that's worthwhile to do. On the other hand, you can also have some very pernicious things happen like governments deciding that they're going to use the Internet as a medium of surveillance. Or they're going to use the Internet to control what people do here or either can say or can't see or hear. So, you see governments in China, for example, working to limit access to some material on the Internet. We have some experience with that at Google. And I'm actually surprised to say that there isn't as much suppression of access as I would have expected.

It comes and goes. Sometimes YouTube gets shut down as it was by accident, if I remember it right, in Pakistan. They wanted to shut down access to YouTube only in Pakistan. And they managed to shut down access to YouTube for a good portion of the world because of the way they implemented their constraints. They basically lied about what the address of YouTube was. So, I'm worried a little bit there about governments that may choose to use the Internet to suppress it or to suppress access to it. So far as I can tell, though, most of the attempts to do that haven't been 100% effective. People can find their ways around some of these constraints. John Gilmore, who is one of the founders of Sun Microsystems, was quoted as saying that the Internet interprets censorship as damage and roots around it. And that's what a lot of people do.

So, I am concerned, to come back to the core of your question, that governments are starting to realize a love-hate relationship with the Internet. They see that it has economic importance to them. They know that they cannot operate in a

global economy without something like that. It's a way of allowing industry to discover each other's goods and services. At the same time, for some governments it's a source of information that they sometimes don't want their populations to have. And there is tension there; it is not resolved. Yes, Ma'am? Let me go here and near. Dr. Cerf, thank you for coming. Call me Vint like everybody else does.

With some historic things that happened here on this campus and we can feel the energy bursting in this room when you talked about what it takes to reproduce that. But I have just a short question for you. Do you feel the singularity is near? Yes. And there is the Singularity University that has just gotten started. Well, let me take your question literally. Ray Kurzweil, as you all know, has written his interesting book called The Singularity is Near. I am simultaneously fascinated by Ray's data collection and analysis and I'm also very skeptical of the conclusions that he reaches. For example, if it's true that the number of transistors in a computer reaches the same number of neurons there are in the human brain, does that instantly suggest to you that the computer will become sentient. And my reaction to this is no. And part of it is based on layman's understanding of how brains work.

The interconnections are extraordinarily complex because you have detritic interactions as well as axon-dendrite transmission of pulses. You have all kinds of inhibitory and excitatory signals. The human brain is a really complex piece of equipment. And I would say that no matter how many transistors you have in the next 10 years or so, the interactions among them is not likely to mirror the same thing that we have in our brains. The other reason that causes me to be skeptical of this notion that somehow a machine will become smarter than we are, and then it will invent the machine that is smarter than all the people that have ever lived, and then we'll upload ourselves into the machines and that's how we'll explore the galaxy because the machines will last longer than we will. And so, we just invented the replacement species, Homo Computational NZes or something. The other reason that I am somewhat skeptical of the notion of sentience arising in this fashion is that, Terry, here's something that I'm sure you'll be able to comment on because of your own expertise, my sense is that we learn language, we learn to understand the world around us because we physically interact with it. So, when I say the word "water", for you that has meaning because you've experienced water in so many different ways. You've touched it. You've drunk it.

You swam in it. You stood under it while it's raining. You have all these physical experiences of water. And so, you have some notion of what it means. Computers don't interact with the world that way, at least not yet. More and more of them do have the ability to sense their environment. And it may be that as we increase the capacity of computers to experience the world somewhat in the way we do, maybe there will be more likelihood that there is literally a semantic there that would allow a machine to be more like us. But at the moment, I think it's very much a prescribed territory. Would you care to comment on that, Terry? I don't mean to put you on the spot. But I know your work has had a lot of relevance to that.

You know that singularity is always around the corner. I was a graduate student when you were. I think most of the people that lived in the laboratory were pretty sure by the 21st Century that would have happened. It would happen. It's hard. The brain is an amazing thing. OK. You had your hand up. It's Dick Scott. Yeah.

I wanted to follow up on Terry's question. He talked about the possibility of... What about the kinds of commercial threats, whether proprietary controls and so on. What about that? That's a very good question. It's actually a really big question, too. Gosh, there are so many aspects of it. Well, let me pick a few and you can say if I miss the one that you particularly wanted to get at. Let's look at intellectual property for a moment. Software patents are a hot button for a lot of people. I had quite a vigorous debate on this subject with the head of the Intellectual Property Management here at Stanford University.

There was a meeting sponsored by the Marconi Foundation recently. It wasn't a staged debate but it was pretty vigorous. My reaction right now is that software is almost like an apprenticeship. It's a skill thing. You learn from reading other people's code. And I would cite just as a small example of that how quickly Webmaster is developed because browsers allowed you to see the text of the program that produced the Web page that you thought was pretty cool. You could see the HTML. People learn how to copy other people's HTML script. And of course, as you all know we're turning now to XML. But the fact is that we had people learning from other people's work and experience and copying each other.

It's very powerful. Putting software patents into that loop is a really troublesome thing. So, I'm concerned there because I think it's stymies creativity and innovation. Another example of commercial interest that can stymie innovation. There is a great deal of debate about broadband access to the Internet in the United States. It isn't a highly competitive environment at all. Generally, maybe 60% of the US population, these are 2007 data from the FCC, had access to two broadband providers, typically a cable company and a telephone company. Thirty percent had access to one or the other but not both. And 10% had no access to broadband at all. Now, it's 2009.

The data are probably different. But you'll notice that the Obama administration has already allocated something like $7.2 billion to try to stimulate broadband access in the US. If you look elsewhere, you'll find that broadband access is more widely available. In Hong Kong, in South Korea in Seoul, in Sweden, in Tokyo, you can get gigabit-per-second service in Japan for about 8700 yen a month, which is a little over $100. It almost made me want to move to Kyoto. So, we don't do very well in that

regard. One of the worries that these, let's say, under-competed service poses is that the broadband carrier will say, "Well, I've built this facility and now I want to get as much return on the investment as possible, good old American tradition. And so, if there are other people who are competing with my applications, I will find ways of interfering with the competitor." So, if you're a cable company and you make money out of providing Internet service and video and somebody on the Internet service decides to provide video to your customer, you might say, "Well, I'm going to starve that competitor and make sure that their video isn't as good as my video." Now, I'm not saying that they would all do that. In fact, when legislation was threatened, they swore they would never do that. The legislation was called Net Neutrality, which meant nondiscriminatory access to the Internet so that you as a consumer could get access to any service that anyone offered on the Internet independent of the Broadband provider, well, except of course, you're depending on that broadband provider for raw access.

So, they said, "We'll never do that. It would be bad for us to do that." And then, as soon as the threat of legislation went away, a whole bunch of interfering activities surfaced. So, this is not to say that everybody is evil or bad or something like that. But I do believe that in an environment where you don't have enough competition to discipline the market, that you really do need to introduce regulatory constraints. I think we've all learned on the financial markets that the lack of regulatory constraints can lead to some pretty severe consequences. And we're living through that right now. So, that's another example of a commercial concern. It's fair to say that people who are commercializing the Internet need to be watched to make sure that they don't abuse that capability. And this is something which we really are smart to assure that there is some sort of awareness and oversight. Google has just gone through a lot of trouble to build a facility to help people analyze their broadband access to the Internet so they can tell whether or not there may be some constraints on it.

That way, with 10 million eyes or more than that, 150 million eyes looking, that should be a fairly inhibiting influence. Did you have some other specific commercial examples that you wanted to get to or did we cover them? OK. Other questions? Over here. I'll get one or two, OK. So, I was wondering how you got the title Chief Internet Evangelist and what it means. What does it mean? Actually, I didn't ask for the title Chief Internet Evangelist. They said, "What title do you want?" And I said, "How about Archduke?" It sounded pretty good to me. I don't remember who it was, Larry or Erik or Sergey said, "The last archduke was Ferdinand and he was assassinated in 1914 and it started World War I. So, maybe you don't want that title." And then, they said, "Considering what you've been doing for the last 30 years, why don't you be our Chief Internet Evangelist?" So, I thought, "Well, it sounded like a good title." So, on my first day of work, when I showed up at Google, I should have put that picture up, I was wearing the academic robe of the University of Balearic Islands. Now, if you know anything about Spanish academic robes, you know they are highly ecclesiastical in their appearance with lace sleeves and a cape, white gloves.

It's amazing. So, it was an interesting anthropological experiment to show up in this outfit. And I watched the younger people reacted differently. There are some people at Google that are a little older than 25. I mean, I know that the average age went up when I joined the company. And so, anybody who was over about 40 understood the humor in all this. The younger people couldn't figure out what was this all about. Some of them didn't even want to look, "I don't believe this." I'd go over and introduce myself and they would back away. One guy came up and said, "What are you selling?" So, that's how I got my title. I spend a lot of my time in these sorts of public relations activities, which I consider this to be one of, or on university campuses, talking to people about research problems that haven't been solved yet that would benefit everybody.

I'm talking to students about possible dissertation topics. I spend a lot of time visiting the engineering offices of Google. This is the Intellectual Bumblebee Face of the Chief Internet Evangelist job. Recognizing that somebody has a problem here, which is isomorphic to a problem that we solved over there is an interesting challenge, something I enjoy trying to do. And finally, I get an awful lot of proposals coming through the email of a form, "I just patented this. Would you like to license it?" "We just started a company. Would you buy us?" "Can I have a job?" "Do you have room for any apprentice Internet evangelist?" So, that sort of makes a good portion of my day. The last part is policy, whether it's domestic policy, international policy especially regarding Internet, or policy internal to the company. So, that sort of makes up my day. OK.

Did you say that was the last one? OK. Well, then, that's all the time we've got. I'm glad you have the chance to have some lunch.