Stipulated Smartphones for Students: The Requirements of Modern Technology for Academia

Rob McGuinness
University of California, San Diego
La Jolla, California, USA
jrmcguin@cs.ucsd.edu

George Porter
University of California, San Diego
La Jolla, California, USA
gmporter@cs.ucsd.edu

ABSTRACT

Technology has become pervasive in everyday life. College students are required to use internet-enabled computers and smartphones to access and interact with course material, often via browser software. However, the trend of requiring technology disadvantages lower-income students. Additionally, the increased cycle of manufacturing, purchasing, and discarding devices comes with an environmental cost in the form of eWaste. Research has been done examining how long users can and do keep devices before discarding them, but we wish to understand this in the context of requirements for college students.

In this paper, we examine how well online learning platforms function on older browser software and device hardware. We then perform an analysis over several years of data from a learning management system used at the authors' university campus. We find that software and hardware become obsolete within roughly a four-year period, meaning students are likely to be required to purchase a new smartphone or laptop during their college careers. In reality, these "obsolete" devices are just as capable as they were when they were new. We advocate for a student-focused solution and examine possible future lines of research in this area.

KEYWORDS

device longevity, online education, computing sustainability, hardware end-of-life, browser usage

Reference Format:
Rob McGuinness and George Porter. 2021. Stipulated Smartphones for Students: The Requirements of Modern Technology for Academia. In *LIMITS '21: Workshop on Computing within Limits, June 14–15, 2021.*

1 INTRODUCTION

Computing has become an invariable and essential part of everyday life. Individuals are now required to use digital devices in order to carry out daily tasks in both personal and professional settings. Supporting this change is a rapid and constant evolution in hardware and software capabilities. This is evident in the purchase rate of new smartphones, with users replacing them roughly every 20 months on average [22]. Users may replace their device for a large number of reasons, such as wanting higher resolution screens, better cameras, or better software performance.

This inevitable "march of progress" can result in newer, widely used software that that is cumbersome or impossible to use on older devices. Users that require this software for their work, or students requiring this software for their education may find that their older devices no longer work correctly and no longer receive necessary software updates. They may be able to compromise by relying on a degraded interface to these software services (e.g. via a browser instead of a dedicated app). However, as we will show later in this paper, even browser interfaces often become unusuable, forcing users to upgrade their devices to continue accessing the modern web. Those upgraded devices comes with numerous costs, both to users in the form of monetary cost and the world at large in the form of pollution, eWaste, and increased carbon output.

In this work, we examine these costs with respect to undergraduate college students, a sector of the population that now is required to use technology to receive an education in the wake of the COVID-19 pandemic. Even before to the pandemic, learning management systems (LMS) such as Blackboard and Canvas have garnered widespread use [6, 15], forcing students to use internet browsers or mobile apps to access and interact with coursework.

We look at three different types of costs that occur in relation to these educational users:

Monetary Costs: With smartphone costs doubling between 2014 to 2018 [26], and continuing to rise, we note that students are disproportionately affected by the cost of frequently purchasing a new device. A majority of students now receive some form of financial aid to afford a college education [18], and some percentage of students additionally require aid for basic needs like food and shelter. Students that are required to purchase new devices to continue to access basic education materials may have to make compromises against their own personal well-being in order to fulfill class requirements. We argue that a student should never have to make this compromise in order to receive an education, especially since at a component level, their "older" devices are likely released only a few years ago.

eWaste Costs: The other significant cost is to the world at large in the form of the eWaste generated from discarding old devices. In 2016, the world produced over 44 million metric tons of eWaste, with computers (laptops, desktops, smartphones, and tablets) accounting for about a quarter of that total [5]. The vast majority of eWaste is not properly discarded or recycled [4], which results in long-term damaging effects to the environment.

Manufacturing Costs: Additionally, the vast majority of environmental damage from computers and smartphones comes in their manufacturing process [4]. Manufacturing incurs enormous environmental damage, involving mining for minerals, generating greenhouse gasses during transportation and assembly, and processing elements like cobalt, lithium, and mercury. At the same time, once a computer goes into service, it is quite efficient, due to lower power components, flash storage instead of spinning hard drives, and more efficient battery technologies, among other reasons.

Taken together, it is clear that upgrading devices unnecessarily has huge personal and environmental effects on students. And yet because the web-hosted services that students rely on to complete their studies become increasingly complex over time, simply keeping older devices in use for longer is often no longer a feasible option. In this paper we take a critical look at this issue.

Our research in this paper focus on the question: **In the worst case, how long can a user realistically continue to use their device to access online educational resources?** We examine multiple sources of data to answer this question. We both independently test websites with a range of browser software and device hardware, and also perform analysis of real user data obtained from the authors' university, UC San Diego.

As a motivating use case, we look at websites used by an undergraduate computer science student. Using an online platform for browser testing, we survey a sample set of educational websites they might access across a range of legacy operating systems and browsers, dating back to 2012. We find that after about four years, the educational websites become increasingly inaccessible to older devices and software. We carry this information forward into our investigation of user data.

Our user data investigation is done via several years of access logs to a web-based learning management system, Blackboard, used by classes at the authors' college campus. We study the average age of browsers and devices used by students across each year, and study the upgrade pattern of devices to understand how old the software and hardware is "in the wild". We find a reinforcement of the "four year" lifetime implied by our website survey, seeing that there are extremely few cases where a user operates with software or hardware outside of this timespan.

Because a traditional undergraduate degree at a university takes four years to complete (and often more than four years), this implies that a student may be forced to purchase a new device within their career as a student at a university in order to complete their degree. However, we argue that a student should not be required to make this purchase, as their old hardware should be more than suitable.

We acknowledge that there are a number of both technical and non-technical reasons why people dispose of their computers to get new ones. However, performance and functionality are common concerns [4]. In reality, our "obsolete" devices are just as capable today as when they were brand new. There is a growing market of "refurbished devices" showing that users are willing to purchase older hardware, provided it still functions well [24]. We posit that this obsolescence is a function of the evolution of modern apps, websites, and web-based services that have grown increasingly more sophisticated and resource-hungry over time. Since most applications that we use on a day-to-day basis require some Internet-enabled functionality, the evolution of web-based services render our devices prematurely obsolete. In an era of long-distance learning, it is of particular importance to increase device longevity for students who now require computers to complete even baseline academic tasks like attending class.

2 BACKGROUND AND MOTIVATION

2.1 The ubiquity of computing in teaching and learning

Academic platforms are commonplace as a support framework for presenting coursework and managing student submissions for assignments. An academic-focused learning management system (LMS), Blackboard, reported more than 100 million users in 2017 [6]. A recently adopted LMS, Canvas, reported more than 30 million users as of 2019 [15]. Students are expected to use these systems to access and complete coursework, and require computing devices in order to do so. Commonly, these are accessed using internet browsers or dedicated mobile application software.

Additionally, the COVID-19 pandemic has had an extreme effect on the role of technology in teaching by requiring that classes begin to be taught remotely over video conferencing platforms. Instructors need to use online teaching methods for courses [19], which further reinforces the obligation for a student to own a device capable of accessing online learning platforms.

This trend is not particularly surprising when reflecting on the growth of technology across the world as a whole. The proliferation of mobile networks has meant that over 80% of the world's population is now covered by a mobile broadband signal [5]. The number of users has grown dramatically as well, with over 4.2 billion mobile broadband subscriptions active as of 2017 [5], over 50% of the world global population.

2.2 The cost and life cycle of a computer

In order to access the aforementioned online learning platforms, a suitable computing device must be available for each student. We must examine the costs of a computing device to provide context the effects they have on students and eWaste generation.

There are variety of costs observed when examining the lifetime of a computing device. These include monetary costs to the users and the environmental costs of manufacturing and discarding the device. Each of these costs can have disproportionate effects in different ways. Extending the life cycles of devices will offset these costs.

2.2.1 Monetary costs. Purchasing a computer or smartphone is a non-trivial cost to a user. While users have varying upgrade rates depending on a large range of factors, users with less financial assets are more likely to be unable to upgrade their devices as frequently. With the cost of smartphones rising quickly in the past decade [26], it is important to consider how students who operate on less funding can access technologies required to complete their coursework.

It is an unfortunate reality that many students require financial assistance. At the authors' university campus, a report [18] states that over 60% of students require some form of financial aid. Additionally, the on-campus food pantry provided food to thousands of students in a single quarter, with that number expected to climb in the following years. A separate report on community colleges [12] reported that two in three students were food insecure, and 13 to 14 percent of students were homeless.

With this in mind, we believe that the monetary costs for students to provide their own devices to access educational opportunities must be reasonable. Unfortunately, no programs to widely provide remote access technologies existed at the time of the shift to remote learning due to the COVID-19 pandemic at the authors' campus, which is a relatively large institution (over 40,000 enrolled students in the 2020 academic year [30]). Smaller and less funded institutions struggled disproportionately during the pandemic [19].

2.2.2 Environmental costs. The environmental cost of a device comes in the form of eWaste. eWaste is defined as any device with a plug, electrical cord, or battery that is no longer used and thus has reached the end of its useful life [4, 29]. eWaste is divided into six categories: temperature exchange equipment, screens and monitors, lamps, large equipment, small equipment, and small IT and telecommunications equipment. In this paper, we will use the term *devices* or *computers* to refer to the two categories of "screens and monitors" as well as "small IT and telecommunications equipment", which includes laptops, desktops, smartphones, and tablets. In 2016, the world produced 44.7 million metric tons of eWaste [29], with screens and computers accounting for about a quarter of that total volume. And while eWaste only accounts for about 2% of the total waste volume in landfills, it represents about 70% of the volume of hazardous waste that makes its way back into the ecosystem [29]. This waste is rich in precious, heavy, and rare-earth materials, with computers often consisting of up to 60 different elements from the periodic table [29].

To understand the environmental impact of eWaste, it is important to understand the entire lifecycle of modern computers, which can be broken down into four phases [4]: Manufacturing, use, eWaste generation, and eWaste disposal. Each of these phases has a different environmental impact:

Phase 1: Manufacturing. During manufacturing, materials are brought together to create integrated circuits, flash memory, screens, and other components. This process is very resource intensive, relying on a number of materials including gold, silver, copper, platinum, and aluminum, and heavy metals such as mercury, cobalt, iridium, cadmium, lead, and lithium. Mining these materials results in significant environmental damage, and transporting components from their origins to be integrated and delivered to their ultimate destinations incurs a significant carbon footprint as well.

Phase 2: Use. In this phase, the device is put into use, either by the primary owner, or subsequent owners in secondary markets. Here the environmental impact of the device is primarily due to its energy demands (e.g. to recharge its internal batteries).

Phase 3: eWaste generation. At the end of a device's usable life, it is no longer used and becomes eWaste. Unfortunately, as we'll highlight below, the replacement cycle for devices has become more rapid over time. Users report replacing devices to keep current with the most advanced models, obtaining new manufacturer warranties, and supporting increasingly complex and resource-demanding applications and apps [5]. Manufacturers have encouraged short replacement lifecycles through planned obsolescence, subsidized replacement programs, and making it difficult or impossible to maintain and repair devices past their planned "end of life" and warranty period [5].

Phase 4: eWaste disposal. There are two ways to dispose of eWaste. The first approach includes official eWaste recycling programs which safely recycle and reclaim materials before disposing of the devices. This method is preferred and has a negative environmental impact. The second method includes sending devices to landfills, incinerators, unregulated reuse and reclamation channels, and other untracked disposal methods. This latter case greatly impacts the environment as the aforementioned materials used in manufacturing cause damage when eWaste is removed in this manner. Unfortunately, only 20% of eWaste is properly recycled, and the remaining 80% ends up in other channels [4].

2.2.3 Is recycling the answer? Recycling, by itself, is not able to fully address the scale and scope of these challenges. As mentioned above, only about 20% of eWaste is currently recycled. But even increasing that ratio to 100% would not solve the problem for two primary reasons. First, the internal components in modern computers are increasingly monolithic. For example, modern CPUs typically include not only compute cores, but external graphics support and GPUs, and even networking and wireless LAN support (e.g. the Atom x3). It simply isn't possible to recover the underlying elements from these chips and devices, in the same way that you can't unbake a cake to recover its underlying flour and sugar.

Even if components could be recovered, recycling by itself is not sufficient to address these problems. Research has shown that most of the impact of creating computing devices resides in the Phase 1 manufacturing step. Bakker et al. performed a lifecycle assessment (LCA) that evaluated laptops both in 1990 and in 2010, assuming they are used for only one year. They found that the total environmental impact of the manufacturing step rose from 68% in 1990 to 78% in 2010. However, the impact of the per-year use of the device fell from 31% in 1990 to 21% in 2010. Lastly, transportation accounted for a negligible percentage [4]. Despite improvements to laptops in other respects, Kasulatis et al. found that between 1999 to 2008, the impact of the manufacturing phase of laptops did not decrease at all [16].

Further, a metastudy by Suckling and Lee found that, similar to laptops, the manufacturing phase of smartphones represents the majority of the device's environmental impact [27]. In fact, they found that for a set of smartphones manufactured after 2010, the manufacturing phase accounted for an average of almost 75% of the total impact, with the use phase making up the majority of the remaining 25%.

2.2.4 How long should devices last? In recent years, use periods of devices have decreased, raising the impact of device manufacturing and disposal on the environment. Bakker et al. found that the Phase 2 (Use) period of laptops decreased from 4.3 years in 2000 to 4.1 years in 2005 [4]. A study in 2016 [22] found that the lifespan of smartphones varied from under 18 months to just under 2 years.

Given the environmental impacts described above, Bakker et al. suggest that the optimal replacement point of a laptop is at least after seven years after manufacturing, based on the increased operational efficiency of laptops during their use [4]. Further, Suckling and Lee suggest that smartphones need to be usable for approximately 5 years before impacts are amortized [27].

This information shows us that devices are used, on average, for approximately half as long as they would need to be in order to offset the costs of manufacturing and eWaste disposal have on the environment. Unfortunately, as we discuss below, there are significant challenges that prevent users from simply using their devices for longer periods of time.

2.2.5 Extending a computer's life. Given the oversized role that cost and manufacturing plays in a computer's impact, a key to reducing that impact is extending the usable life of computing devices. This idea is referred to as a *Circular Economy* and aims to keep devices in circulation to avoid them become eWaste and to reduce the ownership burden of the device across its lifespan. This concept is also called *Product Lifetime Extension (PLE).* Lifetime extension is sustainable, economical, good for the environment, and helps to address the "technology gap". As shown above, aiming to double the usable life of computing equipment would contribute greatly to solving economic and environmental challenges. Surveys of consumers find that they do wish that their devices did last longer, and are unsatisfied with their typical short lifespan [8, 9, 31]. If we are able to extend the usable lifespan of computers, we can address the eWaste problem while reducing the monetary and environmental cost of computing devices.

Why do devices get outdated? An inevitable scenario that forces a user to discard their older hardware is when the device becomes unable to readily access online resources. In this paper we focus on web browsers, given their importance, especially for students. Internet browsers are an extremely elaborate and rapidly-evolving software domain. Major browsers release updates regularly in order to improve security and performance for their users. Updates also include additional features that websites can use to present and render more complex content. This can come at the cost of removing compatibility for older browser versions. Since many websites deprecate support for older browsers, users have needed to continuously upgrade their devices to stay apace with modern websites.

At a higher level, mobile devices can become outdated at a hardware level if manufacturer of the device chooses to stop supporting it. Qualcomm officially supports updates for their mobile chipsets for only three years [1]. Apple does not support iPhones more than five years out of date [3]. The official "end of life" declaration from a smartphone manufacturer effectively eliminates any further usefulness it has for a typical consumer, who will not or cannot go out of their way to install alternative system software to extend the phone's usefulness.

Does Moore's Law help? The ending of Moore's law in the early 2000s fundamentally changed computing, and at first glance seems like it might make computers usable for longer, since CPU frequencies no longer increase at the rate they did prior. Unfortunately, available evidence shows the opposite, with the usable life of laptops decreasing from 2000 to 2005 [4]. Further, although CPU frequencies have not continued to increase, other resources such as memory, the number of CPU cores, flash storage capacity, and the prevalence of GPU units have provided computers with increased capabilities. Further, newer operating systems and device hardware offers new security primitives.

What about upgrades? What about repairing, upgrading, and extending computing devices in the field? In other arenas, repair is common, such as replacing a car's flat tire with a new tire instead of buying an entirely new car. Unfortunately, several trends in computing make repair and upgrade more infeasible. As previously mentioned, laptops, smartphones, and tablet computers rely on increasingly monolithic integrated subsystems and "systems on chip (SoC)" designs, which combine compute, graphics, storage, and even networking into a single chip that can only be replaced, not repaired. This is in contrast to pre-2000 era desktops, which consisted of a number of discrete components like sound cards, RAM modules, network cards, etc., which could all be independently upgraded or replaced. But even then, machines were often replaced rather than piecemeal upgraded over time. For these reasons, we largely rule out upgrading computing devices' hardware directly.

What about repairs? As mentioned above with respect to automobile tires, repairing damaged or defective components of very expensive equipment is commonplace. But for computers, it is unlikely that repair by itself will significantly address the needs we outline, since there is little evidence that durability is a significant factor in replacement decisions. The above-referenced studies show that the majority of devices are replaced either for non-technical reasons or to obtain new features/capabilities. When devices are eventually replaced, they are typically just as performant as the day they were manufactured. Further, as computers become more integrated, several sources of device failure are simply removed. For example, flash storage has largely replaced spinning hard drives, batteries are less likely to leak, and some laptops no longer rely on spinning fans for heat management. Smartphone screen repair remains popular, but nearly any other damage to a smartphone requires replacing the entire device.

2.3 Summary

The increasing sophistication and resource requirements of modern websites, especially in the education technology space, render consumer computing devices prematurely obsolete, resulting in significant contributions to the eWaste problem and to financial burdens for students. If we can extend the usable life of these devices, we will be able to reduce the financial and environmental impact of these devices.

3 BROWSER OBSOLESCENCE

To understand the ability of older devices to use modern websites, we now study their compatibility to various "eras" of hardware and software, with configurations representative of a given year. Each configuration uses a device, operating system, and browser version that was released in the appropriate year we wish to investigate. Comparison between years reveals the trends of obsolescence for different devices over time. To focus our study, we target representative websites used by students in an undergraduate computer science program.

The websites we consider are Google Drive, Canvas (a learning management system), StackOverflow, Jupyter Notebook, and Piazza (a message board), which are all platforms frequently used by computer science undergraduate students. We wished to test against Blackboard, the learning management system we analyze data from in Section 4, but it unfortunately has been decommissioned at the authors' campus.

We leverage the BrowserStack[7] online browser sandboxing platform to access each of these websites. BrowserStack provides a wide selection of smartphone hardware and desktop/laptop browser

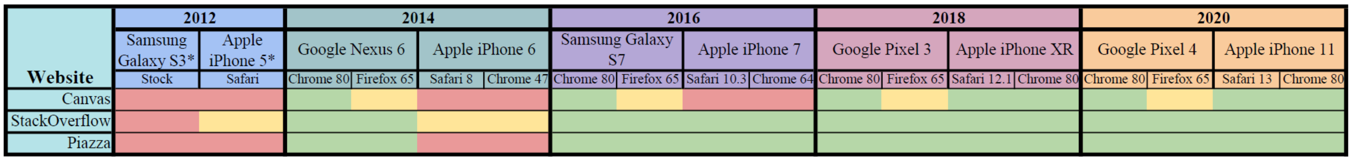

Figure 1: Mobile device/browser compatibility results via BrowserStack. '*' marks devices run via emulation. Browser versions used are the newest available for the given mobile OS.

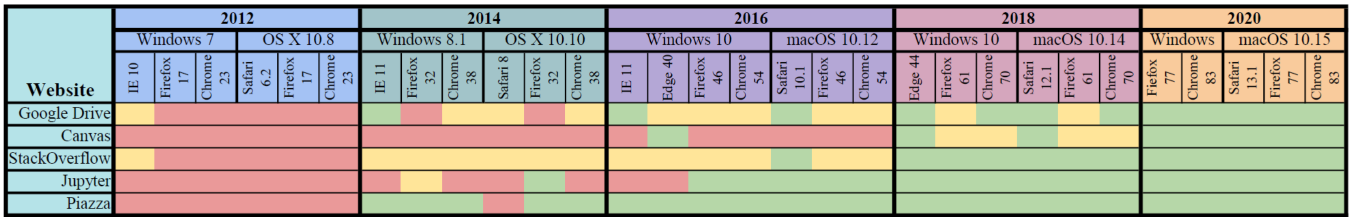

Figure 2: Emulated desktop browser compatibility results via BrowserStack.

versions to the user dating back to more than eight years. Smartphone browsers are not emulated, and are run on live devices. The client is required to select the specific device to use. We bin configurations into "eras" in time. Each era is represented by one of five fixed years, from 2012 to 2020, where we select device, operating system, and browser versions representative of what was up to date at the time. For each test, we complete a basic task on the website in question. We then determine the quality of the user experience, assigning it one of three outcomes: good (green), okay (yellow), and unacceptable (red). *Good* outcomes represent the website working as intended, *Okay* indicates minor issues such as incorrect CSS rendering or unsupported version error messages, and *Unacceptable* represents the website being unusable, due to intentional deprecation, HTTPS issues, or JavaScript errors.

3.1 Mobile Browsers

We evaluate both Android and iOS devices on BrowserStack, omitting apps for this study, though we acknowledge that apps do have a large user base (the Google Play store reports that Google Drive has 5 billion installations[13], and our results in Section 4 show many app users as well). The mobile results are shown in Figure 1.

We see that there are definite limitations placed on users who are using older smartphones. Devices from the 2012 era are unusable in almost all cases. In these cases, HTTPS errors occur on StackOverflow, and the other two websites fail to render. Perhaps the most shocking result is that an iPhone 6 and iPhone 7 cannot access Canvas on nearly up-to-date browsers: the main window is non-functional. Older Android devices receive much better compatibility, but we do begin to see some issues with Canvas.

3.2 Desktop Browsers

The results for our study of desktop browsers are shown in Figure 2. These results solely look at the incompatibility of browser versions, which helps us understand what failure modes occur in browsers.

Our results reveal a very apparent trend that older browser software often has difficulties supporting these websites, with 2012-era systems almost always failing. Even 2016-era software have significant issues, which is alarming given they are only four years outdated. Except for one case, every yellow cell is caused by either a CSS rendering error or a notification being given to the user that their browser is no longer supported. Google Drive and StackOverflow have frequent CSS issues that do not render the site unusable, but clearly impact the rendered appearance. Canvas gives an unsupported browser notification on every page that temporarily covers page elements for versions it deems are outdated. For 2012-era browsers, half of the errors are due to HTTPS negotiation failures; the browser software does not support the newer encryption schemes used by the site. For 2014-era and 2016-era browsers, cases of various JavaScript or rendering errors make websites unusable. These browsers are essentially inoperative across the spectrum of requirements a student may have if they wished to use these sets of websites in a class.

3.2.1 Browser Version Usage. It's clear from the above results that there are issues with using older browser versions on modern websites. However, it is unclear whether a user would be able to resolve these issues without needing to upgrade their hardware or needing an edge-cloud solution. One point of insight that we have been able to gather was to examine the browser version usage statistics gathered by StatCounter[25]. We analyzed usage data for the month of May 2020, and found that a non-trivial 8.6% percent of users are using browser versions released on or before 2016. We compare against this baseline when studying browser usage data of of real students in the following section.

4 STUDENT BROWSERS AND DEVICES

Given our baseline of how older browsers and the devices that run them operate on websites used for coursework, we next move to analysis of real user data. We analyze seven years worth of access logs to Blackboard, a learning management system (LMS) commonly used on the authors' university campus for both undergraduate and graduate courses. An LMS provides functionality for both instructors and students to post and review course material,

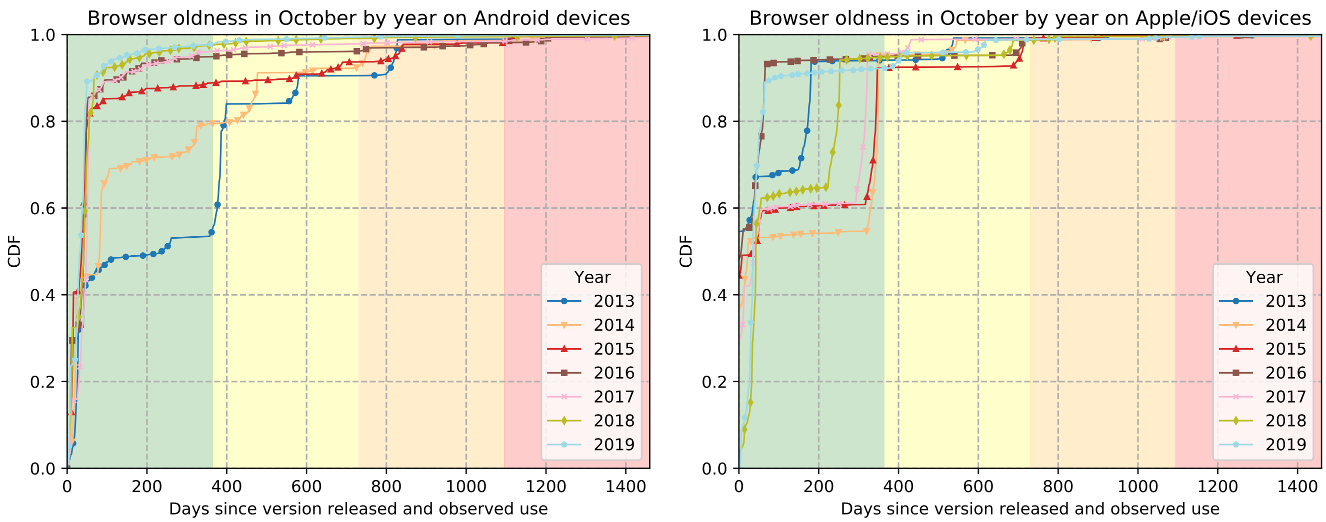

Figure 3: Browser age by year for Android (left) and iOS (right). Background colors denote yearly boundaries.

submit and grade assignments, and more. Instructors that choose to leverage an LMS require students to use a device in order to access material and assignments hosted there.

Log frequency of our LMS is non-uniform due to a varying degree of usage between these dates, beginning in the fall of 2013 and ending in the spring of 2020. However, we are still able to observe some definite trends that reveal the limitations of software and the devices that run them among the student population. Teaching and administration staff are included in our dataset, but are not the vast majority of users. Our dataset consists of over 165 million requests to the LMS web server. The use of this dataset was approved by the authors' campus for this study.

Every user in our dataset is first anonymized with a unique identifier. The information we base our study on are this identifier, access time, event type, and the user agent string given to the web server. Beginning in 2015 on iOS devices, and 2017 on Android devices, we begin to see user agent strings for the Blackboard mobile app, which includes a unique device identification string for a single app installation, and a more specific device model identifier. This allows for wider analysis of device upgrade rates among users.

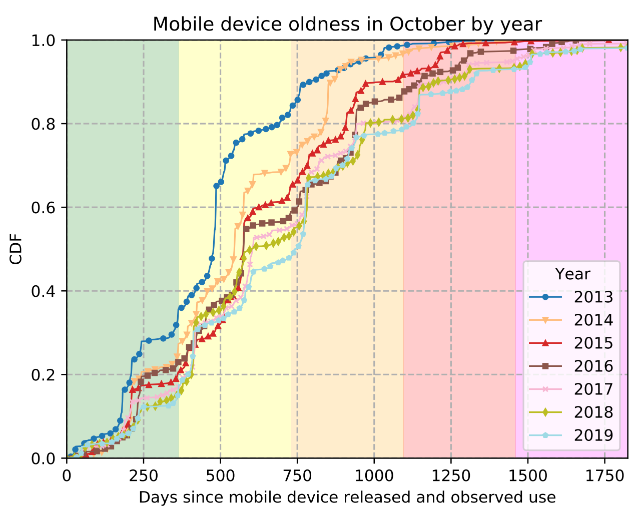

Figure 4: Device age at time of use across various years.

For our purposes, we filter event types in our dataset to only include successful login and logout events to capture the range of usage of a device, and to filter out devices attempting to use the LMS unsuccessfully. Because our dataset is based around user agent strings provided to a web server, we cannot claim 100% accuracy as users may modify their user agent however they see fit (the authors have directly observed a nonzero but negligible number of user agents that were modified by a user).

All data is processed by user agent parsing software, and then inspected by a variety of assertions manually crafted by the authors. Assertion failures trigger a fallback routine that sends the user agent to a large-scale subscription user agent parsing service in order to provide verification of results. Mobile app agent strings are manually parsed using a key-value scheme given as part of the agent string.

4.1 Browser usage

The most obvious and accurate information we can obtain from browser user agent strings is the versions of browsers themselves. Using the access date of the user and recording the browser and operating system information provided, we are able to obtain how dated a given access to the browser is for a given record. To more accurately survey our dataset, we group accesses by user and browser identifiers, and count only the latest entry from that user/browser pair. This means that we record the point where the browser was at its "oldest" and still used by the client to access the LMS.

We create an "oldness" metric representing the number of days between when the version string of the browser given was released

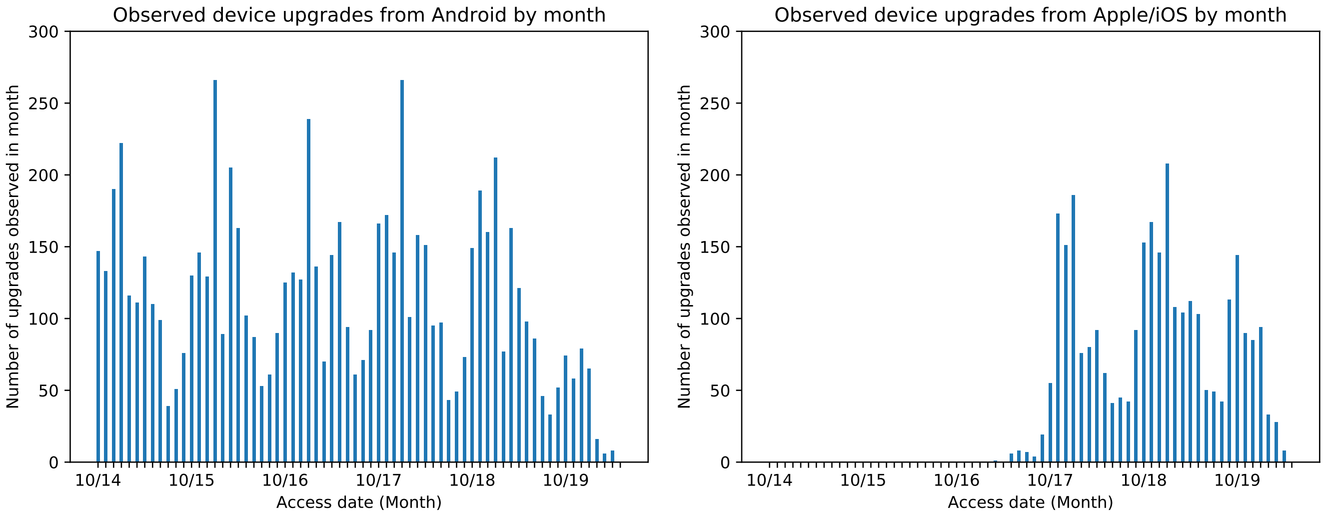

Figure 5: Device upgrades by month over time for Android (left) and iOS (right). iOS is limited to app installs, which began in 2017.

and when it was used by the client. We record this "oldness" metric over the period of each month in our dataset, i.e. each record indicates the oldest point that a browser was used by a given client within that one-month period. We plot this metric for a single month, October, over the different years of data that we have available. We select October as it is near the beginning of the academic year, and patterns across subsequent months within an academic year were similar.

The oldness data of browsers used to access the LMS is given in Figure 3. We observe that effectively nearly 99% of all users access the LMS with a browser that has been released within a four year window. Only mobile results are shown, but the results for desktop devices (Windows and Mac) are similar.

There is no observable trend of browser age by year, except on Android, where there is a clear trend of browser versions becoming newer over time. We do not investigate why this is the case, but recognize that there is an obvious push for Android users to run newer software over time that has proven successful.

From this data we can conclude that nearly all users in most cases tend to use a browser that was released well within the past year. However, there is a non-trivial number of users that wish to run older software on their devices. Despite this, they still all typically fall within a four-year window, a trend that follows from our previous section.

Comparing this to the browser data we studied previously in Section 3.2.1, there is a significantly lower percentage of students in our dataset that use a browser version four or more years out of date than on the internet at large. This may imply that our LMS system has upgrade requirements more aggressive than is typical for other web-based applications.

4.2 Device oldness

Our dataset also provides significant insight into the trends of smartphone and tablet models used by students to access the LMS. Android devices frequently provided device information directly present in the user age or had characteristics unique to an individual model. Additionally, agents given from the the LMS mobile app provided device model information as well as a unique "device identifier" generated at the time of app installation. This identifier provided additional insight as to when a user changed which device they leveraged to access the LMS.

To support this data, we used multiple services to acquire and store device release dates, brands, and full model names. Device naming schemes ultimately proved to be extremely ephemeral and volatile, particularly for less popular devices such as low-cost and rebranded aftermarket devices. Manual inspection, human data entry, and secondary/tertiary verification of the dataset proved invaluable to solving these challenges, but the authors still cannot claim 100% accuracy for every device present in the the LMS dataset.

We repeat our "oldness" metric where we track the last time the device was used by a client within a month's time. The results are shown in Figure 4. We use a similar presentation as in the browser data shown previously.

We once again see that the majority of devices used were released within the past four years from when it was used. There is a more significant tail of devices (<5%) from users that have smartphones up to five to six years years old. The devices in the long tail shown are both from Android and Apple users; there is no clear distinction between device family for the long tail of users.

There is an observable trend that devices seem to grow slightly older over time. However, we posit that this is a result of the number of devices used on the LMS also growing over time, as more and more students became required to use it in order to interact with coursework. The number of entries in our dataset between 2013 and 2017 grows by approximately 1.6x.

4.3 Device upgrades

The last statistic we gather from our dataset is understanding how users upgrade their devices. Because these devices are used to access an essential academic service, they represent a case where a user is *required* to acquire a newer device. Of course, users may (and as shown, do) upgrade their devices well before they are outdated.

The LMS user data we studied provides a significant amount of information, but does not give a direct signal for when a user upgrades their device. A user may simply use multiple devices, or get a tablet or other device they use "on the side" of their primary smartphone. We create a fixed set of criteria we use when detecting an upgrade. In order for a device to count as an "upgrade" for a user, the following criteria must apply:

- The new device was released within the past year, *or* was newer than the old device by at least two years.
- The new device was released at least six months after the release date of the old device.
- The new device was first used within 90 days of the time the old device was last used.
- The old device stopped being used within 90 days from the first time the new device was used.
- The new device was last used at least a month after the old device stopped being used.
- To remove duplicate upgrade events (e.g. a user purchases two new devices), additional devices added or removed within two months of the first upgrade event are not counted.

We do not use a rolling monthly window as in the previous two subsections, but instead use the first time the new device is seen as the "time" of the upgrade event. Each user's entire set of devices is gathered and computed against the above requirements in order to detect eligible upgrade events. Device identifiers are used where possible, and our duplication filtering prevents counting two separate devices in the case where a user both installs the mobile app and accessed the LMS via a traditional web browser.

The number of upgrade events detected is shown in Figure 5. Each graph shows the number of times an Android or Apple device was discarded in the monthly window. We break down upgrade events into separate Apple and Android graphs, as Apple devices hide their model identifier in web browser user agents, whereas the mobile app provides an exact identifier of the device. Therefore identifiers for Apple devices begin only in the summer of 2017.

Interestingly, the pattern for the two device families is different, with Apple upgrades occurring towards the end of the year and Android occurring somewhat after the academic year begins, but with a spike at the start of the calendar year. We do not have insights into whether this coincides with other events, such as the release of a new flagship smartphone. However, we wish to primarily note that we see a fairly consistent pattern of upgrades across all years across device types (with the exception of the, 2019-2020 academic year, which has far less events than the other years due to the introduction and use of a new learning management system).

Next, we break down each upgrade event into percentiles by the age of the old device when it was upgraded. This will allow us to understand how different types of users upgrade their devices- for example, users who upgrade their device within a year of it first releasing are perhaps more likely to purchase a brand new device, rather than a device that may be somewhat dated. We then plot the age of the device these users upgraded to as a CDF. We also include the raw thresholds of old devices ages.

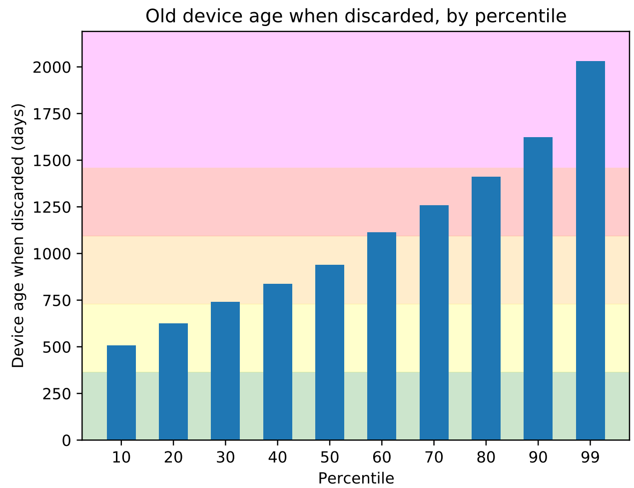

Figure 6: Ages of devices at time of upgrade, by percentile.

The percentile thresholds are shown in Figure 6. We immediately see that very few users actually upgrade their phone within a year-well under 10%. Additionally, there is a fairly linear growth between the most vigorous of upgrades and the least. Our results show that almost all users upgrade their phone once it is four years out of date, although once again we observe a decent tail end of users (<10%) hold onto their devices for roughly five years instead.

The upgrade newness CDF is shown in Figure 7. There is a clear trend that the users who discard their devices sooner are more likely to purchase a newer device. Note that the users who discard their phone within three years of its original release date all have roughly equivalent behavior when choosing a new device to purchase. It can be conjectured that the majority of users fall within a two to three year life cycle for their device, which is in line with information we presented in Section 2.

However, the users from the 60th percentile onward clearly tend to purchase older devices, likely from a second-hand market. We observe that these users often discard their device when it is four years old, and then purchase a device that is three years old to replace it. Because of our previous results on device ages showing that phones are discarded after four years, these users are likely purchasing hardware that they will only find useful for just over a year before discarding it. This is a frightening trend: these users are purchasing a second-hand device that is then likely discarded within a year.

This implies that the second-hand market is not sufficiently extending the lifetime of devices: if it was, we would hopefully see some non-trivial number of students still leveraging these older phones. The reasoning for this is not obvious, but we can conclude that these users either cannot or do not wish to purchase newer devices.

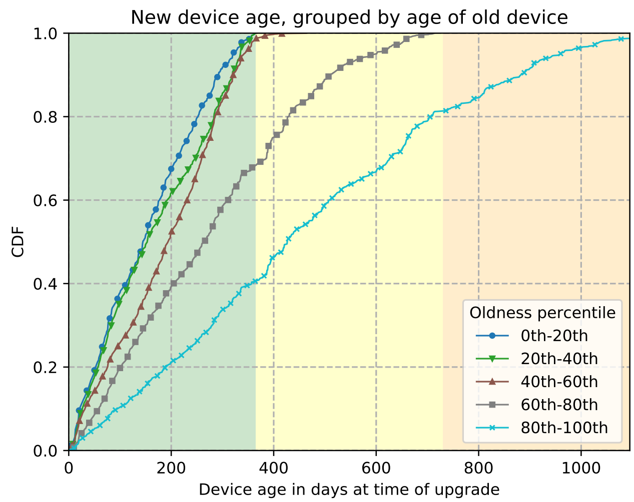

Figure 7: Age of new devices users upgraded to at time of first use, by oldness percentiles.

5 BEATING THE FOUR YEAR LIFE CYCLE

The common trend from all of our data is that roughly every four years, both software and hardware become obsolete. The need for extending device lifetimes to both reduce the financial burden on low income students and lower the total e-Waste generated suggests that this lifetime should be extended. However, it is not immediately obvious how to achieve this goal.

The cyclical nature of device manufacturing, use, and discard is a function of both consumer and producer. There are a vast variety of solutions that can increase the lifetime of computing devices, reaching far beyond the scope of our work here. It is difficult to say whether focusing on producers or consumers to extend device longevity would be more impactful. There are recent efforts [1] to develop longer-lasting devices from the producer side, and there are software solutions from consumers [17] to support devices past this four year duration.

For our purposes, we believe that a more universal and transparent approach that involves neither the consumer or the producer could have the most impact. In this section, we describe our proposed approach to extending the lifespan of consumer computing devices through cloud offload. We start by outlining the benefits of cloud-offloaded applications. We then discuss previous cloud-offloaded web browsers intended for other purposes, finishing by proposing the requirements and functionality of a possible offload architecture for browsers with the goal of enabling students to access the web with legacy devices.

5.1 Cloud offload for applications

Cloud computing at the edge for radio networks has been in development for several years [28], and with the recent advent of large scale 5G network deployments, it has become an attractive target for cloud-based applications that require low latency. Interactive applications that run on a user's device are a great target for cloud offload. However, the requirements of offloading are different from application to application.

A more blunt solution for an outdated device may be to run remote desktop to a more powerful machine. However, we think this is a poor solution. Remote desktop clients do not provide offline access to applications or files to the user, and do not take advantage of the client's hardware capabilities past simple video processing and I/O input. Legacy client devices still can easily run applications such as text processors, presentation software, etc. for periods longer than four years, making them still suitable for many student needs.

For students, the core target application for cloud offloading is a web browser. Academic software with high compute requirements is an interesting target, but solutions for offloading typically either already exist (e.g. remote compilation for programming projects), or are better suited to remote desktop environments even on the newest of student devices (e.g. complex engineering software such as Matlab).

5.2 Cloud-backed browsers

We continue to assume the environment of a student accessing educational resources remotely with a legacy device that is no longer supported. They need to access a set of educational websites and tools, as described in Section 3. In this context, any solution to achieve computer lifetime extension must meet the following requirements.

The user must be able to access modern websites and web-based applications. This includes the most recent versions of JavaScript, CSS-support, etc. While functionality is important, the user's experience must be performant, similar to the experience they would have received on modern hardware. Lastly, a typical student relies on a number of resources when doing school work, and so any solution must not consume an overly significant amount of resources on the target device.

5.2.1 Existing cloud-backed browsers. Development of a split-browser architecture has been explored for multiple decades [10, 11]. Recent efforts using virtual machines in cloud environments exist, but they primarily focus on providing client security rather than extending device lifetimes [23]. They do provide insight into how a modern split-browser architecture on the cloud may be achieved.

One example is the recently developed browser isolation system from Cloudflare [20], which uses a modern WebAssembly framework to provide a "remote browser" from the client to a cloud VM. The cloud VM performs all browser functions, and the client simply receives a network stream from the VM to interact with the rendered webpage. There is no offline functionality, which does not satisfy our requirements.

Other examples of recent cloud-backed browsers include Amazon's "Silk" browser [2] and Opera's "mini" browser [21]. Silk targets performance as a primary objective and is more in line with our requirements, but still requires browsers to run all of the end-result web content locally, which does not solve our compatibility issue. Opera Mini solely performs webpage compression to save on networking overheads.

5.2.2 A cloud browser for legacy devices. Extending previous efforts of split browser architectures is likely the best way forward to create a cloud browser that fits our requirements. Using a remote VM in the cloud to create a fully-featured browser frame that the client interacts with is a strong solution that offloads all security and compatibility requirements off of the client, and has the potential to be useful for a large range of legacy devices.

In order to meet the offline requirement that is absent in the Cloudflare solution, there is a rather large development effort needed in order to translate page content into a safe and offline viewable format for each client. Original split browser designs from Fox et al. [11] could be useful in order to "distill" webpage information into a proper format.

A significant challenge with legacy devices using modern browsers lies in memory utilization. Browsers now use gigabytes of memory when a large number of tabs are simultaneously opened, and legacy devices may not be able to support this. A distilled webpage may allow a separate type of process to view the content, which may save on memory. Additionally, pages that are not currently active when the user is online can be cached in the cloud VM, creating further savings in memory used.

5.2.3 The challenges of 5G networking. Our solution relies on the dual wins of 5G networking by providing students with access to a low-latency, high-bandwidth link to a nearby server in order to support a cloud-backed browser. However, 5G is a new and rapidly developing technology that has its own challenges to overcome to make this a reality. The promises of modern 5G networks have yet to come to fruition.

One concern is that the environmental impact of 5G networking will offset the benefits that green-focused proposals like ours receive from it. There is significant effort in making new 5G deployments focused on reducing environmental impact at the power and antenna level [14], which alleviates some of this concern. However, it remains to be seen if the local edge datacenter deployments for these networks will receive a similar focus on reducing carbon emissions.

Additionally, while 5G networks are likely to become available at large university campuses and their students, it is unknown to what extent availability will benefit those in more remote areas, and if students on remote learning platforms will have access.

We believe solutions targeting 5G networks are still beneficial despite these issues. However, it is imperative that additional solutions for device longevity will need to be explored and implemented in order to achieve a zero-carbon future.

5.3 Future work

The approaches and concerns listed above are certainly not exhaustive, and there is an extremely large range of possible solutions for elongating the lifetime of computers. We hope that future work can examine and implement many such ideas and introduce new systems that promote a circular economy and aid in reducing eWaste.

In this work we primarily target one limited, specific type of user. There has already been work targeting other types of applications, such as those targeting GPU-related applications for videogame platforms. We expect to see additional work also examine many more types of consumers around the world.

Lastly, while work can be done individually by producers and consumers in order to increase device lifetimes, we believe that there needs to be work on how to change societal patterns as a whole regarding smartphones. The increasingly rapid cycle of device replacement cannot solely be attributed to consumers' or producers' lack of ability to provide software or hardware solutions that increase device longevity. In order to create a holistic solution for a zero-carbon future, it will become necessary to create pressure via policy or market demand to change how the idea of "new" smartphones and other computing devices are viewed within the public consciousness.

6 CONCLUSION

The cycle of device obsolescence has created worrying trends for students, who have in recent years have been required to use technology to access educational resources. The data we have analyzed shows that students are not exempt from the increased rates of software and hardware obsolescence. We observe an overall trend that roughly every four years, a student is required to upgrade their hardware and/or software. While many students can afford to purchase new devices, many experience financial hardship and cannot easily do so. Additionally, the frequent manufacturing and discarding of new devices increases the amount of generated eWaste in the world.

The observed trends in the LMS data we analyze point to the need for a solution for students to access educational resources without the need of purchasing a new device. The class of users that wishes to upgrade their devices less regularly and not purchase a brand new device shows that there is a definite need for a system that extends the lifetime of devices within a teaching environment. The previous solutions for cloud-based browser offloading present an interesting solution space that we believe should be explored in order to provide sufficient extensions to product lifetimes for undergraduate students.

7 ACKNOWLEDGEMENTS

We would like to thank the UC San Diego Campus IT Research team for coordinating with us to provide the Blackboard access data for this study, including Cyd Burrows-Schilling, Harry Zhou, and Claire Mizumoto. This work is supported by the Department of Energy through grant ARPA-E DE-AR000084, and by the National Science Foundation through grants CNS-1564185, CNS-1553490, CSR-1629973, and CNS-1911104. We would like to thank our shepherd, Jay Chen, as well as the reviewers for their useful feedback.

REFERENCES

[1] R. Amadeo. Fairphone suggests qualcomm is the biggest barrier to long-term android support. https://arstechnica.com/gadgets/2021/03/the-fairphone-2-hits-five-years-of-updates-with-some-help-from-lineageos/, Mar 2021.

[2] Amazon. Amazon silk documentation. https://docs.aws.amazon.com/silk/index.html, 2020.

[3] Apple Inc. Supported iphone models. https://support.apple.com/guide/iphone/supported-iphone-models-iphe3fa5df43/ios, 2021.

[4] C. Bakker, C. Schuit, et al. The long view: Exploring product lifetime extension. https://www.oneplanetnetwork.org/resource/long-view-exploring-product-lifetime-extension, 2017.

[5] C. Baldé, V. Forti, V. Gray, R. Kuehr, and P. Stegmann. The Global E-waste Monitor–2017, United Nations University (UNU), International Telecommunication Union (ITU) & International Solid Waste Association (ISWA), Bonn/Geneva/Vienna. *ISBN Electronic Version*, pages 978–92, 2017.

[6] Blackboard Inc. Blackboard delivers worldwide growth. https://www.prnewswire.com/news-releases/blackboard-delivers-worldwide-growth-300398129.html, Jun 2018.

[7] BrowserStack. Browserstack. https://www.browserstack.com/, 2011.

[8] T. Cooper. Inadequate Life? Evidence of Consumer Attitudes to Product Obsolescence. *J. Consum Police*, 27:421–449, 2004.

[9] F. Echegaray. Consumers' reactions to product obsolescence in emerging markets: the case of Brazil. *Journal of Cleaner Production*, 134:191–203, 2016.

[10] A. Fox, I. Goldberg, S. D. Gribble, D. C. Lee, A. Polito, and E. A. Brewer. Experience with Top Gun Wingman: a proxy-based graphical web browser for the 3Com PalmPilot. In *Middleware 98*, pages 407–424. Springer, 1998.

[11] A. Fox, S. D. Gribble, E. A. Brewer, and E. Amir. Adapting to client variability via on-demand dynamic distillation. In *Proc. of the 7th ACM Inter. Conference on Architectural support for Programming Languages and Operating Systems*, 1996.

[12] S. Goldrick-Rab, J. Richardson, and A. Hernandez. Hungry and homeless in college: Results from a national study of basic needs insecurity in higher education. *Wisconson HOPE Lab*, 2017.

[13] Google. Google drive app. https://play.google.com/store/apps/details?id=com.google.android.apps.docs, 2020.

[14] C.-L. I, C. Rowell, S. Han, Z. Xu, G. Li, and Z. Pan. Toward green and soft: a 5G perspective. *IEEE Communications Magazine*, 52(2):66–73, 2014.

[15] Instructure Inc. Our story. https://www.instructure.com/about/our-story, 2021.

[16] B. V. Kasulaitis, C. W. Babbitt, R. Kahhat, E. Williams, and E. G. Ryen. Evolving materials, attributes, and functionality in consumer electronics: Case study of laptop computers. *Resources, conservation and recycling*, 100:1–10, 2015.

[17] LineageOS. https://wiki.lineageos.org/devices/, 2021.

[18] P. Mahaffey, D. Suvonnasupa, H. Weddle, and K. Hosch. Report on food and housing insecurity at uc san diego. *UC San Diego Basic Needs Insecurity Committee*, Jul 2016.

[19] L. Mishra, T. Gupta, and A. Shree. Online teaching-learning in higher education during lockdown period of covid-19 pandemic. *International Journal of Educational Research Open*, 1:100012, 2020.

[20] T. Obezuk. Introducing cloudflare browser isolation beta. https://blog.cloudflare.com/browser-beta/, Jan 2021.

[21] Opera. Opera mini. https://www.opera.com/mobile/mini.

[22] K. W. Panel. Kantar WorldPanel: Double Digit Smarphone Market Growth is over. https://www.kantarworldpanel.com/global/News/Double-Digit-Smartphone-Market-Growth-is-over.

[23] M. Patel. Gartner report on remote browser isolation: Menlo security. https://www.menlosecurity.com/blog/gartner-report-on-remote-browser-isolation-menlo-securitys-continued-validation, Jan 2020.

[24] S. Phantratanamongkol, F. Casalin, G. Pang, and J. Sanderson. The price-volume relationship for new and remanufactured smartphones. *International Journal of Production Economics*, 199:78–94, 2018.

[25] StatCounter. Statcounter global stats. https://gs.statcounter.com/.

[26] P. Suciu. Consumers balk at premium smartphone prices. https://www.technewsworld.com/story/85981.html, Apr 2019.

[27] J. Suckling and J. Lee. Redefining scope: the true environmental impact of smartphones? *The International Journal of Life Cycle Assessment*, 20(8):1181–1196, 2015.

[28] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella. On multi-access edge computing: A survey of the emerging 5g network edge cloud architecture and orchestration. *IEEE Communications Surveys Tutorials*, 19(3):1657–1681, 2017.

[29] P. UNEP, I. ITU, and U. UNIDO. A new circular vision for electronics time for a global reboot. https://www.weforum.org/reports/a-new-circular-vision-for-electronics-time-for-a-global-reboot, 2019.

[30] University of California, San Diego. UC San Diego campus profile. https://ucpa.ucsd.edu/campus-profile/, 2021.

[31] H. Wieser, N. Tröger, and R. Hübner. The consumers' desired and expected product lifetimes. *Product Lifetimes And The Environment*, 2015.