Interviewee: Douglas Comer

Interviewer: Fan Yuanyuan from Oral History of the Internet project

Date: March 24, 2023

OHI: Okay, so today is March 24. We are glad to have our second interview with Professor Comer. So, let's start today's interview with your role in Internet Protocols.

DC: Okay. What would you like to know?

OHI: So, what's your role on the developing of the protocols? And can you describe your early involvement with TCP/IP?

DC: Okay, so TCP and IP, the basics had been defined in 1978 when I joined the project. And the questions that remained were all of the details: what happens in this special case? What happens in that special case? I know it's hard to imagine, but when you're writing down protocol specification, you have to write down not only the format of the messages that get exchanged. We have two computers talking. They're going to send a message. Here's the message format. Here's what all the fields mean. But you have to define what happens in every case. What happens if the packet doesn't get there? What happens if lightning strikes a near telephone wire and the packet has some of the bits okay, but some are not? So, all those special cases. And then there were special cases for routing protocols. What happens if two routing protocols start to change routes, and it turns out that they are doing conflicting things? How should you respond when an error occurs? All those questions were still being answered when I first joined the project.

So, I came along, and I had written a proposal to NSF. We were going to use TCP/IP. That was the agreement. And I had to study the protocols. And when I wrote the proposal, I wrote down, I'm going to use TCP/IP, and here's what I'm going to do. I'm going to tunnel it over X25, because we don't have enough money to actually hook up a brand new big wide network in the United States. So, all we were going to do was tunnel it. And I had to understand the protocols.

When the reviews came back from the NSF proposal, one of the reviewers wrote about my section. “It won't work.” Didn't give any reason. Just said it won't work. I mean, there was a little bit more to it than that, but basically no explanation. And the reviews are anonymous. So, I was an assistant professor at the time. I didn't have tenure yet. And of course, assistant professors are very worried about their reputation. Can you imagine what would happen if I got this warning, it doesn't work. I proceeded to do it. It didn't work. And then they came to judge me for tenure, and they say, well, not only did his stuff not work, but he was warned. So, as an assistant professor, I was very nervous. I better do this right.

So, one of the things that I did, I decided I should implement the protocols, try them out, figure out how all those special cases were supposed to work, figure out new special cases that hadn't been thought of. So, I didn't define the protocols, but I helped sort of debug them and hone them. And I kept asking Jon Postel, who had done most of the work on IP and ICMP and some of those protocols, I would ask him questions. Suppose, you know, a router has incorrect routes in it and we send a packet and it does this. What should be the correct action? What should be the correct response? Suppose you get an error message. What do you do with it? So, we kept going back and forth, talking about those things. As I said, I was very worried, so that's why I implemented it. I wanted to make sure that I understood it completely and understood why it wouldn't work.

Well, the good news is that it did work. The reviewers were wrong. It did work and it worked very well. We used it for many years until faster, better networking came along. But the good news for me was I learned a lot about TCP/IP protocols. I learned all the special little things that could go wrong. Of course, I had to learn how it was supposed to work in the correct case, but I learned a lot. So, that reviewer threatening me, telling me it doesn't work, actually caused me to dig in and learn a lot about the protocols.

And then, you know, years later, I would go to meetings and they would introduce me as the world's leading expert on TCP/IP. And the first time they did it, I said, no, I'm not really the world's expert. I'm sure some of these other people who helped define the protocols know more. Jon Postel knows more about IP and I'm sure Dave Clark knows a lot more about TCP. But after about the 10th time that I got introduced as the world's leading expert, I realized maybe I didn't know everything about any one protocol, but I knew a lot about all the protocols. And that's why I got the reputation. It wasn't just one of the protocols that I had worked on. I worked on a lot and how they work together.

OHI: So, who did you work with on TCP/IP?

DC: Well, I mentioned the late Jon Postel. He did work on IP, ICMP, on routing. Dave Clark did a lot of the work on TCP. Of course, Vint Cerf was the project manager. Who else? Paul Mockapetris did DNS. There were other people around. Those were the sort of major contributors to the protocols. Then there were sort of, I guess I would call them the politicians. You can't build a network unless you get government approval, unless you get rights to use things. And there were people around who took it upon themselves to work with the government, to work with DoD, to work out how we are going to build a network and how we're going to use some government money and how we're not going to use government money for other parts. There are certain restrictions. For example, the DoD, the Department of Defense, was funding a lot of the work. And they had what I consider to be some silly restrictions, but it was the Department of Defense. First of all, in the early days of the Internet, you were only allowed to use the ARPANET and the other backbone networks, the DoD-sponsored networks, for government business. You weren't allowed to use it for anything commercial. You weren't even allowed to use it for academic research. You had to justify everything.

And then there are some fun stories. One day, I was asked to talk to a person in the military about DNS, the Domain Name System. You know, that's the system where you put in google.com and you get back an IP address. Then you need that. So, every time you go to the Internet and look up a name, you get back an IP address. And that's what the protocols really use, the address, not the name. So, how should this work? Well, the military had made a rule that there had to be at least a backup for every domain name server. And that makes sense. It's always good to have redundancy. So, we have a domain name server, for example, at Purdue. Purdue.edu has its own domain name server. And we had to have a backup server. Not a problem. I believe that was fine. But then the military said, you need to have two backup servers. OK, it's a lot of work. Do you really think that one server is going to be so unreliable? Do you really think two of them are going to be really unreliable? So, you get two of them down at the same time? All right, maybe. And then they said, they need to be served power from different parts of the power grid. I said, well, I don't understand. I don't even know where what part of the power grid my university is on. What are you talking about? Well, what they were talking about very gently and gingerly was war. What if someone bombs, let's say, the power grid near Purdue University, then the server would be down. But what if the backup is, you know, 10 miles away, not too far away? It has the same power grid. So, if the power grid went out, then that whole area would go out. So, they made the rule that you had to have servers far away. So, I made agreements with other universities and we had we hosted their server and they hosted ours and we had backups far away.

But I was talking to the person from the military and I said Do you understand what the domain name system really does? All it really does is takes a name and changes it to an IP address. -Yes. So, suppose there's this catastrophic event. In fact, the military man had said catastrophic event with impact on population. He never said, you know, a bomb goes off very, very gently. And I said, suppose there is a catastrophic event with impact on population. Do you really think that if there's a catastrophic event near Purdue University, that people still need to be able to look up the name Purdue.edu and get an IP address? And I thought I had sort of proven the folly of what he was saying. Obviously, if Purdue isn't there, no one needs to look up Purdue.edu anymore. And he looked at me very gently. I asked my question. “Do we still need to be able to look up Purdue.edu, even if there's a catastrophic event and Purdue doesn't happen to be around?” He very, very seriously looked at me and said, “Yes.”

So. So, you know, dealing with the military was an interesting idea. It was something I wasn't used to. I was used to living in the academic world where you make logical arguments and people can have discussion, but they always have logical reasons for everything they're saying. And it was just a military rule. We're done. You don't have to understand it.

So, there were people in the Internet Project who talked to the military to the regulatory commissions. One of the other things is now, of course, everybody believes in data networking, but all the regulations for telecommunication have been set up for voice networks. And governments around the world had set up regulations, the US, Europe, all over, and they each had slightly different regulations. And now we propose to hook up a data network. We're going to have two computers sending data back and forth. And when we first started to do this, the regulatory people said, that doesn't make sense. Where are the voice channels? Where are you sending telephone voice? We said we're not sending telephone voice. So. It was a whole new world. How should it be regulated? And I know it sounds funny, but all the regulations were set up for a person on one end of the telephone. “Hello.” Talking to a person on the other end of the telephone. So. There were people who worked on that. Getting new regulations that would let data networking happen. And it sounds trivial, but it's a whole another case of regulatory stuff is not logical. It's been built up over many, many years. By the time we got there, it had been going for decades and they put in more regulations and more regulations and they had all these silly regulations about who could own the wires and where you could string wires. You know, if you're going to run cable or fiber to everyone's house to do data, you have to do it somehow on utility poles or underground or... And all the regulations were set up for phone companies. And now along came ISPs and they wanted to set up Data lines. They didn't have permission. So, there were people that worked on those problems. There were lots of people who contributed to the Internet in various ways to the commercial successful Internet.

OHI: So, that's kind of Internet governance as you described.

DC: It's related to Internet governance. Who owns the right to change the Internet? There were all sorts of legal and regulatory and unbelievable questions that came up. Who has the right to change the standards? Well, in the beginning, it was just, you know, a handful of people writing RFCs. And suddenly, in order to make it commercial, we're going to have commercial companies doing this. But wait, we know what happens when commercial companies come along. The companies with the most money have the most influence. They can volunteer people to go to the committees and write documents. They can pay their salary because they have a lot of money. Well, we didn't want that. How can we set up a system so that we don't just have the companies with the most money sending the most people and getting the standards changed to favor their products? Again, there were people who volunteered to take on that job of figuring out how to set up, in this case, it became the IETF. And the IETF organization had to make rules for how to how to do standards, had to expand it from a handful of people to thousands of people going to an IETF meeting, trying to vote on standards. Phill Gross did a lot of that work.

OHI: And so, for how long time did you work with TCP/IP team?

DC: for how long. Well, the first time I met them was in late 78 early 79 I don't know which. And things went along. we expanded a bit, we branched out a bit, but it was pretty much the same until 1989. In 1989, the internet had started to become commercial. And that's when the big changes started to occur. In 1989, The IAB the Internet Architecture Board was reorganized and Internet Society came along, the IETF became a separate entity doing the standards. And they decided that we had to have commercial representation. All these companies that were building what became the commercial internet had to have a way to influence things. So, the reorganization in 1989 changed it. In my opinion, that's when it stopped becoming research and switched over to production. Now when they first announced the change they came to an IAB meeting, and they said we're going to do this. We're going to make the change, what do you think? And each of us there had to say what we thought and I said, I think we need 10 years. If we had 10 more years of research, we could get these protocols really tuned up, we could get them just right. Everything would be just right. And in response, Phill Gross came across the room, pointed to me and said, Doug, you don't understand. Now there's money and money changes everything. And when he said it, I didn't quite get what he meant but within a year I did. The internet had taken off, and it grew so fast that it was hard to even keep track of what was happening.

In 1988, when I wrote my first book on TCP/IP protocols, I did a map for the cover. It was a map of the US, and we plotted every internet site. That cover had every internet site on the map, except for a couple of sites outside the US. By 1989, things started to grow. And by `92, when I was doing a revision of the book, next edition, I realized I could not possibly plot all the sites on the map. First of all, I'd have to make them so small that they'd be little tiny one pixel dots, because some of them were so close together, there'd be three or four. And if you have just a map on a cover of a book, there's not much geography there. They'd be so small, they'd just be lots and lots of dots. And that all happened just in a few years.

So, I still maintain contact with everybody, but I think that was sort of the end of the research era.

OHI: So, during the research era, I mean, before 1989, before it's commercialized, so what's the environment of the handful people, the research team?

DC: Let's see how to characterize them. First of all, the people who worked on it were incredibly smart. They were hardworking, but they were humble. They were never the kind of people who go around bragging, look what I did, I'm so great. They would always say, look, I've got it to go to this much. Now we need to work on the following. They would always look ahead and see if they could expand it. There was always a challenge.

Vint Cerf would always lay down the next challenge. At one point in time, he laid down a challenge very early, actually, he said, can we send video over the internet protocols? And when he said it, you have to understand the backbone network in the country was running at 50 kilobits per second. 50 kilobits per second. How could you send video over that? But that was the backbone. It had to be shared with all internet communication. It seems, you know, what an unbelievable challenge. But whenever we would achieve something, someone on the team would say, “I'm able to do this.” Vint would move the goalposts. “Now, can you do it more reliably?” Wait a minute, you know, the guy would say, “I just did it.” “Do it better.” Now, have a new application. What would happen if we changed the networks underneath? Suppose we add this new kind of network. Suppose we add a satellite network. It always had one more thing to work on. So, it was incredibly good people working very hard.

The other thing I'll say about the researchers is none of them said this is impossible, even though lots of other people said it was impossible. When TCP was first described, there were two or three people in the CS community who said, that's impossible, you can't do it. It seemed impossible at the time. If you had been back there, you have to remember it was a time when no one really knew what computer networking was all about. Everyone was a beginner. And you looked at some problems and given everything we know, we know how to do really slow networking. How do you do high speed networking? Nothing we've done in the past applies. The people who had studied telephone voice networks do regular standard old analog telephones, they looked at this and said, “no one will ever need it. No one will ever need high speed. And we already know how to do low speed. We've got it done.” And they had honed the voice network. They had gotten a very, very excellent engineering. All for voice. But it didn't work for data. Suddenly, a new world. How do you make things work for data? But the interesting thing about the TCP/IP team is that they never said, “well, in the past, we never did this. We can't do it.” They always said, “we'll see what we can do. We'll look, we'll think about it.” So, it was intimidating to be around these people because they were really good. They never gave up. They worked hard. Sometimes they would be working in the middle of the night.

I had a deal with Dave Clark. I had left the discard port on my VAX here at Purdue open. He was working on TCP. And he would hook up TCP on the machines and then send packets from his machine to my machine and measure things. And sometimes I would see him measuring it in the middle of the night. So, when you take really outstanding people, you give them a challenge, and then you give them the freedom to do it their own way. You can get incredible results.

And by the way, Vint Cerf had a real talent for understanding who to put on the team. If you looked at the people, they had very little in common. They didn't have the same educational background, they weren't from the same part of the country. They weren't all the same age. They weren't, you know, what was it? He had a way of selecting people to work on the project so that he got a really good team.

OHI: So, he is the manager we were talking about yesterday, as a good manager.

DC: Yes. He was an excellent manager. He set high standards. He constantly pushed, but he never pushed people to the point of breaking. He seemed to know how much people could do. So, when I got the tunneling working, his first question was, “it looks like the throughput is low, how are we going to get it higher?” Now, I know there are lots of people who, if they had just, you know, gotten one thing done, they'd want to bask in the glory and they'd want to get a lot of credit. But somehow he knew he could challenge me right away without offending me. And I knew, I knew right away, I knew it was low throughput. But I hadn`t said that. What's the next step? How can we do better buffering? How can we do a bigger window size?

OHI: So, was there any achievements or celebration where you have all those little achievements?

DC: Are there any big achievements? Well, there were certainly milestones along the way. We're now able to do this. We're now able to do that. But when you're living through it, it's incremental, you know. One person is working on routing while another person is working on TCP. So, when the person working on TCP has a breakthrough and now we can do this, the person working on routing is still working along and they will announce a breakthrough later. So, when you're living through it, there were a series of steps. It wasn't as if one day somebody came in and said, now I've solved the whole problem. We're done. It was a series of steps and a series of applications. One of the rules of TCP/IP is you should be able to use any application over TCP/IP, whether it's email or teleconferencing or anything at all. Some physicists sending huge amounts of data from an experiment to a researcher. Whatever it is, you should be able to do it. And gradually we added more and more applications. But you can say that each one of those was some sort of a milestone, but it was continuous.

Meanwhile, the hardware underneath was getting better and better, and there were more types of hardware. And every time a new type of hardware came along, the challenge was, how can we put IP over that hardware? The rule is TCP has to run over any kind of hardware and it has to allow any kind of application to use it. So, it's in the middle. And everything was changing all the time. By the way, while all this was happening, computers were changing. During the 1980s, we moved from having large departmental computers, a single computer for a whole department. So, we had moved from mainframes, a single computer for a whole organization, to departmental computers, the Digital Equipment Corporation VAX, for example, one computer per department. But in the 80s, suddenly, workstations, one computer per person, and a little later, PCs, one computer for everybody at home. So, computing was changing while the internet work was going on. Everything was in flux. The computers, the networks, the applications. So, you can't really expect at any point for someone to come along and say, no, I've solved everything. This is the big change. This is the moment. It's continuous.

OHI: So, you just went on and on to conquer all those problems to make it better, to suit with the new hardware, all the new developments.

DC: Right. So, here's another big change that happened. I was working on the CSNET project, and one of the questions that arose was, suppose we build a new backbone for CSNET. And eventually it became the Science Network, and it became the new internet backbone. But we already had the ARPANET as a backbone, and now we're going to have a new backbone. And the question arose, how do we do routing, IP routing, with two backbones? Now, these days, that sounds like a trivial problem. In fact, undergraduates in the networking course learn how to do that. Not a problem. But in those days, remember, this was all brand new. The internet protocols had grown up entirely around one backbone. And it turns out that putting a second backbone, a parallel backbone in, is a challenge, and we had to change our whole notion of routing. We had a meeting. All the people that were there were trying to figure this out. It took a long time, because it sounds trivial, but it was a major change. So, that was yet another thing happening that went into the mix. How do we do routing with multiple backbones?

OHI: So, in which year did you have to deal with multiple backbones?

DC: That was in, I think, 86, somewhere around 86, when NSF wanted to have a science network and they wanted to build a backbone, what became known as NSFNET. And at first, they were just going to replace the ARPANET. The ARPANET was the old DoD network. They were going to replace that. And then they realized they were going to have to have two backbones. And maybe someone else would build a backbone. Maybe a commercial ISP would build a backbone. So, that's when we started thinking about multiple backbones.

OHI: So, how long did you spend to resolve this problem?

DC: Well, I guess it was about two years before we came up with the first solutions. By 1987, we knew how to do a basic solution, but we needed to really work more to get a more general solution. See, there's also the question, first of all, you start with two backbones. Great. Let's build all the protocols that we need to do two backbones. But what if AT&T builds its own backbone or some other ISP? Now we're going to have more than two. So, we need to generalize. What if we have multiple countries connected to the internet and they each have a backbone? So, you know, it's hard to say we solved the problem. We incrementally solved it. Add this step, add that step, one more step.

OHI: So, what is the major milestones you have in the development of TCP/IP? Major milestone.

DC: Major milestones. Well, there was a major milestone, a flag day, in which the ARPANET and MILNET were divided. And we switched over and decided to use TCP/IP everywhere on the ARPANET. So, until then, there was an old protocol that had been used on the ARPANET. And there were old computers running the old protocol. And TCP/IP got good enough so that they could say, as of this date, we're going to switch over to TCP/IP. It's going to be the production system. So, that was a major milestone.

When we switched from host tables to the domain name system, that was a major milestone.

Well, of course, when we had multiple backbones, that was a major milestone.

I suppose that you could say the first time anybody successfully used a satellite in the Internet, was a major milestone as part of the development.

What else? I suppose the first time we successfully sent video.

The first time that people built high-level search services started with things like Gopher. How do you find people on the Internet? How do you find websites? Or in that day, it was file transfer protocol sites, FTP sites. How do you get to information? So, building a system that could do that.

Here's another milestone that happened fairly early. Hooking up Internet email to other email systems. That was a big thing.

I'm sure I'm leaving something out. But you see what I mean. There were a lot of steps along the way. I know one of my colleagues would say, well, the first time we had an Internet email exchange with, and he would name a country, France, Germany. You know, he kept counting how many new countries. He had a map of all the countries of the world that were on the Internet. And he would count how many countries the Internet reached. You know, at least one site one country.

OHI: So, all those milestones helped the TCP/IP and Internet to grow to what we can see nowadays.

DC: Absolutely. It was all about to become what we see now. If you had been back there at the time, the researchers working on it didn't really imagine what would happen. You know, we were all scientists and engineers. And we imagined people doing science and engineering. But we didn't imagine the sort of things that popped up once the general population had Internet access. Furthermore, in the 1980s, the Internet project was a small research project on the side. The telephone companies argued that it was unimportant. It's not really a network. It's just some people fooling around on the side. After all, the network, and they would say THE network with, you know, uppercase T-H-E, the one and only network, is the telephone voice network. And it was true. Telephone network went around the world. And here we were hooking up so few sites that I could put them on a map on a book cover. So, we were a little, a research project on the side. And I think all of us working on the Internet project felt that way. There were computer companies that had their own networks. IBM, Digital Equipment Corporation. If you wanted to hook together Digital Equipment Corporation and computers, you could buy a network from them. And they sort of believed, in fact, they sort of said it like, we really have the network for the important thing to connect computers together. We have the important network. That thing that they're doing over in the Internet research project, you know, that won't go very far.

So, it's hard to really understand how small the Internet project was. It was off in a corner. The telephone companies were big. The computer companies were big. We were small.

OHI: And with your help, it grew from nothing, from zero to numerous.

DC: Well, I've always said that the real reason it grew, because the people who did the protocols, the people who did the initial work, were so very good that it was an excellent technology. And it could be expanded easily. It could be adapted. It could run over any kind of network. They had done all the work in the beginning so that when the commercial world discovered it and tried to use it, it worked and it worked well. So, in some sense, everything was done before the commercial world discovered the Internet. It was done so well that it was usable, as is.

OHI: So, some Internet researcher has said that the reason why the Internet can grow so large is because it has almost 25 years development without the involvement of government and business world. It had a very deep foundation for the further development.

DC: Well, the 25 years depends on how you count the beginning of the Internet. Of course, there were computer networks before the Internet. The Internet project started when Cerf and Kahn proposed connecting multiple networks together. The Internet isn't just a computer network. It is a whole set of computer networks connected together. And that idea, the Internet idea, rather than just a network idea, didn't start until 73. That's when they wrote the paper proposing it. The work done between 73 and let's say 1990 is when the commercial world started picking it up. That work was the foundational work of the Internet. And it was incredibly short time period, but it was done by incredibly good people. And you're right, part of the part of the beauty of working in those early days was we didn't have any commercial constraints. It was a research project. Figure out how to do it well. Don't worry about who's going to make the most money. In fact, we were at a meeting in 1986. Dave Clark was there. I was there. Paul Mockapetris was there. A bunch of us were there. And Dan Lynch was a guy who was coming along. He was going to start a conference to bring together vendors and people using the Internet. So, it was going to be a commercial venture. And he said, you know, there's got to be money in this. And uniformly, all the researchers, including me, said, no, Dan, you don't understand. The telephone companies have their big network. There's a lot of money there. The computer companies have their networks. There's money there. This is a research project. There's no real money here. You can't make money off of the Internet. We were still living in research land.

OHI: So, how did you approach collaborating with other members of TCP/IP development team? So, what was the process like for designing and implementing new features or changes to the protocols?

DC: In those days, the IAB was really small. The Internet Architecture Board was only a few people. And all you had to do was take your idea there, propose it. And then everybody would either, you know, say thumbs up or thumbs down. And Vint would say, “yeah, why don't you try that?” Sometimes he would say, “well, we don't think it's going to work, but show me, you know, go off, try it, and show me.” It was really easy because it was a small group.

OHI: So, you all had quite a fairly relationship with each other.

DC: Yes, it was all very collegial, very informal. When I had a question about IP, I just send an email off to Jon Postel. I had a question about TCP, send an email off to Dave Clark. Get an answer. They were never too busy to answer. But remember, there's only a handful of people. So, it wasn't like a million people were sending them email questions, how are you doing this in TCP?

OHI: So, it's really a friendly and researching environment for everyone here.

DC: Absolutely, very friendly, very collegial, definitely high scientific standards. But everybody would, after a meeting, they'd all go to dinner together.

OHI: So, can you describe some of the key design decisions you made when you're developing IP or TCP layers? And how did these decisions shape the overall functionality and performance of the products?

DC: Okay, so let me say again, I didn't design TCP or IP. That work had pretty much been done by the time I got there, the overall design, the layers of TCP and IP. Why was it done? Originally, before version 4, the original design was TCP was one protocol. And it was going to be the protocol used across the internet. But then, Jon and others realized, maybe some applications won't want to use a stream-oriented protocol like TCP. Maybe they'll want to use individual messages. So, that's why Jon invented UDP. And once you have these two flavors of transport protocols, a message-oriented and a stream-oriented, it makes sense to have a common carrier underneath. Instead of just building everything into UDP and building everything into TCP, factor out the IP section and have a separate protocol and have both of them run inside IP. So, that's how we got the layering.

1978 is when IP was separated from TCP. So, that we could have IP as a fundamental layer and then the internet layer. And then TCP and UDP on top of it at the next layer. And either of them could use IP. So, the other thing that I'll say about all of the work that was done on TCP/IP and the internet, everything was done by finding a reason, building software, and then writing protocol standards. It's exactly the opposite of the way most standards bodies work. If you go to the phone companies, the phone companies around the world, they have standards body and everybody sits down, everybody writes a document. They say, here's how it ought to be done. And then they go out to build it. So, this was engineers building, testing and proving things before writing the standards. And of course, if you're an engineer, you know that works much, much better. Once the standard was built, you had to get two people beside yourself to implement the protocol according to the standard and then test to see that the implementations were interoperable. That proved that the standard was written well. So, the internet procedures were basically, here's how you do good engineering. You do good engineering by building it first and then writing the standard later, making sure that all the pieces work before you document how it should be done. And two, testing the document when you finally get a standards document, test it by having independent people implement the protocol and then see if they all work together. Because after all, that's the goal of a standards document for networking. We want to be able to have a standard so that everybody can build products according to the standard and then they all are able to communicate.

OHI: So, practice goes first and prove it by others and then it turns out to be a standard.

DC: exactly.

OHI: So, after the TCP/IP has been widely adopted by the world, what kind of impact has TCP/IP had on computer networking and how do you feel about your role played in its development?

DC: Well, let's see, TCP/IP has taken over everything, everything from IoT to desktop systems to enterprises around the world, organizations, governments use TCP/IP, everybody uses TCP/IP. In fact, if you go to modern data centers, cloud computing data centers, and look at what they're using inside the data center, it's all TCP/IP. So, how has it had impact? It has taken over everything. How do I feel about my role? Well, I'm very, very proud to have been part of it, small part of everything that happened. One of the things that I'm most proud of is my books on TCP/IP helped engineers around the world start to build, use, configure internet equipment and build internet ISPs and I constantly meet people who say, I used your book in such and such country.

OHI: So, speaking of textbook, you've written numerous textbooks and books about networking and the protocols. So, you've mentioned that your taking class of English has helped you to write the books. So, what's your advice to people who want to write those technical books?

DC: All right, well, I think it's actually fairly straightforward. First you have to learn how to write prose. Most engineers and scientists never learn how to write prose. They learn how to write mathematical proofs. They learn how to write equations. They learn how to write computer programs. But they never learn how to write prose. So, you have to learn something about writing prose. And it doesn't have to be that much. You just have to get the basics, understand how to put together a paragraph, how to put together sentences in logical order that explain something. Once you know how to do the basic, how to basically write prose, then what you do is this: 1, become an absolute expert on a subject. Lots of people make the mistake of thinking they can write a textbook by going out to Google and searching for whatever they need at the time. There are some very horrible textbooks. Lots of people try and write textbooks without really becoming an expert. So, step one, become an expert on the subject. And I don't mean just read somebody else's work or read somebody else's book. I mean, dig in and learn it really well, in my case, well enough to build the protocols. Understand what's going on. Number 2, think about your audience. One of the mistakes that people make with a textbook is they think it's important to write down every fact they know about a subject. And that's absolutely not true. What you have to do is imagine someone sitting in front of you and you are explaining the world to that person. So, pick a typical person.

For example, you might pick computer science or a computer engineering undergraduate student. They know a lot of mathematics; they know basics of computer programming. So, you can assume they know those things and you can use those as examples.

Pick somebody else, pick a grad student. Grad students want to know what's hard and what's easy, what research problems are left. So, when you're talking to grad students, you have to put down enough of the facts so that you've explained the situation, but then you have to write down the problem. Now how does this expand to a world of billions of people? How do you build a network that expands that far? Or suppose we have this much data to process. How can you possibly build a network to handle that much data if a given network link has a limit of how many bits per second you can send? So, you pose the problem so that they can understand the research issues.

General public. Maybe you want to pick the general public. I wrote a book called The Internet Book. It's aimed at non-scientists, non-scientists, non-engineers. They need to understand something about computer networking. Why do they need to understand it? Because I think it's important for the general public to understand what's going on so they don't get scammed. It's really easy to get scammed. There was a salesman who came to my door one day, didn't know who I was. This was years and years ago. “I'm selling internet service.” Oh really? “Oh yeah. If you buy our internet service, everything you do will be 100 times faster.” Now he didn't even ask me what internet service I had. But he started in with his claim. I said, OK, does that mean when I send an email, I press send and instantaneously it goes to anybody that I sent it to? Oh yes, he said. “Oh yes, there's no delay at all.” “When I am shopping and I click on a catalog page, does it mean that it comes instantly?” “Absolutely, absolutely, there's no delay at all, never.” Every question I asked this guy, he got wrong. I closed the door and I asked myself, what if you didn't know anything about the internet? What if you didn't know how it operated? What if you didn't know what this guy was selling? In particular, all the guy was selling was a higher speed connection from your house to an ISP. One link, nothing to do with sending things across the world. There are other links involved in that. But if you didn't know that, you wouldn't be able to fend off these horrible people who are selling you things that aren't true. Maybe you know that at one point, a group of people in a city was given gigabit internet. And when they were given gigabit internet, internet speeds were very slow where they were. And along came an ISP and said, we'll give you free gigabit internet. They were so excited. You should have heard the interviews. We're going to get this gigabit internet; it's going to be so fast. And then you should have heard the interviews after they got it. The local news station interviewed them. “What do you think of gigabit internet?” And remember, this is a long time ago. “It's not as fast as we thought. Sometimes when you click on a web page, it doesn't come up right away. Something's wrong.” All they got was gigabit from their house to an ISP. Of course, there are still other links in the internet. They didn't know that. So, anyway, that's the sort of thing that I think everyone needs to understand. And I wrote a book aimed at the general public. And I had to explain things without resorting to any equations, no mathematics, no, not even any real numerical stuff at all. Just analogies. So, think of an ISP selling you a high speed service. They're selling you more lanes on your driveway. But once you get out to the road, you still may find that the roads are congested during rush hour.

So, that's my summary. Be an expert. Pick an audience. And then write to that audience.

OHI: So, you should really know your audience before you pick up your pencil to write it.

DC: Yes.

OHI: And so you are a writer, a professor, and engineer, and some other roles. So, if you divide your career into phases, so how many will there be? And what role they are?

DC: So, what are they? Let's see. Starting at the top, of course, I'm a professor, so I do teach classes. And that's a big part of what I do. I write textbooks. I do research. And I do consulting.

I have spent a lot of time writing books because my books are successful, and they get used around the world. And I get positive feedback from lots of places. It's very nice to have someone tell me that they've just found my book and they're using it and it's helpful.

Consulting for many years, I spent lots of time consulting, explaining mostly to engineers and engineering companies, but also to enterprises how to use TCP. I explained to engineers how it works, how to build products. I've explained to companies how to architect their networks, what makes a good enterprise network. I've done other consulting for various groups. Now that I'm getting a little older, I've been trying to cut down the consulting.

You know, once you get to an age where people say, why aren't you retired yet? You know that you shouldn't be doing multiple jobs. So, I've been cutting down consulting a bit. Still writing books, still teaching, and I'm still doing research.

OHI: So, if you divide your career into phases, I mean, this year you've done this and that year you've done other things. So, how many phases there were?

DC: Oh, let's see. I guess you'd say the first part of my career was exploring. I got to Purdue, I was trained as a theoretician, but I did papers. I did research and papers on compilers, software engineering, programming languages, databases, operating systems, computer networks. I explored lots of areas. So, I'll put that down to the exploring phase, the early years.

Within 10 years after I started, it became obvious to me that computer networking was really going to become important, and particularly the internet is going to become important. And I wrote my first book. I had written a book on operating systems, but I wrote the first book on TCP/IP protocols. And that really was a transition. Suddenly I became famous in Silicon Valley and every place wanted to talk to me. Every company wanted to have me come visit. One day I was visiting Stanford and I went down to the Stanford bookstore and I was just looking around to see what was there, to see if they happened to have my book, brand new book. And then the bookstore fellow said to me, what are you looking for? And I said, oh, I was just looking to see if you happen to have any books on TCP/IP and the internet protocols. And he said, "oh, we have that book, but it's out of stock right now." "What book? I assume you have more than one book." Oh, he said, "yeah, we've got a little other book back there, but you don't want that one. You want to get the book by Comer." "Normally," he said, "I have hundreds of copies right here on the table by the cash register. People just come in and get them. I can't keep them in stock." And I was so elated. I was just I walked out of there on air. Can you imagine? writing a book, and by the way, the Stanford bookstore was Stacey's bookstore was the bookstore in Silicon Valley. Everybody went there. So, normally, my book would be sitting on the counter, but he couldn't keep them in stock. And that really started a new phase for me. Lots of consulting. I had done some consulting before, but suddenly everybody wanted me to come consult. So, that was a that was a big phase. I was giving tutorials on TCP/IP. People would call me in and I would give a two day explanation of for engineers. What the protocols were, how they work, how to implement them, answer questions. And let's say that lasted through the mid 90s or so.

And then I started sort of branching out again. Started exploring some new things. How can I do something better? There was a research topic that came along in the early 2000s. Software defined networking. So, I started doing research on that. I had done some research on ATM to show that TCP/IP didn't work well with ATM. Just exploring some new ideas.

And then companies came along and wanted me to come work for them. One year, three companies all invited me to come be a researcher at their company. And I chose Cisco. I thought I'll just try it. So, I went out and explored what it would be like to be VP of Research at a company. And that phase lasted through, I don't know, 2010. And now I'm back at Purdue and I've explored some new ideas, worked on some cloud computing, cloud networks, what kind of network would you use in cloud data center. I would say that I've done a little more exploring in the last years. One of the luxuries of being a professor, you get to move on to new things.

OHI: So, how is it like working in Cisco? It's a completely different environment with the university.

DC: It's a completely different environment. I didn't understand it at all going in. I had never worked in industry. It took me a while to sort of figure out what was going on. One of the differences has to do with corporate politics versus academic politics. In academia, people argue all the time and they don't agree on things, but usually if you can give a logical argument, you can win. If you can present enough data, enough evidence to support your case, you can win.

And I was at Cisco for a short time and I saw two people arguing. And I thought, wow, it was before a meeting started and they were in the corner, but everybody could hear them and they were arguing, no, we should do it this way, no, we should do it that way, we should do it this way. And I thought, you know, it's a lot like a university. We've got these two guys, they have different views of how things should be done. And then something happened that changed my view of industry forever. One of them said, last quarter I brought in $35 million. How many dollars did you bring in? And the other guy said, all right, I didn't bring in that much. We'll do it your way.

So, suddenly it was a different world. It wasn't who had the logical argument, who had the data to support their arguments. It was who had made the most money. So, it was indeed a foreign environment for me.

OHI: So, for all those projects you have worked through your career, what were some of the most significant findings or contributions that emerged from those projects?

DC: Oh, wow. Well, I think from operating systems, I already told you a few of the things, but understanding the relationship between hardware interrupts and processes in an operating system, the process abstraction was a significant contribution.

For the Internet, we had a meeting. In fact, the reason I wrote the first book was the meeting I told you about where we were trying to do multiple backbones. No one there seemed to know what to do. And during that meeting, it was a two-day meeting, I kept writing down principles that I saw. And I thought maybe I should write a paper, research paper, and put down the principles. But on the plane coming home, I got on the plane and I thought, you know, pretty soon everybody's going to want to know about the Internet. If I write a research paper, no one will read it. I should write a book. So, I came home and, the first chapter I wrote down was about routing. That became volume one of the TCP/IP books. So, what was significant about it? The first theorem that I had written down that I was going to write the research paper about was, in an Internet, if all routers contain a default route, there exists a routing loop. Now, if you take that theorem and you look at it, I'll bet you within a half an hour, most people who know anything about routing could come up with what they think of as a counterexample, because it doesn't seem to be true. But that is true. Now, what's significant about that theorem? These days, at the center of the Internet, we have what's known as a default free zone. The reason we have it, is because of theorem one that I never published as a paper. You must have at least one router know all routes explicitly. And I think that was a significant contribution that I made to the Internet.

OHI: So, I believe that's almost all questions I prepared for the interview. Do you think there are some things you should share with us that I failed to ask you?

DC: I think you've done a pretty good job of going through most things. Of course, there are little things along the way that I could point out, but you get the idea.

The Internet, in my opinion, is probably the most significant research project of the 20th century in computer science. It has affected more people, has changed the world in more significant ways than anything else. And I'm very, very thankful that I got to be part of it.

OHI: OK, there are still some questions about China. So, have you ever been to China before?

DC: Yes, I've been to China several times. I've given talks, mostly at universities, Tsinghua, other universities in China. I have some interesting stories about China. One of them is I was giving a talk on network processors at Tsinghua, and some Chinese engineers came up to me after my talk. And they didn't speak English. They had a translator, and the translator said, these are engineers who came from Shanghai to hear your talk. And I said, "hello" And I asked them, "are you working with network processors?" Because that was what the talk was about. It was talking about network processors. They said, no, we don't work with network processors. We use Xinu. And we wanted to come meet the person who invented it. We were so pleased to work with it. We wanted to come tell you we enjoyed it. And I said, that's very nice. I said, interestingly, I had visited Shanghai, and I had just flown up from Shanghai to Beijing a couple of days earlier. And they said, "oh, we didn't fly up. Oh, we couldn't afford to fly up. What we did was we rented a van, and after work, we took off. We took turns driving all night so that we wouldn't have to take much time off from work. And we wanted to hear your talk." And I asked them, "does your company have a branch office here in Beijing?" "No. So, you came just to hear me?" "Yes. Just to meet you. And now we're going to get back in the car, and we're going to drive all the way home. And then we're going to have to work at night to make up for it." I was incredibly impressed that they would go to that extent to come hear my talk.

OHI: So, they were coming to meet their hero.

DC: Yes.

OHI: So, even you don't have a hero of your own, but you've already a hero to others.

DC: I suppose I am, yes.

OHI: Okay. So, what's your impression on the development of the Internet in China?

DC: Well, interestingly, China sort of jumped from zero to infinity very fast. My impression is, of course, I'm not talking about the very rural areas. I'm talking about the major population areas. There was a time when people from China couldn't even spell Internet. They didn't know anything about it.

And then suddenly there were universities looking at how to do advanced things. I was surprised in the early 2000s when Chinese universities were already working on IPv6, you know, as a research project. So, they moved very fast. My impression is, they moved very fast from not being involved to suddenly…

OHI: And so what are your suggestions to China's Internet development?

DC: Oh, I haven't studied the Internet environment that much. I know that most countries, not just China, but most countries, including the U.S., have the situation where now commercial companies are taking care of most of the high-density population areas. But there are rural areas that are not served as well. In the U.S., there's constantly there's all these farms out in the middle of nowhere. And, of course, no ISP wants to put in miles and miles and miles and miles of fiber to reach a farm and have only one farm at the end of the fiber. The cost is prohibitive. So, I suspect that soon governments will do what they did with telephones. In the early days of telephones, many governments made the rule that we want everybody to have a telephone. And to do that, they passed special legislation, they did some taxation, they did various things to make sure that everybody was served with the telephone. I think governments around the world will start to do that with Internet service. They will find a way to incentivize companies because it is a huge loss of revenue to put in one long line for one subscriber. So, governments have to step in to make that happen. And I think that will happen everywhere.

OHI: And when will that happen?

DC: Well, the easy answer is that I don't do government things. I don't understand governments. Sometimes they move faster than I expect. Most of the times they move much slower than I expect. And things that I think ought to be easy, some governments find very hard to do. You know, it's not really, it's not about technology. It's not about how do you put in wires. It's all about convincing the population that this is a worthwhile thing to do. You know, convincing other government agencies that taking money for this is better than taking money for what they want. That's always hard.

OHI: And so in your opinion, what are some of the most pressing issues and challenges facing the Internet and the computer networking today? And how do you think they should be addressed?

DC: Well, you know, if you had asked me this in 1980, I would have told you a long list of technical things. I don't think there are many technical challenges left, really. Well, we can go faster. We can make it more reliable. We can add more redundancy. But all in all, the Internet works and it works really well. So, I don't think that the challenges are technical. I think they are social, legal, and I guess ethical. There are real questions about Internet companies holding information about individuals. How much information about you should a company be allowed to hold? Should they be allowed to sell it? Who should they be allowed to sell it to? You know, everybody knows that there are bad people in the world. There are always going to be criminals. Should criminals be allowed to buy detailed information about individuals? Would the criminals use it to extort money or threaten people? Or would political parties use it to hurt their opponents? You know, if you find out, if an Internet company can deduce how someone voted, and they can do that by knowing something about the votes in an area and how many people live in the area, and they've got all these crazy ways to guess who voted what way, and they brag about it, they say, “we can now guess the ways people vote based on what they buy in the store, what they, you know,” should that information be sold? Those are real questions. Those are hard questions. How can we stop the bad guys from doing bad things? I think it won't be long before we have cyber warfare kinds of things going on. I think that there will be, you know, if you're going to have war, you're going to have all sorts of technologies being used. And in the 20th century, it was all about dropping bombs. But in the 21st century, there are ways that you can really hurt a population without killing people. At least not without killing them outright. All you have to do is take out the power grid. So, there are already people studying those kinds of things. And I think the Internet has made those kinds of attacks much more possible and easy. Question, can we somehow stop the bad guys from using the Internet without stopping everyone from using it? So, I think those are the big questions that we're facing. Social, legal, ethical.

OHI: Yeah. So, these questions seem to be all that might be facing for those techniques after a while of development. Like this is kind of a question facing by Internet now, but later on, it might be a question facing by AI, especially after the development of chatGPT. There will be the same questions.

DC: One of the interesting things about AI and the Internet is when you're on the Internet, it's much harder to find out whether you're talking to a real person or an AI. So, that makes it easier to scam people. Yes, there are lots of questions that we're going to have to decide on going forward.

OHI: Yes, maybe you are talking to AI right now.

DC: Maybe you are.

OHI: Yeah. And so for the one last question, what kind of advice do you have for students and young professionals interested in pursuing a career in computer science and computer networking? And what kind of qualities do you think are the most important ones for success in the field?

DC: Let's see. First and foremost, the ability to solve problems. I know that sounds sort of simplistic, but let me say it a different way. The ability to figure things out is much more important than the ability to memorize. Of course, one has to memorize some basic facts, but the real quality you need to write a good computer program, to do a good research project, to understand how to build a complex computer network, the real quality that one needs is the ability to figure it out. Because most of these answers aren't found written down in pages of a textbook or written down in a document somewhere or even on the Internet. You don't just get the answer. It's a problem that has to be figured out and it has its own characteristics. It has its own parameters. And you have to look at the problem, ask very carefully, what are the possible approaches? What approach can be used here? What approach is optimal here? Under what circumstances does each approach work? And then you can build the thing that you're trying to build. And that ability to stand back and analyze and ask the question rather than just memorize the answer is the key. Students who do best can figure out new things. When they have a bug in a program, you write a computer program, almost no computer program starts out as perfect. You write it and you make mistakes. And then you have to figure out when it doesn't work. And that debugging process, that process of figuring out where the bugs are in the program, yet another example of you have to have the ability to go in and look at all the evidence and figure out what was going on. You can't find it in a book. You can't search the Internet and say, my program does the following crazy things. What's wrong with it? You won't get a good answer. You have to have that ability to ask yourself, let's see. Okay, let me try this or let me see if it does this. You know, what output does it get for this input? Oh, that's wrong. What output does it get for this input? That's right. What's the difference? Why would it give correct output here, incorrect there? Figure it out.

OHI: So, thank you very much for your advice. And so we'll put our interview online so everyone who is interested about it may have a chat about it.