DIRTY BITS

An Environmental History of Computing

Nathan Ensmenger
School of Informatics and Computing
Indiana University

March 18, 2016

NOTE: a version of this paper has now been published in the journal Technology & Culture

Ensmenger, Nathan. "The Environmental History of Computing." Technology and Culture, vol. 59 no. 5, 2018, pp. S7-S33. Project MUSE, doi:10.1353/tech.2018.0148

It can be found at https://muse.jhu.edu/article/712112

The Materiality of the Virtual

This project grew out of a course I taught several years ago on "information and computer ethics" for undergraduate majors in our Informatics & Computing program. Most of the topics that I covered in class were already at least vaguely familiar to my students: privacy, intellectual property, cyber-crime, professional ethics, and the ethics of design.

There was one set of topics, however, that proved almost universally novel and disturbing, even to this reasonably well-informed (and generally very jaded) audience. This was a series of lectures in which I explored the environmental consequences of electronic digital computing. As we traced the global life-cycle of a typical laptop computer or cellphone from its material origins in rare earth element mines in Africa and South America to its manufacture and assembly in the factory cities of China through its transportation and distribution to retail stores and households across America and finally to its eventual disposal in places like the slums of Agbogbloshie, Ghana, the students discovered that the computer industry is built on more than just abstractions, algorithms, and information. Whether it was studying the toxic byproducts of semiconductor manufacture (there are twenty-nine EPA Superfund sites in Silicon Valley alone, the largest concentration in the nation), or the enormous amounts of energy and water consumed daily by massive Google and Facebook server-farms, or the use of child labor in the "computer graveyards" of the developing world, they were forced to confront the fact that computer power comes at a cost, and that the physical infrastructure that enables their virtual interactions are resource-intensive, pollution-producing, and potentially damaging to the environment. For many of these aspiring computer professionals, this was a sobering reality.

The point of these lectures was not to be alarmist, but to make an important point about the material underpinnings of the information revolution. The public conversation about the about the social implications of computer technology must consider its costs as well as its benefits. But it occurred to me as I was delivering my grim message that there were also valuable lessons to be learned for historians of computing. By focusing on the materiality of the electronic digital computer as constructed technology, I was forced to think about tangible objects, specific social contexts, and particular times and places in ways that I had not done previously.

And so this paper is an attempt to develop what I am calling the "an environmental history of computing." As most of my readers will be aware, environmental history is about more than just the study of the environmental "impacts" of human activity. Environmental historians are interested both how humans shape their environment and are shaped by it. Environmental history emphasizes the active role that nature plays in influencing human affairs. And finally, it is concerned with the changing ways that humans perceive and understand the natural world.

For historians of technology, environmental history has proven an enormously productive tool for thinking with, a way to expand the scope of our field to include new actors, places, and questions. I am particularly interested in ways in which focusing the material underpinnings of the digital economy allows me to explore a global history of computing that encompasses surprising and heretofore neglected participants in the digital economy.

My goal is to develop a history of computing that moves beyond inventors, algorithms, and information. The phrase "computer power" is more than just a metaphor; when we look beyond the consumption of the digital, and focus on the physical infrastructure that makes our online interactions possible, we realize that, just as with more traditional forms of technological and industrial development, computer power comes as a cost. From Bitcoin "mines" to server "farms" to data "warehouses", virtual commodities can have surprisingly material dimensions. They can be resource-intensive, pollution-producing, and potentially damaging to the environment.

In this chapter, I will outline several approaches to integrating the methods and insights of environmental history into the history of computing.

But since high-level, historiographical discussions tend to be abstract and difficult to follow, I am going to organize this introduction around a central narrative, which I like to call…

A Bitminer's Blues

On October 31, 2008, a person or group of persons operating under the pseudonym Satoshi Nakamoto posted a white paper to a cryptography mailing list that proposed a "new electronic cash system that's fully peer-to-peer, with no trusted third party." He (or she, or they) called their new currency Bitcoin. The central innovation of Bitcoin was a distributed virtual ledgerbook called a blockchain. In January of 2009, Nakamoto made the first entry in this blockchain, known as the Genesis Block, and in the process of doing so "mined" the first Bitcoin (more on that word mining later).

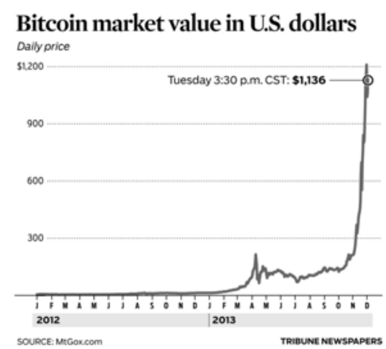

Figure 1: The market capitalization of Bitcoin peaked in late 2013.

Within a few months of its invention, Bitcoin was being actively traded on Internet forums. By early 2010 it was exchanged for the first time for a real-world commodity (two Domino's pizzas). In early 2011 the value of a Bitcoin reached parity with the US dollar.

At this point interest in Bitcoin was nevertheless still largely confined to cryptography experts, computer enthusiasts, financial speculators, and techno-libertarians. In 2010, Bitcoin had become the currency of choice of the Silk Road, an online exchange of illegal drugs, child pornography, and other illicit goods, but this was, at best, a very mixed blessing.

But in 2013 interest in Bitcoin grew exponentially. It began attracting serious investors, including the Winklevoss twins (of *Social Network* notoriety). More importantly, it began being used in China as means of avoiding currency export restrictions. For a brief period in 2013 the value of a single Bitcoin rivaled that of an ounce of gold (that other currency beloved of libertarians) , and the total market value of the Bitcoin economy reached almost $14 billion. Today it is possible to use this virtual currency to purchase goods from more than 12,000 retailers, including such giants as Dell, Microsoft, Amazon, and Victoria's Secret. There are even Bitcoin ATMs!

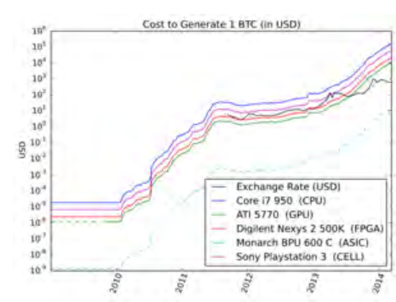

Figure 2: The cost of Bitcoin mining (USD).

There are many interesting political, economic, and technological issues raised by the existence of a thriving international Bitcoin market, including important questions about legal jurisdiction, government regulation, and the ontology of money.

What interests me about Bitcoin, however, is that it makes tangible the relationship between a virtual digital resource and the real-world material inputs required to make that product a "reality." In this case, the input is electricity — and by extension, the coal, oil, gas, water, and/or uranium required to produce that electricity.

To give you an example of the scale of this relationship, on September 30, 2013, there were 4000 Bitcoins mined and 44,494 Bitcoin transactions processed. The total energy required on this day was 18,465.02 megawatts. The total cost of this electricity was $2,769,752.25.

To provide some context for these numbers, this puts the Bitcoin network somewhere between the nation states of Iceland and Ireland in terms of annual energy consumption. The carbon footprint of the Bitcoin network rivals that of the entire island of Cyprus.[1] And both the energy use and greenhouse gas emissions of the Bitcoin network are expected to — and, in fact, are designed to — increase directly with the market value of the digital currency.

I should say that there is some controversy over these figures, and I have seen much lower estimates. But whether the daily energy requirements of Bitcoin are 5,000 megawatt/hrs or 15,000, the true figure will still end up being considerably more than zero.

For most people the fact that virtual commodities consume any material resources comes as a complete surprise. In the era of the Internet, we have become accustomed to thinking of information as being "free". We might pay a monthly fee for an Internet connection, but most of the services that we access online cost us nothing. As more and more goods and services shift from bricks-and-mortar providers into the invisible realm of Cyberspace, its seems to many that we are moving to a less materially-oriented and environmentally impactful economy. This is particularly true as more of our virtual services move to the so-called "Cloud."

What is most shocking about the case Bitcoin is that its massive energy consumption is in fact a feature, and not a bug. In order to circumvent the need for a central authority to create and regulate its currency, the Bitcoin network relies on a shared, distributed record of transactions known as a *blockchain*. In order to contribute to the Blockchain — either to validate a transaction or to create new currency — participants must demonstrate what cryptographers call "proof of work."

In essence, mining a Bitcoin requires the solution to an arbitrary (and meaningless) cryptographic puzzle that can only be solved via brute-force approaches. And to make matters worse, the difficulty of that puzzle is being constantly adjusted. The computational work required to mine a Bitcoin increases exponentially with every leading bit that is required for validation. The required length is adjusted approximately every two weeks, meaning that

[1] Based on those numbers, the total carbon footprint of Bitcoin mining would translate to 8.25 megatons of CO2 per year. That's 0.03 percent of the world's total greenhouse gas output, or equivalent to that of the nation of Cyprus. And that number would scale directly with the market value of the Bitcoin. If Bitcoin's value reaches $100,000, that impact will reach 3 percent of the world's total, or that of Germany. At $1 million – which seems farcical but which may not be out of the realm of possibility given the artificially limited bitcoin supply – this impact rises to 8.25 gigatonnes, or 30 percent of today's global output, and equivalent to that of China and Japan combined.

Bitcoin mining is an increasingly expensive (and competitive) endeavor. No matter how much computing power is thrown at the Bitcoin puzzle, mining a Bitcoin is designed to require about 10 minutes of processing power. It is an infinite sink for computing power — and material resources.

And the size of the Bitcoin energy sink is enormous: as of today, Bitcoin's worldwide computational output is closing on on 200 exaflops — or 8 times the combined capacity of the top 500 supercomputers in the world. All of this power devoted to the production of an entirely virtual commodity whose real-world value is, at best, highly dubious…

The Death of the Digital Wildcatter

For a short period after the introduction of the currency, it was possible to mine Bitcoins in a cost-effective manner using a standard personal computer. (This was, in fact, one of its original appeals: make money at home — literally!) Very quickly, however, the energy costs required to process the hashing algorithm began to outweigh the market value of the Bitcoins produced. Miners quickly switched to running their algorithms on high-end video graphics cards, which for another short period was cost-effective.[2]

Today, it is impossible to mine Bitcoins profitably without a significant investment in high-value capital equipment. The vast majority of Bitcoin mining is being done by a small number of specialized operators. This is indeed one of the tragic ironies of Bitcoin, the melancholy refrain of what I like to call the Bitminer's Blues: that a technology which began as a techno-libertarian fantasy of freedom from large corporations and centralized banks has now become the exclusive province of large-scale, capital-intensive industrial conglomerates.

In their vision of yet another California Gold Rush, the Bitcoiner's imagined themselves as heroic '49ers; what they ended up with was the Empire Mining Corporation.

Because its value is so directly tied to both the US dollar and the market price of electricity, the case of Bitcoin provides a particularly compelling example of the relationship between virtual things and material reality. But a similar relationship could be revealed for every digital activity and commodity, from the video games that we play to the email that we send to the websites currently open in our browsers. Everything we do online ties us to a global network of resources firmly grounded in material substance and physical environments. Cyberspace does not exist outside or above our planet, but is rather firmly intertwined with its most precious resources, its most vulnerable populations, and its most seemingly isolated places.

[2] The amount of electrical power that digital devices require varies dramatically with the computational work that they are being required to perform. A laptop computer running idle (meaning with no significant applications running) might pull 25 watts. A desktop computer running a graphics-intensive video game (or the Bitcoin algorithm) might easily draw 500 watts.

Information as Infrastructure

I grew up in a steel town in Pennsylvania. The presence of the factory in our community was visceral and omnipresent: the plant stretched like a smoking dragon for five miles along the river that ran through the center of our town. We could read its moods in the sounds and rhythms of its daily operations. We went to bed to the glow of its furnaces reflected in the night sky; in the morning we would wipe its dust from our windows and automobiles. The smell of it was always on the wind. Most of us knew at least one person who had been injured or killed working there. And yet for the most part we were willing to live with the environmental hazards of the steel industry because we recognized that it was part of a larger give-and-take between technology and society: in exchange for inconvenience and risk we received jobs and security, parks and schools, and the promise, at least, of a middle-class American life-style.

Figure 3: Bethlehem Steel Works, circa 1908.

But even as I was growing up the United States was already transitioning into what the sociologist Daniel Bell termed a "post-industrial" society. Post-industrial societies, Bell argued, were characterized by a shift from manufacturing to services, from manual to cognitive labor, from production to consumption. Citizens in the post-industrial society would be increasingly reliant on technology, but not on the traditional technologies of industrial modernity — roads and bridges, assembly lines and automobiles, factories and farm equipment — but rather on information and communications technologies.

Figure 4: US Army Computing Division, circa 1919.

Indeed, the very meaning of the word technology has shifted enormously in the past several decades. In most contemporary contexts, the word technology has come to mean computer and information technology. When educators advocate for more technology in the classroom, medical practitioners for more technology in the hospital, and economists for the development of a more technology-proficient workforce, they are not talking about filing cabinets, stethoscopes, or drill-press operators: what they are calling for is more computers, computer-based diagnostic systems, and computer-savvy technicians.

Figure 5: Sears Roebuck Data Division, circa 1913.

For many middle-class, white-collar, eco-sensitive Americans, the shift from industrial to post-industrial society seems to be a largely positive development. While the dirty and dangerous work of industrial manufacturing might not have disappeared entirely, at least it is no longer visible. It happens in other parts of the world, and by (and to) other kinds of people. And where so many of humankind's other great technological accomplishments have been compromised by war, disease, pollution, and other unintended, and undesirable consequences, information technology does appear to be clean, safe, and of relatively low impact on the environment. Indeed, the seemingly inexorable march of Moore's Law towards smaller, faster, and more powerful computers serves for many American the last remaining remnant of our long tradition of technology-driven utopianism.

If we focus solely on the consumption-side of computing and information technology, this model of a radical, discontinuous break with our industrial history does seem believable. For the most part, we experience only the positive benefits of information technology, in large part because we experience them primarily as consumers. Our digital devices provide us with useful information and connections, they amuse and entertain us, they are our helpmates and constant companions. We don't just utilize our technologies, we have relationships with them.

But if we look at the production-side of computer technology, at the vast web of wires, cables, towers, generators, and other physical equipment that underlies the apparently virtual realm of Cyberspace, the digital present does not seem quite so discontinuous with our industrial past.

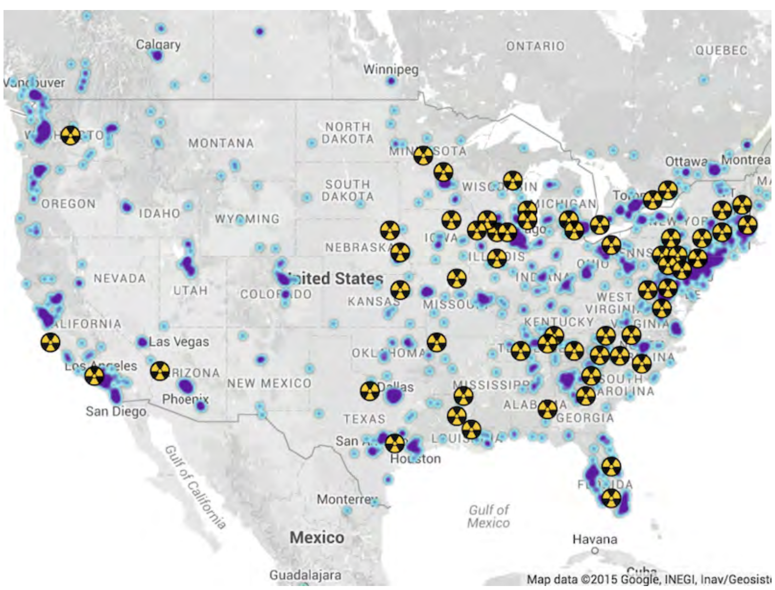

Figure 6: Bitcoin activity (purple) mapped against the location of nuclear power plants in the United States

Let us consider, for a moment, the infrastructure of the Bitcoin network. Since the key input to the Bitcoin mining process is electrons, we should not be surprised to see the virtual map of Bitcoin activity corresponds closely to physical infrastructure of the electrical power grid.

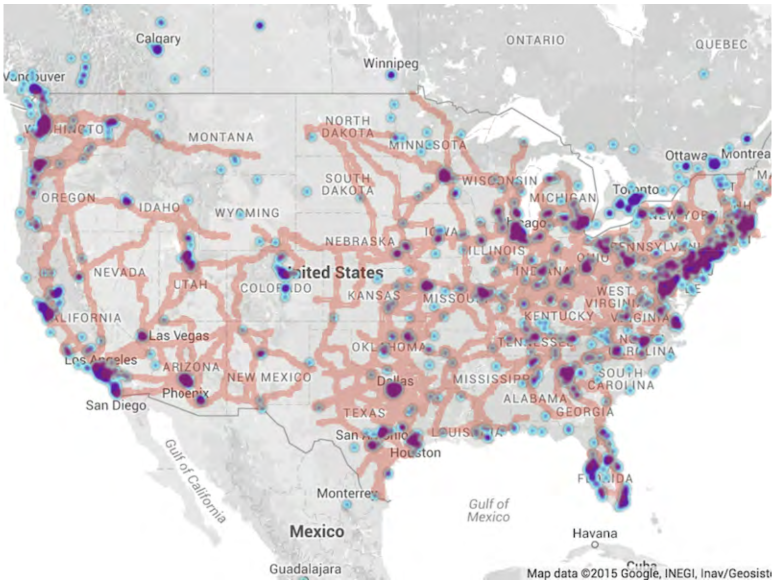

Figure 7: We see a similar correspondence between the Bitcoin network and the power grid distribution infrastructure (red). Cyberspace is **not** evenly distributed!

This map should look familiar: all of you have read your William Cronon, and so are at least vaguely familiar with he story of the American railroad system and how it helped knit together — physically, economically, and symbolically — the increasingly United States in the mid-19th century.

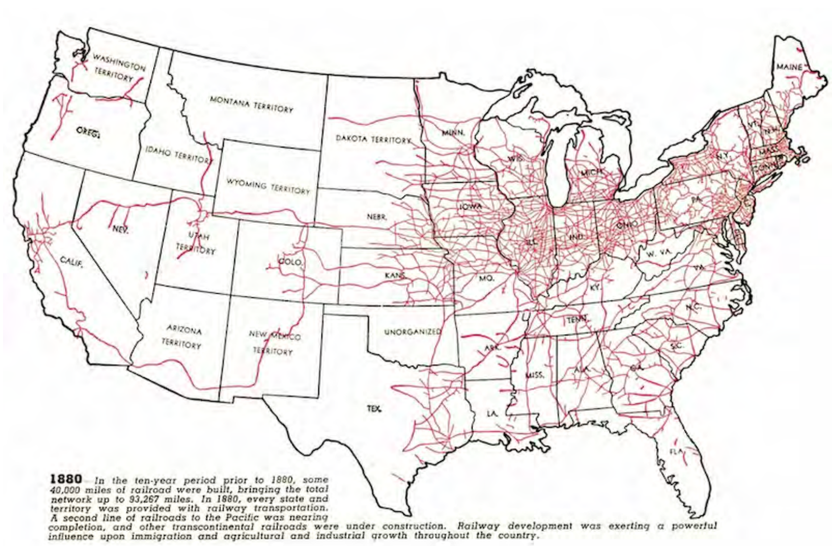

We often tell the story of the railroad in terms of the history of transportation technology.

But equally important is the history of information technology.

The railroad network was enabled by another kind of network, which was the telegraph. The two technologies grew up together; in many ways they were co-constructed.

The railroad needed the telegraph for the purposes of communication and control. A series of nasty railroad accidents in the early 19th century highlighted the need for careful coordination of traffic, and by the end of the 1850s most railroads had adopted the telegraph as a mechanism for signaling and traffic control.

Figure 8: The railroad and the telegraph networks were built on top of one another.

At the same time, the railroad served as the perfect foundation for the growth of the telegraph system. Railroads were flat, cleared pathways along which is was easy to construct and maintain the long strings of copper wire required by a communications networks.

Figure 9: John Gast, *American Progress* (1872)

The two systems almost always were constructed simultaneously. Collectively they enabled the 19th century annihilation of space and time that made possible the exploration — and exploitation — of vast amounts of previously unclaimed or inaccessible resources. Consider these classic images from the era of Manifest destiny: Lady Liberty might be leading the people, but she was trailing a telegraph wire and the railroad followed closely at her heels. From the American West to Southern Africa, the link between information and empire was always apparent.

Figure 10: Cecil Rhodes and his Transcontinental Telegraph

This is a map of the Southern Pacific Railroad network from its heyday in the early 20th century:

This is another famous map, this time of the early ARPAnet.

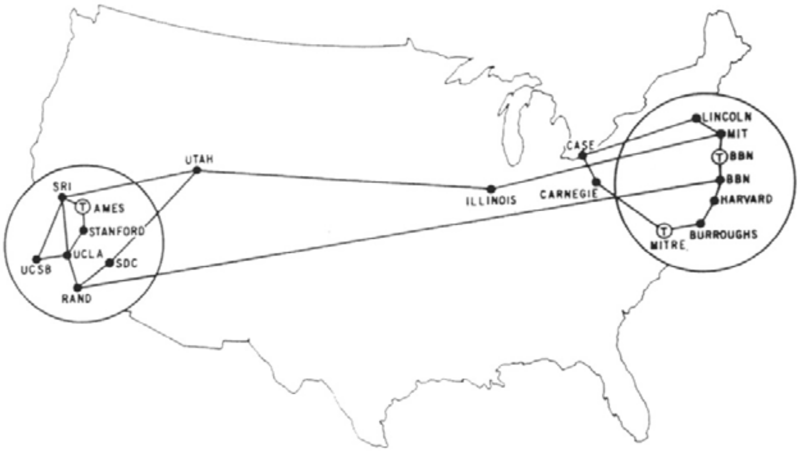

MAP 4 September 1971

Although one is a physical map, and the other is symbolic, you can see their similarity.

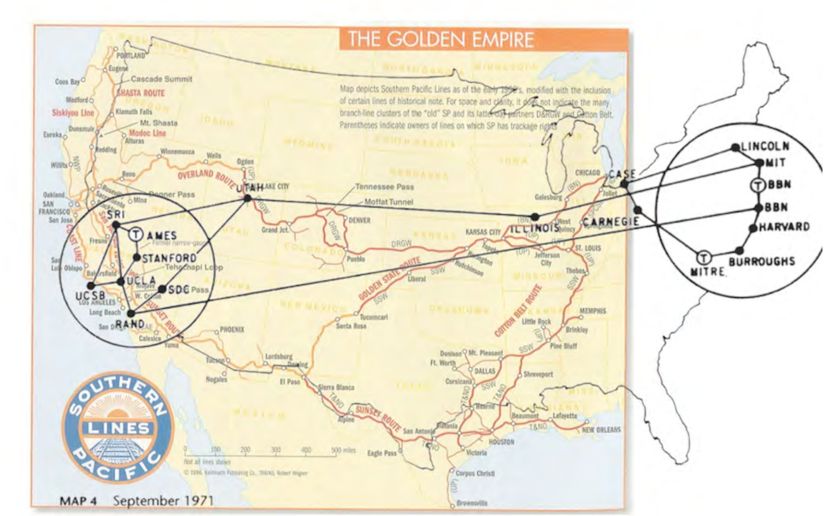

Here is an even better example: this time using the NSFNet from the early 1990s. Again, see how closely the infrastructure of the computer network aligns with the traditional geography of the railroad network.

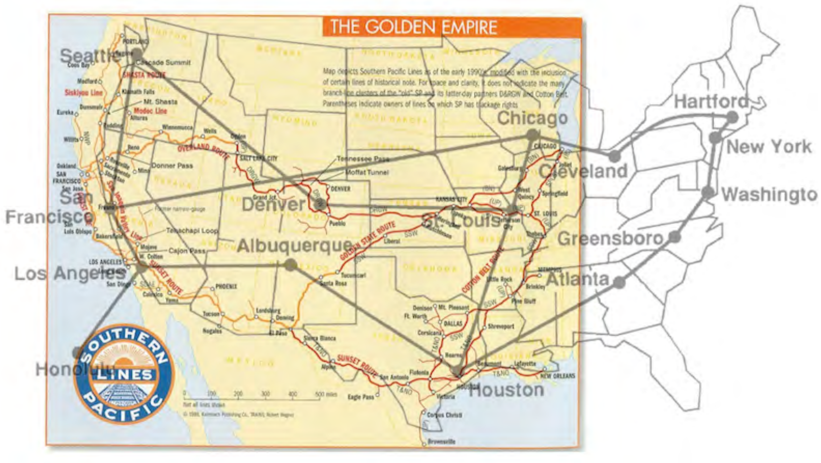

This is not a coincidence:

The Southern Pacific Communications Company, a unit of the larger railroad company, maintained a series of microwave towers along the right-of-way provided by its railroad lines, and in the early 1970s had started laying down fiber optic cable. Eventually this unit would spin-off and rename itself the Southern Pacific Railroad Internal Network Telecommunications, or SPRINT. Today SPRINT is the 3rd largest telecommunications carrier in the United States, and is a Tier-1 Internet Service provider, operating major segments of the national Internet backbone infrastructure.

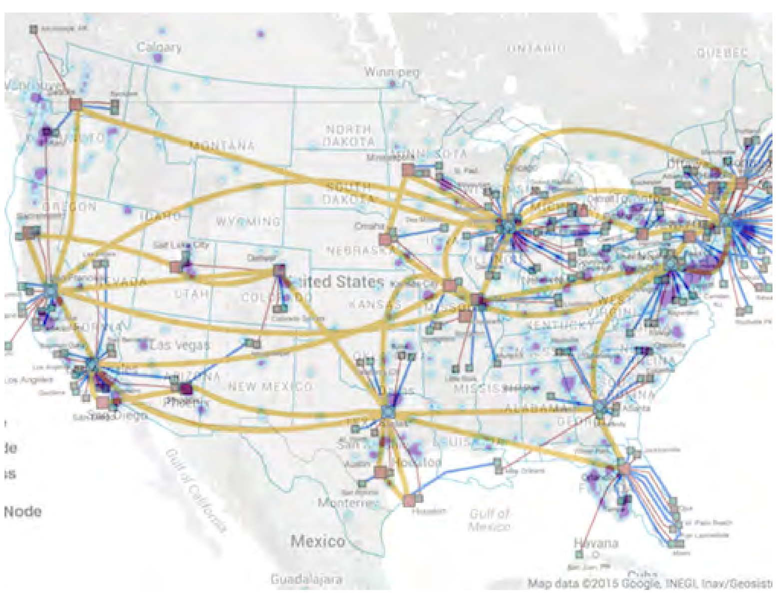

Figure 11: Yet another illustration of the close correspondence between multiple layers of information and communications infrastructure.

We have become accustomed to thinking about information technology as making place irrelevant. Via our digital devices, we can connect to anyone, anywhere, from anywhere. The global Internet is both omnipresent and invisible. It is the ideal infrastructure; it transcends traditional geography; it is everywhere and nowhere.

In fact, many of most significant social and economic nodes of the Information Society sit at the intersection of traditional, material infrastructures like railroads, power grids, and river systems. Geography shapes technology, and vice versa.

The Information Infrastructure of the 21st Century is built around the bones of the 19th century transportation and communication network.

When we look closely at the flows of material that make the virtual possible, we discover that many of most significant social and economic nodes of the Information Society sit at the intersection of traditional, material infrastructures like railroads, power grids, and river systems. Geography shapes technology, and vice versa.

Of course, Bitcoin is not the only virtual activity that requires a close connection to the physical environment. In fact, most Bitcoin "mines" are co-located with other centralized computational activities. We call these centers of activity data warehouses, data farms, server farms, or, more recently, the Cloud.

The Cloud is a Factory

The Cloud is a brilliant and wickedly misleading metaphor. It implies both ubiquity and ethereally. What tech services companies want us to know about the cloud is that it is always available, largely transparent its users, and never needs much thinking about. In other words, the Cloud is the perfect infrastructure. But unlike traditional infrastructure like roads and bridges and sewer systems, the Cloud requires no violence to the physical environment. It floats above, silent and unobtrusive, a force of nature rather than a human-built technology.

Recently Google has made available a series of images of its Cloud facilities. For the most part, they are beautiful and quieting images, a post-industrial reinterpretation the visual genre that David Nye has referred to as the "technological sublime." [Although I should note that those of you familiar with Paul Edward's work on the Closed World might find these sterile and uninhabited landscapes also a little disquieting]

But as the media historian and social theorist John Durham Peters has suggested, although the rhetoric of the Cloud has mobilized by the Internet giants to invoke images of ethereal other-worldliness, Google is a fire-god, and not a spirit of the air. In 2011, for example, Google data centers used more than 2.3 billion kilowatt-hours of electricity, which represented about 2 percent of the annual electricity consumption of the entire United States. That same year, Facebook consumed an additional 532 million kilowatt hours. The collective global demand for power for digital data centers accounts for the output of roughly 30 nuclear power plants, according to a recent article in the New York Times, with server farms in the United States accounting for as much as 1/3 of this total load.

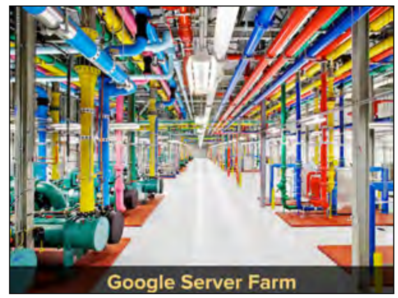

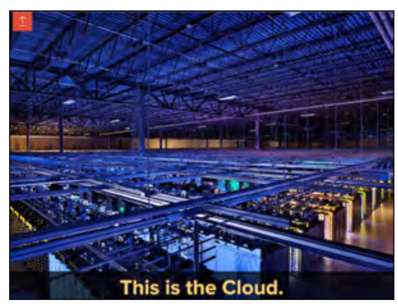

Of course, where energy is used, heat is created. Cooling even a medium-size high-density server farm can require as much as 360,000 gallons of water a day. The new NSA Intelligence Community Comprehensive National Cybersecurity Initiative Data Center in Bluffdale, Utah will consume 1.7 million gallons every day. Such consumption patterns stretch the limits of almost any municipal water supply, and given the looming global shortage of clean water, water scarcities represent one of the many unanticipated consequences of computing whose implications are only just beginning to be realized.

One of the great myths of the global Internet is that it makes place irrelevant. Users, servers, and resources can be located anywhere in the world, and still be equally accessible — or so the story goes.

Across the United States, and presumably the world, the invisible infrastructure of the Internet follows the contours of geography and human settlement. When big data providers like Google and Microsoft locate their server farms, they have to take into consideration the same factors that guided more traditional manufacturers in the Industrial Era: transportation networks, water supplies, power grids, labor markets, and a local political climate amenable to development. When Microsoft chose to locate a new data warehouse in the small farm town of Quincy, WA, for example, it was pursing not only the availability of cheap hydroelectric power, but also lucrative tax breaks. Like any factory moving into a community, they promised local residents jobs, tax revenue, and the ability to maintain their rural lifestyle in the face of a decades long decline in the local agricultural economy.

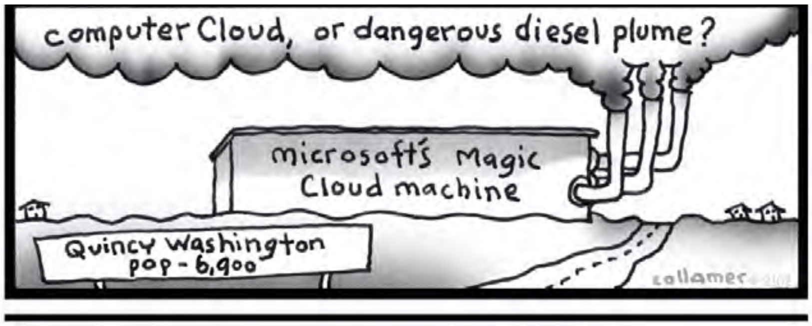

What the Quincy residents did not expect was pressure from Microsoft to invest municipal funds into the further expansion of its electrical grid, or the 40 giant diesel generators that Microsoft installed to provide backup power to its data facility. These generators, which were located near an elementary school, are a source of diesel particulate pollution, which is a potential carcinogen. A similar Microsoft facility located in Santa Clara had recently been identified as one of the largest polluters in the San Francisco Bay Area. Because the performance of data centers is measured in terms of continuous "up-time" (rather than in terms of cost-effectiveness, for example, or energy efficiency) the generators at Quincy were being run almost continuously, rather than for the limited periods that local residents had been lead to expect. To add insult to injury, Microsoft, as part of a related dispute with the local government over its electricity rates, began using power at its facility in what it acknowledged was an "unnecessarily wasteful" way until the fines that had been imposed on it were substantially reduced.

If the Cloud is indeed a factory, and the Internet a form of infrastructure, what does this get us, in terms of understanding the larger connections between computing and the environment?

To begin with, by focusing on the places where large-scale computation happens, on the people who work in those places, and the embeddedness of those sites in specific social and geographical landscapes, we can re-situate the history of computing in the larger history of the American industrial development, and ask new and important questions about labor, capital, politics, and power (in a variety of meanings of that word).

I have elsewhere argued that that computerization of modern society can only be understood in the context of a much longer process of the industrialization of information processing. Like the term "cloud", the term "post-industrial" can conceals more than it illuminates.

Even if you believe, as I do, that Bitcoin is at heart an absurd techno-libertarian fantasy that fails even on its own terms as a digital currency and is unsustainable on almost every level, the more fundamental technology that enables Bitcoin, known as the blockchain, is likely to survive and prosper. This is because the blockchain is seen to have value not in terms of any one application, but as an infrastructure for enabling decentralized networks of trust relationships. The are many promising applications of such trust networks, as this recent powerpoint pitch from a venture capitalist illustrates…

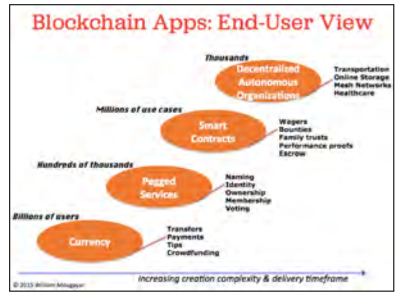

Thinking about the Bitcoin/blockchain network as a form of infrastructure is revealing. Infrastructure are an extreme example what the historian of technology Thomas Parke Hughes famously described as Large Technological Systems. Unlike individual inventions, Large Technological Systems cannot exist in isolation, but are inextricably linked to other technological, social, political, and economic actors, networks, and processes. The first high resistance filament incandescent light bulb was an invention, in the Hughesian taxonomy; the vast and interconnected network of electrical generation and distribution that are required to light up an entire city is a Large Technological System.

The largest of the Large Technological Systems we often call infrastructure: the electrical grid, the sewer system, the interstate highways, the AT&T network. Infrastructures are critical enabling technologies; their primary purpose is to make other technological and commercial activities possible. As a result, as Susan Leigh Star and Karen Ruhleder have reminded us, infrastructures are intended not to be seen. Technologies become infrastructure only after they are perfected to the point of being routine. We notice them only when they fail.

The global Internet is in that respect the perfect infrastructure: it is omnipresent and invisible; everywhere and nowhere. Using it we can connect to anyone, anywhere, from anywhere, but it does not otherwise intrude on our material reality. In our post-industrial society, information technologies make place and space increasingly irrelevant. Distance is dead, the world is flat; the singularity is near. Except, of course, when it is not.

The World in a Machine

One of the most widely-read histories of technology published in the past several decades tells the tale of the 18th-century clock-maker John Harrison, who according to book's self-description, was the "Lone Genius Who Solved the Greatest Scientific Problem of His Time."

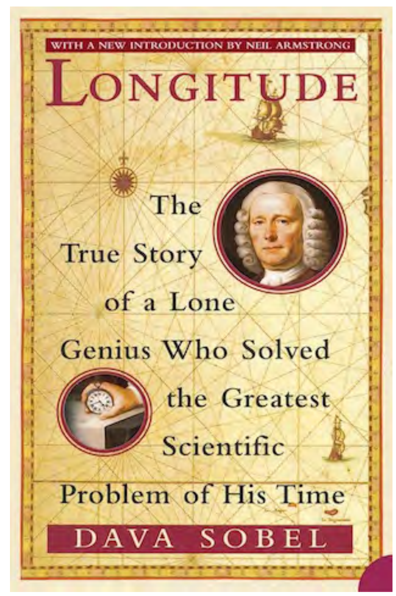

The scientific problem in question was the determination of longitude, and the solution was the invention of an accurate chronometer. What was at stake was the ability of the British empire to effectively manage and protect its growing fleet of military and commercial vessels, a fleet that would ultimately allow this tiny island nation to control over one-quarter of the land mass of the entire planet.

In Davel Sobel's heroic narrative, the working-class, self-educated Harrison triumphs against all odds over the entrenched interests of a the scientific and political establishment, as embodied in the figure of Nevil Maskelyne, the Astronomer Royal.

But as Mary Croaken, in her less celebrated but more balanced history, it was in fact Maskelyne who provided the solution to the longitude problem that was actually used in most maritime vessels well into the 19th century.

Harrison's chronometers, while accurate, were, at £200, far too expensive to be widely deployed; Maskelyne's Nautical Almanac, on the other hand, cost 2 shillings 6 pence, and when combined with a £8 Hadley's quadrant and an equally inexpensive book of mathematical tables, provided the cheap and effective means of determining longitude that would continue to dominate navigation for the next half-century.

I tell you this here because at the heart of Maskelyne's longitude solution was his development of effective computational techniques. Although these techniques were pre-industrial, and as such did not require much in terms of mechanization, they bear all the hallmarks of what would soon develop into the computational approach to data processing and management. His distributed network of human computers, whose work was carefully coordinated, monitored, and compiled, anticipated the "information factory" approach to computation that would dominate the 19th century and beyond.

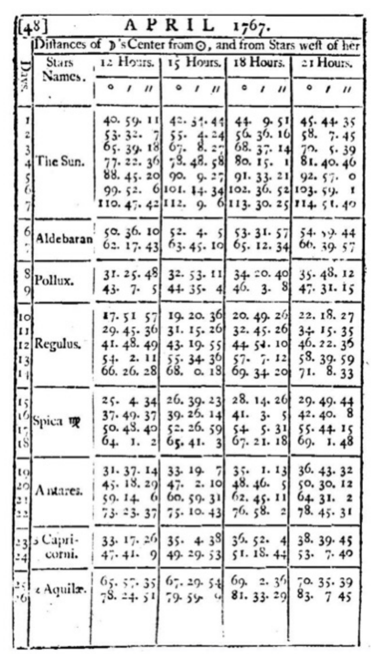

[48] APRIL 1767.

Diſtances of ☽'s Center from ☉, and from Stars weſt of her

Days.	Stars Names.	12 Hours. ° ′ ″	15 Hours. ° ′ ″	18 Hours. ° ′ ″	21 Hours. ° ′ ″
1	The Sun.	40. 59. 11	42. 34. 44	44. 9. 51	45. 44. 35
2		53. 32. 7	55. 4. 24	56. 36. 16	58. 7. 45
3		65. 39. 18	67. 8. 27	68. 37. 14	70. 5. 39
4		77. 22. 36	78. 48. 58	80. 15. 1	81. 40. 46
5		88. 45. 20	90. 9. 27	91. 33. 21	92. 57. 0
6		99. 52. 6	101. 14. 34	102. 36. 52	103. 59. 1
7		110. 47. 42	112. 9. 6	113. 30. 25	114. 51. 40
6	Aldebaran	50. 36. 10	52. 4. 5	53. 31. 57	54. 59. 44
7		62. 17. 43	63. 45. 10	65. 12. 34	66. 39. 57
8	Pollux.	31. 25. 48	32. 53. 11	34. 20. 40	35. 48. 12
9		43. 7. 5	44. 35. 4	46. 3. 8	47. 31. 15
10	Regulus.	17. 51. 57	19. 20. 36	20. 49. 26	22. 18. 27
11		29. 45. 36	31. 15. 26	32. 45. 26	34. 15. 35
12		41. 48. 49	43. 19. 55	44. 54. 10	46. 22. 36
13		54. 2. 11	55. 34. 36	57. 7. 12	58. 39. 59
14		66. 26. 28	68. 0. 18	69. 34. 20	71. 8. 33
15	Spica ♍	25. 4. 34	26. 39. 23	28. 14. 26	29. 49. 44
16		37. 49. 37	39. 26. 14	41. 3. 5	42. 40. 8
17		50. 48. 40	52. 26. 59	54. 5. 31	55. 44. 15
18		64. 1. 2	65. 41. 3	67. 21. 18	69. 1. 48
19	Antares.	31. 37. 14	33. 19. 7	35. 1. 13	36. 43. 32
20		45. 18. 29	47. 2. 10	48. 46. 5	50. 30. 12
21		59. 14. 6	60. 59. 31	62. 45. 11	64. 31. 2
22		73. 23. 37	75. 10. 43	76. 58. 2	78. 45. 31
23	♑ Capricorni.	33. 17. 26	35. 4. 38	36. 52. 4	38. 39. 45
24		47. 41. 9	49. 29. 53	51. 18. 44	53. 7. 40
25	ε Aquilæ.	65. 57. 35	67. 29. 54	69. 2. 36	70. 35. 39
26		78. 24. 51	79. 59. [illegible]	81. 33. 29	83. 7. 45

Figure 12: A computed table from the 1767 Nautical Almanac.

In fact, we could tell a compelling history of modern computing simply by following the problem of knowing exactly where in the world you were, and how to get where you wanted to go next. For example, consider Charles Babbage, often identified as the great-grandfather of modern computing. Like Maskelyne, Babbage was also an astronomer; like Maskelyne, Babbage was charged by the British government with producing an accurate version of the Nautical Alamanac to be used in the pursuit of Imperial ambitions. He and his supporters hint at grave dangers associated with the current unreliability of the Nautical Almanac:

> A yet undetected error in a logarithmic table is like a sunken rock at sea yet undiscovered, on which it is impossible to say what wrecks may have yet taken place.

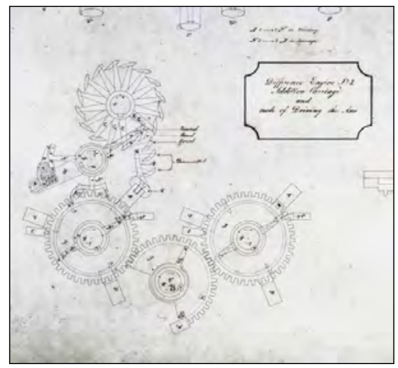

Figure 13: Blueprints for the never-constructed Babbage Difference Engine (1822)

The solution that Babbage proposed also computational, but was constructed around machines rather than people. Babbage's Difference Engine, the basis of his popular status as the great-grandfather of the modern electronic digital computing, was a mechanical implementation of Maskelyne's human-powered method of differences.

In the early 20th century, yet another British astronomer, Leslie John Comrie, would bring to the Nautical Alamanac the American-born technology of the Hollerith punch-card tabulating machine, developed by the company that would eventually become the International Business Machines Company. Comrie is often referred to as the father of scientific computing. Among other things, he implemented an algorithmic solution for applying Fourier synthesis to compute the principal terms in the motion of the Moon. He founded the world's first private company for scientific computing, incorporated as Scientific Computing Service, and during the Second World War headed a team of 30 scientists devoted to the calculation of ballistics tables.

Figure 14: One of many maps of the natural world included in the 1880 census.

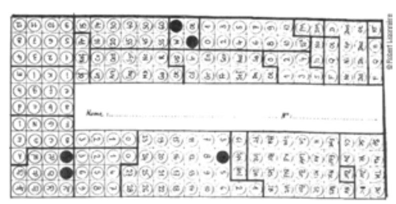

Figure 15: A Hollerith Punch-Card, circa 1880

There is an argument to be made that many of the most important moments in the history of computing have involved changes in the ways in which human beings perceive and relate to their environment — one of the traditional themes of environmental history.

Consider, for example, the Hollerith tabulating technology that John Leslie Comrie used to "computerize" scientific data processing. The punch-card technology that Herman Hollerith invented (and IBM later perfected) was itself invented in direct response to another challenge associated with "knowing" the world and the place of humankind within it.

In this case, the massive geographical and demographic expansion of the United States in the mid-19th century had stretched the ability of the US Census Bureau to perceive and measure the nation-state. Not only was it responsible for enumerating and documenting a rapidly growing population, but the government was asking of its more and more detailed questions about the characteristics of that population — and about the natural world on which that population was dependent. Among the 21,000 pages of the Hollerith machine-produced 1880 census were hundreds of maps describing the physical, social, and economic geographic of the increasingly united United States. The digital data encoded on a Hollerith punch card would become an essential tool for making the world legible to governments and corporations. The mother/inventor of the Hollerith tabulating machine was the necessity of a decennial census mandated by the constitution and greatly complicated by the data requirements of the modern bureaucratic nation-state. Several decades later, a similar dilemma would face the IBM Corporation, the direct successor of the Hollerith firm, when it took on the job of administering the newly established Social Security program.

From the Census Bureau to the National Security Agency, the American government has increasingly perceived its citizens through the lens of an electronic digital computer. These efforts are not simply an extension of the modernist imperative outlined most famously by James Scott in *Seeing like a State*. Computational technologies represented a new viewing and organizing the world, and one which has only been lightly touched upon by historians of computing.[3]

[3] A notable exception is Jon Agar's brilliant and quirky *The Government Machine*, which argues that Alan Turing's inspiration for the Universal Turing Machine was the administrative bureaucracy/technology of the British Civil Service.

The need for governments to "see" the world extended to the environment. As William Aspray, Rick Nebeker, and James Fleming (among others) have convincingly argued, the need of the United States government to perceive and predict the weather, largely for the purposes of making war, was an important motivation for its early investments in large-scale computational technologies. And as Sharon Kingsland (and others) have shown us, the modern science of ecology is a product of computational and cybernetic technologies, as are most of the tools used for environmental impact assessment, natural resource management, and petrochemical exploration and extraction.

From the problem of longitude to simulations of global climate to the computer systems that enable computerized shipping, computers have fundamentally altered the ways in which human beings understand, navigate, and manipulate the world. From the 3D computer models that make possible new extractive processes like hydraulic fracturing, to the more local use of economic models in the FORPLAN system (which is used by the national Forestry Service to allocate natural resources) I am interested in understanding and revealing the ways in which the techniques, assumptions, and politics of computer experts are built into — and then concealed within — the algorithms that shape our world.

Being Digital

In 1995, a professor at the influential Media Lab at MIT named Nicholas Negroponte published a book called *Being Digital*, in which he laid out a vision of the techno-utopian future that was both imminent and inevitable.

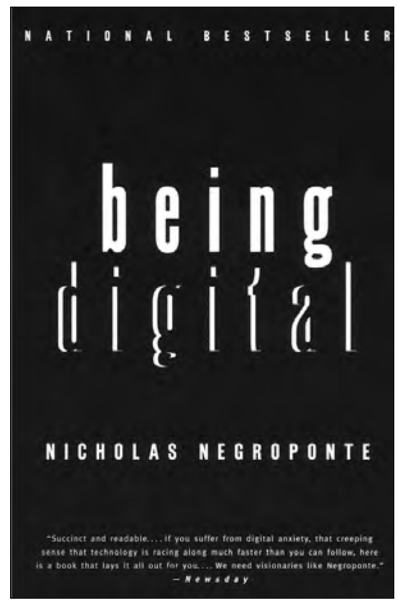

The book opens with an extended discussion of the difference between bits and atoms; atoms represent the material, bits the virtual; the "rapid exponential shift from atoms to bits" in the digital era is, according to Negroponte, both "inevitable and unstoppable"

In this final section of my paper, I will **not** be talking about Nicholas Negroponte.

Instead, the title of this section is a play on the work of the historian of science and technology Gabrielle Hecht. In her recent book *Being Nuclear*, Hecht details the many ways in which the struggle for control over uranium, the core component of both atomic weapons and nuclear power plants, has shaped African political, social, and economic life, as well as the lives and health of thousands, if not millions, of African workers and citizens. She shows how the relationships between the "nuclear powers" (often former colonial powers) and the "developing nations" (often former colonies), were constructed around previously established networks of political influence, commercial exchange, and the movements of goods and people. Her work demonstrates the necessity (and value) of adopting a global perspective, even when dealing with technologies that are predominantly associated with Western industrial nations; in her analysis, impoverished African miners are as much participants in the history of the nuclear economy as are scientists, engineers, and politicians.

Historians of computing have much to learn from *Being Nuclear*. Among other things, control over the supply of key minerals is as important to the computer industry as it is to the nuclear nation-state.

From Conflict Minerals to Death Metal(s)

In the early 16th century, Bolivian silver provided much of the wealth of the Spanish Empire. The city of Potosi, founded in 1545, was for a time the largest city in the New World, with a population that exceeded 150,000. Today, the Potosi region is the single largest supplier of lithium carbonate, the central component of the batteries used in everything from the Apple iPad to the Tesla Model S.

Almost half of the world's lithium is found in Uyuni region of Bolivia.

The primary concern among Bolivians is not human or environmental safety; lithium is extracted from brine is a relatively low-impact process.

Figure 16: Dead Man's Salt Flats (Argentina)

The bigger concern is foreign exploitation: "The previous imperialist model of exploitation of our natural resources will never be repeated in Bolivia," said Saúl Villegas, head of the government run-organization that oversees lithium extraction.

Similar debates are occurring in nearby Chile and Argentina, also large suppliers of lithium. Once again the politics of mineral extraction, foreign investment, and neocolonialism have emerged as central concerns of South and Central American government.

"Of course, lithium is the mineral that will lead us to the post-petroleum era," argues Marcelo Castro, a Bolivian mining engineer: "But in order to go down that road, we must raise the revolutionary consciousness of our people, starting on the floor of this very factory."

You can imagine what this kind of rhetoric sounds like to the heads of major American high-tech firms. The electric car company Tesla, for example, is on track to becoming the largest consumer of lithium-ion batteries in the world. The have recently negotiated with Nevada to locate a "gigafactory" in that state; in exchange, Nevada has committed to developing the Chemetall-Foote lithium mine. Nevada taxpayers will be asked to construct a connector between the Tesla Gigafactory and Highway 50 ("the loneliest road in America") whose sole purpose is to connect the Chemetall-Foote mine with the proposed Tesla Gigafactory.

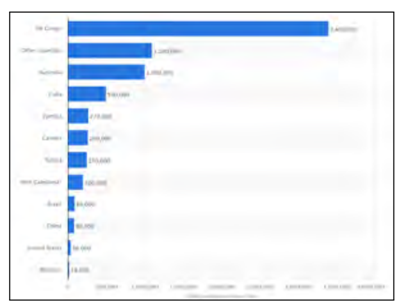

Figure 17: The world's cobalt supplies are concentrated in the DRC

Lithium is not the only mineral component of modern electronics for which there is great and growing demand. Tesla batteries also require cobalt, which is conflict mineral (the major supplies are located in the Democratic Republic of Congo).

If we had time, we could follow the chain of materials required to construct digital devices around the world and across the periodic table. In many of these cases and places, we would find depressingly familiar stories of environmental degradation, human rights violations, and a complete disregard for health and safety.

Let me just highlight one last specific example:

Rare earth elements are essential components of digital devices. There are seventeen rare earth elements, of which five are particularly rare and valuable to the electronics industry. Among them are **yttrium**, which is used in LCD screens and fuel cells; **europium**, a key component of compact fluorescent lights, computer monitors, and iPhone screens; and **neodymium**, **terbium** and **dysprosium**, all essential ingredients in the magnets of wind turbines and computer hard drives

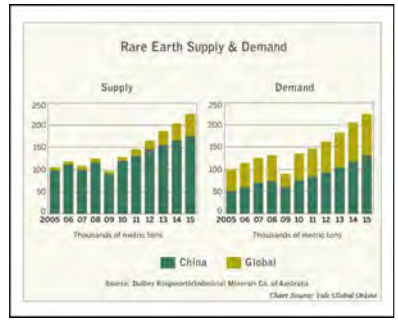

Global demand for rare earth elements has more than doubled in the last decade, largely as a result of the high-tech industry. With continued global growth of the middle class, especially in China, India and Africa, demand will continue to grow.

Here's the kicker: Aside from the a small amount recovered during recycling, the United States is 100% reliant on external sources of rare earth elements.

At the moment, China currently supplies 97 percent of global rare earth metal demand, and 100 percent of certain specific rare earth metals such as terbium and dysprosium. Already their monopoly over certain rare elements has been used by the Chinese government to "encourage" high-tech companies to locate their factories in China, and many analysts are predicting that the struggle for control over rare earth resources will only grow more contentious. As Deng Xiaopeng famously declared in a 1997, "The Middle East has oil, we have rare earth"

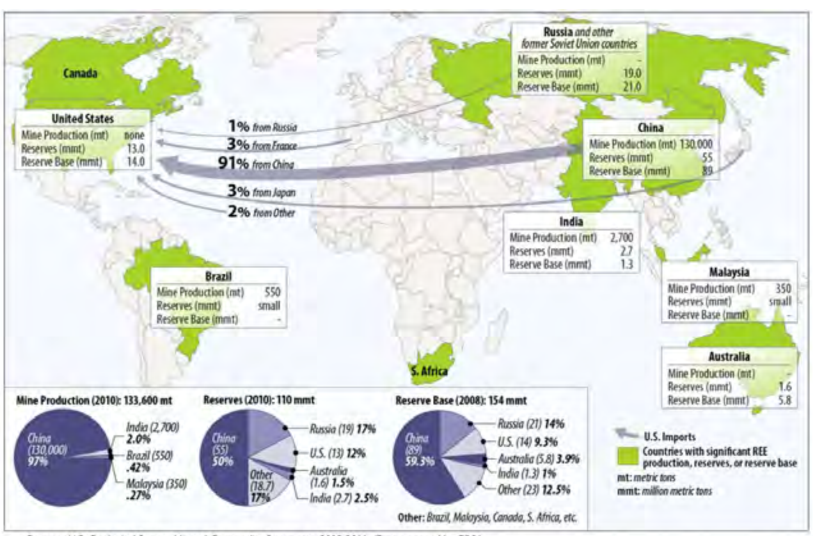

Source: U.S. Geological Survey, Mineral Commodity Summaries, 2008-2011. (Figure created by CRS.)

And you thought our dependency on foreign oil made for interesting geo-politics!

Also, because of the particular chemistry and geology of rare earth elements, almost all of the current supplies are produced as a byproduct of more traditional mining operations. For example, tellurium is produced as part of the copper mining process. Germanium is found alongside of zinc. What this means is that the supply of these materials is inelastic. You can't just add dig another tellurium mine when you need more tellurium.

Also, also: most of the heavy metals and rare earth elements that are essential to the digital economy are also what are called Energy Critical Elements — which means that they are also in demand by the alternative fuel industry. The use of solar and wind power at data centers simply shifts the costs from one type of environmentally destructive mining operation to another.

The Global Lifecycle of Digital Goods

One of my goals for the larger project is to trace the lifecycle of digital products as they circle around the globe, with an emphasis on the ways in which every stage of this sequence, from design to production to consumption to destruction, in influenced by and impacts upon the physical world. But I am not only interested in the ways in which the movement of materials is essential to the construction and operation of information technologies, but in how information technology, by fundamentally changing the way humans perceive and navigate the world, has dramatically expanded our ability to manipulate and exploit the environment.

For example, one of the most revolutionary recent developments in the history of global capitalism was the invention, in the late 1950s, of the standardized metal shipping container. By dramatically reducing the time required to unload cargo — instead of moving materials from one container to another, the entire container would be lifted from the ship and placed directly on a train or truck — containerized shipping increased throughput and reduced cost simultaneously. In the half-century following the introduction of the containerized shipping, global trade has increased 27 times over. By the year 2000, 300 million 20-foot containers were moved by sea each year.

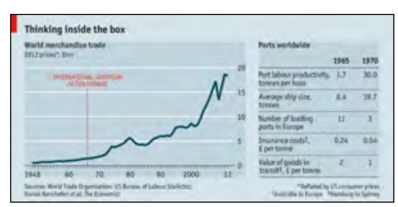

The story of containerized shipping is, of course, in large part the story of the container. But it is also a story of computerization. Computers were used to track, organize, and direct the millions of identical containers — as well as to control the automated cranes that were used to unload them. Today, computer controlled and GPS guided drone ships are poised to further revolutionize the industry. We often think that digital technologies *eliminate* the need for the physical movement of materials and people. The case of containerized shipping suggests one of the many ways in which information technologies *facilitate the expansion* of such movement.

Most of the manufacturing of these products has shifted to countries like China, whose environmental and worker-safety regulations are notoriously lax. The recent controversy about labor conditions in Apple Computer's Foxconn facility — where as many as 400,000 workers inhabit a 1.16 square mile walled "campus" — made many Americans aware the hazards and pollution associated with computer manufacturer have not disappeared, but have only been shifted abroad. A recent United Nations study estimated that the production of just one desktop computer required 240 kilograms of fossil fuels, 22 kilograms of chemicals and 1,500 kilograms of water — and that does not include the human labor involved. Each one of these resources and resource-chains represents a set of stories to be told about global politics, international trade, worker safety, and environmental consequences.This highlights another virtue of exploring the lifecycle of information technology: this is an essentially international story, one which necessarily shifts the focus from (typically American, or at least Western) users towards a broader range of workers and production sites.

All of this I imagine being a central component of the larger project. But for the time being, let me focus for a moment not on how digital goods are born, or how they move about the planet, but on where and how they die …

Digital Residues

It is tempting to embrace the apparent immateriality of the information age: doing so absolve us of responsibility for its less visible and desirable side-effects. But what information technology does is not eliminate but **conceal** the materiality of the so-called "new" economy. It externalizes the costs, and centralizes the benefits. This is particularly true in the case of the environmental pollution associated with both the production and disposal of electronic goods. Both problems have been shifted to parts of the world — India, China, Africa — where environmental and worker-safety regulations are relatively lax. Nowhere is this more true than in the context of the disposal of digital refuse.

The average life-span of a digital device is less than two years. The economics of the semi-conductor industry, along with our own seemingly insatiable appetite for what Michael Lewis famously described as the "new, new thing" have created a system of deliberate and desirable obsolesce. In many respects the high-tech economy is critically reliant on — one might even argue enslaved in bondage to — and endless cycle of "perpetual innovation."

The rapid pace of change in information technology creates an enormous amount of digital refuse. Each year almost 40 million tons of electronic waste is disposed of, including almost 3.3 million tons from the United States alone. In the developed world this so-called “e-waste” accounts for as much as 8% of the total amount of municipal garbage.

But while the majority of e-waste is generated in Western industrial nations, it is disposed of in the poorer Global South, often in places in which environmental protections and regulations are almost entirely non-existent. For example, in a recent expose of a “computer graveyard” in Agbogbloshie, Ghana, journalists with the BBC discovered that more than 50 tons of illegal e-waste was being transported into the area each year. Of this illegal waste, only 10% was recycled; the other 90%, which included lead, dioxin, and other toxins and carcinogens, was dumped directly into primitive landfills, where it quickly contaminated the water supply.

Figure 18: Children burn electronic equipment over open fire pits in Agbogbloshie, Ghana, in order to recover precious minerals.

Even the materials that were recycled were harmful to the environment: over open fires fueled with equally hazardous materials, workers as young as nine years old melted down components to extract valuables such as copper, aluminum, and mercury. Both the smoke from the fire and the materials they reclaimed represent personal and environmental dangers. In a 2008 study, researchers at Greenpeace discovered high levels of lead, cadmium, antimony, PCBs, and chlorinated dioxins in the soils in Agbogbloshie.

This is, of course, the end of the story, the final destination of my imagined global lifecycle of digital technologies. In between are a series of transnational movement and exchanges that encompass the developed and the developing world, a diverse range of sites of production, and an ever-expanding cast of characters.

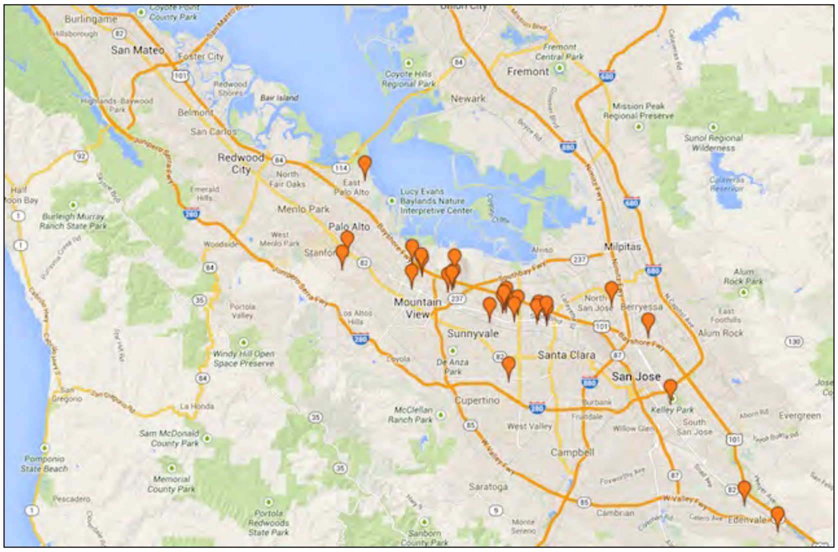

Figure 19: A map of the twenty-nine Superfund sites in Silicon Valley.

It is important to note that the problem of digital pollution is not only a problem of the developing world. Many Americans are unaware, for example, that the single largest concentration of Superfund sites (that is to say, locations designated by the EPA as being particularly polluted and in need of immediate remediation) are located in Silicon Valley. In the roughly 10 by 40 mile strip of land that comprises Santa Clara County, California, there are 23 Superfund sites, most of them contaminated by the by-products of semiconductor manufacturing, including such highly toxic chemicals as trichloroethylene, Freon, and poly-chlorinated biphenyls (PCBs). These chemicals have been linked to elevated rates of miscarriages, birth defects, and cancer. So far more than $200 million has been spent on cleaning up soil and ground-water pollution in the area, and the extent of the problem is only just starting to be addressed. Most of the well-educated and well-paid engineers and scientists who live in the area are unaware of the environmental dangers posed by their seemingly "clean" post-industrial information industry.[4]

[4] Even in tiny little Bloomington, Indiana (pop. 60,000) where I live and work and raise my family, there are three EPA superfund sites associated with the production of consumer electronics. For decades the Westinghouse Corporation dumped waste from its capacitor plant into unregulated landfills and quarries in the Bloomington area, and in the 1980s the Bloomington local government distributed this contaminated soil as part of a municipal gardening program. Today, Bloomington has the dubious distinction of being the "PCP capital" of the United States. This is not a good thing…

Why the Dirty Bits Matter…

Silicon Valley is an appropriate place to wrap up this introduction, because in part this is where the larger project began. This is not simply the obvious association of Silicon Valley with the history of computing, but rather because Silicon Valley has, for me, always been the most striking example of the paradox of place that sits at the center of the Information Society. Silicon Valley makes the technologies that are supposed to make place irrelevant, and yet its own significance as a place remains remarkably persistent. Despite the fact that it contains some of the most expensive real estate in the world — and is both over-crowded and polluted — high-tech and Internet firms still flock to Silicon Valley, and despite many attempts to replicate its magic in industrial and university cities around the globe, no one has ever succeeded. As Annalee Saxenian and others have amply demonstrated, geography still matters, even in Silicon Valley.

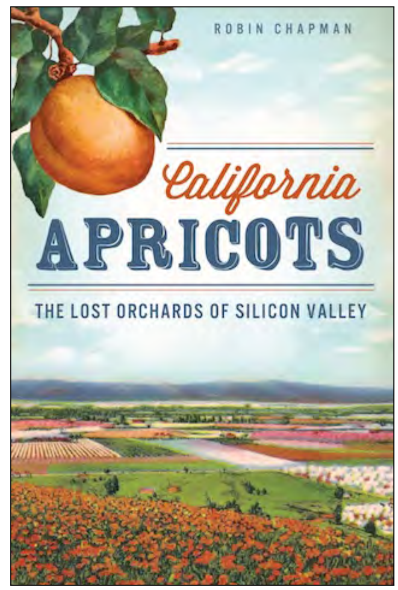

Figure 20: Santa Clara contains some of the richest farmlands in North America. It was once best known for its apricots.

But Silicon Valley also illustrates the ways in which both history and environment matter as well. We also can peel back the layers of the Santa Clara county landscape to reveal multiple sediment stratum of industrial activity and infrastructure. In the convention histories of computing the prehistory of Silicon Valley is that of an idealized agrarian economy — Santa Clara county contains some of the most fertile soil in the United States, and indeed was formerly known as the "The Valley of Dreams," "The Valley of the Heart's Delight," "The Fruit Bowl of America," and the "Garden of the World." But the extensive but undocumented underground water system that carried the pollution of the semiconductor industry were laid down during the era of industrial agriculture, and the canneries of Santa Clara county were once the heaviest polluters of the San Francisco Bay. And prior to that, quicksilver mines in New Alamaden and other towns radically transformed the landscape using destructive hydraulic mining techniques. Both industries relied heavily on immigrant labor migrating back and forth between California and Mexico — this same migrant laborer force was essential to the growth of the early semiconductor industry. There is a racial and ethnic geography from this period that still persists in Silicon Valley — and over which maps a similarly segregated map of environmental pollution and human health effects.

And yet Silicon Valley has also worked hard — and largely successfully — to erase erase both its own past and its connections to the material world.

By distancing themselves from the past, the computer industry can claim the positive benefits of technological progress without bearing any of the burdens of the larger technological history of which it is only the most recent iteration. Automobiles pollute, nuclear energy produces toxic waste, industrial agriculture is giving us all cancer — but computers keep getting faster, smaller, better. By making the physical world increasingly irrelevant, information technologies allow us to avoid the consequences of our actions on the environment. We assume that by putting things online we are removing them from realm of the physical, and that going digital means going green. In some cases this may be true; in others it is clearly not.

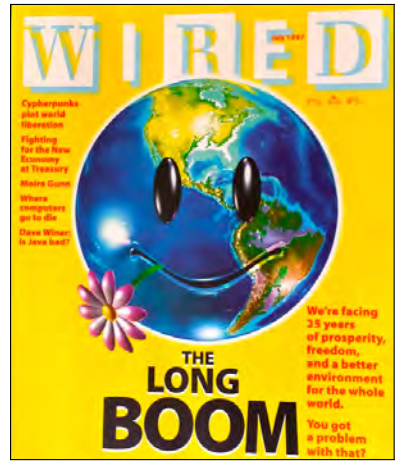

And so, in addition to contributing to the literature in the history of computing and environmental history, I hope that this project will also lead to the development of a new environmental ethic of design in information technology. Bitcoin is an ill-conceived, entirely unnecessary, environmental disaster, but the "proof of work" requirement that makes it so is not an essential requirement of a digital currency. Once we accept that the provision of "computer power," like other forms of industrial activity, is necessarily resource-intensive, pollution-producing, and potentially damaging to the environment, we can make more informed choices about when, how, and for what purposes we employ such power. There is no such thing as a free lunch, even in the virtual domain of cyberspace, but there are meals that are less expensive and more sustainable than others.

A closing note on the title of the paper: the term "dirty bit" is, of course, a reference to the fact that even virtual goods can be tainted by their association with human and environmental degradation, but it is also a term of art in computer science. A dirty bit is a bit that is associated with a block of computer memory and indicates whether or not the corresponding block of memory has been modified. Dirty bits are to "mark segments of data that need to be processed or have yet to be processed." This double meaning seems to me to be significant. I am working on ways to draw out the analogy. Suggestions welcome.