

Macroscopically Sustainable Networking: An Internet Quine

Barath Raghavan and Shaddi Hasan
ICSI and UC Berkeley

ABSTRACT

The Internet stands atop an unseen industrial system required for its continued growth, operation, and maintenance. While its scale could not have been achieved without this reliance, its dependencies—ranging from sophisticated manufacturing facilities to limited raw materials—make it vulnerable to disruptions. To achieve independence requires an Internet quine—a set of devices, protocols, manufacturing facilities, software tools, and other related components that is self-bootstrapping and capable of being used by engineers to reproduce itself and all the needed components of the Internet. In this paper, we study how such an Internet quine could be built. We also attempt to identify a minimal set of such tools and facilities, and how small and inexpensive they can be made.

1. INTRODUCTION

The Internet is arguably the largest and most successful system humanity has ever created, and is composed of a complex array of hardware and software assembled around the world with materials, energy, skills, and designs also from a global resource base. The fact that the Internet has spread to all regions of the world despite this complexity is a testament to its tremendous utility and the ingenuity of those involved in the design, manufacture, distribution, and operation of its constituent components.

With this complexity comes vulnerability to threats: disruptions in global supply chains, energy shortages, societal instability, and government intervention into its operation. The last two decades have been relatively placid compared with the previous eight or previous eighty, and have provided an environment in which the Internet has been able to grow at a phenomenal pace. Recent political instability in the Middle East has dramatically demonstrated the fragility of the Internet on a regional scale. The effects of widespread climate change are predicted to have significant effects on the global economy, in turn threatening to the disrupt global supply chains upon which the Internet depends.

Achieving structural resilience against such challenges, be they at a global, regional, or national scale, calls for a set of networking technologies and components that do not carry such heavy external dependencies. For lack of a better term, we call such a set of technologies required to create a self-sustainable Internet an *Internet quine*, and in this paper explore how to build an Internet quine. In particular, we consider quines with minimal dependencies, adopting as much as possible the principles of “appropriate technology” [7, 13], namely an Internet quine that is a) simple, b) locally reproducible, c) composed of local materials and resources, d) easily repairable, e) affordable, and f) easily recyclable.

Such a quine would have a variety of useful applications. In resource-constrained environments (such as parts of the developing world, or in areas isolated by natural disasters), a functional Internet quine could be a link to the rest of the world. It could also serve as the basis for a non-government controlled communications infrastructure for people living under an authoritarian regime that practices widespread censorship. The process of building an Internet quine could also have great pedagogical value, allowing students to “build the Internet from scratch.” Moreover, many of the objectives we explore are part of the networking canon: we aim to build scalable networks that can be started and tested at a small scale with modest resources, resilient networks that remain useful under changing conditions and respond to pressures, modular networks separated into distinct elements that can be replaced at different scales and technological levels, and open networks that do not demand a certain system or set of components to function.

What an Internet quine would look like is highly dependent on the particular context in which it would be used. As such, we don’t propose a single way forward, but weigh what knowledge and resources can be assumed to exist outside the quine and what services it must support. Thus, in this paper we explore the manufacturing and operational dependencies of supporting a common use case of the Internet today: end to end communication over long distances. We then consider possible alterna-

tive architectures for supporting this use case. Our hope is to provide a small nudge towards sustainability.

2. RELATED WORK

The goal of independence of the sort we study here is not a new one—its roots date to the work of the economist E.F. Schumacher, who introduced the concept of *intermediate technology*, now known as *appropriate technology*—technology that is appropriate for a region rather than being created as a global commodity [13]. Schumacher’s idea, one that we might well heed today, is that independence leads to better human systems. Similarly, the structure, and downsides, of a complex dependent system such as the Internet was studied in the abstract by historian of technology Lewis Mumford throughout the 20th century [8, 9].

More recently, a range of scientists and engineers, from ecologists to computer scientists, have explored building sets of self-sufficient machines. The Open Source Ecology project aims to provide a complete set of designs for appropriate technology that meets all basic human needs [10]; among their tools are several means of manufacturing, including a 3D printer, RepRap [12]. While the RepRap is nearly able to manufacture itself, it still requires significant material inputs. The Fablab project has similar abilities and ambitions and recently has aided in setting up networks in Afghanistan [5].

Many of these efforts stem from the information and communications technology for development (ICTD) community, which has also studied many of the issues we consider here. Since resource-limited environments break design assumptions that the Internet’s infrastructure relies upon, they complicate the deployment and evaluation of networked systems [4]. Understanding the challenges that a system’s economic, technical, and social contexts presents is essential for achieving its long-term sustainable operation [3, 14]. We observe that such an understanding of the present Internet infrastructure is surprisingly difficult to come by, further motivating our current study.

3. DISASSEMBLY

We cannot begin to explore the Internet’s potential for independence without knowing its dependencies, so we first attempt to decompose and stratify its components. A simple way of viewing this process is to think of it hierarchically, where each layer depends upon the layer below it. However, we will soon see that such a clean hierarchy breaks down quickly.

3.1 Use case

What is the set of actions and uses that define the Internet? The Internet has myriad uses and thus it is difficult to say which of these uses is most intrinsic or es-

sential. For concreteness, here we consider the use case of end to end communication between two Internet users who are geographically distant.¹ Curiously, though the example we begin with is simple and natural, upon study of its dependencies we uncover a great deal of complexity and interdependence. At the outset, we note that in piecemeal this analysis has been done by many in the research literature and industry, but to our knowledge a semi-comprehensive picture of the Internet’s dependencies has never been painted.

3.2 Dependencies

We depict the dependencies of the various components of the Internet in Figure 1. Each box denotes an abstract or concrete component, resource, or function of the Internet or one of its recursive dependencies. Arrows denote dependency; dashed arrows denote optional dependency—that is, a “one of these” relationship. At the top is the use case. We quickly move through the higher layers—which represent mostly functional and abstract pieces that we recognize as pieces of the Internet—and on to the lower layers which reside in the realm of hardware. As the Internet currently uses the latest IC manufacturing, we consider the consequent dependencies of that and similar processes. Once we move beyond hardware manufacturing we enter the realm of chemical compounds and natural resources that are required for many of the relevant manufacturing processes, from making fiber optics to printing circuit boards. We terminate with special nodes that represent ores or otherwise naturally occurring resources.

We do not depict operational or deployment dependencies since they tend to involve a small number of processes—notably, power generation and transportation—which ultimately have to do with energy production. These dependencies are still important as energy production is known to be a significant global challenge for this century.

3.3 Why consider these dependencies?

At first glance the dependencies we find take us far afield of what we think of as networking. As a result, it raises a natural question: why consider such esoteric and deeply embedded dependencies of the Internet? After all, in one view the resources we build upon and the dependencies we consider are commodities, and the problems of sourcing the inputs for these are the responsibilities of their respective designers.

We argue that this is not the case and that networking is unique in that it has yielded a system — the Internet — that is far greater than the sum of its parts. Were

¹We assume that the components they use are standard—the routers, data centers, transmission systems, operating systems, and so forth that are in use today.




Figure 1: Dependencies of the end-to-end communication use case.




Figure 2: Locations of necessary natural resources.

some specific peripheral or piece of hardware to become unavailable due to dependency failure, the world would scarcely notice, unless of course it affected the functioning of the Internet in a material way. While we are not considering the process by which such dependency failures could occur, knowing what they are down to irreducible roots gives us the ability to secure the whole by divesting from the parts.

3.4 Where do components come from?

Where the dependencies come from is just as important as what they are. Once again, the idea of global commodities and international trade are so firmly entrenched in our thinking today that we take for granted that a resource in another country can be tapped if the price or force is right. However, this is a modern notion—one whose history is scarcely four decades old—and the increasing cost of energy and increased resource competition might complicate the picture increasingly in the decades to come.²

Thus we consider the geographic locations from which the root nodes in our dependency analysis are typically sourced. In all, we found that our simple Internet use case depended upon over a dozen natural resources sourced from numerous countries, as seen in Figure 2. While some resources are both abundant and widespread, key resources are not. Bauxite, for instance, a principal ore for aluminum and rare metals such as gallium (an important semiconductor), is found in sizeable deposits only in Australia, Brazil, Kazakhstan, and Venezuela, with smaller deposits in West Africa and the Balkans.

4. REASSEMBLY

So far we have identified the deep dependencies of the Internet. In doing so, we are reminded of how dependent the Internet is, on materials and processes that we, and likely most networking researchers, scarcely recognize, let alone understand.³ We would like to shrink this dependency set as we design an Internet quine. At the outset, we should observe that the entire global industrial system is an Internet quine, just not a very compact one; our goal is to identify a smaller Internet quine.

How do we design for self-sufficiency given our knowledge of dependencies? It seems a profitable way to begin is to eliminate dependencies as high up on the chart as possible, as this enables us to replace a whole subtree of dependencies. This necessarily raises the complex issue of “what is the set of essential functions and/or guarantees of the Internet?” We address this next.

²Consider, for example, recent actions by China to control its domestic supply of rare-earth elements [2].

³While this was for just one use case, we believe that it covers the preponderance of possible use cases and their dependencies.

4.1 What is an Internet quine?

Given the complex dependency graph for supporting our use case in the current Internet, is it possible to construct an Internet-like system that will both support the use case and have a simple dependency graph? Point-to-point communication between two geographically distant parties does not fundamentally carry a large set of dependencies. Trivially, such a use case could be satisfied by a physical messenger (indeed, trained pigeons once carried out this task).⁴ Large-scale communication networks existed before many of the technologies used in the current Internet were developed as well, so clearly the set of dependencies we describe above is not strictly necessary to achieve our goal.

For the purposes of our discussion, we consider the primary service that the Internet provides to be routing datagrams between independent administrative domains for automated processing at endpoints. We require only *functional equivalence*, not interoperability, with the existing Internet. In addition, such a network should provide a service interface that is sufficiently generic for multiple application-level services to use it: the system should support layering of protocols, and should not enable only a particular class of applications (assuming the network is capable of meeting an application’s performance requirements; inability to stream high-definition video does not prevent a system from meeting our definition). While the postal system meets many of these requirements and coupled with storage media provides functional equivalence, ideally we would also maintain an additional function: near real-time communication.

In addition to defining the minimum services an Internet quine must provide, we must also define what we mean by self-sufficiency. A minimal Internet quine is one that achieves our definition of an Internet-like system, as well as any application-level requirements, while requiring as few resources as possible for its construction and sustainable operation. What resources are available, and what the application-level requirements of the system are, will of course vary depending on the context in which the quine exists.⁵ We require that the resources in question be locally plentiful for some long period of time, but how long this duration must be leads to different designs. They need not be what we would consider “natural resources” if this duration is fairly short (on the

⁴A modern variant of this that implements a delay-tolerant network might include birds with microSD cards strapped to their legs, which in a recent instance proved more efficient than using conventional network infrastructure [1].

⁵Since the global industrial system was developed over the course of a few hundred years using solely raw resources, in theory any combination of those resources and the knowledge required to build all the needed systems for the Internet could replicate it. We don’t consider quines on this timescale.

order of a few decades).⁶ Indeed, the first quine we consider uses a large supply of salvaged parts.

4.2 A Salvage Internet

One way to reduce the set of dependencies would be to create a networking infrastructure from only salvaged hardware, thus *replacing* many of the dependencies in Figure 1 with pre-manufactured, but widely available, components. Thanks to the proliferation of current computing technologies, such equipment is widely available in many places, including still-functional equipment which has been discarded and is no longer in use. Moreover, local expertise for extending the useful life of computers, cell phones, and other devices already exists in many places. Not only does re-using and re-purposing existing equipment avoid the need for most manufacturing, it also dramatically reduces transportation needs, as most equipment can be found locally.

For our example use case, the salvage requirements are surprisingly modest. Commodity 802.11 WiFi equipment, running open-source drivers and management software, has been successfully used to establish radio links hundreds of kilometers in length, even in very resource-poor environments [11]. Thus we can replace the router node of our dependency graph with such software-based WiFi routers, eliminate cabling, and replace peripherals and laptops with salvaged devices. Assuming access to a working 802.11 radio, the necessary antennas can be fabricated easily by unskilled labor from sheet metal, another common commodity. Working computers are plentiful worldwide, with a vibrant repair market already in existence. Many cellular phones, the single most widely deployed communications technology in the world, have WiFi capabilities as well, making integration with such a system trivial. Indeed, the power needs of this Internet quine could be satisfied using the same power infrastructure currently used to recharge cellular phones. With this equipment, one could easily imagine a multihop network (though perhaps a low-throughput one) transmitting data across considerable distances.

This Internet quine can be constructed using only common, locally available components; as such, it carries a much simpler set of dependencies. It would also be directly interoperable with the existing Internet.

4.3 Low-tech Internet

A salvage Internet, while having substantially fewer dependencies than the current Internet infrastructure, still makes significant assumptions about physical resources available outside the quine: namely, a surplus of working or repairable computing equipment that could be repurposed for use as a communications infrastructure.

⁶If the duration is longer than any non-renewable input can't be depended upon.

Clearly, this assumption may not always hold, nor could it be sustained in the long-term.

The fundamental challenge we encounter when no longer using salvage is the replacement of integrated circuits. As our dependency analysis shows, the set of dependencies and processes is significantly complicated by IC manufacturing and the fabrication techniques modern ICs require, such as photolithography.

Manufacturing radios from simple components using older technology is nothing new. Amateur radio has a long history of utilizing operator ingenuity to create sophisticated communications systems using components and manufacturing techniques that are accessible to hobbyists. Manuals describe how to design and construct fully-working radios using no ICs [6], so neither have the knowledge nor the resources for pre-IC long-distance communication been lost. In such a design, these low-tech radios would provide essentially the backhaul that 802.11 provided in our previous design. We imagine that local networks could still be composed using physical cabling using copper. It is unclear, however, what the end hosts would be or look like in this scenario.⁷

4.4 Summary

These two Internet quines both provide the set of services that one would expect the Internet to provide, and meet basic and reasonable “self-sufficiency” requirements. Given our modest application requirements, we find that an Internet quine can be surprisingly minimal.

5. CONCLUSIONS

Our dependency analysis reveals several interesting properties about the nature of both the Internet's present physical infrastructure as well as its architecture. It also provides us a reminder of the inherent value in reducing the Internet's dependencies.

5.1 Lessons

Our foray into this subject is new, so our lessons aren't deep. Nevertheless, we've never seen (or done) an analysis like the one we did in this paper, and as a result have a few general observations that we feel are worth making.

The Internet is highly dependent. Most obviously, we observe that the present Internet is heavily dependent on many industries not typically associated with network equipment. In particular, several of these dependencies include processes which rely on limited natural resources. Many others require highly-skilled workers and specialized manufacturing processes.

⁷Clearly a non-IC based computer, either mechanical or using tubes, could be designed to interface with this low-tech Internet, but these options are far more degenerate than we'd like.

This situation is ironic given the Internet’s widespread and deserved reputation as a remarkably resilient system. It achieves this resilience from a functionality perspective by ensuring that failed links do not break the entire system: traffic can easily be routed through substitute paths in these events. From a structural perspective, however, substituting for a failed dependency is extremely difficult.

We also note that many of the dependencies we identified have no clear alternative that can meet the performance demands of the current Internet. For example, reasons of physics make fiber-optic cables perhaps the only transport medium that can satisfy today’s requirements for data volume and transmission distance. As a result, disruptions in the availability of these dependencies will require users to adjust their expectations about the Internet’s performance characteristics.

Fundamental dependencies are limited. We can provide a reasonable equivalent to the basic service model of the Internet with substantially fewer resource dependencies than needed today. In this sense, the architecture of the Internet itself does not *inherently* require the complex dependency graph shown in Figure 1. Given appropriate performance expectations, an Internet quine may be very small, requiring only resources that are both plentiful and widely distributed.

Knowledge is a cheap resource. Information is the cheapest commodity on the Internet. As a result, any opportunities to substitute physical resource dependencies with information-based dependencies can dramatically simplify the dependency graph. In general, we can decrease the size of an Internet quine by substituting information resources for specialized dependencies. Where this is not possible, general purpose tools should be favored over special purpose machines. Here we use the definition of tool and machine due to Mumford [8]: the user of a tool imparts the tool’s purpose and role, whereas those qualities of a machine are typically statically determined in its design. Relying on users’ knowledge to operate general purpose tools, rather than on purpose-built machines, can simplify an Internet quine’s dependency graph. Such tools can serve multiple roles in the dependency graph, and likely have fewer, simpler dependencies themselves.

5.2 Benefits of Simplicity

We’ve already explored the primary benefit of simplicity — building a (more minimal) Internet quine. Independently, it also has two other nice benefits.

Networking education. Networking as it is taught today is parochial. Those who are interested in systems discuss the distributed systems that compose the Internet. Others discuss the network architecture and how and why it

makes sense. Others still examine it using the language of queueing theory. These strands of networking pedagogy all have their place and contribute to our understanding, but they do not connect with the real world — the operation of a real-world system. Part of the reason for this is that it is difficult to teach such a big thing.

However, an Internet quine might be something that could be sourced and built from scratch for pedagogical purposes. This could be a valuable undergraduate course or course series: “Building the Internet from Scratch.” While universities with large engineering departments probably already have many of the resources needed today, almost all have the resources for a salvage Internet.

Inching towards sustainability. In decreasing the resource base required to achieve the Internet’s functionality, there is a good chance we also increase the Internet’s sustainability in every sense of the word. That is, an Internet whose functionality can be replicated independently in many places is one that will not stop functioning due to a single common-cause failure. Similarly, as a result of decreasing the complexity of the Internet’s requirements, it is possible (though difficult to verify) that the components yielded are more environmentally sustainable, in that they either re-use or do without components that would have otherwise been manufactured anew.

6. REFERENCES

- [1] SA pigeon “faster than broadband”. *BBC News*, Sept. 10, 2009.
- [2] K. Bradsher. China said to widen its embargo of minerals. *New York Times*, Oct. 19, 2010.
- [3] P. Brand and A. Schwittay. The Missing Piece: Human-Driven Design and Research in ICT and Development. In *Proceedings of the International Conference on Information and Communication Technologies and Development*, 2006.
- [4] E. Brewer, M. Demmer, M. Ho, R. Honicky, J. Pal, M. Plauch, and S. Surana. The Challenges of Technology Research for Developing Regions. *IEEE Pervasive Computing*, 5(2), 2006.
- [5] Fab Lab Project. <http://fab.cba.mit.edu/>.
- [6] H. Friedrichs. *Instruments of Amplification: Fun with Homemade Tubes, Transistors, and More*. Lindsay Publications, 2003.
- [7] B. Hazeltine and C. Bull. *Appropriate Technology: Tools, Choices, and Implications*. Academic Press, 1998.
- [8] L. Mumford. *Technics and civilization*. Harcourt, Brace & Company, 1934.
- [9] L. Mumford. *The myth of the machine*. Harcourt Brace Jovanovich, 1967.
- [10] Open Source Ecology Project. <http://opensourceecology.org/>.
- [11] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and implementation of High Performance WiFi Based Long Distance Networks. In *Proceedings of USENIX/ACM NSDI*, 2007.
- [12] RepRap Project. <http://reprap.org/>.
- [13] E. Schumacher. *Small is Beautiful: A study of economics as if people mattered*. Abacus, 1974.
- [14] S. Surana, R. Patra, S. Nedeveschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *Proceedings of USENIX/ACM NSDI*, 2008.