

this effort is the fundamental assumption that such content is in the form of *static content* such as a book or a musical score.

Today, an increasing fraction of the world's intellectual output is in the form of *executable content*. These include simulation models, tutoring systems, expert systems, data visualization tools, and so on. Even content that could be static (such as a company's product information Web page) is often dynamically generated through execution of code that customizes the content and appearance at runtime. What does it mean to preserve such content with perfect fidelity over time? There is no static object or collection of static objects that can serve as shared context to freeze and preserve. Rather, we need to *freeze and precisely reproduce the execution that dynamically produces the content*.

2. Execution Fidelity

Precise reproduction of software execution, which we call *execution fidelity*, is a complex problem in which many moving parts must all be perfectly aligned for a solution. Preserving this alignment over space and time is difficult. Many things can change: the hardware, the operating system, dynamically linked libraries, configuration and user preferences, specifications, geographic location, execution timing, and so on. Even a single change may hurt fidelity or completely break execution.

The need for execution fidelity is not unique to the digital library community. In the domain of commercial software, it is a problem faced by every software vendor. There is a strong desire to package software in a manner that installs and works with minimal tinkering so that the *time to value* for a customer is as short as possible. This desire to avoid extensive installation effort argues for rigid constraints on the software. At the same time, vendors wish to target the largest possible market. This argues for rich flexibility and configurability of software at installation time, to address the many unique needs and constraints of different market niches and diverse individual customers. Reconciling these contradictory goals is a difficult challenge.

Execution fidelity is also valuable when collaboration across space involves executable content. When help desk technicians assist remote customers, their ability to observe software behavior jointly in real time is likely to be more effective than having customers describe what they are seeing on their screens. Today, the most common class of tools for this purpose involves *remote desktop technology*. Examples include RDP in Windows [2] and VNC [3] in Linux, as well as Web-based services such as GoToMyPC [4].

Unfortunately, the available mechanisms for ensuring execution fidelity are weak. Most software distribution today takes the form of *install packages*, typically

in binary form but sometimes in source form. The act of installing a package involves checking for a wide range of dependencies, discovering missing components, and ensuring that the transitive closure of dependencies involving these components is addressed. Tools can simplify and partially automate these steps. However, the process still involves considerable skill and knowledge, remains failure-prone, and typically results in a long time-to-value metric for a customer.

These limitations directly affect collaboration across time. The install packages themselves are static content, and can be archived with existing mechanisms. However, the chances of successfully installing and executing this software in the distant future are low. In addition to the software installation challenges mentioned above, there is the additional difficulty that the passage of time makes hardware and software environments obsolete. The chances of finding compatible hardware and operating system on which to even attempt an install become vanishingly small over time scales of decades. These challenges have long stymied efforts by the digital library community to archive executable content [5–8].

For collaboration across space, today's use of remote desktop technology works reasonably well but faces fundamental challenges. For highly interactive applications such as games and data visualizations, the end-to-end network latency between collaborating parties limits the perception of perfect fidelity of execution. There is growing evidence that although Internet bandwidth will continue to improve over time, there are few economic incentives to keep end-to-end latency low. As discussed elsewhere [9, 10], this limits the applicability of the remote desktop approach at Internet scale to a fairly narrow class of applications.

Is there a better way to ensure execution fidelity? Ideally, such a mechanism would simultaneously address the challenges of software distribution as well as collaboration with executable content over space and time. We present such a mechanism in the next section.

3. Virtual Machine Technology

A *virtual machine (VM)* is a computer architecture and instruction set emulator of such high accuracy that neither applications nor the operating system is able to detect its presence. In other words, emulated




Figure 1. Basic VM Concepts

execution is indistinguishable from execution on genuine hardware. Since the emulated hardware is usually of the same type as the bare hardware on which emulation is performed, we are effectively cloning a physical machine into multiple VMs.

Figure 1 illustrates the key abstractions in the realization of a VM. The *guest* layer represents an unmodified legacy computing environment, including one or more applications and the operating system. The *host* layer is responsible for virtualizing the bare hardware, including devices such as disk, network, display and so on. The host layer multiplexes the concurrent execution of multiple VMs, and ensures isolation. Each VM has execution context represented in a *VM monitor (VMM)*, as shown in Figure 1.

VMs were invented by IBM in the mid-1960s as a timesharing approach that enabled concurrent, cleanly-isolated system software development on a single mainframe by multiple programmers. Since mainframe hardware of that era was expensive, VMs were a cost-effective approach to enhancing programmer productivity by providing a private “mainframe” for each developer rather than reserving dedicated time on real hardware. Accuracy of hardware emulation was paramount because the system software developed on a VM was intended for use in a mainframe operating system. That software had to work with negligible changes on the real hardware. The techniques for efficient and accurate hardware virtualization that were developed in this era have proved to be of lasting value. The accuracy of virtualization is a key reason for VMs being a potential solution to the problem of execution fidelity today.

By the early 1970s, VMs were being used for purposes well beyond those that motivated their invention. A commercial product, the VM/370 timesharing system [11], was highly successful and was in widespread use for well over a decade. Its descendants continue to be in use on IBM mainframes as z/VM today [12]. However, the emergence of the personal computer led to the eclipse of the VM as a computing abstraction by the late 1980s. Cheap PCs with a rigidly-controlled operating system (MS-DOS initially, and various versions of Windows later) displaced VMs as the most cost-effective way to deliver computing with high execution fidelity. Events have now come full circle. The emergence of cloud computing has restored lustre to the VM abstraction.² Large commercial investments are now being made in VM technology and VM-based systems.

²The VMs in cloud computing are typically based on the Intel x86 architecture rather than the IBM 370 architecture, but that is a minor detail relative to the big picture.




Figure 2. Legacy-compatible VM Ecosystem

Figure 2 illustrates the value proposition of the VM abstraction today. A large legacy world of software, including operating system software, represents a substantial intellectual and financial investment that has been already been made on existing computer hardware. This legacy world can be completely closed-source: there is no requirement for availability of source code, nor a requirement for recompilation or relinking. All that is needed is standard software installation, just as on real hardware. Beneath the implementation of the emulated hardware interface, there can be extensive innovation along many dimensions. This innovation is totally isolated from the idiosyncrasies of the legacy world above. In cloud computing, the freedom to innovate beneath the hardware interface allows the creation of valuable system management functionality such as server consolidation for energy savings and elastic computing for bursty server loads. More generally, it allows an enterprise to virtualize and outsource its entire data center to a remote third party such as Amazon Web Services [13].

Certain specific attributes of the VM abstraction accounts for its longevity and success. The hourglass shape in Figure 2 could depict the interface specification of any abstraction (thin waist), applications dependent on this interface (above), and implementations of this interface (below). More specifically, it could represent an execution engine such as the Java Virtual Machine (JVM) [14] or the Dalvik Virtual Machine for the Android platform [15]. While these alternatives (collectively referred to as *software virtualization*) have successful niches, they do not approach the VM abstraction (also referred to as *hardware virtualization*) in terms of longevity, widespread usage and real-world impact.

Why is hardware virtualization so much more successful? First, the interface presented by the VM abstraction is compatible with legacy operating systems and their valuable ecosystems of applications. The ability to sustain these ecosystems without code modifications is a powerful advantage of VMs. The ecosystems supported by software virtualization tend to be much smaller. For example, a JVM is only valuable in supporting applications written in specific languages

such as Java. In contrast, a VM is language-agnostic and OS-agnostic. In fact, a JVM can be part of the ecosystem supported by a VM. Hardware virtualization can thus subsume software virtualization. Second, a VM interface is *narrow* and *stable* relative to typical software interfaces. In combination, these two attributes ensure adequate return on investments in the layer below. By definition, a narrow interface imposes less constraints on the layer below and thus provides greater freedom for innovation. The stability of a VM interface arises from the fact that the hardware it emulates evolves very slowly and almost always in an upward-compatible manner. In contrast, the mutability of software results in more rapid evolution and obsolescence of interfaces. Keeping up with the changes requires high maintenance effort. Mutability also leads to widening of narrow interfaces over time because it is difficult to resist the temptation to extend an interface for the benefit of a key application. Over time, the burden of sustaining a wide interface constrains innovation below the interface.

The importance of the narrowness and stability of an interface can be seen in the contrasting fortunes of *process migration* and *VM migration*, which are essentially the same concept applied at different levels of abstraction. Process migration is an operating system capability that allows a running process to be paused, relocated to another machine, and continued there. It has been a research focus of many experimental operating systems built in the past 20 years. Examples include Demos [16], V [17], Mach [18], Sprite [19], Charlotte [20], and Condor [21]. These independent validation efforts have shown beyond reasonable doubt that process migration can indeed be implemented with acceptable efficiency. Yet, in spite of its research popularity, no operating system in widespread use today (proprietary or open source) supports process migration as a standard facility. The reason for this paradox is that a typical implementation of process migration involves such a wide interface that it is easily rendered incompatible by a modest external change such as an operating system upgrade. Long-term maintenance of machines with support for process migration involves too much effort relative to the benefits it provides. The same concept of pausing an executing entity, relocating it, and resuming execution can be applied to an entire VM rather than a single process. This is *VM migration*, a capability that is in widespread use in cloud computing systems today. Guest OS changes do not affect an implementation of VM migration. This insulation from change, embodied in the narrow and stable VM interface, is crucial to the real-world success of VM migration.

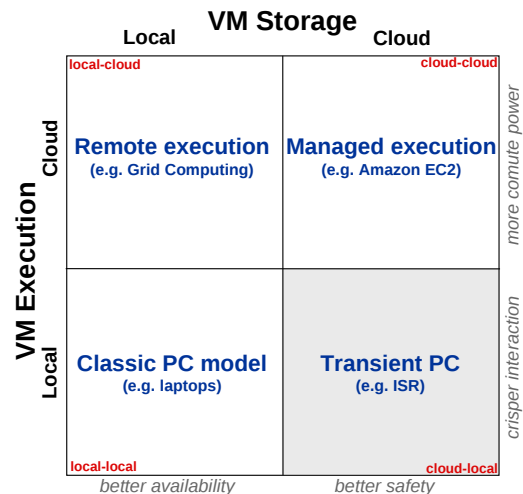


Figure 3. VM-based Cloud Computing

4. Transporting VMs across Space and Time

Modern implementations of the VM abstraction provide support for an executing VM to be *suspended*. This is conceptually similar to what happens when you close the lid of a laptop computer. An exact record of the execution state of the VM at suspend is captured and saved as files in the host file system. There are typically two distinct representations of the VM's persistent state (i.e., its virtual disks) and its volatile state (i.e. memory, registers, etc.). These are referred to as a *disk image* and a *memory image* respectively, and their combination is called a *VM image*. Execution can be *resumed* from a VM image. This is equivalent to opening the lid of a suspended laptop, and seamlessly continuing work. If there is no memory image, resuming a VM is equivalent to powering up hardware, resulting in a fresh boot of the guest OS. An executable VM that is created from a VM image is called a *VM instance*. Some use cases may restrict the number of concurrent VM instances that can be created from a single VM image.

A VM image is static content that can be transmitted across the Internet for sharing, or archived for later retrieval. If you have a VM image, resuming the suspended VM only requires a compatible VMM and host hardware. Everything else needed for perfect execution fidelity is contained within the VM image. The use of VMs thus transforms the difficult problem of collaborating with executable content into the more tractable problem of efficient transmission, storage and retrieval of VM images.

Separating image storage from instance execution leads to a rich design space for VM-based cloud computing. As Figure 3 shows, this design space has two dimensions. One dimension is the location of the definitive copy of the VM image (ignoring cache copies): local to the execution site, or in the cloud. The other dimension is where a VM instance executes: locally (at

the edges of the Internet) or within the cloud. This results in the four quadrants depicted in Figure 3:

- The local-local quadrant is equivalent to the classic unvirtualized model of desktops and laptops today.
- The cloud-cloud quadrant corresponds to situations where both VM image storage and execution occur within the cloud. This is the model supported by public clouds such as Amazon Web Services, and is often what is meant by the term “cloud computing.”
- The local-cloud quadrant corresponds to situations where a local VM image is instantiated in a remote cloud, typically for use of more powerful computing resources or for proximity to large data sets. This is essentially the metaphor of *grid computing*. Use cases in which large data sets are shared by a community of researchers using this mechanism can be viewed as a form of collaboration across space.
- The cloud-local quadrant is of particular interest to the topic of this paper and is discussed in detail below. This quadrant corresponds to a VM image archived in the cloud being instantiated and executed at an edge node of the Internet.³ We label this quadrant *Transient PC* because it evokes the vision of a transient or disposable computing instance that is brought into existence anywhere on the Internet, used and then discarded. If its final state is valuable, that could be archived as a new image in the cloud.

The Transient PC model is well suited to supporting collaboration across time. Just as we retrieve a pdf file today and view it on a desktop, laptop or smartphone, we can envision a VM corresponding to executable content being dynamically instantiated and executed on an edge device. There are many details to get right, but in principle this approach offers better execution fidelity than any of the alternatives available today. This is because the VM contains the entire transitive closure of dependencies (including the operating system and dynamically linked libraries) along with the application of interest.

With additional machinery, the Transient PC model could also support near real-time collaboration across space between multiple users. This would consist of a VM instance executing at each user’s site, along with mechanisms to propagate changes resulting from each user’s interactions to the other VM instances. In

effect, a collection of VM instances dispersed across the edges of the Internet are kept in sync as they interact with their users. This achieves the effect of the remote desktop approach mentioned in Section 2, but with much better execution fidelity since it is less sensitive to wide-area Internet latency. There is an extensive body of research on logging and replay of VMs that could be leveraged for this purpose [22–24]. In addition, there would need to be an appropriate model of concurrency control to ensure that the divergence of VM instances remains within acceptable bounds. The simplest concurrency control approach would be a token-passing scheme that allows at most one user at a time to be the “writer.” More sophisticated models of collaboration will require more complex concurrency control approaches. For example, multi-user graphics-intensive computer gaming at Internet scale may benefit from this approach if game-specific concurrency control can be efficiently supported.

The superior execution fidelity of VMs comes at a price. A typical VM image is many gigabytes in size, possibly tens of gigabytes. In contrast, a typical pdf document is tens of kilobytes to a few megabytes in size. While storage capacity in the cloud for VM images is of some concern, a more serious concern is their efficient transmission over the Internet. Even at 100 Mbps, full transfer of a 10 GB image will take at least 800 seconds (over 13 minutes). Over the wide-area Internet, with end-to-end bandwidths approaching 10 Mbps at well-connected sites, the transfer will take over 8000 seconds (over two hours). These are clearly unacceptably long periods for a user to wait before executable content can be viewed. Ideally, the startup delay experienced by a user at a well-connected Internet site should be no more than a few seconds to tens of seconds. While Internet bandwidth will continue to improve over time, it is also likely that VM images will grow larger.

The obvious solution is to adopt the approach of *streaming* used by video sites such as YouTube. This would allow a user to begin viewing content as soon as a modest prefix has been transferred. Unfortunately, the problem is more complex than video streaming because a VM instance is unlikely to access its VM image in simple linear order. It is the complex runtime behavior of a VM instance (including user interactions, data dependencies, temporal locality and spatial locality) that determines the reference pattern to its VM image. To address this problem, we have shown how *demand-paged execution* of a VM instance can be implemented. Using this approach resume latencies of tens of seconds are achievable today at well-connected Internet sites. This work has been done in the context of the *Internet Suspend/Resume (ISR)* system [25–27]. Recently, Abe et al [28] have shown that history-based prefetching can lower the resume latency of VM-based systems on last-mile networks such as 4G wireless networks.

³The private cloud of an enterprise can be viewed as “an edge node of the Internet” for this discussion.



Figure 4. Olive as a Concept at Multiple Levels

5. The Olive Vision

Our initial exploration with ISR suggested that VMs can indeed archive executable content. To realize the full potential of this capability we envision *Olive*, a VM-based Internet ecosystem for executable content. At the heart of Olive is open-source software for contributing VM images, curating them in Internet repositories so that licensing and intellectual property constraints are enforced, and efficiently retrieving and executing them at a cloud provider or at an edge device (i.e., the cloud-cloud and cloud-local quadrants of Figure 3).

As shown in Figure 4, Olive is a concept that applies at multiple levels of the Internet. At the global or planetary level, we envision a relatively small number of archival sites (less than 100, more likely 1–10) where VM images are curated. Olive also has a presence at the level of a public cloud provider such as Amazon Web Services, or within the private cloud of an enterprise (most likely totaling a few hundred to a few thousand clouds). At this level, VM images may be cached in their entirety from global Olive repositories. They may also have customizations that are cloud-specific or enterprise-specific. Finally, Olive’s presence at any edge device on the Internet enables ubiquitous streamed execution of VM images over wired and wireless networks. For scalability, VM images may be streamed to an edge device from a cloud provider rather than directly from a global Olive repository.

The ability to customize VM images within a cloud is an important Olive capability because it helps VM image contributors to honor software licensing constraints. Executable software content created by one person almost always depends on other software content. Take for example someone who wants to contribute the popular computer game Pac-Man to a

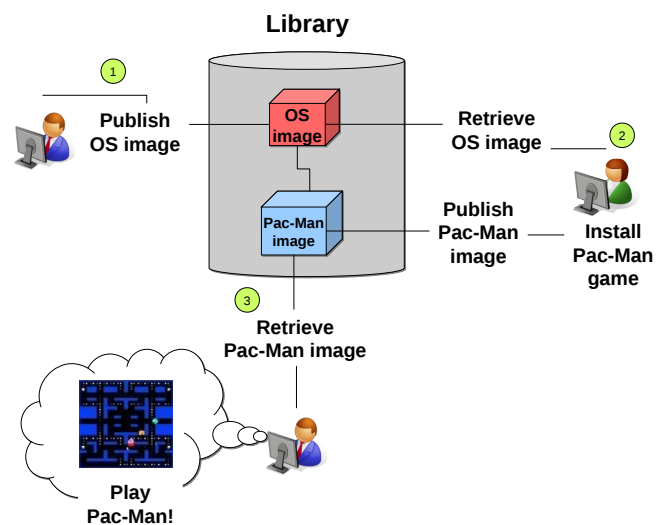


Figure 5. Extending an Olive VM Image

global Olive repository. This person may have the rights to contribute the Pac-Man application, but not the rights to the operating system and other software that must exist on the VM inside which Pac-Man has to run. If the contribution to Olive included those other software components, it would require negotiation of re-distribution rights on all those other components with their corresponding owners. This is not an issue with a totally open source software stack, but it can be a laborious and expensive task when proprietary software is involved. This legal hurdle and its associated litigation and business risks could discourage many from making valuable Olive contributions.

Olive could solve this problem by allowing one VM to extend another VM, thereby creating a new VM. Figure 5 illustrates one way this can work. In step 1, a person or organization that wants to contribute an operating system to Olive, does so as a VM that includes a fully installed and tested image of that operating system. A different person or organization can then retrieve this VM, start it on their own laptop, and install the Pac-Man application. The resulting VM, which includes both the operating system and the Pac-Man application, can then be published back into Olive. Olive technology will recognize that this new VM is an extension of a virtual machine that already exists in its library. It will automatically compute the delta content between the two (which is the contribution from the Pac-Man application), and store it internally. Thus, the Pac-Man contributor only contributes bits that he owns to Olive and does not violate any licensing restrictions. This process is reversed during the customization step within a cloud: the Pac-Man delta is applied to the original VM to produce a merged VM. In this context, the *dynamic VM synthesis* approach explored by Ha et al [29] is important related work.

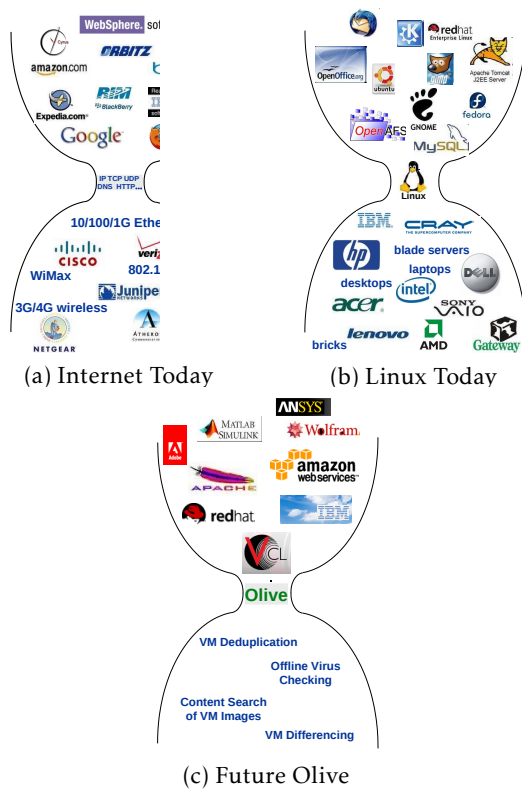


Figure 6. The Power of Open Ecosystems

We envision Olive to be an ecosystem built around open standards and interfaces rather than a closed, proprietary system. The vibrant Internet and Linux ecosystems shown in Figures 6(a) and 6(b) are testimonials to the power of open standards. In both cases, a relatively small collection of open components at the waistline sustains large business investments above and below that waistline. Figures 6(c) illustrates the hypothetical Olive ecosystem of the future. At the waistline are a small set of open components such as ISR. Above are publishers of open or proprietary executable content. Below are open or proprietary tools and mechanisms. We believe that such an open ecosystem has the potential to be far more influential and viable than a proprietary ecosystem. This is especially true for an archival system that has to survive in sound operating condition over time periods of decades and (optimistically) centuries.

6. Olive and Digital Libraries

Historically across the disciplines, academic libraries have maintained the scholarly record, which fixes credit for innovation, invention, and progress. For the stone tablet, papyrus, vellum, and paper environments, successful retention has been challenging but attainable. The realm of static digital resources has proven more challenging [30], but standards and best practices have led several libraries through successful migrations of large corpuses. Libraries are also the logical stewards

for executable content as for other scholarly content. Academic libraries have universally accepted their faculty's research and publication products, which may range from cooking, to chemistry, to computing. The last now complain that they cannot fully appreciate and understand work if it is just represented by object code divorced from source code. Computer scientists need to see what a program does when it runs to be fully informed and enriched by their tradition. As Olive is realized, libraries, working with their computing colleagues, will be able to commit to the long term preservation of diverse forms of executable content. Examples include tutoring systems (such as Loop Tutor, Hypertutor, and iCarnegie), games (such as Rogue, MUDs, Doom, and Duke Nukem), databases (such as Clueset), and obsolete versions of major commercial operating systems and applications.

Trying to maintain the equipment to run these older works is not an elegant preservation solution. The Joint Information Systems Committee (JISC) in "Creative Archiving at Michigan and Leeds: Emulating the Old on the New(CAMiLEON)" looked at the feasibility and costs of using emulation for software[31]. In 2011, The Baden-Wurttemberg Functional Longterm Archiving and Access project, building on earlier work done in the European Union, announced a two year state-sponsored project to develop scalable cloud-based emulation as a service [32]. The Internet Archive is currently using JSMESS to present early productivity software, such as VisiCalc and Wordstar, on its web site while archivists at Stanford University are experimenting with Xenserver and Xenclient. None of these alternatives meets our needs for performance and long term sustainability by an open source community.

Managing the metadata around the items in the collection offers additional challenges. Dr. Jerome McDonough from the Preserving Virtual Worlds project manages metadata for the Olive project. McDonough employed OAI-ORE and METS while examining further work, such as the Carolina Digital Repository's Curator Workbench tool. We are also considering other approaches, such as Trustworthy Online Technical Environment Metadata (TOTEM) [33] and XSEAD, an NSF-supported online platform supporting networks of creativity and innovation across science, engineering, art and design [34]. We will be doing additional research in this area.

Many of the first choices of works to be archived may, of necessity, be those with less commercial value. The JSTOR database has demonstrated that rights owners can be persuaded to give permission for the use of less valuable content (journal articles over five years old). Nevertheless, intellectual property restrictions will be a significant challenge.

Another significant challenge will be sustaining the archive for academic products long term. Many models




Figure 7. Microsoft Office 6.0 on Windows 3.1

exist: JSTOR and HathiTrust (subscription), the dark archives Portico and LOCKSS (membership), and NIH's PubMedCentral and Bibliotheque Nationale de France's DeepArc (government funding). The Million Books Project model (multinational governments, subscription, and university computer science department and library support) has proven to be less effective [35]. The Google Books Library partnership, a commercial venture, is perhaps the most successful in the breadth of its accomplishments.

7. Olive Status

With funding from the Sloan Foundation and the Institute for Museum and Library Services, Carnegie Mellon University has created an initial proof of concept implementation of Olive. Because of software licensing constraints, this system is currently available only to small number of research collaborators. In the near future, we hope to obtain the necessary legal permissions to allow use of Olive content by a broader user community. Olive contains over 15 VMs today, and continues to grow. A summary of the collection is shown at the end of the paper (Figure 11). It includes operating systems and applications dating back to the late 1980s. For brevity, we describe only four of these VMs below.

Microsoft Office 6.0. Figure 7 shows a screenshot of this VM, containing Word, Excel and PowerPoint for Windows 3.1 from the 1993-94 time period.

Great American History Machine. This application was originally created in the late 1980s by Professor David Miller of Carnegie Mellon University. It was used by him and by many professors at other universities nationwide to teach 19th century and early 20th century American History. As the screenshot in Figure 8 shows, this educational software used census and election data to teach students important historical

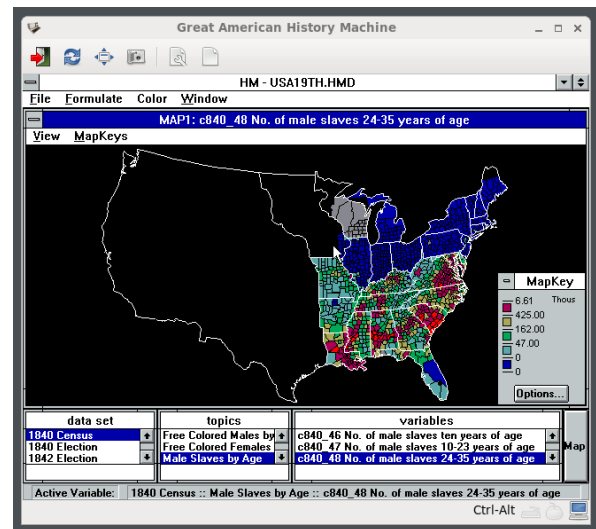


Figure 8. Great American History Machine on Windows 3.1

concepts such as the origins of the Civil War. The Windows 3.1 version of this software was created in collaboration with the University of Maryland. Because of lack of financial resources to port the software to newer Windows platforms, it fell into disuse over time. No modern equivalent of this software exists today.

TurboTax 1997. This application for Windows 3.1 and Windows 95 was used by millions of Americans to prepare their 1997 tax returns. Figure 9 shows a screenshot of this application. Since TurboTax is updated each year to reflect the current tax laws, a suite of TurboTax VMs from consecutive years can offer unique historical value. Imagine a class in political science, public policy or economics assigning students a project based on TurboTax versions that are ten years apart. By calculating the tax returns for hypothetical families with different income profiles, students can see for themselves the impact of tax code changes over time. Such active learning can transform the abstract topic of tax law into a source of valuable real-world insights.

NCSA Mosaic. As the world's first widely-used web browser dating back to 1992-93, Mosaic has a unique historical status. This VM, whose screenshot is shown in Figure 10, is also interesting for a second reason. The version of Mosaic that it encapsulates was written for the Apple MacOS 7.5 operating system on Motorola 68040 hardware. The VM also encapsulates Basilisk II, an open source hardware emulator for Motorola 68040 on modern Intel x86 hardware running Linux. The bootable disk image of MacOS 7.5 with Mosaic is stored as a file in the virtual file system of the outer Linux guest. In spite of two levels of virtualization, performance is acceptable because modern hardware is so much faster than the original Apple hardware. Pointing the Mosaic browser at modern web sites is instructive. Since Mosaic predates web technologies such as Javascript, Cascading Style Sheets, and HTML5,



Figure 9. TurboTax 1997 on Windows 3.1

this old browser is unable to render content from modern web sites. It can, however, render static web pages from some older sites that are still on the Internet.

8. Conclusion

Rich shared context facilitates human collaboration. Face to face interaction is the gold standard for shared context, since it supports the full range of visual, aural and tactile real-time interactions involved in conveying procedural and declarative knowledge. Even in the 21st century, people travel to meetings, workshops and conferences for this highest quality of shared context.

Extending collaboration across space and time to approach this gold standard has been a human quest since the dawn of civilization. The invention of writing, and the emergence of libraries profoundly changed collaboration across time. Over centuries, this quest has led to other transformative inventions such as the printing press, telephony, the motion picture, television, the Internet, and social networks such as Facebook. Yet, in spite of these impressive advances, the shared context supported by them has always been static in nature. Today's Flickr images and YouTube videos may only bear a distant relationship to writing on clay tablets, but they both represent static content from the viewpoint of archiving. Sharing executable content has been theoretically possible since the invention of the computer. But only now have the foundational technologies for this capability matured to the point where we can build upon them.

Procedural artifacts ranging from scientific simulation models to interactive games play an increasingly important role in our lives. The ability to archive these artifacts for posterity would be an important transformative step. Imagine being able to reach back across time to execute the simulation model of a long-dead scientist on new data that you have just acquired. What do the results suggest? Would they have changed the




Figure 10. Mosaic Browser on MacOS 7.5

conclusions of that scientist? Although you aren't quite bringing the scientist back to life, you are collaborating with that person in a way that was not possible until now. As we have shown in this paper, there is now a convergence of the key technologies needed for such scenarios: VM technology, cloud computing and high bandwidth wide-area networks. From these roots, the Olive vision emerges as a natural next step for supporting collaborations across space and time.

Acknowledgements. We acknowledge the extensive contributions of Benjamin Gilbert, who has been the primary contributor to ISR since 2006. Earlier contributors to ISR include Michael Kozuch, Casey Helfrich, Dave O'Hallaron, Partho Nath, and Matt Toups. We also acknowledge the many people who contributed to the Olive vision presented here: Benjamin Gilbert, Jan Harkes, Glenn Ammons, Catherine Zhang, Mladen Vouk, Peng Ning and many other contributors at Carnegie Mellon, IBM, and North Carolina State University, as well as the participants of the Olive workshop held at IBM in April 2011. We thank Dan Ryan for his contributions to this paper, and David Baisley for providing us with the original install disks for the operating systems shown in Figure 7 through Figure 10. The creation of ISR was supported by the National Science Foundation (NSF) under grant number CNS-0509004, an IBM Open Collaborative Research grant, and Intel. The creation of Olive was supported by the Sloan foundation and the Institute for Museum and Library Services. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily represent the views of their employers or the funding agencies. Internet Suspend/Resume and OpenISR are registered trademarks of Carnegie Mellon University.

References

- [1] Oxford Dictionaries, <http://oxforddictionaries.com>. Accessed on May 2, 2014.
- [2] Remote Desktop Protocol, http://en.wikipedia.org/wiki/Remote_Desktop_Protocol. Accessed on May 2, 2014.

- [3] RICHARDSON, T., STAFFORD-FRASER, Q., WOOD, K.R. and HOPPER, A. (1998) Virtual Network Computing. *IEEE Internet Computing* 2(1).
- [4] GoToMyPC: Access Your Mac or PC from Anywhere, http://www.gotomypc.com/remote_access/remote_access. Accessed on May 2, 2014.
- [5] CONWAY, P. (1996), Preservation in the Digital World, <http://www.clir.org/pubs/reports/conway2/>. Accessed on May 2, 2014.
- [6] CONWAY, P. (2010) Preservation in the Age of Google: Digitization, Digital Preservation, and Dilemmas. *Library Quarterly* 80(1).
- [7] HEDSTROM, M., LEE, C., OLSON, J. and LAMPE, C. (2006) 'The Old Version Flickers More': Digital Preservation from the User's Perspective. *The American Archivist* 69: 159–187.
- [8] MATTHEWS, B., SHAON, A., BICARREGUIL, J. and JONES, C. (2010) A Framework for Software Preservation. *The International Journal of Digital Curation* 5(1).
- [9] TOLIA, N., ANDERSEN, D. and SATYANARAYANAN, M. (2006) Quantifying Interactive Experience on Thin Clients. *IEEE Computer* 39(3).
- [10] SATYANARAYANAN, M., BAHL, V., CACERES, R. and DAVIES, N. (2009) The Case for VM-based Cloudlets in Mobile Computing. *IEEE Pervasive Computing* 8(4).
- [11] CREASY, R. (1981) The Origin of the VM/370 Timesharing System. *IBM Journal of Research and Development* 25(5).
- [12] VM (operating system), [http://en.wikipedia.org/wiki/VM\(operatingsystem\)](http://en.wikipedia.org/wiki/VM(operatingsystem)). Accessed on May 2, 2014.
- [13] Amazon Web Services, <http://aws.amazon.com/>. Accessed on May 2, 2014.
- [14] LINDHOLM, T. and YELLIN, F. (1999) *The Java Virtual Machine Specification (2nd Edition)* (Prentice Hall).
- [15] Dalvik (software), http://en.wikipedia.org/wiki/Dalvik_%28software%29. Accessed on May 2, 2014.
- [16] POWELL, M. and MILLER, B. (1983) Process Migration in DEMOS/MP. In *Proceedings of the 9th ACM Symposium on Operating Systems Principles* (Bretton Woods, NH).
- [17] THEIMER, M., LANTZ, K. and CHERITON, D. (1985) Preemptable Remote Execution Facilities for the V-System. In *Proceedings of the 10th Symposium on Operating System Principles* (Orcas Island, WA).
- [18] ZAYAS, E. (1987) Attacking the Process Migration Bottleneck. In *Proceedings of the 11th ACM Symposium on Operating System Principles* (Austin, TX).
- [19] DOUGLIS, F. and OUSTERHOUT, J. (1991) Transparent Process Migration: Design Alternatives and the Sprite Implementation. *Software Practice and Experience* 21(8).
- [20] ARTSY, Y. and FINKEL, R. (1989) Designing a Process Migration Facility: The Charlotte Experience. *IEEE Computer* 22(9).
- [21] ZANDY, V., MILLER, B. and LIVNY, M. (1999) Process Hijacking. In *8th International Symposium on High Performance Distributed Computing* (Redondo Beach, CA).
- [22] DUNLAP, G.W., KING, S.T., CINAR, S., BASRAI, M. and CHEN, P.M. (2002) ReVirt: Enabling Intrusion Analysis through Virtual-Machine Logging and Replay. In *Proceedings of the 2002 Symposium on Operating Systems Design and Implementation (OSDI)*.
- [23] KING, S.T., DUNLAP, G.W. and CHEN, P.M. (2005) Debugging operating systems with time-traveling virtual machines. In *Proceedings of the 2005 Annual USENIX Technical Conference*.
- [24] VEERARAGHAVAN, K., LEE, D., WESTER, B., OUYANG, J., CHEN, P.M., FLINN, J. and NARAYANASAMY, S. (2011) DoublePlay: Parallelizing sequential logging and replay. In *Proceedings of the 2011 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [25] Internet Suspend/Resume, <http://isr.cmu.edu/>. Accessed on May 2, 2014.
- [26] SATYANARAYANAN, M., KOZUCH, M., HELFRICH, C. and O'HALLARON, D.R. (2005) Towards Seamless Mobility on Pervasive Hardware. *Pervasive and Mobile Computing* 1(2).
- [27] SATYANARAYANAN, M., GILBERT, B., TOUPS, M., TOLIA, N., SURIE, A., O'HALLARON, D.R., WOLBACH, A. *et al.* (2007) Pervasive Personal Computing in an Internet Suspend/Resume System. *IEEE Internet Computing* 11(2).
- [28] ABE, Y., GEAMBASU, R., JOSHI, K., LAGAR-CAVILLA, A. and SATYANARAYANAN, M. (2013) vTube: Efficient Streaming of Virtual Appliances Over Last-Mile Networks. In *Proceedings of the ACM Symposium on Cloud Computing* (Santa Clara, CA).
- [29] HA, K., PILLAI, P., RICHTER, W., ABE, Y. and SATYANARAYANAN, M. (2013) Just-in-time Provisioning for Cyber Foraging. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys 2013)* (Taipei, Taiwan).
- [30] McDONOUGH, J. ET AL (2010), Preserving Virtual Worlds, <http://hdl.handle.net/2142/17097>. Accessed on May 2, 2014; also see 'Preserving Virtual Worlds' project website: <http://pvw.illinois.edu/pvw/>.
- [31] CAMiLEON Project : Creative Archiving at Michigan and Leeds Emulating the Old on the New, <http://www.webarchive.org.uk/ukwa/target/113954/source/alpha>. Accessed on May 2, 2014.
- [32] bwFLA – Emulation as a Service, <http://bw-fla.uni-freiburg.de/>. Accessed on May 2, 2014.
- [33] TOTEM, The Trustworthy Online Technical Environment Metadata Database – TOTEM, <http://keep-totem.co.uk/>. Accessed on May 2, 2014.
- [34] Welcome to XSEAD: A community platform for those working across disciplines., <http://xsead.org>. Accessed on May 2, 2014.
- [35] ST. CLAIR, G. (2010), Challenges in Sustaining the Million Book Project, a project supported by the National Science Foundation, In 'The Selected Works of Gloriana St. Clair' at http://works.bepress.com/gloriana_stclair/19.

	Timeframe	Operating System	Application	Description
1	Late 1980s to early 1990s	Microsoft MS-DOS	Preferred orientation package - Los Alamos	Texture analysis software package that provides a comprehensive treatment of material texture analysis.
2			Air Stripper Design and Costing (ASDC)	ASDC enables rapid generation and evaluation of alternative air stripper designs for removal of volatile organic compounds (VOCs) from water.
3			Amortizer Plus 3.01	Amortizer calculates loan amortization in a user-friendly, character-cell interface
4			Wanderer	Wanderer is a game similar to the old "Boulderdash" or "Repton" games, and was originally written to run under UNIX on text terminals (TVI910 and Wyse60)
5			DOOM for DOS	The original DOOM First Person Shooter game
6	Early to mid-1990s	Microsoft Windows 3.1	Microsoft Office 6.0	Microsoft Word 6.0c, Excel 5.0, and PowerPoint 4.0
7			WordPerfect 6.1	Widely used word-processing software before Microsoft Word became dominant
8			Electronic Anesthesiology Library 1991-95	A compilation in multimedia format of four journals (Anesthesiology, Anesthesia and Analgesia, British Journal of Anaesthesia, and The Canadian Journal of Anaesthesia). Includes the Knowledge Finder software for searching the journals.
9			TurboTax 1997	Tax-preparation software for tax year 1997
10			Great American History Machine	Visualization software to explore historical American census and election data from the 19th and early 20th centuries. The original version was created in the late 1980s for Carnegie Mellon University's Andrew system.
11	Mid-1990s	Apple Macintosh 7.5	Oregon Trail 1.1	A game designed to teach school children about the realities of 19th century pioneer life on the Oregon Trail. The original pre-Mac version was conceived in 1971 and produced by the Minnesota Educational Computing Consortium (MECC) in 1974.
12			HyperCard 2.4.1	The last version of Apple's HyperCard multimedia authoring system for Macintosh.
13			NCSA Mosaic 1.0	A very early web browser that triggered world-wide awareness of the Internet
14			Netscape Navigator 1.12	One of the earliest versions of the first commercial web browser.
15	Early 2000s	Microsoft Windows XP		Basic operating system environment
16	2012	Scientific Linux 6.4		A Linux release put together by Fermilab, CERN, and various other labs and universities around the world. Its primary purpose is to reduce duplicated effort of the labs, and to have a common install base for the various experimenters.
17			ChemCollective	A collection of virtual labs, scenario-based learning activities, tutorials, and concept tests to teach and learn chemistry.

Figure 11. VM Collection in Olive as of April 2014