BUILDING A BETTER INTERNET

By

Zain Humayun

B.S. Computer Science
Lahore University of Management Sciences

SUBMITTED TO THE PROGRAM IN COMPARATIVE MEDIA STUDIES/WRITING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTERS OF SCIENCE IN SCIENCE WRITING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

FEBRUARY 2022

Signature of Author: _____

Zain Humayun
Department of Comparative Media Studies/Writing
September 6, 2021

Certified and Accepted by: _____

Tim De Chant
Department of Comparative Media Studies/Writing
Thesis Advisor and Graduate Program of Science Writing Lecturer
September 6, 2021

_____

Seth Mnookin
Department of Comparative Media Studies/Writing
Graduate Program of Science Writing Director
September 6, 2021

BUILDING A BETTER INTERNET

By

Zain Humayun

Submitted to the Program in Comparative Media Studies/Writing on September 6, 2021 in Partial Fulfillment of the Requirements for the Degree of Masters of Science in Science Writing

ABSTRACT

In the summer of 2016, dozens of computer scientists gathered at a former church in San Francisco. For a long time, they had all been worrying about the same thing: the future of the Internet.

Some of those present were dissatisfied with an online world where websites could go down, and the information on them lost forever; others were alarmed by an Internet dominated by a few powerful tech companies, running opaque algorithms and surveilling millions of people. Some believed that the Internet's nuts and bolts presented easy targets to hackers and enemy states, while others were more concerned about countries censoring the Internet within their own borders.

Those issues seemed to be unrelated, but each reflected a part of the Internet where control was growing concentrated in the hands of a few companies, governments, or infrastructural services: an overarching trend that computer scientists refer to as centralization.

The meeting in California was a call to arms for those trying to organize a counter-movement: a mission to save the Internet, to build a better one, and to realize many of its early, unfulfilled promises. In the five years since, has the movement to decentralize the Internet come any closer to success?

Thesis Advisor: Tim De Chant

Title: Lecturer

The onslaught began at 7:10 a.m.[1] The earliest reports came from the East Coast, where early risers, trying to read the *New York Times*, watch Netflix, or scroll through Twitter, stared in confusion at their computer screens.[2] The Internet was under attack.

The attackers had hit Dyn, an internet infrastructure company based in New Hampshire. On the morning of October 21, 2016, the company's servers were hit by a type of cyberattack known as a distributed denial-of service, or DDoS. A barrage of junk traffic from millions of internet addresses flooded Dyn's computers, crippling their ability to reply to legitimate requests.

Dyn's servers acted as one of the Internet's address books. When a person entered a name like google.com into their browser, Dyn provided a number, called an IP address, needed for the browser to find that website. The DDoS attack that overwhelmed Dyn's computers did not bring down websites like Twitter and Netflix, but it did prevent users from being able to find them.

Around 9:30 a.m., Dyn's engineers said they had thwarted the assault. Or so they thought. Two hours later there was another salvo, this one spreading to users on the West Coast. At 5 p.m., there was a third round. By 6 p.m., Dyn had weathered the largest DDoS ever recorded.[3]

Some had seen it coming. A month earlier, a cybersecurity expert named Bruce Schneier claimed that someone was testing the defenses of companies running important parts of the Internet. "We don't know who is doing this, but it feels like a large nation-state," he wrote online. "China and Russia would be my first guesses."[4]

Schneier wasn't the only one worrying about the Internet. Three months earlier, in June, an eclectic group of computer scientists, lawyers, policy-makers, and activists, all gathered in San Francisco to discuss their growing concerns about the Internet and its vulnerabilities.[5] Some were dissatisfied with an Internet where websites could go down, and information lost forever; others were alarmed by an Internet dominated by just a few powerful tech companies, running opaque algorithms and collecting vast amounts of data on their users. Some believed that the Internet's nuts and bolts — like its DNS system — presented weaknesses for hackers or enemy

---

1. "Update Regarding DDoS Event Against Dyn Managed DNS on October 21, 2016," Dyn Status Updates, October 29, 2016, https://www.dynstatus.com/incidents/5r9mppc1kb77.
2. Nicole Perlroth, "Hackers Used New Weapons to Disrupt Major Websites Across U.S," *The New York Times,* October 21, 2016, https://www.nytimes.com/2016/10/22/business/internet-problems-attack.html.
3. Nick Woolf, "DDoS attack that disrupted internet was largest of its kind in history, experts say." *The Guardian*, October 26, 2016, https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet.
4. Bruce Schneier, "Someone Is Learning How to Take Down the Internet," Schneier on Security, September 13, 2016, https://www.schneier.com/blog/archives/2016/09/someone_is_lear.html.
5. Quentin Hardy, "The Web's Creator Looks to Reinvent It," *The New York Times*, June 07, 2016, https://www.nytimes.com/2016/06/08/technology/the-webs-creator-looks-to-reinvent-it.html.

states to target, while others were more concerned with countries like China, censoring the Internet within their own borders.

While those issues seemed to be unrelated, each reflected an aspect of the Internet where control was growing concentrated in the hands of a few companies, governments, or infrastructural services: an overarching trend that computer scientists call centralization.

Many in the technology community believe that the counter-movement, called decentralization, is as much about a way forward as it is a return to an original promise.

"When the Internet started, it was full of optimism," says Wendy Hanamura, director at a non-profit called the Internet Archive. "Things like the Snowden revelations, Cambridge Analytica, just show us how far from the ideals we've strayed. Those are wake up calls, that the web we use is failing us."[6]

The foundations of today's Internet

The Internet's early days offer both lessons and inspiration for those trying to shape its future. Many of the programmers who are trying to design new kinds of Internet applications, for example, draw on the work of a Cold War engineer named Paul Baran.[7]

In 1959, at the heart of the Cold War, Baran began working at a US military think tank. He was concerned that the country lacked a communications system capable of withstanding nuclear attack, so he began designing one.[8]

Baran knew that the system would have to be different from existing ones. In a telephone network, for example, all the calls from one neighborhood were collected by the local telephone office. The problem was that although each household was connected to only one telephone office, that office served many households. In a war, if the enemy attacked a single telephone office, it would disconnect an entire neighborhood.[9]

To work around this, Baran envisioned a system of "nodes" arranged in a lattice, so that every node was connected to several others. If a link connecting two nodes was destroyed, there were still a way to send a message from one to the other, through other nodes.[10] But this

6. Wendy Hanamura, interview with author, April 15, 2021.
7. David Pierce, "The Centralized Internet is Coming", Protocol Source Code, January 17, 2021. https://www.protocol.com/newsletters/sourcecode/decentralized-blockchain-internet?rebelltitem=9#rebelltitem9.
8. Janet Abbate. *Inventing the Internet* (Cambridge, MA: MIT Press, 2000), 10;
"Paul Baran and the Origins of the Internet," RAND Corporation, https://www.rand.org/about/history/baran.html;
Katie Hafner and Matthew Lyon. *Where Wizards Stay Up Late: The Origins of the Internet* (New York: Simon & Schuster, 1999), 54, Kindle.
9. Abbate, 11.
10. Abbate, 11; Hafner and Lyon, 58.

raised a new question: with multiple paths to choose from, which one should a message take?

Baran drew inspiration from the postal system, where a letter makes its way from post office to post office, and each post office keeps the letter until it decides which one to send it to next. It also offered Baran the resilience required for a nuclear-proof communications system. If an enemy attack destroyed a node, the network could adapt: each node could make a real-time decision about the best path to take, based on the connections available.[11]

Baran designed his system so that, unlike a telephone network, it could transmit not only sound, but any data that could be represented in binary numbers.[12] He also proposed a novel method to make this process economical and mathematically efficient.[13] At the sender's end, a message would be split up into equal-sized chunks.[14] Each chunk would be labelled with the sending and receiving address, like a letter; then, guided by the network's nodes, each chunk would make its own way to the recipient. There, the chunks would be stripped of their address labels and combined to form the original message.[15]

The idea of breaking a message into chunks, sending them along independent paths, and piecing them together at the destination would become prevalent in the field of computer networking.[16] The chunks would come to be called packets: the overall design, "packet switching".[17]

In 1964, Baran published his design in an 11-volume work, titled *On Distributed Communication.*[18] Remarkably, across the Atlantic Ocean, an engineer named Donald Davies would soon come up with packet-switching himself[19] – unaware of Baran's existence[20], and motivated by an entirely different challenge.[21]

***

In the early 1960s, computers were large, expensive machines, so much so that scientists had to share them.[22] They would feed a computer all their programs, and it would process them one after the other.[23] This was efficient because it ensured that the computer didn't waste any

11. Abbate, 13.
12. Abbate, 17.
13. Abbate, 20.
14. Abbate, 17; Hafner and Lyon, 59.
15. Abbate, 17.
16. Abbate, 18.
17. Abbate, 18.
18. "Paul Baran and the Origins of the Internet," RAND Corporation, https://www.rand.org/about/history/baran.html.
19. "Donald Davies," Internet Hall of Fame, https://www.internethalloffame.org/inductees/donald-davies.
20. Hafner and Lyon, 64.
21. Hafner and Lyon, 66.
22. Abbate, 23.
23. Abbate, 24.

time sitting idle, but it was inconvenient for programmers because they couldn't get an immediate response from a computer. Often, they'd have to wait for hours.[24]

At a research center in Britain called the National Physics Laboratory, Donald Davies was exploring a new way of using computers. What if, instead of completing one program before starting another, a computer could quickly cycle through all of them, spending a fraction of a second on each? Someone using this method, called time-sharing, would feel like the machine was responding instantly.[25] But there was a problem.

When several people were using a time-sharing computer together, each would enter their commands using a keyboard-and-monitor-type device called a terminal. Usually, these would all be in the same room as the computer itself. Soon, however, people began finding it convenient to have their terminal somewhere else, further away from the computer.[26] Computer scientists like Davies found themselves facing the same challenge Baran had encountered: finding an efficient and economical way of sending data from one device to another[27]. Davies, drawing on the telegraph system, proposed a network of nodes that would divide a message into packets, carrying them to and from time-sharing computers.[28]

When Davies publicly shared his design, he was surprised to learn that Baran had beaten him to it.[29] Both had converged on the same idea — packet-switching — at around the same time.[30] The world was growing ready for the Internet.[31]

***

But there was still work to do. The institution that would actually try to build something using Baran and Davies's ideas was a research center of the US military called the Advanced Research Projects Agency, or ARPA.[32] The US government had created ARPA in response to the Soviet Union's successful launch of Sputnik I – humanity's first satellite – hoping that the agency would fuel America's own technological progress.[33] It did that, and then some.

When ARPA was formed in 1962, computer science wasn't yet an established discipline. But when it started funding research at universities like MIT and Carnegie Mellon, the field grew rapidly.[34] Soon, ARPA had funded several computing research centers across the United

24. Abbate, 24.
25. Abbate, 24.
26. Abbate, 26.
27. Abbate, 26.
28. Abbate, 26.
29. Abbate, 27.
30. Hafner and Lyon, 52.
31. Abbate, 27.
32. Abbate, 36.
33. Abbate, 36.
34. Abbate, 36.

States.[35] The machines at these sites were highly specialized, and ARPA realized that if there was a way for scientists to use them all, it would benefit their research.[36] So ARPA's leadership began brainstorming a network that could connect the country's computers.[37]

To build this network, which would be called the ARPANET, ARPA contracted different parts of the project to researchers at various institutions.[38] Like the decentralization movement today, the ARPANET offered computer scientists from different backgrounds an opportunity to come together, sharing their ideas for an ambitious, unprecedented technology.[39]

ARPANET's engineers could leverage the powerful packet-switching ideas of Paul Baran and Donald Davies, but they still had other challenges to overcome. The major one was that they were trying to connect computers which used vastly different hardware and software, and were therefore incompatible with each other.[40] Users working with more than one machine, for example, wasted valuable time reformatting their code and data.[41]

To solve this problem, ARPANET's engineers decided that each computer would agree to follow certain rules, called protocols.[42] These would act as lingua francas, allowing the computers to communicate with each other. As the engineers designed the network, they separated the protocols into conceptual layers.[43] The "lower" layers would deal with more concrete tasks, like sending electrons through wires; the "higher" layers would handle abstract tasks, like interpreting user commands.[44] Higher layers would build on the layers below, but they would also be independent: an engineer could focus on designing and debugging one layer without worrying about how the other ones worked.[45] First there were two layers, then three; more would be used after the ARPANET was finished.[46] Soon, the idea of a protocol "stack" became a central idea in the design of computer networks, and is frequently used by those trying to decentralize the Internet today.[47]

On the ARPANET, the lowest layer was responsible for getting packets from one computer to another.[48] (This was done by minicomputers that assisted the main ones, assembling and disassembling packets, and guiding them to their destination.)[49] The next layer had a protocol

35. Abbate, 37.
36. Abbate, 46.
37. Abbate, 37.
38. Abbate, 54.
39. Abbate, 54.
40. Abbate, 44; Abbate, 48.
41. Abbate, 48.
42. Abbate, 49; Hafner and Lyon, 145.
43. Abbate, 53.
44. Abbate, 53.
45. Jonathan L. Zittrain, *The Future of the Internet and How to Stop It* (Yale University Press: 2008), Chapter 4.
46. Abbate, 53.
47. "Origin Story," Get DWeb, accessed September 5, 2021, https://getdweb.net/origin-story/.
48. Abbate, 53.
49. Abbate, 61.

for setting up communications between two computers — the initial "hello" between them.[50] The highest layer, called the applications layer, included protocols for the first two applications of the ARPANET. The first allowed a user to access a computer remotely, such as a machine at another university; the second allowed a user to send files from one computer to another.[51]

In September 1969, ARPANET's first computer went online.[52] By the end of the year, ARPANET spanned four institutions;[53] by the end of 1971, it connected fifteen.[54] The next year, ARPA demonstrated the network to the public. At the International Conference of Computer Communications, held at the Hilton in Washington, D.C.[55], a thousand people stared in awe at the network that let them connect to distant computers, and on them, applications as diverse as a weather modelling system, a chess player, and an air traffic simulator.[56]

But the feature that became the ARPANET's killer app[57] took everyone by surprise — including its inventors.[58]

***

ARPA built the ARPANET so scientists could share the powerful computers scattered across the country. This did happen, to some extent: computer scientists collaborated by sharing files over the ARPANET, and they studied the network itself, by analyzing its traffic.[59] Some other kinds of researchers – like physicists, chemists, seismologists, and climate scientists – also used computers on the network, as did military personnel.[60]

But there were problems. The ARPANET didn't come with instructions, and it was difficult for users to learn what resources were available on the network, or how to use them.[61] In addition, ARPANET had been meant to offer an economical way of connecting the country's few large computers. But computers were shrinking in size and cost,[62] and it was becoming easier for scientists to just buy their own machines, instead of having to share them with others.[63]

Instead, ARPANET's users realized that the most useful information they could send over

50. Abbate, 67.
51. Abbate, 67.
52. Abbate, 64.
53. Abbate, 64; Hafner and Lyon, 155.
54. Abbate, 78.
55. Hafner and Lyon, 177.
56. Abbate, 79.
57. Hafner and Lyon, 186.
58. Abbate, 108; Hafner and Lyon, 188.
59. Abbate, 100.
60. Abbate, 102-103.
61. Abbate, 86.
62. "1970s - Early 1980s," Computer History Museum, accessed September 5, 2021, https://www.computerhistory.org/brochures/1970s-early1980s/.
63. Abbate, 105.

the network wasn't scientific data, but the text message. Soon, electronic mail — or email — became the ARPANET's most popular application.[64] It was better than real mail because it was instantaneous, and it was better than the telephone because it was asynchronous: a sender and a receiver didn't have to be online at the same time.[65] The ARPANET had been designed to connect computers, but no one had foreseen how successful it would be at connecting people.[66]

***

Around the same time as the ARPANET, engineers in France and the UK had also been trying to build packet-switching networks.[67] Like the ARPANET, these had their own protocols. First, the challenge for computer scientists had been to connect multiple computers into a single network. Now, the task was to connect entire networks together. Over the next few years, ARPANET's engineers collaborated with researchers in other countries[68] to come up with a design for connecting different networks, or inter-networking.[69]

Eventually, they reached a consensus: a practical and efficient inter-network would use two protocols.[70] The first, called the internetwork protocol, or IP, moved packets from machine to machine. The other, the transmission control protocol, or TCP, protected against errors. It made sure that packets got to their destination and retransmitted them if they got damaged in transit (these errors didn't happen on the ARPANET, but other networks were less reliable).[71]

The new design kept things simple. Most processing would take place on computers, and not the network, which would just serve as a set of pipes for information to flow through.[72] In the 1980s, as personal computers grew ever more affordable, and the two protocols (together known as TCP/IP) became an unofficial standard,[73] the ARPANET evolved from a single network into the Internet, a network of networks.[74]

But even though more and more people were connecting to the Internet, they weren't using it to do much. "If you think back to the 1980s, for somebody to say they were on the Internet, what that really meant is that they had an email account," says David Clark, a computer scientist at MIT who led the Internet's development during the decade.[75] Whether as an email

64. Abbate, 107.
65. Abbate, 107.
66. Abbate, 111.
67. M. Campbell-Kelly. "Data Communications at the National Physical Laboratory (1965-1975)," *IEEE Annals of the History of Computing*, vol. 9, no. 03, pp. 221-247 (1987), doi: 10.1109/MAHC.1987.10023; James Gillies and Robert Cailliau, *How the Web Was Born: The Story of the World Wide Web* (Oxford: Oxford University Press, 2007), 36.
68. Hafner and Lyon, 221.
69. Abbate, 123.
70. Abbate, 130.
71. Abbate, 128; Hafner and Lyon, 224.
72. Abbate, 130, Hafner and Lyon, 226.
73. Abbate, 147.
74. Abbate, 186; Hafner and Lyon, 254
75. David Clark, interview with author, January 23, 2020.

or a file, the Internet was only really useful for sending information from one computer to another: there was no way to navigate the totality of information that existed across the network.

This would soon change, however, when an English scientist named Tim Berners-Lee, growing frustrated with the state of the Internet, would decide to take matters into his own hands.

***

In 1989, Berners-Lee was working as a software engineer at CERN, the particle physics laboratory in Geneva, Switzerland. [76] Scientists from all over the world came to CERN, and they brought with them different kinds of computers, most of which stored data in incompatible formats. Berners-Lee saw that even with the Internet, scientists like him had no easy way to share information with each other. If he wanted to take advantage of the various computers at his disposal, he would have to go through the hassle of writing new programs, or converting data from one format to another.

Berners-Lee realized that the data on the lab's different computers could be thought to be virtual documents. If he could link those documents to each other, this would create the "worldwide web" of information that he envisioned for the Internet. To do this, he came up with the idea of web "pages", which would display information; addresses for finding those pages, called URLs; and an application-layer protocol for sending and receiving those pages over the Internet, called hypertext transfer protocol, or http.

"All of a sudden, the web was something that lots of people could understand," Clark says. "It was a very compelling idea for people who were not computer science professionals."[77]

Over the next few years, as the World Wide Web flourished, an American computer scientist named Brewster Kahle watched with excitement, but also with mounting concern.

The dangers of a centralized Internet, and the path toward decentralization

Kahle realized that although the Web was making the Internet usable for millions of people, it fell short of the resilient network envisioned by people like Paul Baran. The Web was transient: webpages were updated frequently, and if a site ever went down, its content was lost forever.[78]

---

76. TED, "Tim Berners-Lee: The next Web of open, linked data," YouTube video, 16:51, Mar 13, 2009, https://www.youtube.com/watch?v=OM6XIICm_qo.
77. David Clark, interview with author, January 23, 2020.
78. Brewster Kahle, "Siteless Website Possible? If bittorrent is a fileserver without a server, what about a website without a site?" Internet Archive Blogs, published October 22, 2012, http://blog.archive.org/2012/10/22/siteless-website-possible-if-bittorrent-is-a-fileserver-without-a-server-what-about-a-website-without-a-site.

"I said, 'We need a better system here,'" Kahle recalls. "'If this is going to be how we're going to build out contemporary culture, based on this technology, it can't be this shaky.'"[79]

In 1996, Kahle launched a non-profit digital library called the Internet Archive. The Archive's computers had an ambitious mission: they would try to create a backup of the entire Web, visiting its pages and taking snapshots across time. This way, anyone could see what a website looked like in the past.[80]

Over the next two decades, the Archive expanded its storage and added new features. But the Web was growing swiftly too, and Kahle was beginning to realize that the Archive's backup endeavor was simply a band-aid solution.[81] A permanent solution would need an entirely new set of technologies.

In 2014, some friends introduced Kahle to a German entrepreneur named Joachim Lohkamp.[82] A tech-minded businessman, Lohkamp had grown frustrated with the way information was locked up between companies, and even within them. "The mindset was 'information is power' and you really want to protect it," Lohkamp says. "It was not so much about sharing."[83]

Inspired by applications like BitTorrent – called "peer-to-peer" because they rely on many interconnected computers rather than one central server – Lohkamp was designing an application that could serve as a digital fingerprint. The app, called Jolocom, would offer users a decentralized identity, in contrast to identity documents like passports and credit cards, which were issued by "central" authorities like governments and banks.

In Germany, Lohkamp had already formed a community of programmers working on similar projects; in San Francisco, he was trying to expand the conversation. The discussions with Lohkamp struck a chord with Kahle, who was inspired to flesh out his own ideas.[84] "What we need is a Next Generation Web," Kahle wrote on his blog soon after.[85] Two decades had passed since the invention of the original Web, and computer scientists now possessed a technical arsenal that included faster computers, encryption, and blockchain technology. "I've seen each of these pieces work independently, but never pulled together into a new Web," he wrote on his website. "That is what I am challenging us to do."

Kahle started speaking publicly about his ideas, and meeting like-minded individuals. One

---

79. Brewster Kahle, interview with author, May 4, 2021.
80. "About the Internet Archive," Internet Archive, https://archive.org/about.
81. Brewster Kahle, interview with author, May 4, 2021.
82. Wendy Hanamura, interview with author, April 15, 2021; Joachim Lohkamp, email to author, August 25, 2021.
83. Joachim Lohkamp, interview with author, April 30, 2021.
84. Wendy Hanamura, interview with author, April 15, 2021.
85. Brewster Kahle, "Locking the Web Open, a Call for a Distributed Web," published April 22, 2015, http://brewster.kahle.org/2015/04/22/locking-the-web-open-a-call-for-a-distributed-web.

of them was a young computer scientist named Juan Benet, who was working on a problem of immediate relevance to Kahle and the Internet Archive.[86]

According to the Web's protocols, information is identified by its location — *where* it is, represented by its URL. If the link is inaccessible (perhaps because a government has blocked it, or the hosting computer has gone down) then the information at that link is lost forever. Benet was designing a new protocol — an alternative to the Web's http — called Inter-Planetary File System, or IPFS. This identified information based on *what* it was, rather than where it was located. Instead of existing at one link, a website could exist in multiple places on the Internet, and the protocol would allow a user to find all of them.[87] Kahle realized that such a protocol could be a building block for an Internet where information was never lost.

While Benet was working on decentralized storage, and Lohkamp on decentralized identity, a computer scientist named Muneeb Ali was exploring whether it was possible to hit several birds with one stone.

During the first decade of the Web, users would travel from website to website, and blog to blog. In the aughts, this changed. With the rise of social media applications, large number of users began to congregate on the same platform, like Facebook.[88] By attracting users with free, convenient services like applications, content, and storage — all in one, central place — these platforms collected the power to control how hundreds of millions of people interacted with the Internet. Like many others in the technology community, Ali believed that centralized platforms like Facebook were gatekeeping a digital world that was meant to be democratic.[89]

In 2008, a pseudonymous individual published the design for a digital currency called Bitcoin.[90] What made Bitcoin an unusual currency was that instead of relying on a bank to arbitrate financial transactions and prevent fraud, the design relied on cryptography. Ali, who was a graduate student at Princeton in 2014, when Kahle met Lohkamp, believed that the technology underlying Bitcoin, called blockchain, could be useful for much more than finance. If blockchain could allow people to manage their money without banks, perhaps it could also let them use the Internet without depending on companies like Facebook and Dyn. Designing a blockchain-powered, decentralized Internet, Ali joined the ranks of computer scientists like Kahle and Benet. Meanwhile, just a few hours away at MIT, the inventor of the World Wide Web was pursuing a vision of his own.

Tim Berners-Lee's original vision for the Web was of a world where all information was

---

86. Wendy Hanamura, interview with author, April 15, 2021.
87. "What is IPFS?" IPFS, updated June 22, 2021, https://docs.ipfs.io/concepts/what-is-ipfs/#decentralization.
88. Gili Vidan, conversation with author at Web 3.0 conference, Cambridge, Massachusetts, March 3 – 5, 2020.
89. TEDx Talks "Welcome to the new internet | Muneeb Ali | TEDxNewYork," YouTube video, 10:31, Dec 1, 2016, https://www.youtube.com/watch?v=qtOIh93Hvuw.
90. "Who is Sakoshi Nakamuto," *The Economist*, Nov 02, 2015, https://www.economist.com/the-economist-explains/2015/11/02/who-is-satoshi-nakamoto.

inter-connected.[91] Users would have complete control over their data, and that data would be easy to use with different applications. Like others in the decentralization community, Berners-Lee was alarmed that the contemporary Web offered neither: the average user's data was scattered across the Web in incompatible fragments, owned by companies like Facebook, Google, and Twitter.

In response, Berners-Lee was designing Solid, a conceptual rethink of the Web where data would exist independently of applications.[92] According to Solid's design, for example, a user would have complete control over all their social media data: they could share whichever parts they liked, and however long they wanted, with apps such as Facebook or Twitter. And if they wanted to switch platforms, they could easily take their data with them.

***

As Brewster Kahle met the many people who were independently working toward a decentralized Internet, he decided to bring them together.[93] In June of 2016, Kahle and his colleagues at the Internet Archive hosted the Decentralized Web Summit in San Francisco.[94] On the first day, at the Archive's headquarters, more than seventy technologists gathered to share their ideas and discuss the path forward.[95]

"It was so exciting because they had never been together before," says Wendy Hanamura, who helped organize the conference at the Archive. "To bring people from all over the world working on little bits and pieces of this together was mesmerizing."[96]

Young computer scientists like Juan Benet and Muneeb Ali mingled with veterans like Berners-Lee, who invented the Web, and Vint Cerf, who co-designed the Internet's TCP/IP protocols. But the event was meant to be more than just an exclusive meeting of the technological elite. Over the next two days, the participants presented their ideas to hundreds[97] of attendees from the public, explaining how a decentralized Internet would protect their privacy and freedom.[98]

By engaging the public, the event was meant to honor a tenet of decentralization. "Technology is always better when the people who are going to be using it are sitting at the table with you," Hanamura says.[99]

---

91. TED, "Tim Berners-Lee: The next Web of open, linked data."
92. The Solid Project, https://solid.mit.edu.
93. Wendy Hanamura, interview with author, April 15, 2021.
94. Decentralized Web Summit, accessed September 5, 2021, https://2016.decentralizedweb.net
95. "Origin Story," Get DWeb, accessed September 5, 2021, https://getdweb.net/origin-story/.
96. Wendy Hanamura, interview with author, April 15, 2021.
97. "Origin Story," https://getdweb.net/origin-story/.
98. Decentralized Web Summit, https://2016.decentralizedweb.net.
99. Wendy Hanamura, interview with author, April 15, 2021.

But for one attendee, a writer and activist named Mai Ishikawa Sutton, the summit lacked diversity in age, gender, race, and class. “It felt too similar to the same type of Silicon Valley culture [that led to centralization],” she says.[100] In 2018, when the Archive held a second event on decentralization, Sutton felt the conversation still wasn’t inclusive enough. She joined the Archive herself, determined to make a difference at the following year’s event.

“We don't want it to just be tech people talking to other tech people,” Sutton says. “We want to make sure that journalists and media, academics and artists, and students and activists are part of the conversation.” She believes that while there’s still room for improvement, the decentralization space “has got a lot better” with regards to diversity.[101]

Those involved in decentralization generally agree that the movement has also grown in scale[102], although it’s difficult to come up with exact numbers.[103] “Over the past five years, what we've really seen is the result of the influx of a lot of money into the ecosystem,” says Karissa McKelvey, a researcher who specializes in decentralization. “There’s a lot more people involved now, and a lot more diversity in projects, which I think is really cool.”[104]

Many of the projects that were presented at the first decentralization summit have come far. The coronavirus pandemic disrupted plans for a conference in 2020[105] so, in March of 2021, the Internet Archive organized a virtual event for several projects to share updates with the public.[106]

The first presentation was from Filecoin, Juan Benet’s follow-up project to the IPFS protocol. Filecoin is a storage app, meant to offer users a decentralized alternative to services like Apple’s iCloud and Google’s Drive. It uses IPFS in combination with blockchain technology, allowing people to rent out unused space on their hard drives, and to store their own data securely on others’, all over the world.[107]

A representative from Jolocom, Joachim Lohkamp’s decentralized identity app, shared that the company was working with the German government to develop use cases spanning finance, education, and governance.[108] It’s possible that soon, a tourist visiting Germany won’t have to worry about entrusting their passport to a local hotel — they could just identify themselves

---

100. Mai Ishikawa Sutton, interview with author, April 26, 2021.
101. Mai Ishikawa Sutton, interview with author, April 26, 2021.
102. Joachim Lohkamp, interview with author, April 30, 2021; Juan Ortiz Freuler, interview with author, April 13, 2021; Karissa McKelvey, interview with author, April 26, 2021; Wendy Hanamura, interview with author, April 15, 2021.
103. Juan Ortiz Freuler, interview with author, April 13, 2021; Ruben Verbogh, interview with author, April 28, 2021;
104. Karissa McKelvey, interview with author, April 26, 2021.
105. Wendy Hanamura, interview with author, April 15, 2021
106. Mai Ishikawa Sutton, "DWeb Meetup March 2021: Latest in the DWeb Ecosystem," Internet Archive Blogs, published April 22, 2021, http://blog.archive.org/2021/04/22/dweb-meetup-march-2021-latest-in-the-dweb-ecosystem.
107. FileCoin, https://filecoin.io.
108. Sutton, "DWeb Meetup March 2021."

with the Jolocom app.

Muneeb Ali shared that his company, Hiro Systems, had developed a set of tools that programmers could use to build all kinds of applications relying on the Bitcoin blockchain — everything from decentralized DNS to decentralized WhatsApp.[109]

Meanwhile, Ruben Verbogh, an ambassador for the Solid project, says that the UK's NHS is experimenting with the technology to offer citizens control of their health data, while allowing them to share it easily with doctors; and the Flemish government in Belgium is considering using Solid in governance.[110]

It's taken time for all these applications to mature, largely because designing decentralized software is hard. "It turns out that these technologies take more work than we were hoping, to get to a certain point," says Paul Frazee, a programmer who has worked on several projects.[111]

As they develop, decentralized systems must be in conversation with each other. Projects like Filecoin, Hiro, and Jolocom are all working on different pieces of the puzzle. To offer the average Internet user a compelling alternative to existing centralized applications, those pieces must fit together seamlessly. "That's really a very important point for decentralization," says Joachim Lohkamp. "You want to think not in terms of creating dependencies, but having a more modular and open architecture where the pieces are interchangeable."[112]

By adopting that philosophy, the decentralization movement is drawing inspiration from the original Internet. "When Tim Berners-Lee wanted to invent the Web, he didn't have to talk to Vint Cerf, or David Clark, or anyone else," says Juan Ortiz Freuler, a research affiliate at Harvard's Berkman Klein Center.[113]

If Berners-Lee could easily build the Web on top of the Internet he inherited, it was because the underlying protocols, like TCP/IP, had been designed with flexibility in mind.[114] In the same way, decentralization projects are slowly laying the building blocks for the Internet to come.[115] And by informing each other, those pieces are maximizing the movement's chances of success. "A lot of these systems connect to each other in ways that allow them to scale much better," Freuler says.[116]

***

Although the decentralization community has been able to build foundational infrastructure

109. Sutton, "DWeb Meetup March 2021."
110. Ruben Verbogh, interview with author, April 28, 2021.
111. Paul Frazee, interview with author, April 22, 2021.
112. Joachim Lohkamp, interview with author, April 30, 2021.
113. Juan Ortiz Freuler, interview with author, April 13, 2021.
114. David Clark, interview with author, January 23, 2020;
115. Ross Schulman, interview with author, April 21, 2021.
116. Juan Ortiz Freuler, interview with author, April 13, 2021.

since the first summit in 2016[117], some major challenges remain. Not all the projects from the early days are still around — many have fizzled out.[118]

In 2017, Paul Frazee launched Beaker, a peer-to-peer Web browser (like BitTorrent, but for websites). The project presented several "novel pains in the butt," he says, like the fact that he could only get it to work on desktop browsers, not mobile platforms like Android and iOS. Now, Frazee has moved on to CTZN, a decentralized social media application that will resemble Twitter, but will allow users to control the app's content moderation and curation algorithms. Through trial and error, Frazee has learned sobering lessons about creating decentralized technology. "The people that were involved at the beginning [of the movement] that are still involved have gotten less starry-eyed," he says. "Some of the expectations of success have shifted."[119]

The most visible decentralization projects tend to be the ones that have successfully secured large amounts of funding. [120] Most of those, like Hiro and Filecoin, rely on blockchain[121], a technology others believe has yet to deliver on its promises. [122] "A lot of the momentum that blockchain is experiencing is connected to a place of speculative trends and forces," [123] says Joshua Tan, a computer scientist at Stanford's Digital Civil Society Lab, referring to the unpredictable fluctuations of cryptocurrencies like Bitcoin.

There's some way to go before decentralized apps replace existing ones. Most available today are too slow for users to want to use, says Karissa McKelvey.[124] Ross Schulman, a researcher at the New America think tank, says that over the next few years, the movement has to develop apps that people will use "not because they're decentralized, but because they're the best tool for the job."[125]

---

117. Adin Schmahmann, interview with author, May 5, 2021; Juan Ortiz Freuler, interview with author, April 13, 2021; Karissa McKelvey, interview with author, April 26, 2021; Nicola Greco, interview with author, April 21, 2021; Paul Frazee, interview with author, April 22, 2021; Ross Schulman, interview with author, April 21, 2021.
118. Paul Frazee, interview with author, April 22, 2021; Janus Kopfstein, "The Mission to Decentralize the Internet," *New Yorker*, https://www.newyorker.com/tech/annals-of-technology/the-mission-to-decentralize-the-internet; Morgen Peck, "Pumpers, Dumpers, and Shills: The Skycoin Saga," *New Yorker*, https://www.newyorker.com/tech/annals-of-technology/pumpers-dumpers-and-shills-the-skycoin-saga.
119. Paul Frazee, interview with author, April 22, 2021.
120. Karissa McKelvey, interview with author, April 26, 2021; Paul Frazee, interview with author, April 22, 2021.
121. Juan Ortiz Freuler, interview with author, April 13, 2021; Karissa McKelvey, interview with author, April 26, 2021; Paul Frazee, interview with author, April 22, 2021; Wendy Hanamura, interview with author, April 15, 2021.
122. Joshua Tan, interview with author, April 14, 2021; Paul Frazee, interview with author, April 22, 2021; Ross Schulman, interview with author, April 21, 2021.
123. Joshua Tan, interview with author, April 14, 2021.
124. Karissa McKelvey, interview with author, April 26, 2021.
125. Ross Schulman, interview with author, April 21, 2021.

To get there, computer scientists have to solve some thorny technical problems. Nicola Greco, a computer scientist who helped build Filecoin,[126] says one of the difficulties of offering users a more private Web is that computer scientists still don't have an efficient way to do useful things with data when it's encrypted (and looks like a stream of gibberish).[127]

Others, like Gina Abrams at Hiro, say a major challenge is designing the best experience for users.[128] Some prefer an interface that hides the complicated innards of a decentralized app; others want access to that complexity so they can customize apps for their needs. "There's been a tension between tradeoffs," Abrams says, "What does the decentralized community value?"[129]

For some, the decentralization movement isn't just about inventing better technology; it's about developing business models that are more innovative than the prevailing one, where companies compete ruthlessly for data they can sell to the highest bidder.[130] While some decentralized projects like Hiro and Filecoin are directly taking on the likes of Facebook and Google, those working on Solid believe in the merits of a collaborative approach.[131] For example, says Ruben Verbogh, instead of a world where people are forced to choose between Fitbit and Apple devices to monitor their health metrics, Solid advocates for one where a user can easily share their data with both companies — and with access to more data, those companies can offer their users better analytics. "We look at use cases where there's a win on both sides," Verbogh says. "We look at cases where companies can be more effective and efficient, but people can also take control of their data."[132]

***

From conversations with dozens[133] of individuals working on decentralization, one thing becomes clear: those people care deeply about technology, and the world they are building it

126. Juan Benet, Nicola Greco, et al. "FileCoin Whitepaper," https://filecoin.io/filecoin.pdf.
127. Nicola Greco, interview with author, April 21, 2021.
128. Gina Abrams, interview with author, May 6, 2021; Paul Frazee, interview with author, April 22, 2021.
129. Gina Abrams, interview with author, May 6, 2021.
130. Lawrence Lessig, Web 3.0 conference, Cambridge, Massachusetts, March 5, 2020; Ruben Verbogh, interview with author, April 28, 2021.
131. Ruben Verbogh, interview with author, April 28, 2021.
132. Ruben Verbogh, interview with author, April 28, 2021.
133. Allen Gunn, interview with author, April 23, 2021; Brewster Kahle, interview with author, May 4, 2021; Gina Abrams, interview with author, May 6, 2021; Joachim Lohkamp, interview with author, April 30, 2021; Joshua Tan, interview with author, April 14, 2021; Juan Ortiz Freuler, interview with author, April 13, 2021; Justin Bingham, interview with author, May 6, 2021; Karissa McKelvey, interview with author, April 26, 2021; Mai Ishikawa Sutton, interview with author, April 26, 2021; Paul Frazee, interview with author, April 22, 2021; Ross Schulman, interview with author, April 21, 2021; Wendy Hanamura, interview with author, April 15, 2021; multiple participants at the Web 3.0 conference in Cambridge, Massachusetts, March 3-5, 2020.

for. They agree on the destination: an Internet that's freer, safer, and more egalitarian. There is less consensus, however, on the best way to get there. Some, like Joshua Tan, think that even calling decentralization a movement is too much; he says he prefers the word 'community.'[134]

Allen Gunn, a technology consultant who helped the Internet Archive conduct the first summit in 2016, holds a similar view. "Right now, [the summits are] kind of like a music festival that will let any kind of music play on stage," he says. "The downside of eclecticism is it's a much harder thing to forge focused identity or focus purpose around."[135]

Others disagree, arguing that a tolerance for difference is exactly what they are trying to encode into the Internet of the future. "The movement is comprised by people who have very different backgrounds, and very different political ideologies," says Juan Freuler at the Berkman Klein Center. "Having a space where all of these people can meet, give each other reasons to stick together, it's extremely valuable."[136]

Brewster Kahle, the founder of the Archive, echoes that sentiment. "We don't all have to agree," he says. "Let's build a system with many winners." There are challenges ahead, but Kahle is determined not to overlook the substantial progress that's been made. For him, decentralization can't happen soon enough. After all, he points out, he's already been trying to fix the Internet for twenty-five years. Kahle throws up his arms. "What's taking so long?"[137]

134. Joshua Tan, interview with author, April 14, 2021.
135. Allen Gunn, interview with author, April 23, 2021.
136. Juan Ortiz Freuler, interview with author, April 13, 2021.
137. Brewster Kahle, interview with author, May 4, 2021.