# A Peer-to-Peer Mobile Storage System

*a report by*

**Dr Randolph Wang, Nitin Garg, Yilei Shao, Elisha Ziskind, Sumeet Sobti, Fengzhou Zheng, Junwen Lai** *and* **Arvind Krishnamurthy**

*Princeton University*

Dr Randolph Wang is an assistant professor of computer science at Princeton University.

Nitin Garg is a graduate student in Computer Science at Princeton University.

Yilei Shao is a PhD student in Computer Science at Princeton University.

Elisha Ziskind is a graduate student in Computer Science at Princeton University.

Sumeet Sobti is a graduate student in Computer Science at Princeton University.

Fengzhou Zheng is a PhD student in Computer Science at Princeton University.

Junwen Lai is a PhD student in Computer Science at Princeton University.

Arvind Krishnamurthy is an assistant professor of computer science at Yale University.

## Introduction

As the cost, form factor and capacity of stable storage continues to improve dramatically, one consequence is the emergence of highly compact secondary storage technologies that can be integrated seamlessly into devices of all shapes and forms. Today, these devices are largely disjointed and users are expected to hoard, propagate and back up data on individual devices manually. As these devices proliferate rapidly in our surroundings, we are faced with an increasingly difficult challenge of managing this chaotic sea of 'invisible bits'.

This article argues that an effective way of managing this data is by applying the peer-to-peer philosophy: instead of 'powering' all these devices with an omnipresent external networked storage utility, these devices are, by themselves, able peer components of a mobile storage system, and what is needed is a piece of system software that ties all these disjointed devices into a coherent whole. However, wide area connectivity alone is not always sufficient for the purpose of co-ordinating all these devices. To compensate for this inadequacy, we introduce a small portable storage device that is equipped with several connectivity technologies. This device leverages ad hoc peer-to-peer connectivity and an embedded storage element to overcome the wide area connectivity bottleneck and accomplish its role as a co-ordinator of the other devices. From a user's point of view, this portable device follows them wherever they go (just like a BlackBerry e-mail device would today). As long as the user has the device with them, they can:

- access their own data transparently, regardless of where the bytes are physically stored in relation to the user;
- read other users' data transparently when access has been granted; and
- use the portable device as a storage 'adaptor' for other appliances so that their special purpose data can be readily integrated into the entire storage system.

One important difference between this system and some existing application-specific solutions is that it is designed to be a storage-level solution that can support most existing file systems and applications transparently. The system must address three fundamental problems:

1. How does the system locate data that can be stored on any device?
2. How does the system ensure consistency across multiple devices as old data on these devices becomes obsolete?
3. How does the system ensure a consistent image across all devices for the purpose of back-up and sharing?

The system solves the first problem using a location and topology-sensitive multicast-like solution. The advantage of this solution is that it minimises global state, allows for autonomous data movement decisions, and can effectively exploit locality. The system solves the second problem using lazy peer-to-peer propagation of invalidation information. As a result of decoupling from data propagation, it can quickly bring weakly connected devices up to date. The system solves the third problem using a distributed snapshot mechanism. The advantage of this solution is that it allows a user to continue to modify the storage system without interfering with either back-up or sharing. A Linux-based prototype of the system has been implemented in which the role of the portable storage device is played by a Compaq iPAQ equipped with storage and connectivity accessories. One goal of this design is to minimise the amount of distributed state and global co-ordination. This design goal has resulted in a relatively simple and robust implementation that has little complexity in its crash recovery mechanism, for example. Initial experience with the prototype leads to the conclusion that a peer-to-peer model, when aided by a mobile device equipped with embedded storage and ad hoc connectivity, can effectively overcome wide area connectivity weakness and enable mobile users to access and share data ubiquitously.

## Two Naive Approaches

There are two broad categories of existing solutions to providing ubiquitously available mobile storage. One is relying on ubiquitous connectivity to a 'storage utility' infrastructure; the second is relying exclusively on a small portable storage device.

### *Connectivity Limitations*

High-performance universal network connectivity remains an elusive goal. Even so-called 'broadband' home digital subscriber line (DSL) users typically only have access to an uplink capacity of around 100 kilobits per second (Kb/s) today. The much-anticipated third-generation (3G) wireless networks are designed to achieve 384Kb/s ultimately, but industry observers agree that wide availability of such speeds is many years away. Today, 3G users in the US can realistically expect data speeds of somewhere between 40–80Kb/s, a far cry from the hypothetical speeds of 144Kb/s and 192Kb/s.[1] At any instant, only a small number of devices may be strongly connected to each other, and a mobile storage user cannot always rely on omnipresent high-quality connectivity to a centralised storage service.

### *Limitations of Portable Storage Devices*

Due to the difficulty of accessing data across a mobile wide area network (WAN), we tend to resort to carrying bits with us. Today, when some of us travel, we are apprehensive about leaving our laptops behind, not necessarily because we fear that our travel destinations lack computers, or because our personal machines have any special capabilities, instead, what makes a person's machine personal is the data stored on it. Carrying a laptop for this purpose, however, is cumbersome. Among its many faults, one of the most serious is that the form factor of a generic computing device such as a laptop is unlikely to improve in terms of portability due to user interface considerations such as the size of a screen or a keyboard.

Recognising this inconvenience, manufacturers have started to offer a wide array of mobile storage devices. The form factor of these devices can be as small as a key-ring ornament.[2] It is hoped that as storage density continues to increase, a day may come when all a user would have to carry is such a small device.

Computers should be cheap, widely available and generic. The only thing that a user needs in order to take advantage of these widely available generic computing devices is an easily portable storage device that houses their personal data, something that makes a user unique. However, relying exclusively on such a device for all data storage needs has its pitfalls. First, despite the capacity improvement of storage devices, the nature of new applications' appetite for storage is such that the capacity of a single portable device is unlikely to be sufficient for all of a user's storage needs. The capacity of mobile devices is likely to continue to lag behind that of their stationary counterparts, and it is expected that much data will continue to reside on these stationary devices. Second, mobile storage devices tend to have poorer performance compared with desktop versions due to energy consumption and form factor considerations. Third, mobile storage devices tend to be less reliable due to environmental and human factors such as exposure to shock, moisture and theft. Last, mobile storage devices, by themselves, provide little support for transparent data-sharing between collaborating users.

## The Skunk-based Peer-to-Peer Approach

The portable device referred to here is known as the Skunk device (patent-pending) or simply the Skunk, and the entire peer-to-peer storage system will be referred to as the Skunk system. The Skunk-based approach has two key elements as follows:

1. instead of relying on omnipresent access to a storage utility, the system is based on the co-ordination of users' existing peer devices and their embedded storage elements; and

2. instead of relying exclusively on a portable storage device, the system employs such a device as a co-ordinator of the other peer elements, and exploiting this portable storage element and ad hoc peer-to-peer connectivity becomes a powerful means of overcoming the handicaps of wide area connectivity.

### *Peer-to-Peer Components*

The storage managed by the Skunk system principally resides with the devices already owned by users. These may include personal computers at home, at work or on the road, and various consumer electronic appliances already equipped with storage elements. In many cases, these devices are sufficiently powerful by themselves and are the natural 'homes' of certain data, so a peer-to-peer system consisting of these existing storage elements makes good sense. One of the main

1. *J Wagner, "Getting to Know Your 3G", http://www.internetnews.com/wireless/article/0,,10692_964581,00.html, January 2002.*
2. *Sony hangs its memory on a key-ring, http://www.cnn.com/2001/TECH/ptech/11/13/sony.keychain.idg, November 2001.*

roles of the Skunk system is to co-ordinate these peer devices to form a single coherent name space.

### The Skunk Device

The task of co-ordinating these existing disjointed devices, however, may be difficult with existing connectivity capabilities. The role of the Skunk device is to facilitate this co-ordination across a potentially wide area. Some of the tasks performed by the Skunk include:

- storing and propagating metadata that is used to drive other devices towards eventual consistency;
- caching and propagating data to improve performance;
- providing wireless, peer-to-peer, ad hoc, short-range connectivity between peer devices; and
- providing wireless wide area connectivity between peer devices.

Physically, a Skunk consists of a processor, a storage element, an ad hoc connectivity interface and a WAN connectivity interface.

It is anticipated that an industrial-strength version of the Skunk can be packaged in a form factor that is not much larger than a wristwatch, and, as with a wristwatch or the larger BlackBerry e-mail device, the Skunk is a personal device and accompanies the user at all times. As long as the user has a Skunk with them, they can always access and share the storage system.

### Use Cases

A user works on an office computer with a Skunk nearby, which communicates with the office computer via an ad hoc wireless interface. The user 'sees' a single Skunk-backed file system with data that may be spread physically across any other of their computers. As new data is created on the office computer, some of the new data may be stored on the Skunk, and some of it may be pushed to other computers in the background. When the user leaves the office at the end of the day, carrying only the Skunk home, the same coherent Skunk file system is seen on their home computer, which happens to have only a weak DSL uplink. As the user reads data from their home computer, the Skunk system ensures that they always read a fresh copy of the data, which may reside physically on their home computer, Skunk device, office computer or any other devices that the user may own. The next morning, the user carries the Skunk back to the office and the cycle repeats. The Skunk carries the metadata for ensuring consistency and caches much of the data to avoid overstressing the weak wide area link between home and office.

Two colleagues bring their Skunks on a business trip. The night before their scheduled presentation, they need to collaborate on their slides as they work on the hotel computers. The Skunk system allows each colleague to have read-only access to some of the other colleague's data, which, again, may be stored physically on any of their devices. The system uses this communication pattern (known as Skunkast) to route requests from one colleague to the closest device of the other colleague that houses the desired replica. The ad hoc wireless interfaces on their Skunks allow the two users to share data directly and quickly without resorting to a wide area connection that may be either weak or non-existent on the hotel computers. If the collaborators are separated by a large geographical distance, however, the route choices available to the Skunkast may be more diverse and complex. At any instant, the Skunk system routes the user's request to the closest replica.

The data management needs addressed by the Skunk system are by no means limited to traditional desktop applications. As the data management functionalities are separated from cumbersome generic computing devices, and as these functionalities are encapsulated in modular small form factor devices (such as the Skunk), these devices can then readily interact with other consumer electronic devices in interesting and useful ways. For example, as a camera is coupled with a Skunk, the images captured can be made instantly shareable among family and friends; as an MP3 player is coupled with a Skunk, friends can swap music as ad hoc fan groups form; and as an e-mail device is coupled with a Skunk, file-sharing can be integrated into e-mails. In these examples, the Skunk devices essentially serve as storage 'adaptors' for other appliances so that their special-purpose data becomes accessible on a 'backplane' connecting all peer-to-peer Skunk devices.

### Conclusion

A peer-to-peer storage system has been designed and implemented that caters to the needs of mobile users. The system co-ordinates many of its user's existing devices to form a coherent whole so that it may ease the user's increasingly difficult data management chores. By judiciously introducing and enlisting the aid of ad hoc peer-to-peer connectivity and portable storage elements, the system can overcome the inadequacy of wide area connectivity. A combination of the system's features, including the Skunkast data location mechanism, the lazy invalidation log propagation and the distributed snapshot scheme, are uniquely suited for the decentralised, autonomous and personalised environments it targets.