

Note: Many of these questions are now answered in much more depth in my book, [Weaving the Web](#)

Frequently asked questions



I feel that after a while if I answer the same questions again, I will start answering rather mechanically, and will forget important steps, and after a while it won't make sense. So I have put a few answers from my outgoing mail in this list to save everyone time. But this list is (c) TBL so don't quote directly in the press without permission. Do feel free to quote for school projects. If you are doing a school project, I have a special [page of questions that people tend to ask for reports](#). Thanks.

- [Roles at MIT, W3C and Southampton?](#)
- [Spam - "Please stop sending it to me"! \(2002/4\)](#)
- [I have this great new idea - changing the world](#)
- [What's happening? \(2000\)](#)
- [What about peer-peer file sharing?](#)
- [General questions 1999](#)
- [General questions 1998](#)
 - [Q: I understand you invented the Internet....](#)
 - [Q: What is the difference between the Net and the Web?](#)
 - [What did you have in mind when you first developed the Web?](#)
- [Examples of early WWW hypertext](#)
 - [What was the first web page?](#)
- [Physics: why and influence](#)
- [W3C and standards \(1996\)](#)
- [What computer do you use?](#)
- [Spelling of WWW](#)
- [Why the //, #, etc?](#)
- [What were the first browsers?](#)
- [What influenced the design of the web?](#)
- [Why is your email address on my screen?](#)
- [Can you tell me more about your personal life?](#)

Roles at W3C, MIT and Southampton?

(2004) When I moved to MIT from CERN in 1994, it was to start the World Wide Web Consortium and act as its Director. Since then, my time has been split between the various tasks that involves, and, once the W3C was running smoothly, also forward-looking research into the future of decentralized systems like the Web and specifically the Web of machine-processable data, the "Semantic Web". In 2002, Steve Bratt joined W3C as Chief Operating Officer which made that part of my life much easier, and made W3C run very much more effectively. In 2004, I also accepted a part-time post at Southampton University in the UK. Southampton is one of the leading sites in Semantic Web research in the UK. While this will take a fairly limited amount of my time, I hope it will help collaboration between MIT and Southampton, and it will allow me to help Southampton and MIT to plan future research directions.

My roles as W3C Director and researcher at CSAIL continue. With Steve in the COO position, I can emphasize the technical side of my work such as that with the W3C Technical Architecture Group.

Spam - "please stop sending it to me!"

This question is one I have started (2002/04) getting more and more frequently. It is (ironically) normally sent automatically by people who are so enraged by spam (unsolicited bulk commercial email) that they try to find some way to protest to someone who will be able to stop the spammers. Most self-respecting Internet Service Providers will terminate their contract with anyone who abuses the service. So it is reasonable to take that approach. So these people generally set up a program to check through the email to find the web page it points to. Spammers are always after people's money, so there is some pointer to a web site which will (indirectly) take it. The plan is basically that these folk search the email message for pointers to web sites, and then search the domain name information to find out who is responsible for that domain. They then try to email someone "upstream" who will cut off the spammer's email access.

If you are one of these people, and you end up mailing me (timbl@w3.org) it is probably because I am one of the contacts for www.w3.org. Why do you find www.w3.org? Because you search the hypertext (HTML) email too simplistically and you found the XML namespace identifier which defines the HTML language. This is NOT a hypertext link. It identifies the specification of the language in which the email is written. The identifier in www.w3.org space is there because the World Wide Web consortium is the body which defines HTML. So w3.org has nothing to do with the sender of the spam. So if you vent your frustration on me, it just shows the software you are using, is broken.

By the way, I don't know whether the technique works. I have a horrible feeling that the spammers will just revel in the feedback they get from this. But I don't know. Check out abuse.net from which I have got some of these. I am *not* mad at you for trying to stop spam. I am mad at those who spammed you. For the record,

- I hate spam.
- I and my staff waste a significant amount of time deleting spam.
- I feel that those who make their living sending spam damage the whole community for the sake of greed.
- The lie "you are only getting this because you have been signed up for it" makes me sick.
- My handling charge for unsolicited bulk email is \$10,000 plus and recovery costs and legal fees.
- I would support legislation which made it illegal to falsify or omit the full identification of those responsible for any commercial mail.
- I believe that the falsification of email headers for one's own gain and other's loss is fraudulent.

See also:

- [W3C webmaster FAQ on this issue](#)

I have a great idea -Changing the world

Q: I have been working for a long time on a very special and new idea which will revolutionize computing. Can I tell you about it?

This is the most difficult answer to have to write. I am sorry to say that I can't give your vision of the future the time it would take to compare it with existing architectures and point out the similarities and dissimilarities. I get quite a few requests like this. What I would humbly suggest (and only suggest) is that you do that comparison piecemeal, and - while keeping your vision in mind -- try to find the first piece to implement in the move toward what you envisage. The world is can only really be changed one piece at a time. The art is picking that piece.

When you have, then use the web to find out who is working in that area. Acquaint yourself with the vocabulary they use for talking about it. Find a way of explaining your novel idea in their terms, after you have understood why it has not already been done your way. Then suggest that change. If it is an idea in computing, then you may want to write the code to show that it works first.

(I didn't find lots of people willing to get excited about the idea of the web. They quite reasonably asked to know why it was different from the past, or other hypertext systems. In retrospect, it was mainly that the decentralized database is removed, allowing the system to scale, but allowing for dangling links. But it took a long time for that to surface as the novelty.)

What's happening? 2000

Q: What sort of technology should the forward-looking geeks in my company be looking at?

You probably have a lot of people using XML by now. You should have someone looking at the next level - [RDF](#). Tell them not to worry about the syntax, but check out the model. This is a question of looking the data your company is storing and transferring, and making sure that it can be represented in that simple circles-and-arrows RDF way. This is very simple. An important trick is that you use URIs to identify the arrows as well as the circles. Doing this homework will ensure that you have a well-defined data model, which will allow you data to be combined, merged with any other RDF-model data. It will mean you will be able to multiply the power of separate application areas by running RDF queries and new RDF-based applications across both areas. It will mean that you will be there with talent which understands the basic model as the [Semantic Web](#) becomes all-important.

Other things to watch: [SVG - Scalable Vector Graphics](#) - at last, graphics which can be rendered optimally on all sizes of device. The user interface world is rapidly becoming competent at [voice input and output](#) an W3C has standards in that area coming along. [XML Signature](#) will let you to digitally sign XML documents - find out how. But in general, always check out the [W3C home page](#) for what's new.

If your company/organization/self is a W3C member, then your Advisory Committee representative has the task of understanding everything which is happening in W3C, and everything in your company, and seeing where they should be introduced.

What do you think of peer-peer file sharing? (2000)

Q: What do you think about the peer-peer file sharing technology which allows people to copy copyrighted information so easily? A: The issue is not simple - so I try to put my thoughts into a few words. In general, the way to make a sane society is to enact and enforce laws rather than to ban a given generic technology. (I would make the exception for things which are specifically designed to harm such as guns and nuclear bombs.) That said, one can make technology which supports our social

and legal frameworks better if one does it deliberately. One of the four domains of the World Wide Web Consortium addresses [Technology and Society](#) for this reason. For example, in this case, I think we really need standards for encoding the broad licensing terms of material so it can be read and handled automatically. Then we can see, when the technology allows one to see whether information is free or for pay, whether there is still a substantial problem of theft. The basic idea of forwarding copies automatically between machines is a technical optimization of the distribution protocol which is very useful and should not of itself be disallowed just because it -- like many powerful things -- can be abused. I'd point out that some ostensibly "peer-peer" systems are centralized system in fact, allowing centralized control and profit by the central server's owners. Other systems are really decentralized, having no central server. These are like internet news groups which have been around for ages and which raised similar issues.

General Questions, 1999

Q: What is your opinion on 'Cyber Squatting' for domain names? (-Lia Kim)

Domain names are a scarce resource - one of the few scarce resources in cyberspace. I have little sympathy for those who scoop these up with the hope of speculating on their value. This is not one of the most helpful activities on the net. There are those who use their energy for the purposes of furthering the technology or the content or the world in some way, but just sitting on a domain name without using it in order to cash in later does not seem to me a constructive .

General Questions, 1998

Q: I understand you invented the Internet....

Sorry, not me! I was lucky enough to invent the Web at the time when the Internet already existed - and had for a decade and a half. If you are looking for fathers of the Internet, try **Vint Cerf** and **Bob Kahn** who defined the "Internet Protocol" (IP) by which packets are sent on from one computer to another until they reach their destination. See:

- ["Cerf's Up" : MCI WorldCom on technology](#) with profile and FAQs by Vint, who currently works for MCI.

Vint explains the timing:

"The DESIGN of Internet was done in 1973 and published in 1974. There

ensued about 10 years of hard work, resulting in the roll out of Internet in 1983. Prior to that, a number of demonstrations were made of the technology - such as the first three-network interconnection demonstrated in November 1977 linking SATNET, PRNET and ARPANET in a path leading from Menlo Park, CA to University College London and back to USC/ISI in Marina del Rey, CA."

David Clark, of MIT's LCS, is another one I can point to who put in the work in the 1970s which made the Web possible in the 1990s.

Vint Cerf and Bob Kahn used, in making IP, the concept of packet switching which had been invented by Paul Barran.

It is also good to mention the Domain Name Service upon which the web relies heavily. The protocols which make the DNS work were pioneered and standardized by [Paul Mockapetris](#).

Q: What is the difference between the Net and the Web?

The Internet ('Net) is a network of networks. Basically it is made from computers and cables. What Vint Cerf and Bob Kahn did was to figure out how this could be used to send around little "packets" of information. As Vint points out, a packet is a bit like a postcard with a simple address on it. If you put the right address on a packet, and gave it to any computer which is connected as part of the Net, each computer would figure out which cable to send it down next so that it would get to its destination. That's what the Internet does. It delivers packets - anywhere in the world, normally well under a second.

Lots of different sort of programs use the Internet: electronic mail, for example, was around long before the global hypertext system I invented and called the World Wide Web ('Web). Now, videoconferencing and streamed audio channels are among other things which, like the Web, encode information in different ways and use different languages between computers ("protocols") to do provide a service.

The Web is an abstract (imaginary) space of information. On the Net, you find computers -- on the Web, you find document, sounds, videos,.... information. On the Net, the connections are cables between computers; on the Web, connections are hypertext links. The Web exists because of programs which communicate between computers on the Net. The Web could not be without the Net. The Web made the net useful because people are really interested in information (not to mention knowledge and wisdom!) and don't really want to have know about computers and cables.

Questions below derived from those asked by Taiwan's *Commonwealth* magazine

Q: What did you have in mind when you first developed the Web?

From [A Short Personal History of the Web](#):

The dream behind the Web is of a common information space in which we communicate by sharing information. Its universality is essential: the fact that a hypertext link can point to anything, be it personal, local or global, be it draft or highly polished. There was a second part of the dream, too, dependent on the Web being so generally used that it became a realistic mirror (or in fact the primary embodiment) of the ways in which we work and play and socialize. That was that once the state of our interactions was on line, we could then use computers to help us analyze it, make sense of what we are doing, where we individually fit in, and how we can better work together.

Q: Do you have had mixed emotions about "cashing in" on the Web?

Not really. It was simply that had the technology been proprietary, and in my total control, it would probably not have taken off. The decision to make the Web an open system was necessary for it to be universal. You can't propose that something be a universal space and at the same time keep control of it.

Q: Are you happy with what the World Wide Web has turned out so far?

That is a big question. I am very happy at the incredible richness of material on the Web, and in the diversity of ways in which it is being used. There are many parts of the original dream which are not yet implemented. For example, very few people have an easy, intuitive tool for putting their thoughts into hypertext. And many of the reasons for, and meaning of, links on the web is lost. But these can and I think will change.

Q: What do you think of the commercial turf wars going on the Web?

There has always been a huge competition to come out with the best Web technology. This has followed from the fact that the standards, being open, allow anyone to experiment with new extensions. This produces the threat of fragmentation into many Webs, and that threat brings the companies to the W3C to agree about how to go forward together. It is the tension of this competition and the need for standards which drives W3C forward at such a speed.

Q: What should they lay person be aware of as the Web evolves?

We should all learn to be information smart: to understand when a Web site, or a piece of software, or an Internet Service provider plan, is giving us biased information. We should learn to distinguish quality information and quality links. As technology evolves, and machine-understandable information on the Web becomes available, we should be aware of the sudden changes which large-scale machine processing might have on our businesses.

Q: How could the Web be a more interactive, creative medium?

Nothing can be perfect, but the Web could be a lot better. It would help is we had easy hypertext editors which let us make links between documents with the mouse. It would help if everyone with Web access also had some space they can write to -- and that is changing nowadays as a lot of ISPs give web space to users. It would help if we had an easy way of controlling access to files on the web so that we could safely use it for private, group, or family information without fear of the wrong people being able to access it.

Metadata

Q: You talked about the need for a metadata language. Can you tell us laymen what it is?

"Meta" is used with anything which is about itself - so a metabook would be a book about books, and metadata is data about data. On the Web, this means all sorts of information about information: its ownership, authorship, distribution rights, privacy policy, and so on. These needs are driving us to make ways of putting information on the web designed for computers to be able to understand. Web pages at the moment in HTML are designed to be read by humans. In the future, some Web pages will be in "RDF" -- Resource Description Framework. This will be read by computer programs which will help us organize ourselves and our data and possibly everything we do.

Privacy

Q: Are you worried about privacy on the Web?

When it comes to privacy, my personal view is that the consumer needs some legal or regulatory protection by default. The W3C has a project called "P3P" for privacy which will allow a user to control if and how information is given away to a Web

server. P3P will allow Web sites to specify their privacy policy and users to automatically be warned about sites whose policies they don't like. See the [P3P](#) project.

ECommerce

Q. Do you shop online? What do you think about the E-Commerce?

Yes, I buy a lot of things online myself. I think that Web shopping as it is is only the tip of a huge larger change which will come when I can find things and compare prices automatically, and when electronic financial instruments are commonplace.

Web and Education

Q: Peter Drucker has predicted that information technology will bring about the demise of the university as currently constituted. Do you share this view? What changes will the Web help bring to education?

I hope that educators will pool their resources and create a huge supply of online materials. I hope much of this will be available freely to those especially in developing countries who may not have access to it any other way. Then I think we will see two things. One will be that keeping that web of material up to date will take a lot of time and effort - it will seem like more effort than creating it in the first place. The other is that we will see how essential people, and their wisdom, and their personal interactions, are to the educational process. A university is a lot more than its library.

The effect of the Web on how we work

Q: How do you see the web shape the new, knowledge-based economy?

The Web is simply a name for all the information you can get online. So it will be the abstract place where the knowledge-based economy happens. Already the W3C staff team works with three international sites, many offices, and several people working from or near home. The Web will open up new forms of business altogether, and make us rethink the way we run existing businesses. It can turn bureaucracy over to machines, and let people get on with the creativity. It will help us see where we each fit, with our own experience, talents and passions, among the millions of other people and theirs. It can help us work together more effectively, remove misunderstanding, and bring about peace and harmony on a global scale. But it can only do these things if we learn to use it wisely, and we think very

carefully about both the technology and the laws we make or change around it.

Examples of early WWW hypertext

Q: What was the first web page?

A: Apart from local "file:" URLs on my machine (which was the first browser as well as the first server), the first http one (end of 1990) was basically

<http://nxoc01.cern.ch/hypertext/WWW/TheProject.html>

An alias was made so that this was later known as

<http://info.cern.ch/hypertext/WWW/TheProject.html>

It is not now (alas) served but a later (1992) copy of the original pages exists at <http://www.w3.org/History/19921103-hypertext/hypertext/WWW/TheProject.html>

Q: Do you have any examples of the early Web which we could compare with the current Web?

A: (1997): I don't have a very early 1990,91 snapshot but there is a snapshot of our web as of November 1992, much of which dates from earlier. (Some pages for some reason don't work with Netscape 3.0 for some reason it doesn't the old HTML for some reason or perhaps it just has a bug. They do work with Internet explorer 4.0)

There is a [list of design issues](#) and a [trip report on the 1990 European Conference on HyperText](#) and a [note on the "state of standardization"](#) (!) and an example of the use of the web as a collaborative tool in some [shared notes on the topology of the web](#) wrote and Jean-Francois Groff annotated .

The pages will look much the same as they did originally, although the actual style sheet I used as a default with the original browser/editor you can see converted approximately into a CSS style sheet if you read my [Style Guide for Online Hypertext](#) with a CSS-compliant browser such as IE 4.0.

Some of the links in the historical stuff have been accidentally saved (much later) incorrect absolute links -- if you really want to follow them you can see where they ought to have gone by stripping of the prefix.

Physics: why and influence

(Based on replies to David Brake, "New Scientist", 1997/9)

Q: Why did you study physics?

A (1997) : My parents are both mathematicians: they actually met while working on the Ferranti Mark I, the first computer sold commercially. My mother has been dubbed the "first commercial computer programmer" as she went with the machine when it was installed on the customer site. So we played with 5 hole paper tape, and learned to enjoy mathematics wherever it cropped up, and learned that it cropped up everywhere.

Later on, my hobby was electronics. When I left school, obviously I was going to do something in maths, science and/or engineering. Emanuel school was programmed to send people to Oxford, where the subjects very narrow. I took physics thinking it would be a sort of compromise between maths and electronics, theory and practice. It turned out not to be that, but to be something special and wonderful in itself. Physics was fun, and in fact a good preparation for creating a global system. In physics, you learn to think up some simple mathematical rule on a microscopic scale, which when scaled will explain the macroscopic behavior. On the Internet, we try to dream up computer protocols which when extrapolated to the macroscopic will produce an information space with properties we would like.

Q: Why didn't you stay on to do a PhD in physics?

A: After undergraduate physics, you have a reasonable training in logical thought and common sense, an ounce of philosophy and not enough maths to study physics. I didn't meet anyone who was actually doing physics research at the postgrad level and was really excited about it. I might have been more tempted to take a PhD if I had had a role model who did have that excitement.

What seemed much more exciting was the possibility of that electronic hobby really taking off. The microprocessor was just hitting the world. I got an early M6800 evaluation kit, and built myself a rack-based 8-bit system. I had already while in college slowly put together a display unit out of an old TV, bits of TTL logic and junk from the Tottenham Court Road. I joined Plessey Data Systems: of the telecom companies doing the "milk round" interviews the Poole (Dorset) site won hands down in terms of the sea and the countryside!

Those who got into designing microprocessor hardware and software then rode the crest of the wave of the deployment of microprocessor technology. Compared with

TTL, a microprocessor gave one that feeling of unbounded opportunity which had everyone excited. Later, the thought of building an abstract information space on top of it all had the same sort of kick.

W3C and standards, 1996

Q: What role does the [W3C](#) play in setting standards?

A: (1996) W3C's mission is to realize the full potential of the web, by bringing its members and others together in a neutral forum. The W3C has to move rapidly (time measured in "web years" = 2.6 months) so it cannot afford to have a traditional Standards process. What has happened to date has been that W3C has, by providing a neutral forum and facilitation, and also with the help of its technically astute staff, got a consensus among the developers about a way to go. Then, this has been all that has been needed: once a common specification has been prepared and a general consensus among the experts is seen, companies have been running with that ball. The specifications have become de facto standards. This has happened with for example HTML TABLES, and PICS. Now in fact we have decided to start using not a full standards process, but a process of formal review by the W3C membership, in order to draw attention to specifications, and to cement their status a little. After review by members, the specifications will be known as W3C process.

(See [process of review](#))

Q: What do you make of the branding attempt of companies, by putting little icons on their home pages saying, "best when viewed with Microsoft Explorer, or Navigator?"

This comes from an anxiousness to use the latest proprietary features which have not been agreed by all companies. It is done either by those who have an interest in pushing a particular company, or it is done by those who are anxious to take the community back to the dark ages of computing when a floppy from a PC wouldn't read on a Mac, and a Wordstar document wouldn't read in Word Perfect, or an EBCDIC file wouldn't read on an ASCII machine. It's fine for individuals whose work is going to be transient and who aren't worried about being read by anyone.

However, corporate IT strategists should think very carefully about committing to the use of features which will bind them into the control of any one company. The web has exploded because it is open. It has developed so rapidly because the creative forces of thousands of companies are building on the same platform. Binding oneself to one company means one is limiting one's future to the innovations that one company can provide.

Q: What role do standards play in today's hyper competitive, and fast-changing marketplace?

A: Common specifications are essential. This competition, which is a great force toward innovation, would not be happening if it were not building on a base of HTTP, URL and HTML standards. These forces are strong. They are the forces which, by their threat to tear the web apart into fragmented incompatible pieces, force companies toward common specifications.

Q: Is it overly ambitious to think standards can be set and adhered to? Are they a relic of a kinder, gentler era?

A: Do you think that incompatibility, the impossibility of transferring information between different machines, companies, operating systems, applications, was "kinder, gentler"? It was a harsh, frustrating era. The Web has brought a kindness and gentleness for users, a confidence in technology which is a balm for IT departments everywhere. It has bought new hope. As a result, great things are happening very fast. So this is a faster, more exciting era.

Companies know that it is only interesting to compete over one feature until everyone can do it. After that, that feature becomes part of the base, and everyone wants to do it in one, standard, way. The smart companies are competing on the implementations: the many other aspects such as functionality, speed, ease of use and support which differentiate products.

June 96

Machinery

Q: What sort of computer do you use?

A: (2002) A titanium G4 Powerbook running OS X and under X11 fink -installed stuff including Amaya. I use a Nokia bluetooth 3670 tri-band GSM phone which has a low-res camera. The OS X operating system is very similar to the NeXTStep operating system on which I developed the WorldWideWeb program originally.

Spelling of WWW

Q: How in fact do you spell World Wide Web?

A: It should be spelled as three separate words, so that its acronym is three separate "W"s. There are no hyphens. Yes, I know that it has in some places been spelled with a hyphen but the official way is without. Yes, I know that "worldwide" is a word in the dictionary, but World Wide Web is three words.

I use "Web" with a capital W to indicate that it is an abbreviation for "World Wide Web". Hence, "What a tangled web he wove on his Web site!".

Often, WWW is written and read as W3, which is quicker to say. In particular, the World Wide Web consortium is W3C, never WWWC.

Q: Why did you call it WWW?

A: Looking for a name for a global hypertext system, an essential element I wanted to stress was its decentralized form allowing anything to link to anything. This form is mathematically a graph, or web. It was designed to be global of course. (I had noticed that projects find it useful to have a signature letter, as the Zebra project at CERN which started all its variables with "Z". In fact by the time I had decided on WWW, I had written enough code using global variables starting with "HT" for hypertext that W wasn't used for that.). Alternatives I considered were "Mine of information" ("Moi", c'est un peu egoiste) and "The Information Mine" ("Tim", even more egocentric!), and "Information Mesh" (too like "Mess" though its ability to describe a mess was a requirement!). Karen Sollins at MIT now has a Mesh project.

Why the //, #, etc?

(2000/09) When I was designing the Web, I tried to use forms which people would recognize from elsewhere.

What is the history of the //?

I wanted the syntax of the URI to separate the bit which the web browser has to know about (`www.example.com`) from the rest (the opaque string which is blindly requested by the client from the server). Within the rest of the URI, slashes (/) were the clear choice to separate parts of a hierarchical system, and I wanted to be able to make a link without having to know the name of the service (`www.example.com`) which was publishing the data. The relative URI syntax is just unix pathname syntax reused without apology. Anyone who had used unix would find it quite obvious. Then I needed an extension to add the service name (hostname). In fact this was similar to the problem Apollo domain system had had when they created a network file system. They had extended the filename syntax to allow `//computername/file/path/as/usual`. So I just copied Apollo. Apollo was a brand of unix workstation. (The

Apollo folks, who invented domain and Apollo's Remote procedure call system later I think went largely to Microsoft, and rumor has it that much of Microsoft's RPC system was).

I have to say that now I regret that the syntax is so clumsy. I would like `http://www.example.com/foo/bar/baz` to be just written `http:com/example/foo/bar/baz` where the client would figure out that `www.example.com` existed and was the server to contact. But it is too late now. It turned out the shorthand `//www.example.com/foo/bar/baz` is rarely used, and so we could dispense with the `//`.

What about the "#"?

So, I needed something to separate the document (resource) from the thing (fragment) within that document (or view of that document). In a snail mail address in the US at least, it is common to use the number sign for an apartment number or suite number within a building. So 12 Acacia Av #12 means "The building at 12 Acacia Av, and then within that the unit known numbered 12". It seemed to be a natural character for the task. Now, `http://www.example.com/foo#bar` means "Within resource `http://www.example.com/foo`, the particular view of it known as bar".

It turned out later that in fact another hypertext project of some sort in IBM, and Doug Englebart's NLS system had both independently use "#" for this purpose. So there is something to choosing a character for the way people think of it.

Ray Tomlinson, who invented email, tells a similar story of many years earlier choosing the "@" for email - it made linguistic sense, as "at" was the english preposition which typically connects a person and their address. Hence `ray@example.com` and so on.

What were the first WWW browsers?

WorldWideWeb

I wrote in 1990 the first GUI browser, and called it "*WorldWideWeb*". It ran on the NeXT computer. (I much later renamed the application *Nexus* to avoid confusion between the first client and the abstract space itself).

- [Screenshot of WorldWideWeb](#) taken for a CACM article. *By this time it had color and inline images. The original 1990 version 1.0 would have looked identical except the book icon and the CERN icon would have been in separate windows - and the whole*

thing (like NeXT at the time) would have been in gray scale! The screen shot shows me making a link from "Atlas" in a list of experiments to some marked page. Look - no typing URLs, no <angle brackets>!

WorldWideWeb was a graphical point-and-click browser with mode-free editing and link creation. It used style sheets, and multiple fonts, sizes, and justification styles. It would download and display linked images, diagrams, sounds animations and movies from anything in the large NeXTStep standard repertoire.

(Some have asked for pointers to the source code. I have found an [archive directory](#) including the [HyperText.m](#) module which was the basis for the hypertext functionality. This code, like all my WWW code and later W3C has always been publicly available. This archive has the code, though the libwww code modules are soft links which no longer work. I haven't tried recompiling and linking it for years - so it is probably of historical interest only)

- [More about the WorldWideWeb application](#)

Viola

Pei Wei, student at U.C. Berkeley (not Stanford, as incorrectly reported earlier in a typo here), then wrote *ViolaWWW* for unix, based on his *Viola* language; some students at Helsinki University of Technology wrote "Erwise" for unix; and Tony Johnson of SLAC wrote "Midas" for unix. Pei Wei has passed though history unnoticed among others whose work is not mentioned in the histories, even though there was a year or so when Viola was the best way to browse the web, was the engine driving the installation of new servers, and the recommended browser at CERN for example.

Many people, incidentally, saw the Web for the first time by telnetting into `info.CERN.ch`, which gave them a crude but functional line mode interface. This was the second browser, a text-based browser, called the "line mode" browser, or "www", and written by CERN student Nicola Pellow. Many people imagined that that was all there was to the web. As one journalist wrote "The Web is a way of finding information by typing numbers" as links were numbered on the page. It was only in the community of people who use NeXT computers that the Web could be seen as a point-and-click space of hypertext.

Where does Mosaic fit in?

As I understand it, Marc Andreessen at NCSA was shown *ViolaWWW* by a colleague (David Thompson?) at NCSA. Marc downloaded Midas and tried it out. He and Eric

Bina then wrote their own browser for unix from scratch. Later, several other folks at NCSA joined the team to port the idea to Mac and PC. As they did, Tom Bruce at Cornell was writing "Cello" for the PC which came out neck-and-neck with Mosaic on the PC.

Marc and Eric did a number of very important things. They made a browser which was easy to install and use. They were the first one to get inline images working - to that point browsers had had varieties of fonts and colors, but pictures were displayed in separate windows. This made web pages much sexier. Most importantly, Marc followed up his and Eric's coding with very fast 24hr customer support, really addressing what it took to make the app easy and natural to use, and trivial to install. Other apps had other things going for them. Viola, for example, was more advanced in many ways, with downloaded applets and animations way back then - very like HotJava was later. But Mosaic was the easiest step onto the Web for a beginner, and so was a critical element of the Web explosion.

Marc marketed Mosaic hard on the net, and NCSA hard elsewhere, trying hard to brand the WWW and "Mosaic": "I saw it on Mosaic" etc. When Marc and Jim Clark first started their start-up they first capitalized on the Mosaic brand, but NCSA fought for it and won. When the "Netscape" brand appeared, people realized the difference between the general "World Wide Web" concept and specific software.

Start of the web: Influences

Q. Have your first ideas in regard to the Web been influenced by any specific work or published paper like Vannevar Bush's "As we may think", a publication of Doug Engelbart or Ted Nelson?

A. There wasn't a direct line. I did come across Ted's work while I was working on the WWW -- after my "Enquire" program (1980) but during my reading up on hypertext - probably between March 89 and September 1990. Not sure.. Of course by 1989 there was hypertext as a common word, hypertext help everywhere, so Ted's basic idea had been (sort of) implemented and I came across it though many indirect routes.

I came across Ted's name first of course. Then I ordered "Literary Machines", and I remember I was late paying him as he didn't take credit cards or Swiss cheques - I paid him in August 1992 cash in person in Sausalito.

I came across Vannevar Bush's article first in the documentation of Digital Equipment Corporation's "Memex" project which became "Linkworks" for VMS. I don't remember when that came out. Great paper.

Doug Englebart's work was the closest to the Web design -- when I saw that the first time I was amazed. He had even used the hash sign as a delimiter for the address within a document (I guess like me by analogy with an apartment number). Doug's stuff is unbelievable. You have best to see the video of him demonstrating it or his demo of a recent smalltalk re-implementation. I saw the latter at the Edinburg Hypertext conference ECHT 94.

Q: Any people who personally helped you get to where you are today?

A: I think the list would be too long to mention. Everyone who was fun and encouraging, starting with my parents. On the professional side, here are a few.

The Maths teacher at Emanuel, Frank Grundy, who conveyed the excitement of the subject with a twinkle of his eye, could make numerical approximations in his head faster than we could work it out longhand, and would throw in a teaser question into his conversation to puzzle anyone who thought they had figured the subject out. And Daffy Pennel who also couldn't contain his excitement for Chemistry and anything related to it.

Unlike most people at Oxford I had one tutor for almost all the work. John Moffat has a vary rare talent for being able to understand not only the physics itself, but also my tangled misguided attempts at it, and then showing me in my terms using my strange symbols and vocabulary where I had gone wrong. Many people can only explain the world from their own point of view.

At CERN, I was recruited by Peggie Rimmer who taught me, among other things, how to write a standards document. Ben Segal was a mentor for my RPC project at CERN, and was a sole evangelist for Internet protocols at CERN long before they were adopted. Ben gave me a lot of moral support in the later WWW days too. A few years later, Mike Sendall was my boss who has a great combination of human warmth and technical depth, and actually allowed me unofficially to write the WWW programs. And then everyone across the Internet who thought the Web was a neat idea and worked on it after hours actually built it.

On collaboration and automatability, Sept 95

The web today is a medium for communication between people, using computers as a largely invisible part of the infrastructure. One of the long-term goals of the consortium is "Automatability", the ability for computers to make some sense of the information and so help us in our task. It has been the goal of mankind for so long that machines should help us in more useful ways than they do at present, help us

solve some of those human problems. Maybe this is one of the many ideas (like hypertext) which the web's great scale will allow to work where it did not achieve critical mass on a small scale before. So there are groups looking at a web of knowledge representation. It could be that some scientific field will be the first to be sufficiently disciplined to input its data not just as cool hypertext, but in a machine-readable form, allowing programs to wander the globe analyzing and surmising.

The W3 Consortium started to address this goal with its recent workshop on Collaboration on the Web. The ability of machines to process data on the web for scientific purposes such as checking a scientist's private experimental data against public databases, require databases to be available not only in a raw machine-readable form, but also labelled in a machine readable way as to what they are.

The knowledge engineering field has to learn how to be global, and the web has to learn knowledge engineering, but in the end this might be a way in which again the scientific field leads the world into something very powerful, and a new paradigm shift.

March 95

Q: How did you come to arrive at the idea of WWW?

A: I arrived at the web because the "Enquire" (E not I) program -- short for Enquire Within Upon Everything, named after a Victorian book of that name full of all sorts of useful advice about anything -- was something I found really useful for keeping track of all the random associations one comes across in Real Life and brains are supposed to be so good at remembering but sometimes mine wouldn't. It was very simple but could track those associations which would sometimes develop into structure as ideas became connected, and different projects become involved with each other.

I was using Enquire myself, and realized that (a) it would fulfill my obligation to the world to describe what I was doing if everyone else could get at the data, and (b) it would make it possible for me to check out the other projects in the lab which I could chose to use or not if only their designers had used Enquire and I had access.

Now, the first version of enquire allowed you to make links between files (on one file system) just as easily as between nodes within one file. (It stored many nodes in one database file). The second version, a port from NORD to PC then VMS, would not allow external links.

This proved to be a debilitating problem. To be constrained into database enclosures was too boring, not powerful enough. The whole point about hypertext was that (unlike most project management and documentation systems) it could model a changing morass of relationships which characterized most real environments I knew (and certainly CERN). Only allowing links within distinct boxes killed that. One had to be able to jump from software documentation to a list of people to a phone book to an organizational chart to whatever .. as you can with the web today. The test rule was that if I persuaded two other projects to use it, and they described their systems with it, and then later at any point a module, person etc., in one project used something from another project, that you would be able to add the link and the two webs would become one with no global change -- no "flag day" involving the merging of two databases into one, no scaling problems as the number of connected things grew. Hence the W3 design.

The same lesson applies now to the webs of trust we will be building with linked certificates.

So the requirement was for "external" links to be just as easy to make as "internal" links. Which meant that links had to be one way.

(There was also a requirement that the web should be really easy to add links to, but though that was true in the prototype we are only now starting to see betas of good commercial web editors now.)

June 94

This was an interview in Internet world by Kris Herbst. His questions are his (c) of course. Slightly edited.

IW: What did you think of the first WWW'94 conference?

TBL: Great! It had a unique atmosphere, as there were people from all walks of life brought together by their excitement about the Web. As it was the first one, they hadn't met before, so it was a bit unique. It was very oversubscribed, as you know, so the next one will have to be a lot bigger.

IW: Can you tell us something about your early life, and

how those

experiences might have influenced you later as you developed WWW?

TBL: That's the first time I've been asked to trace WWW history back

that far! I was born in London, England. My parents met while

developing the Ferranti Mark I, the first computer sold commercially,

and I grew up playing with five-hole paper tape and building

computers out of cardboard boxes. Could that have been an influence?

Later on I studied physics as a kind of compromise between mathematics and engineering. As it turned out, it wasn't that

compromise, but it was something special in its own right. Nevertheless,

afterward I went straight into the IT industry where more things

seemed to be happening. So I can't really call myself a physicist.

But physicists spend a lot of time trying to relate macroscopic behavior

of systems to microscopic laws, and that is the essence of the design of

scalable systems. So physics was probably an influence.

IW: What led you to conceive the WWW?

TBL: I dabbled with a number of programs representing information in

a brain-like way. Some of the earlier programs were too abstract and led

to hopelessly undebuggable tangles. One more practical program was a

hypertext notebook I made for my own personal use when I arrived at

CERN. I found I needed it just to keep track of the -- how shall I say --

flexible? creative? -- way new parts of the system, people and modules

were added on and connected together. The project I'd worked on just before starting WWW was a real-time remote procedure call, so that gave me some networking background. Image Computer Systems did a lot of work with text processing and communications -- I was a director before coming to CERN.

IW: What elements in your background or character helped you to conceive WWW as a way to keep track of what was happening at CERN?

TBL: Elements of character?! Anyone who has lost track of time when using a computer knows the propensity to dream, the urge to make dreams come true, and the tendency to miss lunch. The former two probably helped. I think they are called Attention Deficiency Disorder now. ;-)

IW: Do you have some favorite Web sites for browsing?

TBL: (Sigh) I wish I did, but I hardly spend any time browsing. Historically, I appreciate the people who were first and showed others how things could be -- Franz Hoesel's Vatican Library, of course, Steve Putz's map server, lots more.

IW: How do you feel about the fact that WWW promises to generate large amounts of money for some persons?

TBL: If it's good, people will want to buy it, and money is they way they vote on what they want. I believe that system is the

best one we

have, so if it's right, sure people are going to make money. People will

make money building software, selling information, and more importantly doing all kinds of "real" business, which happens to work

much better because the Web is there to make their work easier.

The web is like paper. It doesn't constrain what you use it for:

you have to be able to use it for all of the information flow of

normal life.

My priority is to see it develop and evolve in a way which will hold us in good stead for a long future.

If I, and CERN, hadn't had that attitude, there probably wouldn't be a web now.

Now, if someone tries to monopolize the Web, for example pushes

proprietary variations on network protocols, then that would

make me unhappy.

More obscure questions...

Rendition of links

Q: I'm a student of visual communications and asked myself why links are blue. I found some answers that might be, for example blue is a color of learning, but I'm not sure what is right. Is there any reason, why links are colored blue?

There is no reason why one should use color, or blue, to signify links: it is just a default. I think the first WWW client (WorldWideWeb I wrote for the NeXT) used just underline to represent link, as it was a spare emphasis form which isn't used much in real documents. Blue came in as browsers went color - I don't remember which was the first to use blue. You can change the defaults in most browsers, and certainly in HTML documents, and of course with [CSS](#) style sheets. There are many examples of style sheets which use different colors.

My guess is that blue is the darkest color and so threatens the legibility least. I used

green whenever I could in the early WWW design, for nature and because it is supposed to be relaxing. Robert Cailliau made the WWW icon in many colors but chose green as he had always seen W in his head as green.

One of the nicest link renditions was Dave Raggett's "Arena" browser which had a textured parchment background and embossed out the words of the link with a square apparently raised area.

Why is your email address on my screen?

Q: I get on my connection screen something like

```

Keyword          Decimal      Description
References
-----
-----
http             80/tcp      World Wide Web HTTP
http             80/udp      World Wide Web HTTP
www              80/tcp      World Wide Web HTTP
www              80/udp      World Wide Web HTTP www-http
80/tcp          World Wide Web HTTP www-http
80/udp          World Wide Web HTTP #   Tim Berners-Lee <timbl@W3.
org>

```

Who are you and why are you there?

A:

Your screen is showing you a list of services on the Internet. Service 80, for example, is HTTP, the protocol which allows a web server and client to talk to each other. A web client opens a TCP connection to a port number 80 on the server. It just happened that I designed HTTP and asked for the port number to be assigned for computers everywhere to be able to use for the web. So someone left my name and email against the entry at the time for the record. The hash (#) tells your computer not to take any notice of that line. It is just historical. I am not hacking your computer!

Can you tell me more about your personal life?

No, I can't - sorry. I like to keep work and personal life separate. What is on the web on this page and my home page is all there is. Please do not email me asking for

more information for school projects, etc. Thank you for your understanding.

TimBL