The Design and Implementation of a RAID-3 Multimedia File Server

Alan J. Chaney, Ian D. Wilson and Andrew Hopper*
Olivetti Research Laboratory
24a, Trumpington Street
Cambridge CB2 1QA
United Kingdom

Tel: +44-1223-343000
Fax: +44-1223-313542

Email: {iwilson,achaney,ahopper}@cam-orl.co.uk

Abstract

The Olivetti Research Laboratory has developed an experimental system based on intelligent peripherals connected directly to an ATM network. As well as multimedia modules (e.g. audio and video) the system also includes a directly connected RAID-3 storage server called the "Disc Brick". This paper describes the architecture of the Disc Brick, and discusses some of the hardware and software issues raised by its design. It also presents measurements taken from a Disc Brick in operation, and discusses how the observations relate to the original design objectives. Finally, the paper attempts to evaluate the Disc Brick as part of ORL's family of directly connected peripherals.

Introduction

An experimental system has been constructed at the Olivetti Research Laboratory which aims to provide a rich multimedia environment for a variety of users [1]. It utilises ATM communications technology as the basic interconnect both for computers and multimedia peripheral modules. Each module is made up of a number of standard component parts: the hardware consists of an ARM CPU, up to 32 Mbytes of memory, and a 100Mbit/s ATM TAXI interface; the low-level software consists of a microkernel called ATMos [2] which provides a mechanism for scheduling processes and controlling low level hardware. In addition there is provision for interoperation with other systems by using protocols such as TCP/IP and XTP, or distributed platforms based on CORBA.

The ATMos framework is used to implement both switches (4x4 and 8x8) and end-point modules (or *direct peripherals*). The modules include ATM video, ATM

---

* University of Cambridge and Olivetti Research Laboratory

audio, ATM LCD tile, ATM TV, ATM frame store, ATM processor farm and the Disc Brick, as well as workstations and PCs (see Figure 1).

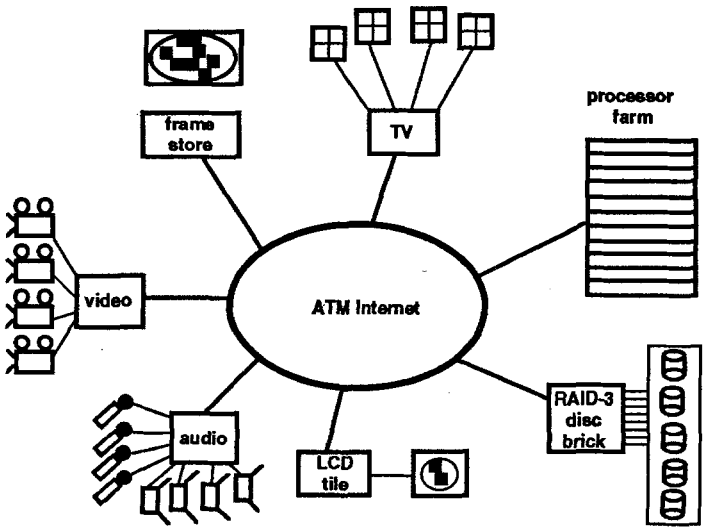

Figure 1

This approach has similarities with the Desk Area Network [3] but also differs from it in important ways. Rather than exploding the workstation, ORL's approach is to treat each direct peripheral module as a first class ATM object; furthermore, most of the traffic from a module will typically flow to points elsewhere in the system and not to a nearby PC or workstation. However, as with the DAN, the approach uses a network as the peripheral interconnect, and hence is scalable.

The system has been deployed in the laboratory and has been in use for about 18 months. Some two hundred modules and switches are available for experimentation. A typical office has four cameras, four speakers, four microphones, a display tile and a workstation as its multimedia infrastructure. In addition, there are five Disc Bricks available on the network for use as multimedia fileservers. Applications in use include a video-mail system, which takes advantage of the Disc Brick's performance by recording several video and audio channels simultaneously. Other applications include a multi-way video-phone which uses four video streams and an audio stream between the corresponding parties.

The rest of this paper concentrates on one particular directly connected ATM peripheral: the RAID-3 Disc Brick.

The Disc Brick Filing System

The Olivetti Research Laboratory has been working in the multimedia area for many years. One of its earlier systems, Pandora [4], relied on a fileserver for the storage of audio and video objects. It is the filing system which was originally designed for the