THE SPIDER'S WEB

A Case for Reforesting the World Wide Web

E. W. Dershowitz
Leland Stanford Junior University

STANFORD UNIVERSITY PRESS
Redwood City, CA, 2026
La Vergne · Toronto · Ingram · N. Yorkshire

CONTENTS

FOREWORD

During the COVID-19 lockdowns, which lasted from early 2020 to about February of 2022 in most municipalities (at least in terms of 'masking mandates'; many at this point had already informally considered the outbreak to be over, as epidemiologically unsound as that idea ever was), I observed as my eldest daughter, Nina, unable to play with her peers as a typical child should, began to grow up on the Internet. As a child of 1988, I was alive (however faintly aware) for the release of Netscape Navigator in 1994 and the sale of NSFNET's assets in 1995. Even still, the early landscape of the Internet, particularly during the Dot-com bubble and before the real emergence of commercialization of the Internet with Web 2.0, was a far cry from the Internet I warily allowed my daughter access to.

Like my daughter, I was raised by a computer scientist. Like my father, I had to (swiftly) weigh the dangers and benefits of a rapidly evolving technology. Though, in the 1990s it was not quite so easy to explore as it is in 2026, when I write this. What you are about to read is a combination of information and opinion. I intend to lead you through a period of time which predates me, starting with its inception in the late 60's, and ending with the current state of the Internet. Afterwards, perhaps you might share my disdain that such a promising expanse of technology has been reduced to little more than a vehicle for a few well-linked sites to advertise to a captive audience.

E. W. Dershowitz, Stanford, 2026

THE ENVELOPE

Because the messenger's mouth was heavy and he couldn't repeat [the message], the Lord of Kulaba patted some clay and put words on it, like a tablet. Until then, there had been no putting words on clay.

– Enmerkar and the Lord of Aratta, c. 1800 BCE

First, do not be one of my 21st century contemporaries. Instead, be a courier in an early civilization which has invented writing. You do not have to remember your message, but you certainly have it in your hands, or in a bag of some kind. If you are literate, you can even read it. You bear a tablet, or a quipu, or even a letter, on paper. You know who you have been sent to find, and you may even know how to get there.

If you don't, you've been given directions to your first stop, and when you get there, you know who to ask to get your next set. When you find your destination, you will present your message, and perhaps bring back a response. You traverse the world by roads, or, perhaps earlier than that, landmarks. The essence of it is, you can walk, or ride, and you can get around by following features of the terrain, from point A, to B, to C, to your destination.

You are fortunate if your way is always paved. A road is the promise of a landmark, it tells you that you know where you are. You may take them as a means of circumventing all of that seeking. But, all of this walking around and looking for your next landmark is

exactly why it was so expensive to have a messenger. When your message is carved into clay, if you wish to be verbose you may have to send multiple. Perhaps they are professionals. Perhaps they are even spies. It's a terrible amount of uncertainty.

The quality of messengers and messaging at your disposal could vary quite a bit. Perhaps the postal service has been invented. Perhaps the messengers are all professionals, and they are fast runners, or they have fast horses. In any case, they must be told where to go, and they must know it. It'd be a much better system if there were a few messengers at every stop, who knew how to get to the places most immediately nearby. That way, they could even make the trips multiple times per day. But then, the destination gets passed along, from mouth to ear, as the tablet passes hand to hand. It would be awfully easy for it to get lost. As you might imagine, writing the delivery address is another fast development.

Even with a robust system, there's another source of risk. More people means more chances for one of them along the way to work for your enemies, in secret. This is an art foiled and unfoiled by keys and boxes, lockpicks, and social engineering, but an art emerges nonetheless. One day, it will resemble what we would call a sealed envelope. The address line is on the outside, and the message is secure inside. If you're sending an awful lot of information, you may have to break it into multiple messages, inside multiple envelopes.

As you might have surmised, I'm not really talking about messengers at the dawn of civilization. I'm talking about the challenges facing the rather clumsy field of computer networking at its

inception. At the dawn of computer networking, researchers knew how to 'join' several computers by means of wired or telecommunications-based networks, consisting of several 'nodes' (computers in this example, cities or recipients in the previous) connected by 'edges' (those would be the wires/telecom, or the messengers).

However, it was exceedingly difficult to communicate necessary information at speed. Moreover, it was vulnerable to attack, disruption, and interference from overusage. Researchers in conjunction with RAND Corporation and the UK's National Physical Laboratory independently invented what we now know as packet switching. The data you wish to send from your device to a connected device is first split into 'packets' if necessary. These resemble envelopes, they have data inside called a 'payload', and an address line called a 'header'. The rest of the devices in the network know how to forward each packet to its destination based on the header. The packet, after all, had already been developed in primordial form at some point in the stone age.

Networks grew with their demand. What started with a concept in two research laboratories became a means to time-share computer equipment, which at that time took an awfully long time to do anything. This saw the establishment of ARPANET in the US and CYCLADES in France. Both operated off of early packet-switching, and had some amount of unreliability in terms of latency and packet arrival. However, the benefit of being able to connect research data to powerful research computers at other institutions led to the continued development of both projects.

Advancements came quickly. The Transmission Control Protocol was invented in 1974, in a paper which saw the first attested use of the phrase 'Internet'. It was developed somewhat in conjunction between the two major networks, and represented a marked improvement over the NCP used by ARPANET at the time. TCP was eventually divided into TCP/IP as specifications improved, with TCP referring to the protocol which 'plots the route', and IP referring to the protocol which 'picks out the landmarks'. Routing consists, in 2026, mainly of routing tables held by Internet Service Providers. Few private users have their internet line connected directly to 'nearby' nodes.

Around the time that public-use networks (most previous were either entirely proprietary or government-sponsored) were beginning to take shape, a model of computer networking was developed to describe some of the concepts I have alluded to in metaphor. The OSI (Open Systems Interconnection) model was published in 1983 by ISO and Hubert Zimmerman, and consists of seven layers. They are, in brief:

1. Physical; the literal wires and telecom connections
2. Data Link; transmission between two devices
3. Network; routing and traffic control within a multi-node network
4. Transport; reliably transporting data from end to end
5. Session; maintaining a continuous communication line
6. Presentation; translating data by means of encoding, compression, etc.
7. Application; resource sharing or remote file access (e.g., HTTP)

In practice, this model is not used. The simpler model proposed as part of TCP/IP is the actual foundation for most of the Internet, especially with the closure of what were known as the Protocol Wars (debates, not wars). This included other debates, including Packet Switching vs. Circuit Switching (in which a 'line' is held open between one endpoint and another to maintain the communication session), are somewhat inconsequential for my model of the Internet that could have been (and perhaps one that still could be).

In 'The Web', I shall explain what the Internet truly looks like, so far removed from the rudiments of message passing and envelopes. In 'The Spider', I shall attempt to convince you how this removal has allowed us to be led astray. And indeed, we have been led astray. In 'Visions of a Forest', I will describe my vision for my model, that dearly-imagined Internet-apparent.

But first, I must lay the groundwork for what I view as the great failure of the commercialized Internet. It remains my opinion that the OSI model will serve well enough as an introductory step to the layers of abstraction which define the Internet that most people are familiar with. It's important that you become acquainted with just how far we've already come from stranding a data-transfer cable from device to device (though you may well still do that, if you're one of my more industrious readers), because that, like so much else, was lost in the abstraction. Please join me now in understanding 'The Web'.

THE WEB

Consider a future device... in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

– As We May Think, Vannevar Bush, 1945

It might have been unthinkable to those early pioneers that one day we might interact, establish sessions, and link computers with individuals and computers we did not know before, in the physical world. In 2026, a staggering proportion of communications are made without so much as the sharing of an IP address. The route has been abstracted away; the client is in communication with the service, and the service (ostensibly) delivers communications intended for another end user via a route that scarcely resembles a direct path along the links between sender and recipient. Far less when the service is more complex than, say, email. The intermediate service, the application, the social networking website, perhaps, has become all-encompassing.

I am not some great opponent of anonymity. While I do not enjoy its comforts (feel free to look me up; I have a 'faculty' page at Stanford University, my first name is Ewa), I see it as a marker of just how far removed we are from the physical layer of Zimmerman's OSI model. I would speculate (perhaps without basis) that the vast majority of internet communications in our day consist between a sender and recipient who may have no knowledge of each others'

physical location, much less IP address, nor the routers and kilometers of fibre-optic cable their packets are routed through.

In 'The Envelope', I described a network comprised of devices connected to other devices by routers, cable, and telecommunications. However, today's Internet bears only a passing resemblance to that to the average user. In the same chapter, I mentioned the *HTTP* in the application layer. The *HyperText Transfer Protocol* is the real backbone of the World Wide Web, which is the user-friendly side of the Internet. Included in the specification of the HTTP was the *hyperlink*, intended to *link* between *hypertext* documents (e.g., sites). The purpose of this protocol was to permit traversal between sites by mouse click, but especially permit traversal without necessarily needing to pre-plan what site you reached.

Nor am I an opponent of the HTTP. There is an undeniable danger in the publicization of users' IP addresses (G-d forbid my daughter's *Minecraft* server-mates knew our home address, though the ISP masks that well enough) and a frankly impractical hassle in the necessity to traverse by knowing your exact destination. However, HTTP represented another divergence from the 'true' network of the Internet, and the creation of a 'second-order' network.

In 1993, there was a single list of every web server on the World Wide Web, hand-indexed by Tim Berners-Lee on the CERN web server. As the World Wide Web became more popular, and increasing demand arose for searching *content* rather than *users*, search engines were invented. A search engine has a triple purpose; most are part 'bot', intended to crawl the Web in order to assemble indices and

ranking information based on what qualities the search engine intends to support. Finally, a search engine permits user searches, attempting to match keywords or other advanced factors. Unlike subjects discussed previously, I'm fairly certain that most readers are familiar with a search engine, or perhaps multiple!

It's worthwhile to note that most modern search engines are assembled by algorithm. As we have seen in recent years, these algorithms range from naïve, to benign, to shockingly well-fit, to (most of them) entirely bent on selling you whatever you least need. RankDex is an early model of search engine algorithm, akin to the algorithm used in Hermes Fulfillment's *Talaria!*. Both algorithms rely on the counting of hyperlinks (though not as naïvely as many attempting 'Search Engine Optimization' may believe). Several engines, including *Talaria!* Successfully adopted a pay-for-placement model, seen as the 'sponsored' or 'suggested' content at the top of most search pages, as early as 1998, originally based on an auction model.

Let's return to the idea of the second-order network. While the real Internet exists as a network of all-too-real connections carried by telecom lines, ground cable, or satellite network, there is another Internet that exists in the form of sites and user devices, connected by hyperlinks. The advent of search engine technology radically changed the topology of this second-order network, rather than traversing along the routers from user to user, or user to web server via an address, perhaps gained from CERN's index, most internet traffic traverses from link to link, some of these resulting in extremely (true network)

inefficient routes that are, to the user, fairly short in connection. For good or ill, the HTTP and search engine were the last two nails in the coffin in terms of abstraction.

It isn't either the HTTP nor the search engine which characterize the Internet of today, of course. Back in 'Foreword', I mentioned Web 2.0, which was not so much a product as a general drift in Internet architecture towards user interactivity, monetary transactions, and multimedia content. By 2006, Web 2.0-based services such as social networking (including the nascent Facebook, created in 2004), blogging, wikis, and media sharing.

Like never before, users of the World Wide Web could meet, interact, collaborate, and converse with other users around the world. User participation, dynamic content, metadata, and of course, tracking information, which quickly grew into a lucrative market. Within the framework of Web 2.0-type entrepreneurial philosophy, the Internet was catapulted in a matter of years into the most remarkable vessel for entertainment on Earth.

Within the paradigm of Web 2.0, users are a major driver of internet content, primarily through content aggregation services, i.e., social networking platforms. As networking platforms, blogs, and vectors of media sharing grew more popular, they gained positions of prominence within search engines, thus reinforcing their impact. Meanwhile, the Internet under this paradigm also becomes a potent tool for marketing. And yet, social networking services, by which I primarily refer to content aggregation services, e.g. *TikTok,* little produce their own content for audiences. Rather, they aggregate

content produced by their users. Some political thinkers have pointed out that many media-sharing sites which utilize advertising as a means of generating revenue or offsetting upkeep costs are exploiting the user participation aspect of the paradigm as a means of free labor. I am inclined to agree, with the small allowance that aggregation is, in fact, a service of value.

However, most worrisome of all the facets of the Web 2.0 paradigm is the level of blending of advertisement and entertainment, a process not wholly unique to the Internet. However, in 2026, the general consensus among avid Internet users is that an ad blocker is necessary for optimal use, not only for preventing slowdowns and unwanted advertisement, but to avoid malicious software infections by way of such ad services. Nowadays, the profit motive is the primary driver of development of the Internet, with a number of large advertising agencies having consolidated control of many major entertainment services and social media distribution platforms.

Advertising algorithms subsist on scraping tracker data nested in webapps and hyperlinks in order to improve the success rate of yielding clicks or interest. I am among the considerable chunk of people which view this not only as an invasion of privacy, but a violation of rights. In recent years, I find I seldom use the internet outside of work. Many, however, can't seem to get enough of it: and indeed, how could they? Web 2.0 makes a compelling offer-- this process of aggregation puts an inexhaustible supply of content and conversation at the user's fingertips. We, the social animals that we are, cannot lament our preoccupation with that.

The monetization of an internetwork of potential consumers, and the reinforcement of schemes with which to advertise to them, was somewhat inevitable. However, I increasingly believe that it is the fault of the naïveté with which we approached this nascent field, which has led to the profit motive so completely subsuming all other hobbyist, egalitarian, non-proprietary, or socially-motivated goals, etc., goals of the Internet.

Furthermore, I strongly believe that the present state of the Internet, i.e. one that profits off of fomenting culture war and political unrest to yield greater clickrates and greater viewrates of advertisements, is not only unethical, but a disgusting distortion of the promises of the technology as set out in the 1960s. I believe that the great modern architects of the Internet willfully abuse the obscurity of their systems, the addictiveness of readily accessed and abundant content, and our own seeming inability to regulate our desires for it. The truth is, we hardly had a chance. Now, it seems, we come to the spider in the web.

THE SPIDER

Oh, what a tangled web we weave,
When first we practise to deceive!

– Marmion, Sir Walter Scott, 1888

When, in 2026, someone is using the Internet, they are typically using a small handful of websites, owned by a smaller handful of service providers. You likely know the ones: *Facebook*, *Reddit*, *Youtube*, *TikTok*. In the late 2010s, *Twitter*. Each of these social networks, which the vast majority of habitual internet users report accessing at least once, and a large proportion thereof report continual use, operate on roughly the same scheme: a continued flow of easy to access content, primarily generated by the users, and curated by model-fitting algorithms, similar to those used by advertising services. These services are funded by advertisement money as well as premium-user packages which unlock certain quality of life features.

In short, the end state of such a content/entertainment ecosystem is one which completely dominates a given user's time online, thus exposing them to the largest possible array of advertisements, use metrics, and other trackers, which self-reinforce in order to keep a user present within the network. This domination is carried out by the services I reference with, in general, frightening efficacy: the old adage goes, *if it bleeds, it leads*, and there is no shortage to the macabre and the sensational online. A consequence-- or, dare I imply, design principle-- of the algorithms which determine the aggregation schemes of each service, is the polarization of communities, in which each

community is driven to fortify its borders and inundate itself in negativity.

I don't need to explain why some psychologists point to this feedback loop as a sort of dopamine addiction, especially when the period of its advent is well associated with a general reduction in attention span, a phenomenon I've noticed both in my students and in my daughter. The user is empowered to obtain a constant supply of novelty, because the aggregation of content means that there is always more to click on just down the page. It is little wonder we have developed our habits of doomscrolling.

While perhaps a less popular scare, the great evil I decry is not quite that businesspeople as they ever have are attempting to entrap potential buyers with newer and shinier distractions, but that this directly leads to an acceptance of a pre-chewed Internet without user input nor control. In our modern world, telecommunications are necessary, and while it chafes, it is possible to get by using the internet solely in the workplace, with some auxiliary use towards convenience and coordination (see, sending a calendar invite over e-mail).

It is the foremost failure of today's cohort of Internet users that so many of them are willing to accept this data-selling scheme packaged as endless user-friendly experiences. It is the duty of the shrewd buyer to ensure they are not being cheated. It is the duty of every user of the Internet to not permit this remarkable technology to be misused in such a way. For my part, this exact issue is the one I seek to impart by thoroughly explaining the mechanisms and basic forms of the Internet. The means by which to create your own

experience are indeed suppressed by the interests who would much rather you take part in a return on their heavy investments in ad revenue and 'walled-garden' type services, but they are immensely rewarding, and only the understanding which comes from experiencing the mechanisms of the internet can afford a true sense of the beauty of its technology.

In short, it is far too easy to be enchanted and led astray by an addiction to clicks, by the doomscroll, and by the promise of infinite content pushed by the entrepreneur-*cum*-server-hosts of the modern Internet. People growing up following the Dot-com bubble, as previously mentioned, or more egregiously, growing up with even limited access to the internet as does my daughter, Nina, and as do many children across the modern world, are extremely likely to be unaware and off-guard against these modern threats.

The technological basis of the Internet as it is, not as these walled-garden services present it, is far more expansive, far less sanitized, codified, or inclined to entrap the user. It is also far more dangerous and far more likely to provide an easy vector for malware. We have always had to balance safety with the freedoms necessary for human life, and I for one wish to impress that I cannot tolerate what I see as the willing sacrifice of the latter in the name of palatability for advertising executives.

This has always been the insidious promise of proprietary software, software which is typically developed closed-source and maintained within-house, which by copyright and intellectual property law prevents free share and modification of the software. This term is

defined in contrast to free and open-source software, which permits and even encourages savvy users to improve upon the software. The great deceit of proprietary and other non-free software is that the rightsholders or publishers are equipped with better knowledge and judgment as to what uses are permitted within the framework of the software, and on a larger level, that users should give up control over what shape the future of the Internet should take. Consider that example of the walled garden: we have trusted the aggregation platforms to determine what we see. They have returned our trust with addiction, polarization, and commodification.

This is not to say there are no contemporary experts or visionaries on the topic of the Internet. At Stanford, I am fortunate to be in contact with some of the brightest minds in computer networking (as well as others). However, it is not the specialists which are making the decisions as to the shape of the Internet, it is the advertising executives, and only really in name. Profitability-- the market-- is making the decisions, which is perhaps the worst idea possible.

More importantly, I believe that such aggressive strains of meritocracy serve only to stifle those which have not yet shown their promise or willingness to contribute. So who should be making the decisions as to the shape of the Internet? Certainly not the specialists alone, we love our creation far too much for that. On this, I find I am something of an absolutist: the only opinion I accept where who ought to be able contribute to the Internet is concerned is 'each and every person who wants to.'

It remains, in my view, among the greatest shames in human history that the singular potential of boundless interconnectivity and communication at speeds and distances never before matched has, at present, become a billboard-ridden closed environment profiting off of the fomenting of the natural rage generated by the economic and political polarization of our societies and the relative listlessness of modern people in the face of back-to-back economic crises. At risk of dallying too long on such grim fronts, I would like to present the function of this drawn-out keynote: my 'Visions of a Forest'.

VISIONS OF A FOREST

When we try to pick out anything by itself we find that it is bound fast by a thousand invisible cords that cannot be broken, to everything in the universe.

– The John Muir Papers, John Muir, 1869

When I set out to write this opinion piece, it was my husband's suggestion that I christen it with a suitably pithy title. Thence comes 'A Case for Reforesting the World Wide Web'. However, in its catchiness, it perhaps loses clarity. While the question, 'how exactly does one reforest a telecommunications network' might serve to get a reader through, with fortune, say, the first fifteen pages, I am an active professor, and my experience in teaching tells me that I am correct in thinking that such a technique is neither intuitive nor suitably didactic to justify itself. Well, you read the title, didn't you? He won; oh well.

Reforesting necessarily must follow deforesting. Otherwise, the apt term is 'foresting'. I quibbled at some length on the topic of these 'walled-garden' services which enrapture the striking majority of habitual Internet users, and that's the kernel of what I mean. I think of the Internet as, perhaps, a forest once fed by seeding from a vibrant ecosystem which was demolished for the construction of a shopping complex (and when I think of the present Zeitgeist-*cum*-specification in website monetization, it makes me happy to think it may one day go the way of the shopping mall: a big, soulless, garish, waste of advertiser money). Since these walled-garden services make up the

bulk of Internet traffic, we may consider the present Internet suitably deforested. The essence of reforesting is simple.

If you are still reading this, then some part of my text has struck a chord with you. Perhaps it angers you in a way you can't quite explain that you have no choice to be pandered to constantly, to watch your friends radicalized or reactionarized, and to argue in an immensely public and embarrassing way about it, and most of all, to be commodified. I hate to be commodified! We should all hate it! This anger is extremely basic, and extremely necessary. As human beings, it is just that we would hate our own exploitation. As good human beings, it is aspirational that we would also hate the exploitation of our fellows.

If any of the above tracks, or if you have free time, I invite you to create a web server. And by G-d, make it ugly and stupid. I said before that the HTTP was not an evil, and I mean that. Hypertext is a basic component of 'web documents', the most basic form of a website, and it is the simplest thing you can do to wrest some control of the Internet away from the technocratic interests which have warped it into something big, soulless, garish, and frankly stupid. That's the first level.

On a more pragmatic level, there are still issues with the above approach. For instance, by perching like spiders in their own rights over the routing tables which allow packet-switched traffic to be directed across the physical anatomy of the Internet, Internet Service Providers may charge and monitor what you choose to do with this freedom. Some of this is necessary for public safety. Most of it is

extortion. We have, for over sixty years, been perfecting the architecture of this backbone, but it is well-attested in mutual aid groups that some things just cannot be trusted to centralization. Therefore, we reach what I really (truly) mean by 'reforesting'.

We return (obliquely) to the OSI model. The 2nd layer, 'Data Link', which holds the key to constructing what some before me have described as the 'Universal Forest' model, devolved, decentralized, and consisting of 'trees'. Topologically, it may be conceived of as a network of networks, where each network is a particular community; a type of devolved 'site'. Individual users act as mirrors, seeding the site instance to other users as relays. One's contribution to the Universal Forest is calculated as a function both of the computational power they afford and the number of people they mirror for, and thereafter serve to mirror for them.

The exact mechanism which permitted the consolidation of the Internet into walled gardens was the divorce of the second-order network from the first-order or real network. A sustainable, forested Internet must be unflinching and bare in terms of its fundamental architecture. Perhaps it is ridiculous of me to ask that everyone who wishes to benefit from the internet understand networking on a basic level, but there must be an answer somewhere between that and 'we let the people who want to sell our data do all the work'.

Such an Internet must carefully balance accessibility with its central mission; but we do not expect our children to be incapable of at the least basic housekeeping chores after a certain age, ergo, accessibility cannot be overfocused. Finally, the Universal Forest must

never compromise on open source principles. Obscurity of protocols, jargon (which I am at times guilty of), and abstraction must be metered as necessary to ensure the user remains the participant.

'Ecologically' sound ranking databases must be specific and democratic in ranking. The domination of market share on the part of services such as *Talaria! i*s a leading contributor towards authoritarian (i.e. a proprietary software foundation of the Internet) and away from a end-user-*cum*-participant based scheme, which is in line with open source principles.

The tight-knitting of a well-advertised internet 'core', the initial state of dominant multiservice platforms such as YouTube or Hermes' webstore, can be prevented by incentivizing 'local' traversal between trees or communities with high affinity. While the relative cost associated with web servers is the primary motivation towards the mirror-based peer-rove traversal of the Universal Forest, the second reason is that when presented with a ranking or routing table, very soon, someone is pointing you where to go.

The second reason for mirroring is that those who mirror multiple trees act as ready-made 'jumps' between the two communities. I will refer to this form of traversal henceforth as proximity traversal, or a 'dive-based' traversal where you rove forth from your own immediate tree of mirrored sites to those connected, viewable in-browser by a 'dive menu', reminiscent of directory menus in file browsers. It is my hope that, with each user acting as a bridge, one is incentivized to take willing part in the technical nature of the Internet.

While much-decried already in this essay, advertising has its place within the Universal Forest model. For one thing, I believe it is necessary for secure and open-source protocols towards resource sharing for the aforementioned mirroring and seeding methods to be effectively carried out by the common denominator of users. A secondary purpose of such protocols is that those seeking to advertise their sites could pay users directly for resource-hours on their devices.

I am not against the protective concealment of artwork or other intellectual creations. Mirrors must have sufficient security protocols to prevent the access of works which the viewer has not compensated. However, the difference arises when we discuss the common 'skeleton' of the Internet, or the Forest, in this case. As I see it, it is wholly unforgivable that an inch of cable or a byte of software architecture of the common, public Internet be maintained in private, i.e. as proprietary; the position of the World Wide Web as a common necessity for work and life and the monetization thereof cannot coexist, and I hope that I have been thoroughly clear which of the two must take precedent.

Loath as I am to say it, we end up with a system nowhere near as exciting as the current World Wide Web. We have deliberately made it unprofitable (though neither is the post office, and the sort of person who thinks it should be stopped reading around the point I suggested commodifying human communication was bad), we have required that people understand the technology they are using, and we have sacrificed a considerable amount of convenience. But this convenience was never free, and its ransom was, societally, astronomical.

This paradigm would require greater trust between agents of a network. Because communication occurs peer-to-peer, the vulnerabilities of the network are exactly the vulnerabilities of its individuals. Trust, on an individual basis, between peers, becomes paramount. The security of each client connection demands a reactive and supportive development team. I, for one, gladly volunteer. Furthermore, while users are expected to have a basic technical understanding, I understand the impracticality of expecting everyone to be able to address its vulnerabilities and intricacies. Specialists would be treated as 'guides'. Because the World Wide Web is a public resource, they could be accredited through local institutional bodies.

In summary, at its inception as a means of time-sharing massive and unwieldy research computers, the Internet held promise for a novel form of communication, and whispers of never-before-seen interconnectivity across the world. However, entrepreneurial motives have at present stymied the organic (open source, free) development of the Internet, leaving it a closed system by which to prey on its own habitual users with the promise of addictive and inexhaustible content.

That is not to say this Internet could be totally free. We must have networking infrastructure and we must have software, and these both must be maintained. In all likelihood, one would continue to pay a fee for connectivity, for the maintenance of physical links, and basic software. But no longer would one pay for *access* to the vaunted hub of the internet from the exclusive cartel of tier 1 networks.

Nor do I have any illusions that we shall (barring some worldwide disaster which destroys all networking infrastructure) have the chance

to build the Web again, only wiser this time, but if we are to escape this trap and benefit from its worldwide interconnectedness both, we must be more mindful of the very means of its operation, down to the level of physical links. Clearly, we cannot trust profit to make the decisions as to the shape of the Internet, and I have little hope that we can trust politicians to understand it. So, we are left facing a great and grueling calling: to determine for ourselves and by ourselves, peer to peer, and in conjunction with a devolved network of specialists ceded some-- but not too much-- control, the very nature of all of our Internet communications.

I propose the auxiliary development of small 'pirate' networks built on mutual aid principles which may eventually be incorporated into the larger body of the Internet, with necessary translation architecture, such that the concept may proliferate more widely. The venture capitalist model is in some sense its own undoing, as in the ability for users to cease use (thus ceasing profitability) by engineering our own alternative, we are able to impress and motivate a wider schema of open-source software with respect to the architecture of the internet. In short, if a sufficient proportion of current heavy-users were to adopt such a paradigm, it would have an impact.

The only mitigating factor is the apathy and emotional exhaustion the present software-capitalist class can manufacture, another deception used to maintain the most profitable state of the proprietary internet. If they cut off our protocols, we will write translators. If they separate us from the internet, we will string our own cables.

This is why I say that the commitment to the Open Source schema and freeware principles at large is the only weapon we have against the construction of walled gardens, which are somewhat intuitive in design and specification when the end goal is the extraction of money from a user class. A necessary step therein is to divorce the act of 'using' from the act of 'participating', even as under the Web 2.0 model the core action of social networking services is participation, i.e., 'posting' original content or otherwise novel material to enrich the service at large. I say that it is entirely critical to the web-entrepreneurial class to maintain ignorance as to the underlying mechanics of the network, as with a savvy population comes the confidence that the web-entrepreneur is, in fact, as unnecessary as I claim them to be.

So, perhaps having found some truth to my words, you will understand why I condemn the World Wide Web in its present state, and why I find pity in the failure of its promise. Perhaps, like me, you will now dedicate yourself to working towards its betterment. If so, warm regards, and good luck-- I look forward to writing the foundations of our new Internet together.

About the Author

E.W. (Ewa Wiktoria) Dershowitz, Ph.D., is a computer scientist, teacher, open source fanatic, small-time hacker, and writer of strong opinions. Most of these things, she picked up along the way, but she's been a computer scientist since her father's machine beat Garry Kasparov in a chess match. She grew up alongside the internet, for better or worse, and when she isn't writing about its many flaws, she can be found singing its praises... within reason, of course.

Officially, Dershowitz lives in Palo Alto with her husband, Mitchell Garber, and their daughters Nina and Tovah. In fact, she lives in her office, 389 Jane Stanford Way, Room E122.