Tim Berners-Lee

Transcription

Let me introduce myself.
I spent ten years working in Geneva at CERN. When I was working there, I was a software engineer. I was not a physicist, although I spent a few years studying physics. I was just there to try and make the place work. At CERN there were a lot of other people – in fact, you may meet people who are physicists there, and they will probably be doing software engineering, or engineering, because it's a huge project, and the number of people really actually doing physics at any one time has to be relatively very small.
I had a practical problem of getting information about what was happening, understanding the place I was in because it was such a heterogeneous place, and I had a practical problem too of disseminating information about my projects, and also working on projects which were really on the sidelines of what was officially sanctioned. Such projects tended to be done by volunteers, so I needed some sort of medium where volunteers could come in, work on a project, contribute to it, and go, without anybody spending time briefing and debriefing them.
The web became a necessity. I wrote a memo about it, suggesting that we needed such a thing, what… it was, 20 years ago this March. Inn March 1989, I wrote a memo suggesting it would be a good idea. A few other people thought it would be a good idea, but CERN had no way of actually letting me do it, because it wasn't a software company, it was… it is a particle physics lab. Some time later, somebody asked me, oh, what happened to that idea? I said "I sent you a memo", and they said "Did you? Could you send it to me again?" So I sent it again, in May 1990.
Eventually what happened – we were talking about innovation – eventually, my boss, Mike Sendall, he let me do it anyway. He let me buy one of those black Next cubes, one of those really neat computers, and he said: "We should experiment with this machine, this is supposed to be really good for developing new programmes, and if it's supposed to be really good for developing new programmes, why don't you try some arbitrary programme to test it on, why don't you try that hypertext thing you were talking about?"
So I fired it up and it said "New Project". It said "what's the name of the project", and I said "World Wide Web" – I'd had some discussions with friends about what I should call it. And so… and, that was how we produced the web.
From then on, it's been a question every day of trying to think what, of all possible things, I should be doing. Should I be trying to enhance the web? Should I be telling people about it? Should I be promoting it? Should I be finding students to work on it? Should I be trying to protect it? Should I be trying to think about its future – well, now more people have got involved, of course now we have organizations that do that.
In 1994, I went to... MIT, to be the base for the World Wide Web Consortium, because there was pressure – so many people felt the web was important, that

we should start an industry-based, industry-funded member consortium, in which lots of people (whether industry or academic) could come together and discuss the future of the web. So, at that point I moved to MIT. But we made sure that it was an international consortium, and that's the perspective I'm coming from.

I've been asked to talk about what the future looked like back then. In our conversation, on the telephone, when I was invited…

In fact, that's not such a silly idea as it may sound. Thinking back, thinking of that time in 1989, I could talk about a network information system of hypertext put on top of the internet and people would say, "Uh-huh yeah, ok, so what's the big deal?" I could show somebody a web page, and click on the page and go onto another web page, and they'd say, "Uh-huh, yeah, okay, so what's the big deal? Why should I take all my precious documents and turn them into your funny language so that I can click? 'Cos, you know I can click on things, I've got CD-ROMs, we have hypertext… hypertext has been and gone, it's sort of a yesterday thing".

So what was the problem? The problem was looking at the future. The future was difficult to imagine. When people said "click", they didn't mean – they didn't understand everything. A teenager today understands that "click" is going to a different place. Potentially going to any piece of information on the planet. People couldn't understand that.

The paradigm shift – I think the concept of paradigm shift was produced for that exact problem: the problem that you can't explain what it's going to be like because back then, as we looked forward, we didn't have the vocabulary for talking about the World Wide Web. People just didn't have the idea that... Except for a few people, nobody could imagine what it would be like. If I click on that link, it could potentially go absolutely everywhere.

What were things like? Well, technically, we were designing things in a very logical way. We were told to do a structured analysis, a structured design. Take the problem space that you have, the programme you've been asked to design, and divide it up. Divide the task into subtasks, just like the way you make a company. Divide it into subdivisions, and as you make those tasks, divide those into further tasks. For each one of those, you create a little group of people to produce the software. They produce modules. And all the modules connect together.

What we had were object-oriented systems, programming languages which allowed you to produce the structured design. And of course, when you do that, you get something which is hierarchical. You get something which has a root. You get somebody, of course, who is in control of the root.

Most of the systems we produced were supposed to be designed in that hierarchical fashion. That was good design; good structured analysis. Structured design was the watchword of the day. The object-oriented systems we produced came out of that.

Looking at the systems I could choose from as a means of storing documentation, for example, these tended to be also hierarchical systems, like file systems, or sometimes matrix-oriented systems. What these systems did is

they competed. At that point they competed for dominance. Each system would be a hierarchical system. It would have its roots, and it would try to compete...
One person would say, “If you want to store your documents on our system, then you have to convert them into SGML,” which is a documentation language, “And you have to store them on a mainframe, and you have to...” You have to arrange them in a hierarchy. If it’s a three-layer hierarchy, you have to give it a project, a category, and a folder. Another person would say “I don’t want to store my documents in SGML, so I made another system. We should all use it. It involves a new thing, networks of Macintoshes. All you have to do is write your document in Microsoft Word, store it on a Macintosh, and then we’ll have a matrix organization.” Now, because the hierarchical organization they have on the mainframe is horrible…
I’m making up these arbitrary examples; there are lots of examples that are like that. When you use the system, they tell you which machine to use, they tell you which software to use, they tell you which format to use, and worse, they tell you how to organize your brain. They tell you how to organize your thoughts. And that didn’t work.
My observation was that if anybody was going to solve this problem, to produce something, it had to be something which was unconstrained. It had to be something which did not require anything. In the end, I was going to ask everybody – for every piece of information that was important – to give it a URI, or a URL as we call it now.
That was asking a lot. If I’m just going to ask that, then I can’t ask for anything else. I can’t ask them to use a particular computer, or to use a particular piece of software or things like that.
The whole point about the web is precisely that: a very important aspect of its design – really important – is its universality. The web is this universal space; it should be usable on any system.
The question was… What was the epiphany? In a way, one of the epiphanies was the move away from hierarchical control – you couldn’t ask somebody to go to a particular centralized spot. The web ushered in an era of decentralized systems: decentralized systems in which you’re not doing a structured design from the top down. You’re designing, instead, a substrate. You’re designing a platform. You’re designing just an environment.
The web is just like the Agar you put in a Petri dish when you want to study a mould or the way a bacterium spreads, so you can model it mathematically. Similarly, the web, in its way… designing the web was like designing the stuff that you put in the Petri dish. The Agar. The stuff which is going to be food for an organism which is going to spread across it. The web is just a platform for creativity.
Fortunately for me, the internet had been designed with exactly that idea. The people who designed the internet protocols, they didn’t make any assumptions about who was going to use them. It’s this wonderful layered architecture.
Here I am today, talking to Telecom Italia. Every time I talk to telecom companies, it’s worth pointing out that layered architecture is still really important.

You guys are delivering bits. Keep delivering bits. Get really good at delivering bits, under all circumstances.

But the great thing about the design of the web, is that the internet would deliver bits from one programme to another. I could write the programme and run it on two machines, and I didn't have to worry about compatibility. Since those days, the number of webpages has gone from about 1 to about 10 to the power of 11. There are about ten billion webpages out there. And all that time, the programme – the way the programme talks to the internet, the way a web browser communicates with the internet – is completely unchanged. From the point of view of how my programme runs, I look at the internet and I see exactly the same thing.

But during that time, the actual way that access has gone, that's changed from going through one of those little acoustic coupling modems that you put on the phone [makes modem connection noises] all those things – remember those things? Those memories of what it was like to communicate? And now, it's going across fibre optics coming into your house. So the speed has gone up through orders of magnitude.. The provision of telecommunications infrastructure has changed hugely and the number of applications on the web has changed hugely. But not the layer in between – or in fact, two layers: there's the web layer, the HTTP layer, which has stayed the same. What it means to make a web server is basically the same; and below that, what it means to send an IP packet across the world is the same.

These two connection points… Keeping those layers independent is what was important… And this was one of the important ideas that was already out there. The whole iso-seven layer of model – I don't know whether any of you remember that? Who remembers the iso-seven model? Well, this model contained the idea of different layers, which is still very important. When we looked into the future, back in 1989, people imagined it would be a layered future. What they didn't realize is that the layering would also apply in the information space. They couldn't imagine the information space.

To recap, a very important thing was not to force things; to make this decentralized.

The reason the web can grow is that you don't have to come and ask me when you want to make a new web server. People were always asking me how many web servers there were… Well, of course, I kept a list, originally, of all the web servers. People would e-mail me. I have a copy of the list from when there were 26 web servers on it, back in 1991. Obviously, the important thing is this: it's something I can never know. No one knows how many there are out there. In fact, the weak point of the web at the moment is that you have to get a domain name. That domain name structure is in fact hierarchical. And so somebody is going to charge you for it. It is not completely decentralized.

Maybe in the future, we'll be making things even more decentralized.

What are the lessons for innovation here? Obviously, the fact that I was left to just go and do it anyway, means that we should make sure that the people working for us now should be given such space. We should provide an environment where people can do that. Where people can just go ahead with an

idea that they've had, and they have some time... Some companies give 10% time or 20% time. I don't know whether Telecom Italia does this, but leaving people in research, having people in research, giving them the time to go and do something which they haven't really thought out enough to explain to anybody, that they haven't thought out to put into a proposal… And even if there was a proposal they were going to write out, there would be absolutely no committee that would accept it. How do we produce this? We have to make a company, if you like, be an agar Petri dish, just like the web and internet has to be a place where people can do things that just seem like a good idea at the time.

What's happened since then? What are the fundamental changes? One of the questions was: "What are the big changes that have happened recently, over the last few years, that affect the way you look at the world?"

One of things now is that we look at the web very differently. I used to look at it just as a web of webpages. I'd think of it as something mostly uniform, like the sea: water molecules, water, all sloshing about. And now, we look at it, and it's big. There are as many webpages out there as there are neurons in your brain. The number of webpages is increasing and the number of neurons in your brain is going down. Sorry! Not only that: when you look at it, it's not uniform. It's got this very interesting structure. The network scientists tell us that it's scale-free; that there is no characteristic scale. When you look out there, there's no such thing as a typical web page. There are some webpages which generate huge amounts of interest, links pouring into them – ridiculous numbers of links. There are some webpages that compete with them. Then there is a very, very large number of webpages with a few links pouring into them, and there's a smooth curve, this thing they call the long tail. The web's got all these really interesting properties.

Now we're looking at the web... The web is a huge thing, like a huge organism. When we look at it now – over the last couple of years, we've realized that nobody has actually been looking at the web, at this big, very, very complex thing – we're looking at it as an object of study. We've started to talk about web science, like a few years ago we started to talk about cognitive science. It's a bit like how we started to look at the brain and study the brain, taking psychology and all those different things like neuroscience and biochemistry – all these different things and the ways we study the brain, we put them together to use them to understand the brain. We now have to do the same thing for the web because it's really important. It's a communications medium. It's no longer, to my mind, a network of webpages. It's a network of people. It's humanity. What makes a link? A person makes a link a reason. When somebody follows a link, they follow a link for a reason. It's psychology that makes you follow the link. We need to understand this. We need to take all kinds of disciplines and put them together to understand the web. We're not doing that yet, so that's what's going to be happening in the future.

The last thing that's happened is the people I work with… We've been thinking about the web as this big important thing. When you think about it as the communication medium for humanity, and suddenly humanity has this huge economic problem, you realize that you need... We need to understand it, we

need to do web science, and we need to make sure that the web serves humanity well. When you're thinking about how the web serves humanity well, you realize that research is not all of it, and development of technology is not all of it. When you talk about humanity, you realize that only 20% of humanity actually uses the web at all. We are talking about planning a Web Foundation to look at the web technology and connect it… look at what it needs to do to connect it to the rest of humanity. I know this is very close to the desire Telecom companies have to see telecommunications – mobile telecommunications in particular – push into the developing world. This is happening very rapidly. We're going to see a huge change in the next two years. We have a lot of hope that it'll have a dramatically positive effect on the development of companies. When we look at the technology, we see something which is designed by the developed world for the developed world, and that could be a problem. We would like this foundation to talk to all the people who understand the problems in developing countries and make sure that there is a good connection and coordination between things happening specifically for development and research, and the deployment of web technology.

That's the plan, and we plan next year (2009) to launch that. It's very exciting talking to people about the early stages of the planning for that.

Anyway, that's my viewpoint, looking back at what it looked like to look forward at that point, and how things have changed over the last 20 – almost! – years. Thanks. [Applause]

Venice Sessions
November 27th, 2008