Bringing the Cloud Down to Earth: Transient PCs Everywhere

Mahadev Satyanarayanan, Stephen Smaldone[†],
Benjamin Gilbert, Jan Harkes, Liviu Iftode[†]

September 2010
CMU-CS-10-143

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

[†]Rutgers University

Abstract

The convergence of cloud computing and mobile computing is the latest chapter in a long-running dialectic between centralization and decentralization in system design. Cloud computing represents a thrust in which the forces of centralization are ascendant. Mobile computing, on the other hand, represents a thrust in which the forces of decentralization dominate. In this paper, we show how use of virtual machine technology can lead to new "sweet spots" in the space of system architectures that try to reconcile the tradeoffs between centralization and decentralization. Notably, it yields the Transient PC computing model which preserves the centralization benefits of cloud computing without sacrificing mobility or usability.

1 The Roots of Cloud Computing

The intersection of cloud computing with mobile computing is a hot topic today. There is a growing sense that at their nexus lies immense potential for the creation of new applications and services that can be delivered to users at any time and any place. Large dollar signs dance before the eyes of investors and entrepreneurs. What is the basis for this excitement? What are the aspects of this technological convergence that are truly new, and what is merely old wine in new bottles? To fully understand the tectonic forces that are driving mobile computing and cloud computing towards each other, it is important to briefly revisit the past.

At the heart of this convergence lies a fundamental tension between *autonomy* and *interdependence* in distributed systems. This can equally well be characterized as a tension between *centralization* and *decentralization.* On the one hand, mobility and decentralization are all about freedom: the ability to do anything, anywhere, anytime in a completely unconstrained and untethered manner. On the other hand, total isolation is rarely desired by a user. Communication across space and time is essential for collaboration and knowledge sharing across users. Even with respect to a single user, there is often a need to access personal information resources (such as email or files) from the past and across many different mobile or static devices. Unless carefully constrained, the totality of these interrelationships can be overwhelming to a user. System-level mechanisms such as distributed file systems, databases, email clients and hypertext systems as well as design primitives such as cache consistency protocols, content-addressable storage, and atomic transactions can be viewed as order-preserving tools that try to simplify the complexity of information space-time for a user. Centralization of storage in a cloud can be also viewed as an order-preserving tool. Centralization is simpler, less expensive and more orderly than decentralization from the viewpoint of system management.

How to preserve order and hence reduce human-visible complexity (both for users and system administrators), while minimally constraining mobility and flexibility (in their broadest sense) across time, space and multiple computing devices is a fundamental challenge. Finding the "sweet spots" in this tradeoff space has challenged the system designers for over a quarter of a century.

One of the earliest efforts to reconcile this tension took place at the dawn of personal computing in the early 1980s. The goal of this effort, called *the Andrew project,* was to create a system in which users enjoyed the load-invariant, high-quality, feature-rich interactive environment of personal computing while preserving the ease of information sharing that had emerged as a valuable side effect of timesharing. Contemporary accounts of this project [11, 16] discuss the design considerations that led to its architecture. The image and
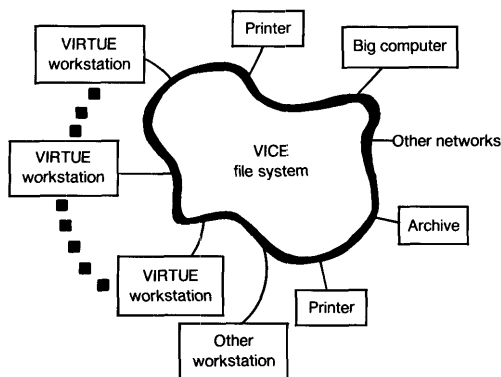
The amoebalike structure in the middle, called VICE, is a collection of communication and computational resources serving as the backbone of a user community. Individual workstations, called VIRTUEs, are attached to VICE and provide users with the computational cycles needed for actual work as well as a sophisticated user–machine interface. (VICE stands for "Vast, Integrated Communications Environment"; VIRTUE, for "Virtue is reached through UNIX® and EMACS.")

(b) Original Caption

(a) Original Image

Figure 1: Cloud Computing *circa* 1986 (Source: Morris et al, 1986 [11])

Figure 2: Cloud-like Attributes *circa* 1990 (Source: Satyanarayanan, 1990 [14])

caption shown in Figure 1 are reproduced from one of those accounts. It is interesting to note the similarity to modern cloud architectures, even though the term "cloud computing" was nearly 20 years in the future. An Andrew user did not know or care where his data was stored: it was just "somewhere in the cloud" and magically appeared at his current location when it was accessed. An Andrew system administrator could move data across servers for load balancing or add new servers without disrupting users. We associate these attributes of mobility and serviceability with cloud computing today. Usage experience in Andrew revealed the importance of these attributes as early as 1990, as shown by the quotation in Figure 2.

2   VM-based Cloud Computing Today

Fast forward to the present day. The buzz around cloud computing directly reflects renewed interest in the attributes shown in Figure 2, in the hope of finding a new "sweet spot" in the centralization-decentralization tradeoff space that reflects current economic reality. As hardware costs continue to plummet, the people cost of system admiminstration looms ever larger. The attention now being paid to the metric called *total cost of ownership* reflects this rise in people cost relative to hardware cost. There is pressure towards centralization because it tends to lower system administration costs.

Virtual machine (VM) technology leads to new "sweet spots" in the tradeoff space that provide greater centralization without loss of user mobility. Figure 3 identifies these "sweet spots" as quadrants of a two-dimensional space. The vendor-neutral and format-independent term *parcel* in this figure refers to encapsulated VM state. A parcel may contain both volatile state (a memory image) and persistent state (a disk image), or just persistent state. The "Parcel Storage" dimension shows where parcels are stored when not in active use; the "Parcel Execution" dimension shows where they run. The top right quadrant, labeled *Managed Execution,* contains approaches such as Amazon's EC2 [1], IBM's Research Compute Cloud (RC2) [2], and the open source Eucalyptus infrastructure [12], where parcel storage and execution both happen within the cloud. The top left quadrant corresponds to *VM-based grid computing,* where a locally-created parcel is remotely executed. The bottom left quadrant, labeled *Classic PC,* is the degenerate case of a standalone PC, laptop or mobile device.

The bottom right quadrant of Figure 3, labeled *Transient PC,* is of particular interest to mobile computing because it makes possible a "carry-nothing" model of mobility. In this case, parcels are stored in the cloud but execute on a computer that is close to the user. That computer could, of course, be a mobile device. Alternatively, mobile computing can be realized by the user taking advantage of whatever hardware is nearby and letting the system magically deliver a relevant parcel upon use. The latter approach directly corresponds to the Andrew architecture of Figure 1, with the important difference that the granularity of location-transparent access is now an entire parcel rather than user files. In other words, the Transient PC model improves upon the Andrew model by ensuring that the operating system, applications and user files are all delivered from the cloud as a cohesive unit. The technique of *hoarding,* originally developed in the context of a distributed file system [6], can be also applied to parcel state. This enables the Transient PC model to support disconnected operation, a critical capability at the intersection of cloud computing and

Parcel Storage

	Local	Cloud
Cloud	local-cloud **Remote execution** (e.g. Grid Computing)	cloud-cloud **Managed execution** (e.g. Amazon EC2)
Local	**Classic PC** (also laptops) local-local	**Transient PC** (e.g. ISR) cloud-local

more compute power

crisper interaction

better availability · *better safety*
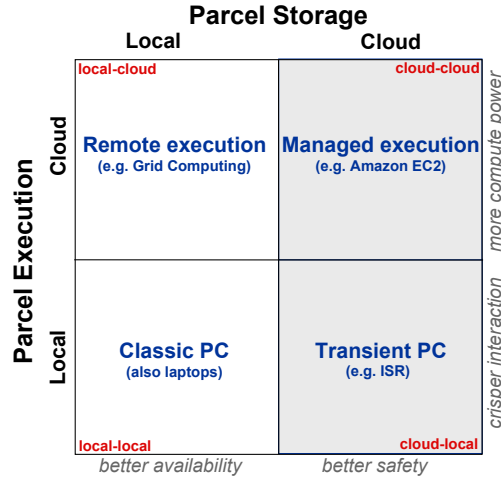
Parcel Execution

Figure 3: Taxonomy of VM-based Cloud Computing

mobile computing.

There are many technical challenges in efficiently implementing the Transient PC model, the most obvious of these being efficient handling of large parcel size. Fortunately, these implementation challenges can be overcome, as shown by our work in the context of the Internet Suspend/Resume® system (ISR) [7, 15] and by related work in the Collective [13, 3] and MokaFive [10].

3 Fully Exploiting Local Hardware Resources in Transient PCs

The Transient PC usage model is quite different from the ubiquitous email, Web access, and social networking capabilities provided by BlackBerries, iPhones, and other mobile devices. The strength of Transient PC systems lies in their ability to precisely, safely and rapidly re-create a user's Windows or Linux desktop environment as a thick client on borrowed hardware at any time and place. When a user chooses to use a Transient PC, he or she implicitly confirms the importance of its strengths over the mobile device alternatives mentioned above. It is therefore critical to take full advantage of the local hardware resources of a Transient PC in order to provide a satisfactory user experience.

The most important property of these local resources is their proximity to the user. For interactive applications, executing the application in a distant cloud and viewing the output locally (as would be the case with the top right quadrant of Figure 3) results in an unsatisfactory user experience. The long WAN latencies to the cloud hurt the crisp interaction that is so critical for smooth and non-disruptive interactions. Humans are acutely sensitive to delay and jitter, and it is very difficult to control these parameters at WAN scale. The work by Lagar-Cavilla et al [8] has shown that latency can negatively impact interactive response even when bandwidth is adequate. Networks with high delay-bandwidth products are therefore disastrous for remote execution of interactive applications; in contrast, local execution on a Transient PC does not suffer from this problem. This is particularly true for graphics-intensive applications such as scientific visualization and games, which can further take advantage of graphics acceleration provided by local hardware. Fitting such applications into the Transient PC model is predicated on the ability to virtualize graphics hardware. This was a questionable assumption for many years because of the lack of standardization in such hardware, but VMM-independent virtualization of graphics hardware is now feasible [9].

Another performance-critical resource is the local disk of a Transient PC. To achieve truly ubiquitous availability of Transient PCs, a zero-install solution would be ideal. This would obviate the need to require pre-installed software on the Transient PC. Instead, the user boots the hardware from a USB storage device to establish the Transient PC environment, and then accesses his cloud-based parcel. The user can thus

Storage Device	Label	Type	Size (GB)	Speed (RPM)	Transfer Rate (MB/sec) Read	Write
PNY Attache USB Flash Drive	SanDrive	USB	16	Flash	30.51 (1.01)	6.65 (0.29)
SanDisk MicroSD Card	MSD	USB	8	Flash	16.03 (0.23)	11.9 (0.2)
Apple iPod	IPOD	USB	20	4200	12.63 (0.18)	12.38 (0.17)
Internal SATA Drive	TransPart	SATA	250	7200	41.78 (2.48)	32.05 (1.14)

Transfer rate results are the mean of 5 measurements. Standard deviations are in parentheses.

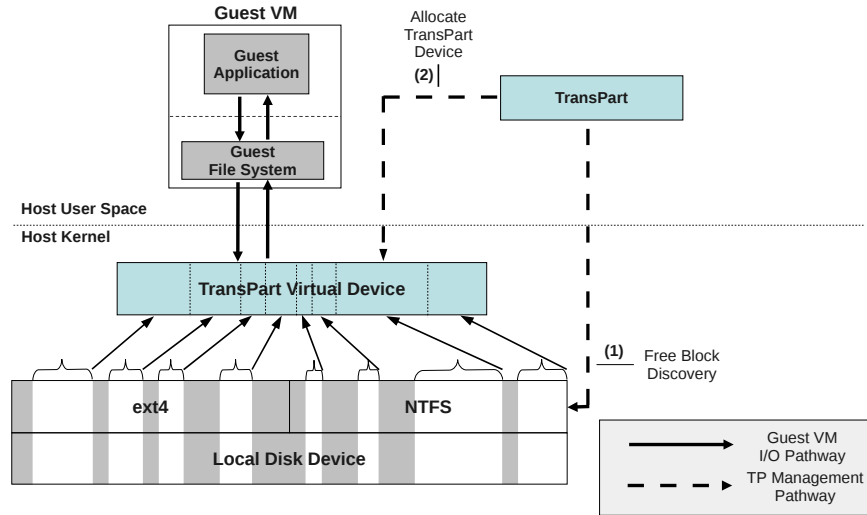Table 1: Portable Storage Device Characteristics.

Figure 4: TransPart Implementation

transform any available hardware with an Internet connection into a Transient PC. Although conceptually simple, this approach runs into a serious performance obstacle. Portable storage devices sacrifice I/O performance in order to obtain the highest capacity and robustness at the lowest cost, size, and weight. As Table 1 shows, the I/O read and write performance of typical USB-attached storage devices is substantially slower that of an internal disk. This severely impacts operating system performance, including basic functionality such as swapping and application launch. Upgrading the interconnect to USB 3.0 will not eliminate this problem, since it is due to internal storage limitations.

To solve this problem, we have developed a mechanism called *TransPart* that constructs a transient virtual disk out of the free disk blocks of local hardware. This mechanism requires no modifications to the software or hardware of a borrowed computer, and thus preserves the zero-install attribute. TransPart currently supports discovery of free disk blocks from ext2/3/4 and NTFS file systems as well as Linux swap partitions.

Figure 4 illustrates the operation of TransPart. During the host boot process from a USB device, TransPart constructs the virtual disk in two phases. In the first phase, TransPart enumerates local devices and then discovers individual storage volumes stored on these devices. Such volumes include physical disk partitions as well as aggregate volumes such as software RAID and Logical Volume Manager volumes. Each storage volume usually contains a single file system or swap partition. In the second phase, TransPart searches through each storage volume to discover the available free blocks. Most modern file systems maintain a set of block allocation tables as meta-data on disk. TransPart utilizes file system on-disk semantics
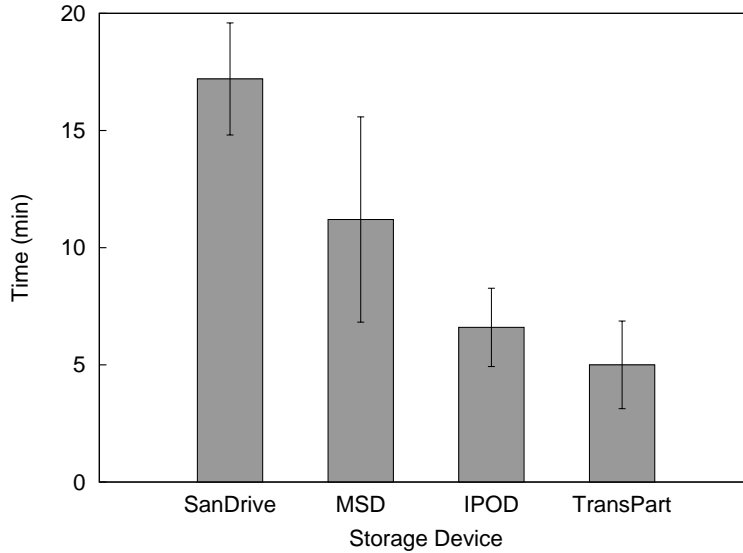
4

Figure 5: Completion time of the Postmark v1.51 benchmark (minutes)

to properly parse the file system metadata and to discover free disk blocks. Once TransPart has discovered free disk blocks, it allocates a TransPart device for the guest VM. An ext4 filesystem is created inside the TransPart device, and this filesystem is used to store parcel data demand-fetched from the cloud. All data written to the TransPart device is encrypted, so that a user never leaves behind residual state that could compromise his privacy.

Our experiments show that TransPart can significantly improve the I/O performance of a Transient PC. For example, Figure 5 shows the performance of the well-known Postmark benchmark [5] on typical desktop hardware. The graph shows that TransPart offers significant performance improvement over the portable storage alternatives shown in Table 1. Figure 6 shows launch latency for a number of common Linux applications on the same Transient PC. Once again, the use of TransPart provides a significant improvement in user-visible performance. Experimental details of the results shown in Figures 5 and 6, as well as additional implementation details on TransPart are available in related technical reports [17, 4].

4   Closing Thoughts

The convergence of cloud computing and mobile computing is the latest chapter in a long-running dialectic between centralization and decentralization in system design. Over the past 50 years, the pendulum has swung back and forth between these extremes. Today, cloud computing represents a thrust in which the forces of centralization are ascendant. Mobile computing, on the other hand, represents a thrust in which the forces of decentralization dominate. Where these tectonic forces meet, there will inevitably be a lot of heat generated and, hopefully, also some light. In this paper, we have put forth the view that use of virtual machine technology can lead to new "sweet spots" in the space of system architectures that try to reconcile the tradeoffs between centralization and decentralization. Notably, it yields the Transient PC computing model which preserves the centralization benefits of cloud computing without sacrificing mobility or usability.
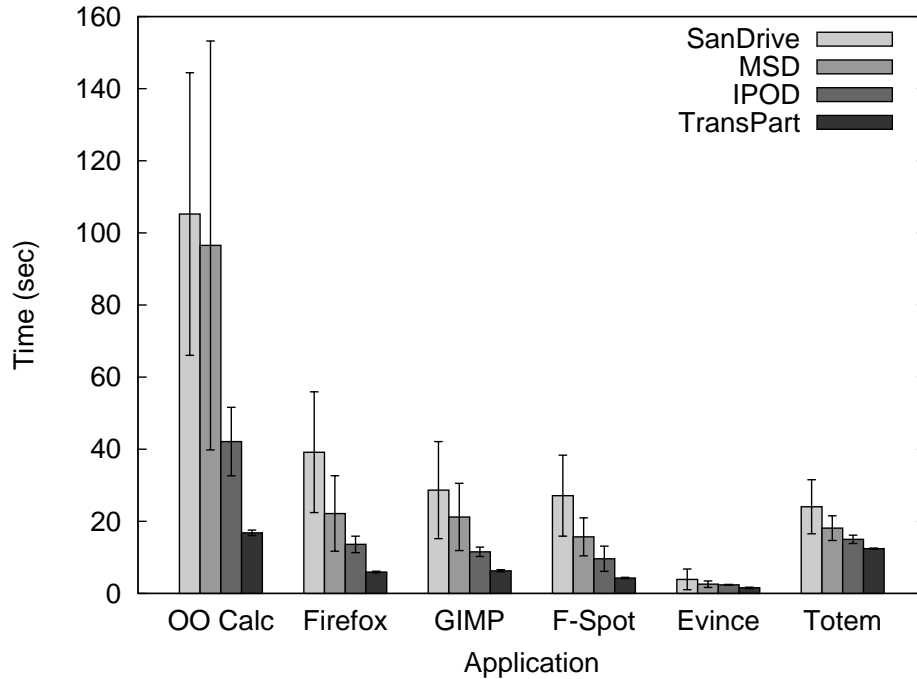
Figure 6: Application launch latency time for 6 common desktop applications (seconds)

References

[1] AMAZON INC. Amazon Web Services - Elastic Compute Cloud. `http://aws.amazon.com/ec2/`.

[2] AMMONS, G., BALA, V., BERGER, S., DA SILVA, D. M., DORAN, J., FRANCO, F., KARVE, A., LEE, H., LINDEMAN, J. A., MOHINDRA, A., OESTERLIN, B., PACIFICI, G., PENDARAKIS, D., REIMER, D., RYU, K. D., SABATH, M., AND ZHANG, X. RC2: A living lab for cloud computing. IBM Research Report RC24947, IBM, 2010.

[3] CHANDRA, R., ZELDOVICH, N., SAPUNTZAKIS, C., AND LAM, M. The Collective: A Cache-Based System Management Architecture. In *Proceedings of the Second Symposium on Networked Systems Design and Implementation* (May 2005).

[4] GILBERT, B., GOODE, A., AND SATYANARAYANAN, M. Pocket ISR: Virtual Machines Anywhere. Tech. Rep. CMU-CS-10-112, School of Computer Science, Carnegie Mellon University, 2010.

[5] KATCHER, J. PostMark: A New File System Benchmark. Tech. Rep. TR3022, Network Appliance, 1997.

[6] KISTLER, J.J., SATYANARAYANAN, M. Disconnected Operation in the Coda File System. *ACM Transactions on Computer Systems 10*, 1 (February 1992).

[7] KOZUCH, M., AND SATYANARAYANAN, M. Internet Suspend/Resume. In *Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications* (Callicoon, NY, June 2002).

[8] LAGAR-CAVILLA, H. A., TOLIA, N., DE LARA, E., SATYANARAYANAN, M., AND O'HALLARON, D. Interactive Resource-Intensive Applications Made Easy. In *Proceedings of Middleware 2007: ACM/IFIP/USENIX 8th International Middlewae Conference* (Newport Beach, CA, November 2007).

[9] LAGAR-CAVILLA, H. A., TOLIA, N., SATYANARAYANAN, M., AND DE LARA, E. VMM-Independent Graphics Acceleration. In *Proceedings of the 3rd ACM SIGPLAN/SIGOPS Conference on Virtual Execution Environments (VEE)* (San Diego, CA, 2007).

[10] MOKAFIVE, INC. MokaFive Home Page. `http://www.mokafive.com`.

[11] MORRIS, J. H., SATYANARAYANAN, M., CONNER, M. H., HOWARD, J. H., ROSENTHAL, D. S., AND SMITH, F. D. Andrew: a Distributed Personal Computing Environment. *Commun. ACM 29*, 3 (March 1986), 184–201.

[12] NURMI, D., WOLSKI, R., GRZEGORCZYK, C., OBERTELLI, G., SOMAN, S., YOUSEFF, L., AND ZAGOROD-NOV, D. The Eucalyptus Open-Source Cloud-Computing System. In *CCGRID '09: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid* (2009).

[13] SAPUNTZAKIS, C., CHANDRA, R., PFAFF, B., CHOW, J., LAM, M., AND ROSENBLUM, M. Optimizing the Migration of Virtual Computers. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation* (Boston, MA, December 2002).

[14] SATYANARAYANAN, M. Scalable, Secure, and Highly Available Distributed File Access. *IEEE Computer 23*, 5 (May 1990), 9–18, 20–21.

[15] SATYANARAYANAN, M., GILBERT, B., TOUPS, M., TOLIA, N., SURIE, A., O'HALLARON, D. R., WOL-BACH, A., HARKES, J., PERRIG, A., FARBER, D. J., KOZUCH, M. A., HELFRICH, C. J., NATH, P., AND LAGAR-CAVILLA, H. A. Pervasive Personal Computing in an Internet Suspend/Resume System. *IEEE Internet Computing 11*, 2 (2007).

[16] SATYANARAYANAN, M., HOWARD, J. H., NICHOLS, D. A., SIDEBOTHAM, R. N., SPECTOR, A. Z., AND WEST, M. J. The ITC Distributed File System: Principles and Design. In *Proceedings of the 10th ACM Symposium on Operating Systems Principles* (Orcas Island, WA, December 1985).

[17] SMALDONE, S., HARKES, J., IFTODE, L., AND SATYANARYANAN, M. Safe Transient Use of Local Storage for VM-based Mobility. Tech. Rep. CMU-CS-10-110, School of Computer Science, Carnegie Mellon University, 2010.