Unplanned Obsolescence: Hardware and Software After Collapse

Esther Jang
University of Washington
infrared@cs.uw.edu

Matthew Johnson
University of Washington
matt9j@cs.uw.edu

Edward Burnell
M.I.T.
eburn@mit.edu

Kurtis Heimerl
University of Washington
kheimerl@cs.washington.edu

ABSTRACT

In a setting of economic and infrastructural collapse, the inability to manufacture and maintain computing resources will be an enormous limitation on the continued use of technology. The concept of "rot" exists for both hardware and software, referring to a slow loss of functionality over time. Given a desire to maintain technological capability, we raise a variety of questions about technology use in such a scenario. How long will current hardware last through repair, robust construction, and good maintenance practices? What would software development and maintenance entail without today's Internet infrastructure? What can be done to keep our software stable and usable for as long as possible in the face of viruses, storage degradation, and other threats? We present rough estimates of the expected longevity of desktop and laptop hardware for various levels of maintenance, and hypothesize that software degradation, not hardware degradation, will be the limiting factor in determining how long devices will remain usable for computing tasks involving any exposure to external files or networks. We propose both physical and social strategies to guard against both modes of degradation.

CCS Concepts

•**Security and privacy** → *Human and societal aspects of security and privacy;* •**Hardware** → *Aging of circuits and systems;* •**Software and its engineering** → *Software creation and management;*

Keywords

longevity; hardware; software; security; malware;

1. INTRODUCTION

Computing resources are integral to the fabric of our modern society. Medical records are stored and accessed electronically, weather is predicted using computational models, and people have access to high-bandwidth long-distance communications infrastructure at their fingertips. In an infrastructural collapse, computing and all of the services which rely on its affordances would be put in jeopardy. In the event that large-scale electronics manufacturing were to suddenly halt, or a region were to be cut off from the global supply chain, computing devices would become precious resources whose functionality would need to be carefully maintained. Furthermore, without reliable power generation and distribution, long distance communication over the Internet as we know it today would likely not exist, even if the hardware itself could be maintained. Lack of connectivity would render all modern network-based services and software maintenance infrastructure defunct. Conservation of existing distributed hardware and software resources would need to continue until we either learned as a society to recreate their functional equivalents in a more sustainable way, or learned to do without them.

1.1 Assumptions

The production of modern computing resources rests on massive technical, social, and economic infrastructures. In this work, we explore a scenario where: 1) the manufacture or acquisition of new integrated circuits (ICs) is prohibitively difficult, perhaps due to a lack of raw materials, accessible production facilities or energy to run them, a disrupted supply chain, or any combination which we believe likely in the event of collapse; and 2) long distance networking and information sharing becomes difficult with the decay of Internet infrastructure due to the same factors as above.

In their work discussing a minimal set of devices and protocols required to reproduce the functionality of the Internet, Raghavan and Hasan detail the extensive network of resource dependencies involved in hardware device manufacture, and recommend reducing these dependencies [41]. However, we assume most communities will not have specially-architected computing devices designed for the loss of present-day manufacturing infrastructure. Most people's only recourse upon failure of hardware components will be to repair them, or procure replacements from those manufactured before collapse.

Hardware alone does not make a modern computing platform; we also anticipate a slew of challenges related to maintaining the correctness of software and user data, especially due to malware infections and "bit rot" on storage media not designed for decades of integrity. Even in the absence of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

global connectivity, we expect malware to remain an issue as long as computers engage in networking and file transfer over any medium. Furthermore, long-term connectivity loss and lack of a centralized trust infrastructure break many fundamental assumptions made in software design, making development, distribution, and verification of new software much more difficult. We hypothesize a dramatic reduction in the authoring and dissemination of software after collapse, to the point where patches and security fixes are no longer widely available.

Finally, we suggest technologies, practices, and social infrastructures yet to be developed that could mitigate the risks collapse imposes on keeping both software and the hardware it runs on functioning in an environment adversarial to users and developers.

2. MITIGATING HARDWARE RISK

2.1 Computing Usage and Environment

We consider two usage scenarios which may characterize either end of a spectrum of computer lifetimes. In scenario one, dedicated computers are set aside for the operation of critical services, such as weather modeling or database accesses, and are kept in a controlled environment such as a clean room to consciously maximize longevity. Scenario two is that of a personal computer, probably a portable laptop, used as is typical today without any special protection from the elements. We use the two scenarios to separately reflect on the inherent effects of computational load and external environmental effects such as impact, water damage, or particle intrusion.

Our motivation for this separation is that many types of damage come from the external environment and can be almost entirely prevented through stringent environmental control and limited device mobility. For example, dust and dirt on electrical components can prevent proper cooling, increasing their chance of failure. Humidity or spilled water can corrode circuits or cause shorts that lead to component damage. Accidental impact due to dropping or jostling during transport may result in mechanical damage to the screen, keyboard, ports, fans, and the chassis, opening additional entry points for dust, dirt, and water. Strict control of the material computing environment and avoidance of machine transport mitigate many of these risks. We argue that environmental control can increase device longevity at the cost of losing some of the social functions of computing permitted by mobility today.

2.2 Computation-limited Components

Computational loads themselves contribute to physical wear on many components, leading to performance degradation and eventual failure with regular use of the computing resource. Storage drives are one such component. A casual study of Internet forums on computer repair suggests that hard drive replacement is one of the most common repairs performed on consumer machines, for a variety of reasons ranging from mobility-related damage to performance deterioration from component wear over time.

The industry standard for manufacturers to provide estimated lifetimes for HDDs has historically been Mean Time Between Failures (MTBF) or Mean Time To Failure (MTTF), measured in hours of uptime. Common MTTF ratings for modern consumer-grade hard drives range from 100,000 to 1,000,000 hours, which represents roughly 100 years of continuous use. In reality, however, real world data has shown that modern consumer HDDs fail at rates of around 2-5% per year, with an observed acceleration to around 10% after the first four to six years [5, 6, 49]. A generous estimate at the original 2-5% puts the half-life of a HDD at 13.5 to 34 years; with the increased failure rate from 5% to 10% after the first 6 years, the half-life is 9.7 years. Furthermore, these empirical failure rates were measured in datacenters, where the drives would have been largely protected from unpredictable power fluctuations and physical damage. Power outages are known to cause "head crashing" in HDDs, where the mechanical disk head snaps back to a starting position upon loss of power and potentially scratches the disk platter [29]. Since HDDs are considered very difficult to repair with common tools, we propose that when worn out or damaged (perhaps every 10-20 years), they will need to be replaced.

SSD manufacturers typically provide lifetimes in terms of number of writes to the drive, since molecular wear occurs with each write on the flash memory gate storing the written value. For example, one 120 GB Samsung SSD has a lifetime of 100 terabytes written. At the typically cited estimated "average" workstation usage of 10 GB per day, this SSD has a lifetime of about 28 years, with lifetime scaling roughly linearly with the size of the drive [1]. Therefore, we propose that a SSD will only need to be replaced every 20-30 years at this stock workload, though performance will decrease steadily throughout the drive's lifetime as cells fail, and may drop below that required by the user. Write intensive workloads will naturally lead to much faster SSD failure depending on the nature of the workload.

Parts with moving components other than HDDs, such as optical drives and fans, are also susceptible to wear over time, but have been less well studied. MTBF values for consumer CPU fans are typically specified in the 30-50,000 hour range, or 3.4-5.7 years, though high-end CPU fans can be found with listed MTBFs of 28 years [37]. However, unlike HDDs, fans are amenable to cleaning, lubrication, and repair, and may not need to be replaced as often with regular maintenance [13].

Finally, some components age over time via chemical processes. One common repair is the replacement of electrolytic capacitors in a power supply unit (PSU) or on a motherboard, due to the slow evaporation of the electrolyte resulting in decreased capacitance. Typical consumer electrolytic capacitors are rated to run for 2000 hours at either 85C or 105C; depending on the type, at a working temperature of 45C they will have a lifetime of around 3.7 or 14.6 years of continuous use, respectively, with the lifetime highly dependent on temperature [15]. Unfortunately, unused electrolytic capacitors have a shelf life of only 2-3 years, due to degradation of the aluminum oxide layer insulating the capacitor foil. They may be usable after "reformation," in which a DC voltage is applied to the capacitor over a period of days or weeks to restore the aluminum oxide layer [42]. A better solution might be to replace the electrolytics with a few smaller but longer-lasting ceramic capacitors (lifetime 100+ years) in parallel and a resistor in series to mimic the properties of the electrolytic capacitor [55].

Also, after just a few years depending on environmental conditions such as temperature, thermal grease applied between a CPU and heatsink may solidify and crack, introducing air gaps that decrease the effectiveness of cooling.

It is unclear from our research exactly when this happens or whether it can be prevented; however, if detected before any damage occurs to the CPU, the hardened grease can be removed with an organic solvent and reapplied. If damage does occur to the CPU, a replacement chip must be procured and substituted, which may be possible or prohibitively difficult depending on whether the CPU is socketed or soldered directly to the motherboard.

2.3 Environmental Management

In order to maintain a longevity-friendly environment for computers in scenario one, the units would ideally be kept in a clean-room-like environment, with air filtering, rigorous entry and exit protocols, low humidity, and cool temperatures to avoid overheating [30]. Regular maintenance, such as cleaning of parts vulnerable to dust such as fans, could also prevent avoidable damage. Finally, one of the most important features of this environment would be a clean, reliable source of power to prevent surges and outages that would damage either the computing devices or the equipment being used to maintain favorable environmental conditions for its survival.

2.3.1 Power Management

Computing will only be possible with some power source, whether via intermittent grid electricity or an off-grid solution. An exploration of the space of power systems that could provide clean, reliable power for computing devices is out of the scope of this paper, but we describe one such minimal, off-grid system to show that it would be feasible to build and maintain.

The following system is based on current solutions for off-grid power used in RVs and boats: A constant-voltage DC power source such as a solar panel charges a 12V battery system, either a 12V car/marine lead-acid battery or pairs of 6V go-kart/motorcycle lead-acid batteries, with a simple low voltage indicator (made from LEDs and resistors, with no IC). A 12V DC car/marine PSU draws power from the batteries, and powers the computer. When the sun is shining, the solar panels charge the batteries up to their "full" voltage via constant-voltage (CV) charging; as the computer runs, it drains the batteries until the low voltage indication, at which point the user should turn the system off until the sun is shining again.

Each part of this system is essential: the solar panels produce power, the batteries handle input dropouts, and the PSU takes the slightly-fluctuating DC input and produces clean power at multiple voltages. Common warranties on modern solar panels guarantee an output of no less than 80% of the rated power over the first 25 years of use. However, with a typical degradation rate of 0.5% a year, the output should not fall below 80% for the first 44.5 years [32]. Typical lead-acid car batteries last 0.5-4 years inside a car depending on usage, but would last longer in more favorable temperatures and avoiding deep discharge while attached to a solar panel [26]. Sealed lead-acid (or VRLA) batteries last up to 10 years without maintenance, and even after sulfation are regularly revived and reused [40]. We expect commonly available DC/DC PSUs to also have electrolytic capacitors, and therefore similar lifetimes to AC/DC PSUs; to extend their lifetime, the same capacitor replacements as described above would be required. Therefore, we conclude that computing would likely not be limited by a lack of mains power; it would be feasible to maintain a power system that would last the lifetime of a computer and inflict minimal damage on its hardware.

If an inverter (with a standard life expectancy 10 years [48]) were added to the system, a standard AC PSU could be used instead of a DC one, and lead-acid batteries could be skipped for an off-the-shelf uninterruptible power supply (UPS) system (with a life expectancy of 3-5 years [51]). It would also be feasible to reconstruct the function of a UPS with a charging circuit, lead-acid battery, and an inverter, which would likely be more robust and have a longer shelf life than a UPS. Many options exist for powering computation according to need and hardware availability at the time.

2.4 Mobility-limited Components

For a baseline failure rate for mobile computing, we refer to a Consumer Reports study in 2015 that claimed modern consumer laptops have a 10-20% chance of failure over the first three years of ownership, with a median of 18% [54]. The median half-life computed from this value is 10.47 years, although as we have explored in previous sections, the annual failure rate of hardware tends to accelerate with age. According to an older study by SquareTrade in 2009 [46], which cited a higher failure rate of 30% over the first three years, about a third of laptop failures were due to accidents as opposed to malfunction. As hardware reliability has improved with SSD proliferation, this proportion has likely risen. Specific repair challenges are detailed below.

Mobile laptops tend to suffer damage from exposure to heat, dust, dirt, and water (especially containing salt). To repair corrosion and/or shorted electronics due to water damage, the corroded metal can be removed using isopropyl alcohol, and the electronics can be replaced by soldering. However, this kind of repair takes considerable care, effort, and expertise, especially with the tight integration and decreasing size of hardware components in modern laptops.

Another limit to longevity is that laptop batteries are consumables with a lifetime of 2-5 years, and need to be replaced for the continuation of mobile use, though said replacement is trivial to perform when the part is available [3].

Repetitive physical handling due to mobility can lead to mechanical constraint wear on case screws, tape, and glue (especially after multiple repair-related disassemblies). Laptop form factors tend to differ significantly between models, so if the chassis is cracked or falls apart due to being handled roughly or dropped, a replacement may have to be fabricated from some renewable material such as wood (which has been proposed for laptop chassis in some renewable designs [21]).

On the other hand, while the chassis and peripherals may be flimsy or complicated to replace, laptops are designed to be compatible with a large variety of spare peripherals such as external monitors and USB mice and keyboards. A laptop may theoretically remain usable for computation long after the chassis has been replaced by a box housing just the motherboard, storage drive, and peripheral connectors.

Ruggedization against foreign particle entry might also help mitigate exposure to the elements. Specifically designed ruggedized computing devices are costly but available according to military specifications [58] for applications such as warfighting or construction. At the most basic level of protective design, a HDD in a laptop can be replaced with a SSD before mobile use in order to avoid mechanical damage to the storage drive, and potential errors in stored data.

Table 1: Summary of Recommended Replacement Parts and Estimated Lifetimes
(H) in the Estimated Life column indicates that the value is a half-life computed from other ratings.

Limited by	Part	Estimated Life (yrs)	Notes
Computation	Ceramic capacitors	100+	(MLCCs) Could replace electrolytic capacitors
	SSD	20–30	120 GB SSD at 10 GBW/day
	HDD	9.7–13.5 (H)	At 5% baseline failure/yr
	Electrolytic capacitors	3.7–14.6	Affects PSU, motherboard, AC/DC adapters
	CPU fans	3.4–5.7	Longer life with cleaning/lubrication/repair
	Thermal grease	2+	Depends on temp and conditions
Mobility	Aggregate of device parts	10.47 (H)	Based on Consumer Reports study
	Peripherals	Variable	Screens/monitors, keyboards, mice
	Li-ion batteries	2–5	Computer technically works without batteries
Power	Solar panels	50+	$\geq$ 80% of rated power output for first 25 yrs
	DC/DC PSU	3.7–14.6	Assumed limited by electrolytic capacitors
	Inverter	10	Standard for solar inverters
	Sealed lead-acid battery	10	Easy to repair with standard tools
	UPS	3–5	May also be built w/ lead-acid battery
	Unsealed lead-acid battery	0.5-4	Easy to repair with standard tools

2.5 Resources for Repair

In order to sustain the repairs mentioned above, replacement parts must either be kept in stock by the device owner, or available through a procurement network. For the mobility-limited scenario, we discussed needing HDDs or SSDs, fans, electrolytic capacitors, PSUs, thermal grease, and possibly CPUs. Additional parts would be desired for the mobile scenario, including Li-ion batteries, screens or external monitors, keyboards, mice and other peripherals, and AC/DC adapters and cords if AC mains power was still available. See Table 1 for a summary of commonly required parts. The question remains whether all of the the replacement components will have shelf lives long enough to be usable for repairs after fifty or a hundred years. For example, SSDs packed for long term storage in a temperate environment, with desiccant, and away from radiation, are likely to remain in good condition after 15 years or more, because integrated circuits are expected to last as long under the same circumstances [31]. However, not much work has been done on measuring their shelf life for longer periods.

Just as important for successful repairs would be human resources with the skills needed to perform them, such as soldering and use of a multimeter. Without intentional teaching and community retention of these skills even in a generation of less computing ubiquity than we have today, the skills could be lost to many communities. Social networks or institutions of people interested in computer repair could be invaluable for sourcing parts and maintaining skills needed to keep computing alive until devices and power are no longer scarce.

3. MITIGATING SOFTWARE RISK

Software degradation is less predictable than hardware degradation and subject to different challenges under collapse. While software does not "wear out" like hardware, it can be slowly corrupted over time, often has external dependencies, and is still subject to contamination from the outside environment.

3.1 Limits on Development and Distribution

In a scenario where power and manufacturing infrastructure have degraded, the Internet would likely cease to exist as well, which poses an enormous number of threats to modern software functionality. Firstly, software distribution would mostly cease, as would the distribution of bug fixes and security patches. Secondly, cloud infrastructure would not be available, which would suspend all web services immediately, and render renewable-license software void with no means to renew at the end of the license period. Finally, software development would slow to a crawl without the current development ecosystem, which has co-evolved with increasing societal connectivity. Unfortunately, software development would be needed as a crucial line of defense against malicious software (malware) infection, another significant threat to computing discussed below.

Rapid innovation in software today depends heavily on web-based tools for easy long-distance discussion, technical search, and software distribution. Without communication tools, online documentation, and cloud-hosted code repositories such as github and npm to facilitate collaboration, developers would have to work and learn individually through time consuming experimentation, likely replicating each other's code [38].

3.1.1 Solutions for Developer Collaboration

Collaboration through distributed version control tools would be possible without the Internet, but would require either co-location of developers or the establishment a highly reliable developer network. From our discussion on hardware above, we believe that tightly integrated mobile computing platforms like modern laptops will fail faster than stationary desktops and servers with easily replaceable components. The eventual depletion of mobile computing resources will make it increasingly difficult to gather people and their computers in a single location. Therefore, it may be crucial to establish communication channels, file sharing practices, and communities for maintaining software engineering knowledge before the breakdown of mobile computing.

These communications could be as simple as broadcast-

ing code over radio, which was done in Finland in 1985 as part of an effort to stimulate interest in computing [28]. Another strategy could be to establish decentralized communication networks over sneakernet with cryptographically assured messaging, or point to point wireless systems as inspired by community networks and ham radio [14]. Regardless of communication medium, person to person networking will be an important part of post-collapse computing, without centralized Internet communities to establish reputation and put developers in initial contact with one another. When remote collaboration becomes infeasible, computing centers could be established to bring software engineers to the same physical location to allow in-person collaboration.

3.2 Data Decay over Time

High barriers to verified file sharing also create challenges to maintaining correct copies of data, including software. All data is vulnerable to subtle faults of the underlying hardware it is stored on, including in-memory bit flips [17, 53] and on-disk file corruption [18]. The widely deployed Windows operating system does not implement error correction codes in its default filesystem, and commodity consumer hardware eschews error correction-enabled memory for lower cost and higher performance. Over time flaws will accumulate; while corruption to non-essential files could be harmless, corruption of key files in the operating system or critical user applications could cause irrecoverable failure of the overall computing resource. In our well-resourced world we can ignore these issues because it is easy to reinstall an application, and software lifetimes are relatively short. However, in a collapse scenario, everyday users must take on the burdens of data management and preservation that are left to data center administrators and archivists today.

3.3 Trust Breakdown

A fundamental but relatively invisible piece of modern software infrastructure is the ubiquitously available public key infrastructure (PKI). Centralized certificate authorities sign and validate website secure socket layer (SSL) credentials, software packages, and system updates to give end users a reasonable way to validate their authenticity. While nothing in the cryptographic principles of PKI requires centralization, it does require a root of trust upon which chains of trust can be built to validate third parties. In an environment with extremely limited connectivity, it will be difficult for content creators to obtain digital signatures that will be trusted by all the end users that content may eventually reach. Most SSL certificates distributed with browsers and operating systems have expiration dates, beyond which key invalid errors will be thrown by the validating software. As seen in Figure 1, all root certificates on a currently up to date system will be invalid in 30 years. While users can continue to rely on expired keys, they will have to override warnings and run the risk of long-held keys being compromised with no way to get replacements.

A systematic breakdown of the current signing infrastructure will further complicate the problem of software authenticity verification and increase the chances that normal users encounter malware through compromised content. Without an understanding of how PKI operates users will have a difficult time handling the remnants of the current implementation and making the right choices with regards to trust and system security that are handled transparently today.

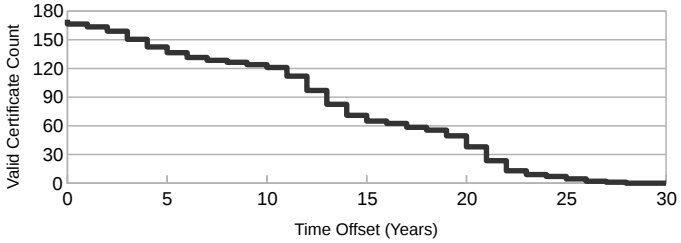

Figure 1: Valid Certificates vs. Time
Measured from expiration dates on installed SSL root certificates on an up to date Ubuntu 16.10 system.

3.4 Malicious Software

Finally, we see malware as potentially the single largest threat to productive computing after collapse given its volatile nature and high risk of harm. Malware can cause varying levels of disruption to a computing system, ranging from passive non-interference up to catastrophic data loss or even irreparable hardware damage [4, 27, 33]. In developing regions today many users cope with systems infected with malware, but at substantial cost to productivity and security [8, 11, 23, 35]. In a collapse scenario where new, trusted copies of data cannot be easily retrieved and systems cannot just be wiped and reinstalled, users may have to cope with the effects of malware infections indefinitely.

While a collapse event significant enough to impact computing capability may also diminish incentives to create malware due to decreased computer usage, in some scenarios they may actually be enhanced. For example, in a collapse triggered by warfare, cyber weapons may be intentionally developed and deployed by opposing factions to harm critical infrastructure[25, 47]. Collateral damage from such weapons could spread unchecked through the software ecosystem if no countermeasures are in place. Malware authors also write for a variety of other personal motivations, such as boredom, which may not disappear after collapse [52]. It only takes one developer to create and release a piece of malware, but containing it requires coordinated effort to update the systems of a large number of vulnerable users. Furthermore, current malware in the wild will not cease to exist, and users will have to contend with any malicious code deployed but not yet patched at the time centralized update services fail.

Many types of malware have the advantage of spreading virally through incidental contact with other systems, while patches, not commonly spread peer to peer, will be slowed by the destruction of centralized distribution channels [10]. Without the Internet, users would have to rely on peer-to-peer file transfers to productively exchange megabytes of information [50]. Direct file transfers provide no way to verify the authenticity of received files before opening, and without updates to malware signature databases users will have no way to identify new malware in received files [16].

An example of the damage malware can cause can be seen today in the computing systems in Internet cafes in developing regions [8]. The authors have also personally witnessed the impact of malware infections on computers in public university computer labs in Ethiopia, where users must resort to unsafe security practices to share data even if they know better. Without reliable connectivity, vulnerable USB drives or direct ad-hoc wireless connections are the file transfer mediums of choice, and without access to official distri-

bution channels the only way to acquire software and media is often through the illegal downloading and sharing [9, 50]. Cracked software and digital rights management stripped media is frequently contaminated and commonly becomes a vector for malware transmission [16]. The contemporary experiences of users in these conditions inform our expectations of a future collapse computing scenario.

3.4.1 Software Recovery

Presently only two main models currently exist for the recovery of systems compromised by malware. The first involves expert security researchers and developers characterizing malware infections, designing a tailored removal tool, and deploying that removal tool to infected users to restore their systems. Experts also generate signatures of the malware to detect and prevent future infections. Severe collapse scenarios preclude usage of this model due to a lack of connectivity for experts to gather malware samples from the broader user base and then distribute fixes. Isolated groups of users will likely not have access to the expert resources and time required to solve problems in this manner.

The other more extreme model, completely wiping and restoring the computer from a new OS image, is often used as a last resort in developed countries against rootkits or sophisticated malware, and as a regular cleansing operation in developing contexts where tailored fixes may not be available [8, 19]. The source image for the new operating system install can come either from a restricted partition on the user's hard drive or from a dedicated piece of external installation media. On new machines commonly provided without disk drives today, the partition approach is favored for most users to decrease costs on the manufacturer and simplify recovery. However, the partition approach presents several notable disadvantages: since the partition is always physically present on the computer sophisticated attacks could bypass OS security measures and modify data on the partition to infect the recovery image. Similarly, since the partition is tied to the same physical disk as the running OS, failure or corruption of that disk could damage the image. Lastly, the image will still be vulnerable to the original exploit and reinstalling it will not prevent future infections.

In a collapse scenario long term maintenance of reliable backup data becomes both much more important for system longevity and much more difficult to achieve with limited resources. Present day solutions rely on software to manage backup images, but secure hardened backup stores grounded in hardware would provide more assurance that software bugs could not be exploited to gain access. Physical switches allowing read, append, or write access to hardware isolated storage would help users take control of their backup data storage reliably and explicitly.

New sophisticated attacks have been recently uncovered that target low level device firmware on the system's hardware itself, persisting across a complete OS level restore [60, 20]. Recovering from these attacks requires either acquiring new hardware or having access to low level firmware flash tools as a part of the recovery process. Without planning for such a contingency prior to collapse, users infected with this type of malware could be unable to restore their systems to working condition [45].

3.4.2 Sustainable Malware Inoculation

While malware has the advantage of self-replication and contact spread, the same principles could be applied by trusted software sources to distribute patches organically. As demonstrated in Ghana by the FlashPatch project [11], it is possible to piggyback antivirus definitions onto regular file transfers over USB, reaching machines otherwise cut off from network connectivity. The same system could be used to transport signed OS packages or core firmware updates which could be incorporated and passed on automatically by end users who trust the original signing authority. Such a system would require a distributed web of trust based on strong cryptography to allow software packages to be validated securely on remote machines that may have never directly communicated with the creating entity. Such a web could be built with primitives that exist today, such as PGP signatures or another certificate framework. Additionally, long term viral distribution of OS patches would require a user-friendly way to manage patch conflicts (imagine two disconnected developers fixing the same bug) and a way to condense layered patches to keep the space required for their distribution in check over time. While storage is relatively large and inexpensive relative to the size of required packages, there is currently no approach for managing patch conflicts at any level higher than the source code. Further research would be required to enable distributed updates in a transparent and user-friendly manner.

3.4.3 Malware-Tolerant Systems

An important aspect of sustainable defense against malware will be not only preventing infection (as it will become increasingly unavoidable), but containing the damage that follows. One approach to increasing system resilience involves sandboxing different parts of the computer system at a low level to provide high assurance of isolation and better user visibility into system behavior. In security-oriented operating systems like Qubes [57], virtualization technology separates small parts of the operating system into isolated zones with well defined communication permissions and protocols between them [44]. This minimizes the attack surface exposed by each component while allowing users to catch anomalies in communication through intelligent monitoring. Action can then be taken to replace compromised zones before the infection spreads to the entire system and user data is compromised. Replacing a single zone of the system is much easier and lower-cost than restoring the entire OS. Additionally, hypervisors enforcing virtualization security policy can be simple and minimalistic enough to be formally verified and guaranteed to meet security specifications [7, 24, 39].

Other resilience models are possible as well, potentially drawing from existing concepts in fault-tolerant computing or the design of secure information systems for classified data [36, 22]. Approximate computing techniques could even be applied where multiple runs of a computation are attempted in corrupted environments and results are combined intelligently to catch and repair introduced errors. The high performance computing community is already exploring such techniques for large scale computing at the limits of error correcting code memory [17, 53]. As long as the "viral load" and corruption introduced into the computation was low enough, useful information could still be extracted from compromised compute resources.

4. DISCUSSION

4.1 Designing Systems for Collapse

Emphasizing longevity and repairability instead of up front cost, maximum initial performance, or low size, weight, and power significantly changes the tradespace in designing a computing machine [43]. Present day conditions without limits incentivize design and construction of machines, which while capable in the present environment, may not be adequate for sustainable computing in an extremely limited collapse environment. Notably, IT professionals often focus on hardware longevity in planning for overall system longevity, assuming the availability of valid software, global connectivity, standardized architectures, and a strong network of software developers. However, in a collapse scenario, access to both replacement hardware and up-to-date, uncorrupted software will become limiting system constraints.

4.1.1 User-Mediated Security

An important fundamental paradigm for collapse computing will be putting control of system security back into the hands of users, with human factors in mind. Cut off from centralized services of security researchers and patches, users will need the tools to take system and network security into their own hands. Permissions based systems, like User/Group permissions in Unix derivatives, provide some security; however, they are often difficult for users to understand and configure correctly deprived of context, and present too many uninformative, ignorable prompts [56, 61]. General purpose monitoring tools like file system monitors, registry watchers, or network traffic classifiers increase system transparency at the risk of overwhelming users with false positive warnings and drowning attack signals in noise from nominal system operation [61]. Ongoing work on privilege elevation triage and system security transparency could make systems better able to detect threats from noise by adapting to expected usage patterns and local states. Once updates cease, malware that works around rigid security paradigms will probably proliferate, but well designed human-in-the-loop security paradigms could continue to function as non-technical end users modify their best practices in response to threats evolving in the wild.

More research could also be done towards establishing strong user data protection in the face of system compromise. Hardware enforced filesystem access could protect critical data stores for keys or recovery images by requiring explicit physical action from the user to enable reads or writes. Hardware could also enforce backup policies, ensuring that recovery copies of data always remain available, or that attempts to destroy or modify backups are brought to the user's attention.

4.2 Social Mechanisms for Maintenance

Per the above analysis, we might expect computing hardware and software to persist in well-maintained environments for several generations, and in mobile forms for approximately one generation. This multigenerational effort relies upon a knowledge and culture of maintenance, and so may fail for cultural reasons; just as we have considered the obsolescence of computing hardware and software, so too must we consider the obsolescence of computing culture, and how it might persist or rot.

History offers many examples of infrastructural maintenance after a collapse, but two interestingly divergent ones are the Chinese and Roman road networks built from around the second century BC to the third century AD, and decaying thereafter. While the Roman network decayed rapidly, contributing to cultural disconnects of the early Middle Ages, the Chinese road network was maintained, albeit reduced from wide roads that could handle drawn carts to narrow ones designed for wheelbarrows [12]. This maintenance was performed by cultural organizations such as the Taoist Yellow Turbans and Buddhist fraternities as a component of their training and service. Perhaps computing could continue similarly after collapse, as public enclaves maintained by semi-ascetic cultural organizations whose primary focus may or may not be computing. Such a situation might lead to a kind of software and hardware monoculture designed for application by non-technical adherents.

For a social model more preserving of technical development effort we can look to the history of early personal computing. As hardware began to enter the mainstream, enthusiast groups maintained and created many of the shared understandings and technologies that allowed individuals to engage with computing [59]. Were post-collapse computing to follow this framework, much of our current technical knowledge, computing heterogeneity, and software development ecosystem might be maintained, but with informal software distribution channels malware could be quite a burden.

Even further back in the history of computing, we recall the development of LISP, whose fundamental lambda calculus was specified in the mathematics literature [2] two decades before it was used for computing [34]. Even as computing collapses, a rich body of computer science literature could survive. New results in encryption, compilers, and other immediately applicable research could be argued mathematically before being input to rare computing resources. Computing could be reserved to polish and finish work already peer-reviewed, maintaining a capable and trusted but highly restricted computing resource for the academic community.

In the discussion of PKI infrastructure above, the importance of trusted transportation was mentioned; historical analogues for this might include the early postal systems of Europe and the Pony Express. Such logistical businesses could of course benefit heavily from computing themselves; one could even imagine overlapping competitive transportation networks offering computing services and software patches from afar, an environment which would fully explore both hardware longevity after collapse and the dangers of malware.

Taken together, these historical examples make it clear that along with the analyses of hardware and software, the roles that computing might take in society are important factors for the continuation of computing after collapse. What groups will have access to what computing resources? Will these resources be captured and centralized by groups with power, or maintained in a decentralized fashion? How will the education and training necessary to fully utilize and adapt computing to new societies be passed down from generation to generation? These questions call for the study and creation of sustainable and resilient modern computing cultures.

5. CONCLUSIONS AND FUTURE WORK

Collapse scenarios present existential challenges to the preservation of computing capability in the post-collapse context. Hardware, software, and user data all face threats to survival in an environment with limited replacement part availability, limited communications and power infrastructure, and limited software development capabilities. While there are challenges to maintaining hardware in such a constrained scenario, they are relatively well-understood. With sufficient replacement parts and care, a commodity computer may be maintained and powered for the duration of a temporary collapse of several decades. Software, however, presents a set of challenges that are harder to mitigate, as the detrimental effects of long term disconnection, software data corruption, and malware are numerous and potentially devastating.

Further research on computing within these limits could directly benefit users in today's collapse scenarios while improving the survivability of computing as a whole. Significant areas for future work include: further investigation into the longevity and care of hardware in use and storage, to improve the overall environmental sustainability of computing; development of flexible user-centric security paradigms so systems can adapt to changing threats without regular software updates; computing systems designed for secure full recovery in the face of malware infection; and design of distribution technologies to allow secure development and deployment of software without a global Internet.

6. ACKNOWLEDGMENTS

We would like to thank our labmates at the University of Washington and the wider ICTD community for their help and inspiration in the creation of this work.

7. REFERENCES

[1] SSD Endurance Test - Live Testing Samsung EVO, SanDisk, Intel and Kingston.

[2] Alonzo Church. *The calculi of lambda-conversion,*. Annals of mathematical studies ; no. 6. Princeton University Press; HMilford, Oxford University Press, Princeton, London, 1941.

[3] Apple Computer. Determining battery cycle count for Mac notebooks.

[4] AVTest. Security Report 2015/2016. Technical report, AV Test, Magdeburg German.

[5] Backblaze. 2016 Hard Drive Failure Rates for 2tb - 8tb Drives, Nov. 2016.

[6] Backblaze. 2016 Hard Drive Reliabilty Benchmark Stats, Jan. 2017.

[7] G. Barthe, G. Betarte, J. D. Campo, and C. Luna. Formally verifying isolation and availability in an idealized model of virtualization. In *International Symposium on Formal Methods*, pages 231–245. Springer, 2011.

[8] P. Bhattacharya and W. Thies. Computer viruses in urban Indian telecenters: Characterizing an unsolved problem. In *Proceedings of the 5th ACM workshop on Networked systems for developing regions*, pages 45–50. ACM, 2011.

[9] J. Chen, M. Paik, and K. McCabe. Exploring Internet Security Perceptions and Practices in Urban Ghana. In *SOUPS*, pages 129–142, 2014.

[10] L.-C. Chen and K. Carley. The Impact of Countermeasure Propagation on the Prevalence of Computer Viruses. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, 34(2):823–833, Apr. 2004.

[11] H. Corrigan-Gibbs and J. Chen. FlashPatch: Spreading Software Updates over Flash Drives in Under-connected Regions. pages 1–10. ACM Press, 2014.

[12] K. De Decker. How to Downsize a Transport Network: The Chinese Wheelbarrow, Dec. 2011.

[13] eBay. How to Repair a CPU Fan, Mar. 2016.

[14] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 27–34. ACM, 2003.

[15] M. Fortunato. Ensure long lifetimes from electrolytic capacitors: A case study in LED light bulbs, Apr. 2013.

[16] J. F. Gantz, P. Soper, T. Vavra, L. Smith, V. Lim, and S. Minton. Unlicensed Software and Cybersecurity Threats, Jan. 2015.

[17] A. Geist. How To Kill A Supercomputer: Dirty Power, Cosmic Rays, and Bad Solder, Feb. 2016.

[18] J. Gray and C. Van Ingen. Empirical measurements of disk failure rates and error rates. *arXiv preprint cs/0701166*, Dec. 2005.

[19] S. Guo, M. H. Falaki, E. A. Oliver, S. Ur Rahman, A. Seth, M. A. Zaharia, and S. Keshav. Very low-cost internet access using KioskNet. *ACM SIGCOMM Computer Communication Review*, 37(5):95–100, 2007.

[20] A. Hern. Lenovo does it again as LSE component removed after security fears. *The Guardian*, Aug. 2015.

[21] S. Hickey, C. Fitzpatrick, P. Maher, J. Ospina, K. Schischke, P. Beigl, I. Vidorreta, M. Yang, I. D. Williams, and E. den Boer. Towards zero waste in industrial networks: A case study of the D4r laptop. In *Proceedings of the Institution of Civil Engineers - Waste and Resource Management*, volume 167, pages 101–108, Aug. 2014.

[22] Information Assurance Directorate. The Community Gold Standard Framework 2.0, June 2014.

[23] IT News Africa. 'Raila Odinga' computer virus routs Malawi's cities, Sept. 2008.

[24] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, and others. seL4: Formal verification of an OS kernel. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pages 207–220. ACM, 2009.

[25] T. Koppel. *Lights Out: A Cyberattack, A Nation Unprepared, Surviving the Aftermath*. Crown, New York, 1st edition edition, Oct. 2015.

[26] B. Lampe. The Average Car Battery Life: When is it Time for a Change?, Mar. 2016.

[27] R. Langner. Stuxnet: Dissecting a Cyberwarfare Weapon. *IEEE Security Privacy*, 9(3):49–51, May 2011.

[28] M. Lasar. Experiments in airborne BASICâĂŤ'buzzing' computer code over FM radio,

Aug. 2012.
[29] J. Lee. The Effects Power Outages Can Have On Your Computer, Sept. 2014.
[30] Liberty Industries. Cleanroom Operating & Maintenance Protocol.
[31] R. R. Madsen. Component Reliability After Long Term Storage. *Texas Instruments*, 2008.
[32] M. A. Maehlum. The Real Lifespan of Solar Panels, May 2014.
[33] Malwarebytes Labs. The State of Malware: 2017. Technical report, Malwarebytes Labs, Feb. 2017.
[34] J. McCarthy. Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I. *Commun. ACM*, 3(4):184–195, Apr. 1960.
[35] C. Michael. Computer viruses slow African expansion. *The Guardian*, Aug. 2009.
[36] National Institute of Standards and Technology. Security requirements for cryptographic modules. *Federal Information Processing Standards Publication Series*, May 2001.
[37] Orion Fans. Life Expectancy, 2017.
[38] B. Penzenstadler, A. Raturi, D. J. Richardson, M. S. Silberman, and B. Tomlinson. Collapse (and other futures) software engineering. *First Monday*, 20(8), 2015.
[39] G. Plouviez, E. Encrenaz, and F. WajsbÃijrt. A Formally Verified Static Hypervisor with Hardware Support for a Many-Core Chip. In *Euro-Par 2013: Parallel Processing Workshops*, pages 801–811. Springer, Berlin, Heidelberg, Aug. 2013.
[40] PowerThru. Lead Acid Battery Working Lifetime Study.
[41] B. Raghavan and S. Hasan. Macroscopically sustainable networking: on internet quines. pages 1–6. ACM Press, 2016.
[42] T. Reese. Strategies to Repair or Replace Old Electrolytic Capacitors.
[43] C. Remy and E. M. Huang. Sustainable Interaction Design: Obsolescence in a Future of Collapse and Resource Scarcity.
[44] J. Rutkowska. Software compartmentalization vs. physical separation, Aug. 2014.
[45] J. Rutkowska. State considered harmful. 2015.
[46] A. Sands and V. Tseng. SquareTrade Laptop Reliability, Nov. 2009.
[47] D. E. Sanger and M. Mazzetti. U.S. Had Cyberattack Plan if Iran Nuclear Dispute Led to Conflict. *The New York Times*, Feb. 2016.
[48] N. Santhanam. What is the Lifetime of Solar Inverters?, Sept. 2015.
[49] B. Schroeder and G. A. Gibson. Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? San Jose, CA, Feb. 2007. USENIX.
[50] T. N. Smyth, S. Kumar, I. Medhi, and K. Toyama. Where there's a will there's a way: mobile media sharing in urban india. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 753–762. ACM, 2010.
[51] J. Solis. Tips to Maximize the Life Expectancy of Your UPS System, Apr. 2015.
[52] E. H. Spafford. Computer Viruses and Ethics. 1991.
[53] V. Sridharan, N. DeBardeleben, S. Blanchard, K. B. Ferreira, J. Stearley, J. Shalf, and S. Gurumurthi. Memory Errors in Modern Systems: The Good, The Bad, and The Ugly. pages 297–310. ACM Press, 2015.
[54] D. Tapellini. Survey Results: The Most Reliable Laptops, Oct. 2015.
[55] TDK. Guide to Replacing an Electrolytic Capacitor with an MLCC | Multilayer Ceramic Chip Capacitors.
[56] The Ponemon Institute. The Cost of Malware Containment. Technical report, The Ponemon Institute, Jan. 2015.
[57] The Qubes OS Project. Qubes OS.
[58] US Department of Defense. Department of Defense Test Method Standard: Environmental Engineering Considerations and Laboratory Tests, Oct. 2008.
[59] S. Wozniak. Homebrew And How The Apple Came To Be.
[60] J. Zaddach. *Implementation and Implications of a Stealth Hard-Drive Backdoor*. 2011. OCLC: 873035573.
[61] M. E. Zurko. User-centered security: Stepping up to the grand challenge. In *Computer Security Applications Conference, 21st Annual*, pages 14–pp. IEEE, 2005.