Operating systems support for the Desk Area Network

Case for Support
March 1992

1 Purpose

Recently Asynchronous Transfer Mode (ATM) has become widely accepted as an important technology for the implementation of future multiservice networks. It has been adopted not only for ubiquitous public wide area service (BISDN), but also as an important local and private networking technology. An important reason for this widespread acceptance has been the offer of a unified transfer mechanism in local and wide area networks providing an end-to-end ATM service.

Our work on host interfaces and multimedia devices has led us to extend ATM principles into the heart of the end system and to use ATM cells as the basic unit of transfer between components of a multimedia workstation; this we have called the Desk Area Network[McAuley 91] (DAN).

The purpose of the proposed work is to investigate an operating system architecture which can manage and control the resources in the DAN so as to allow an applications environment in which it is possible to manipulate real time and continuous media data directly in a general purpose manner.

The nodes of DAN include memory banks, processing elements and IO devices. The processing elements can be either general purpose microprocessors or more specialised signal processing devices (e.g. for compression or pattern recognition), while we expect IO devices to include sources and sinks of various media such as digital video and audio. As the unit of transfer in the DAN is the cell, the *network interface* node, which connects a DAN to a local area ATM network, becomes a simple device, performing cell forwarding similar in many respects to a single port on an ATM switch.

The resources to be managed include the bandwidth available in the switch fabric interconnecting the nodes of the DAN, bandwidth and timing requirements of the different nodes, buffers, processsor cycles (of both general and special purpose processors), and access into the LAN. To provide a general purpose programming environment on such a system we must to schedule the allocation of sets of resources so as to meet the real time and synchronization requirements of applications, as well as being able to provide the graceful degradation of service during overloads which is to be expected of an operating system faced with a dynamic load.

The architectural work in the project will consider the resource model within such a system, while on the practical side, we shall design algorithms and implementation mechanisms to realise this model.

Clearly the implementation phase requires an actual DAN system; however, the design of hardware components for the DAN is not part of this proposal as these have already been developed both directly and indirectly by other projects. The SERC funded 'DAN

Devices' project[1] has developed an ATM camera, and work continues in this project on a display and a compression node. The 'Fairisle' project[2] [Leslie 91], also SERC supported, has provided the ATM networking infrastructure and RISC processor nodes. All these components plug into the standard Fairisle backplane which provides the mechanism for the exchange of ATM cells between nodes.

The starting point for the software within the project is a locally developed microkernel called Wanda. This is already available on the ARM3 based processor nodes and has support for multiprocessor configurations developed for the implementation on the DEC SRC Firefly. The implementation work will involve enhancing the current Wanda resource management system with an implementation of our new architecture.

2 Background

DAN devices

The 'DAN Devices' project is an SERC funded project investigating devices which connect directly into an ATM network. The project aims to provide three devices, camera, framestore and compression. The camera interface pipelines the digitisation and coding of an incoming analogue video signal with transmission into an ATM network. The display interface receives such streams of video as well as graphics requests and renders them into a frame buffer for display. The compression device will use a JPEG device to compress and decompress incoming streams of cels.

The camera interface has been designed and a prototype built. A final version will be available at the start of the proposed project. The framestore is under design. A prototype should be available at the beginning of the project. The compression system will be built during the first year of the project.

Fairisle

Fairisle is an SERC Co-operative Grant with HP Labs Bristol as the co-operating organisation terminating September 1992. It is concerned with the design and implementation of a general topology ATM local area network. Fairisle has developed and built 4 ATM switches which are now in use for experiments investigating issues related to quality of service guarantees. Interswitch links at 100Mbps are provided by standard FDDI components, while the 'Fairisle backplane' provides an 8 bit full duplex data path into the switch fabric which offers 160 Mbps per port.

Pandora

Pandora is a SERC Co-operative project between the Computer Laboratory and Olivetti Research Limited terminating September 1992. It has been concerned with integrating multimedia into a distributed computing environment. Pandora takes the approach that applications orchestrate the establishment and synchronization of real time streams, with the real time data handled by a dedicated peripheral, the Pandora Box.

Our experience with this system has led us to realise the restrictions on possible appli-

[1] The official project title is 'Using ATM to Attach Continuous Media Devices'.

[2] The official project title is 'A Private Fast Packet Switching Network'.

cations caused by isolating the real time data from the application writer. This has led us to consider the more general purpose multimedia workstation and operating system architecture, that is, the DAN and this proposed research respectively.

Wanda

Wanda is a locally developed multiprocessor microkernel developed as a tool for distributed systems research. It has been ported to MIPS, 680x0, ARM and VAX processors (including the DEC SRC Firefly, a shared memory memory multiprocessor composed of 4 uVAX3 devices). The ANSA testbench has been ported to Wanda and allows interaction between Wanda and UNIX based ANSA services.

Wanda is now used as the basis for many of our distributed systems research projects, it is in currently in everyday use on 30 to 40 machines by about 20 researchers. It has served as the platform for a file server project, in the Pandora project for real time network services and on the Fairisle port controller.

Personnel

The main researchers at the Computer Laboratory will be Dr. D.R. McAuley, Dr. I.M. Leslie and Prof. R.M. Needham. All have extensive experience in multiservice network design and distributed operating systems.

3 Programme

DAN Systems

While most of the hardware components will have been designed and built by the start of the project, some straightforward replication of components will be required. It is envisaged that two DAN systems will be built and be in service within a few months of the start of the project. These will be installed on our ATM LAN together with the workstations used for development. 0 **Architecture**

The first phase of the work will concentrate on the design of the resource architecture. One aspect of this will involve investigation and characterisation of the behaviour of possible DAN components, supported by experiments with the available DAN system. Another area of investigation will be to identify the critical resources and to define for each an appropriate degradation policy in the face of overload. Clearly such policies need to be viewed in the scope of the mechanisms and policies of the complete system.

Demonstration

The effectiveness of the architecture will then be investigated by implementation. This will involve the choice, implementation and testing of appropriate algorithms. It is envisaged that in the light of measurements and experience this will led to refinement of the architecture and hence implementation.

4 Resource Requirements

Research Staff

One research associates and one research assistant are required to carry on this work. The research associate will be responsible for the development of the resource management model, the implementation work within the operating system, and a small amount of application work which will be required to demonstrate the utility of the system and approach.

The research assistant will be required to keep the environment running and work under the direction of the research associate. This work will include replication of the DAN hardware at the beginning of the project, keeping the DAN hardware in working order, and providing software support (of the Wanda environment) to the research associate.

Postion	Three Year Cost
Research Associate	67182
Research Assistant	54162
The total cost	£121344

Travel

We are proposing to spend £5K on travel to conferences and workshops. The conferences we wish to attend are SOSP (conference and workshop), SIGCOMM,

Consumables

The replication of DAN hardware will be undertaken by the project. The component cost of 2 DANs is necessarily an estimate.

Ideally each DAN would be built from a 8 by 8 fabric with two camera interfaces, one frame store, a network interface, a memory node, a compression node and two processing elements. In addition cameras and a monitor will be required. However, this can be reduced since we need only ensure that at least one DAN has two processors or two cameras or indeed a compression node. We imagine that a set of devices will be available from the DAN devices project will be available to us. The table below represents the addition items required.

Item	Number Required	Estimate per Unit	Total
fabric	2	1500	3000
rack, power, backplane	2	1000	2000
processor	3	2000	6000
camera and interface	2	1000	2000
framestore and monitor	1	2000	2000
memory system	2	1000	2000
network interface	2	1000	2000
Total			£19000
Total with VAT			22325

The maintainance on workstations for the duration of the project will be £3380 (CHECK)

plus VAT or £3971. Consumables arising from supporting the workstations below are £1050 per year per workstation (including VAT) or £6300 in total.

Grand total for consumables is thus £32596.

Equipment

The research associate and research assistant will each require a workstation. We wish to connect these workstations to the same ATM LAN to which the DAN will be connected. This has led us to choose DECStation 5000/240's with ATM interfaces compatible with the Fairisle network. As one of the main thrusts of the work is concerned with multimedia support, colour workstations are required.

Item	unit cost	total
ATM Interface Cards (from Olivetti Research)	£500	£1000
DECStation 5000 /240 (with discount)	£9639	£19278
Total		£20278
Total with VAT		£23827

Project Total

Staff	121344
Travel	5000
Consumables	32596
Equipment	19278
Total	£178218

5 Collaboration

There is no planned formal collaboration,

6 Exploitation

This work is potentially of major significance. Although the ideas here must still be regarded as speculative and long term, we can forsee a major impact on the way in which workstations and operating systems which support multimedia are designed. We also hope that the resource management model will be extended into application design.

References

[McAuley 91]
"The Desk Area Network"
McAuley D.R. and Hayter M.D.
Operating Systems Review October 1991

[Leslie 91]
"Fairisle : An ATM Network for the Local Area"
Leslie I.M., and McAuley D.R.
ACM Computer Communication Review, Vol 21, No 4 September 1991