A File Migration Architecture for Pervasive Systems

Adam Hudson, Bob Kummerfeld, Aaron Quigley
School of Information Technologies
University of Sydney, Australia
{ahudson, bob, aquigley}@it.usyd.edu.au

ABSTRACT
In this poster, we describe the outline of an architecture that uses caching within a Personal Area Network to increase the availability of data generated by mobile devices, and data migration between these devices and a remote server to increase the amount of reliable storage that is available.

Keywords
Mobile computing, personal area networks, peer-to-peer networking

INTRODUCTION

Modern portable devices are great producers and consumers of data. Digital cameras, mobile phones, music players, PDAs and laptops are some common examples of this. However, as outlined in [1], there are four main constraints experienced by mobile devices when compared to fixed systems:

- They will always be relatively resource-poor.
- They are at higher risk of loss, damage or breaches of security.
- Connectivity is highly variable in terms of availability, reliability and performance.
- All levels of hardware and software must take power efficiency into consideration.

Users need to be able to store the data from these devices to protect it against loss, and to retrieve the data that they wish to access through the devices on demand. However, this can be difficult for an isolated, potentially resource limited device to achieve.

Therefore, when equipped with wireless networking, these devices can join together in an ad hoc manner to form a Personal Area Network (PAN). Communication across this network allows them to share their resources, particularly storage and Internet connectivity, and therefore work around their limitations as mobile devices.

If devices are able to access each others' storage in a coordinated manner, then they can use it to distribute their data around the PAN. Devices in ad hoc networks come and go quite often, usually as they move in and out of range or run out of power. If data is cached in multiple places throughout the PAN, then it remains available to access even if the device that generated it has left. More generally, these caches can be located beyond the PAN to provide further availability.

The need also exists to backup data produced from within a user's PAN to a safer, fixed point, such as a server in the home or office. Likewise, users may wish to access a particular piece of their data while on the move. A dedicated backup server can be reached from within a PAN provided that a connection to the Internet is available. Devices within a PAN may be able to offer this connectivity, such as a connection to a wireless gateway or a GPRS service through a mobile phone. If a direct connection is not available, then a series of store-and-forward caches housed on other portable devices help the data migrate towards its destination until the Internet becomes available. A network of similar caches across the Internet provides the basis for data delivery in both directions between the PAN and the server.

This poster presents the initial design for a file migration system that runs upon a network of portable devices and the fixed Internet. The architecture allows for the devices to share data between them, to distribute data to caches on other systems and to save data to a backup server. A unified view of the distributed data is presented to applications running on portable devices, so that they can easily access the data wherever it resides. The work bears similarity to Internet-scale storage systems such as OceanStore [2], but with a particular focus on the role of portable devices.

NETWORK FORMATION

To support the file distribution operations, a user's portable devices join together in an ad-hoc fashion, and then use a service discovery protocol to identify which other local devices are of use to them. This is determined by the properties that are visible such as current available storage space, battery remaining and connection type and speed. The information is used to establish routing tables for where to migrate the data (as discussed later), with a cost associated with each link.

Distribution between a PAN and the backup server takes place across a peer-to-peer network overlaying the Internet. It initially contains the backup server alone, but other nodes are added as a user's PAN connects to them as forwarding points towards the server.

ARCHITECTURE OVERVIEW

Nodes on the system each have the same layout, whether they are on a portable device, an Internet based system or a backup server. Each contains a file system, a cache, a file manager and a migration queue, in addition to any

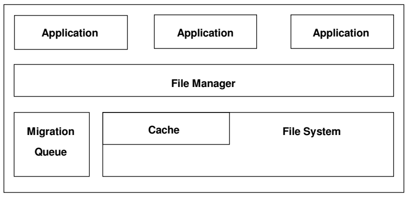

Figure 1 - The layout of a node

applications (see figure 1). Between nodes, data is moved around in segments, rather than complete files, to increase the chance of completing a transfer across an unstable link and to ensure they can fit within caches.

Applications access their local storage through the file manager. When an application wishes to store a file, it passes it to the file manager, along with metadata collected either explicitly from the user or automatically. The file and metadata are stored in the local file system, with pointers to them placed in the migration queue. The migration queue is a data structure which unifies the file manager's outbound communications, housing a series of such pointers in the order that they are to be sent to other nodes. The file manager may also receive data to cache from other nodes, which it stores to the section of the file system reserved for caching and then references in the migration queue.

The file manager provides a view of all the files accessible across the network. If an application wishes to retrieve a remote file, a request is placed on the migration queue, forwarded to the store, and the file manager returns a reference to it once all the segments are received.

DATA MIGRATION

The migration queue on each node holds pointers to file segments inserted by local applications, cached segments that have been received from other nodes and file requests. Each queue entry has associated metadata which aids its migration.

The metadata is used to determine an importance metric, which in turn determines what position the data will take in the queue. When the file manager decides that the importance of the data at the head of the migration queue makes it worthwhile to transfer across an available link, it does so. Data may be transmitted several times to different next hops, to increase the level of redundancy and improve the probability of reaching the destination. Within a PAN, the data may be transferred to another device one hop away, to increase its availability if the originating device ceases to be accessible; it may be sent directly to the backup server over any available connection service; or it may be transferred to a device in another PAN, in the hope that it will be forwarded towards the backup server at some later time. Fixed Internet nodes follow similar rules, although their relative stability means that their focus is on forwarding the data directly towards its destination, rather than to other nodes to increase redundancy.

Changing the compression or dimensions of data before transmission may be worth investigating. For example, a backup of a photo a quarter of its original size is better than none at all, and the smaller size increases the chances of a connection limited portable device completing the transfer.

DATA RETRIEVAL

Once the data is distributed across the network, there needs to be a way for devices to retrieve it. The file manager communicates with its peers in the network to determine the files accessible to its applications, making searches and requests based upon the metadata available. From the point of view of an application on a portable device, there is one integrated file system available to it. However, in reality the data may be stored on the local file system, on another portable device within the PAN, in a cache on some other network node or on the backup server: The metadata values and the location(s) of this data needs to be presented to the user in a way that can inform their usage decisions, yet still permit file access by identical means. This metadata may include such properties as the file size, the speed of the connecting link or the remaining battery life of the remote device.

Once an application has made a request for a particular file, the segments need to be retrieved from across the network. A request for each segment is injected into the migration queue, with its destination set to the location that a file segment should be retrieved from. At the destination, the file manager forwards the segment back to the requesting node, which caches it and returns a reference to the application. While the simple case here is for the data to be returned to the original location, a more sophisticated option under investigation is to build a model of user mobility. This could allow data to be forwarded to where a moving device is predicted to be, rather than where it was when it made the request.

CONCLUSION

The design for this system is still being refined, and the next major step is to build a prototype. Through the implementation of this file migration architecture, it is hoped the availability and reliability of the data held on portable devices can be increased.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Smart Internet Technology CRC in this research.

REFERENCES

1. M. Satyanarayanan, "Fundamental Challenges in Mobile Computing," presented at PODC '96, Pennsylvania, USA, May 1996.
2. J. Kubiatowicz et al, "OceanStore: An Architecture for Global-Scale Persistent Storage," presented at ASPLOS 2000, Cambridge, Massachusetts, USA, November 2000.