

Printing Requisition / Graphic Services

98900

- Please complete unshaded areas on form as applicable.
- Distribute copies as follows: White and Yellow to Graphic Services. Retain Pink Copies for your records.
- On completion of order the Yellow copy will be returned with the printed material.
- Please direct enquiries, quoting requisition number and account number, to extension 3451.

TITLE OR DESCRIPTION

C.S.-88-05

DATE REQUESTED

Nov. 20/89

DATE REQUIRED

ASAP

ACCOUNT NO.

11261894141

REQUISITIONER - PRINT

S. CHRISTODOULAKIS

PHONE

2192

SIGNING AUTHORITY

[Signature]

MAILING INFO -

NAME

S. DEANGELIS

DEPT.

C.S.

BLDG. & ROOM NO.

DC 2314

☐ DELIVER

☐ PICK-UP

Copyright: I hereby agree to assume all responsibility and liability for any infringement of copyrights and/or patent rights which may arise from the processing of, and reproduction of, any of the materials herein requested. I further agree to indemnify and hold blameless the University of Waterloo from any liability which may arise from said processing or reproducing. I also acknowledge that materials processed as a result of this requisition are for educational use only.

NUMBER OF PAGES 14 NUMBER OF COPIES * 20

TYPE OF PAPER STOCK

☒ BOND ☐ NCR ☐ PT. ☒ COVER ☐ BRISTOL ☒ SUPPLIED ☐

PAPER SIZE

☒ 8 1/2 x 11 ☐ 8 1/2 x 14 ☐ 11 x 17 ☐

PAPER COLOUR

☒ WHITE ☐ ☒ BLACK ☐

PRINTING

☐ 1 SIDE ☐ PGS. ☒ 2 SIDES ☐ PGS. FROM ☐ TO

BINDING/FINISHING

☒ COLLATING ☒ STAPLING ☐ HOLE PUNCHED ☐ PLASTIC RING

FOLDING/PADDING

CUTTING SIZE

Special Instructions

Math fronts & backs enclosed

COPY CENTRE

OPER. NO. BLDG. NO. MACH. NO.

DESIGN & PASTE-UP

OPER. NO. TIME LABOUR CODE
☐ ☐ ☐ D 0 1
☐ ☐ ☐ D 0 1
☐ ☐ ☐ D 0 1

TYPESETTING

QUANTITY

P A P 0 0 0 0 0 0 T 0 1
P A P 0 0 0 0 0 0 T 0 1
P A P 0 0 0 0 0 0 T 0 1

PROOF

P R F
P R F
P R F

NEGATIVES

QUANTITY

OPER. NO.

TIME

LABOUR CODE

F L M										C 0 1
F L M										C 0 1
F L M										C 0 1
F L M										C 0 1
F L M										C 0 1

PMT

P M T										C 0 1
P M T										C 0 1
P M T										C 0 1

PLATES

P L T										P 0 1
P L T										P 0 1
P L T										P 0 1

STOCK

										0 0 1
										0 0 1
										0 0 1
										0 0 1

BINDERY

R N G										B 0 1
R N G										B 0 1
R N G										B 0 1
M I S 0 0 0 0 0 0										B 0 1

OUTSIDE SERVICES

COST \$

TAXES - PROVINCIAL ☐ FEDERAL ☐ GRAPHIC SERV. OCT. 85 482-2

**OPTICAL MASS STORAGE SYSTEMS
AND THEIR PERFORMANCE**

S. Christodoulakis, K. Elliott, D. Ford,
K. Hatzilemonias, E. Ledoux, M. Leitch, R. Ng

Research Report CS-88-05

Department of Computer Science
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

Optical Mass Storage Systems and their Performance¹

S. Christodoulakis, K. Elliott, D. Ford, K. Hatzilemonias, E. Ledoux, M. Leitch, R. Ng

Department of Computer Science
University of Waterloo
Waterloo, Ontario N2L 3G1 Canada

1 Introduction

The MINOS project at the University of Waterloo aims at the development of multimedia data base management systems. These systems regularly have as a major resource, information which is not formatted (text, graphics, bitmaps, voice, animation, etc.)

Due to the very large storage requirements of these data types and the archival nature of the applications involved, optical discs are currently the most appropriate storage devices. As part of the MINOS project we study performance aspects of various optical disk based architectures, and we are designing and implementing a high performance server based on optical disk technology. In this paper we summarize our current research activities.

2 ODAS: An Optical Disk Based Archival System

WORM (Write Once Read Many times) storage is secondary memory in which registration of information is permanent (i.e. done once), but retrieval can be done an unlimited number of times. It differs from CD ROM, another, currently popular, type of permanent optical disk storage, in that it has greater storage capacity (atleast twice) and that information is written on the disk by the user, not by the manufacturer during pressing as is the case for CD ROM. It can also differ in the format of the recorded tracks, WORM disks can be either CAV (Constant Angular Velocity, concentric rings) or CLV (Constant Linear Velocity, single spiral track). CD ROM is only available in the CLV format.

The low cost per bit of stored information available when using WORM optical disks, their large

¹to appear in IEEE Database Engineering March 1988

documents. Additionally, index maintenance packages may be interfaced with ODAS to implement content addressability.

The design of ODAS accommodates the peculiarities of the WORM storage and exploits its properties in order to provide a high performance system. The span access capability of optical disks increases the importance of physically clustering data which will be frequently accessed together. Furthermore, if more than one optical disk drive is available, concurrent retrieval can enhance efficiency.

A declarative, high-level script language can be used by application designers to specify rules for data clustering based on properties of document components. Using this facility, information about known access patterns can be utilized to improve retrieval performance for specific applications [10].

Since the retrieval performance of magnetic disks is generally better than that of optical disks, the use of magnetic disk space for temporary buffering of incoming documents improves retrieval speed for documents which tend to be retrieved often soon after they are archived. Periodic flushing of document buffers in batches onto optical disk is done when buffers become full, or during system idle times (to alleviate contention and minimize seeks).

A first version of ODAS was completed at the University of Waterloo [9]. The software consists of approximately 50,000 lines of C source code, developed on Sun 3 workstations (also ported to a VAX 8650). The current implementation of ODAS includes a hardware driver for an Alcatel-Thomson Gigadisc WORM drive [1]. Additionally, two WORM disk simulators of varying complexity and properties have been implemented. Simulated disk drives may be intermixed transparently with real hardware.

3 The Distributed Testbed

The *Distributed Testbed* is a client-server based distributed system. The system features a high performance optical disk based multimedia document server and is intended to support a distributed version of MINOS. More importantly, it provides a testing ground to explore the specific requirements of supporting optical disks, large multimedia documents, and interactive users.

The testbed distributes the basic functionality of *ODAS*, an optical disk based document filing system, across a set of workstations interconnected by a local area network. The testbed consists of two major software subsystems, the *Storage and Retrieval Subsystem* and the *Client Subsystem* (Figure 2). Each subsystem consists of an administrator process and several specialized worker processes [13]. Within a subsystem, the administrator is typically concerned with overall scheduling while the workers perform device and communication management.

The *Storage and Retrieval Subsystem* is the testbed's document server. The server receives requests from clients and replies to them. To retrieve a document, a client sends a retrieval request specifying a document to the server, the request is processed and the document is then returned to the client. To archive a document, the client sends an archival request accompanied by the document. The server batches the document on magnetic disk and the client is free to proceed. The decision to actually archive the document on the optical disk is at the discretion of the server.

The server has two specialized worker processes to manage the optical and magnetic disk storage. The optical disk storage manager interacts with *ODAS* to perform all optical disk requests. The magnetic disk storage manager provides a set of magnetic disk storage caches. The magnetic disk is available for the batching of documents to be archived. It is also available for the prefetching of documents to be retrieved.

The *Client Subsystem* manages user interaction with the testbed server. Each user in the system corresponds to a single instance of the client subsystem, with the potential for multiple clients per workstation. Client workstations may be optionally equipped with magnetic disk storage. A user interacts with the *Client Subsystem* through a set of exported operations through which requests are generated while architecture and communication details are obscured.

Communication within the testbed has severe performance and design implications and may be broken down into communication within subsystems and communication between subsystems. Within a subsystem, a set of reliable interprocess communication primitives is used. They are based upon the blocking send-receive-reply model and are used to relay all information between the processes.

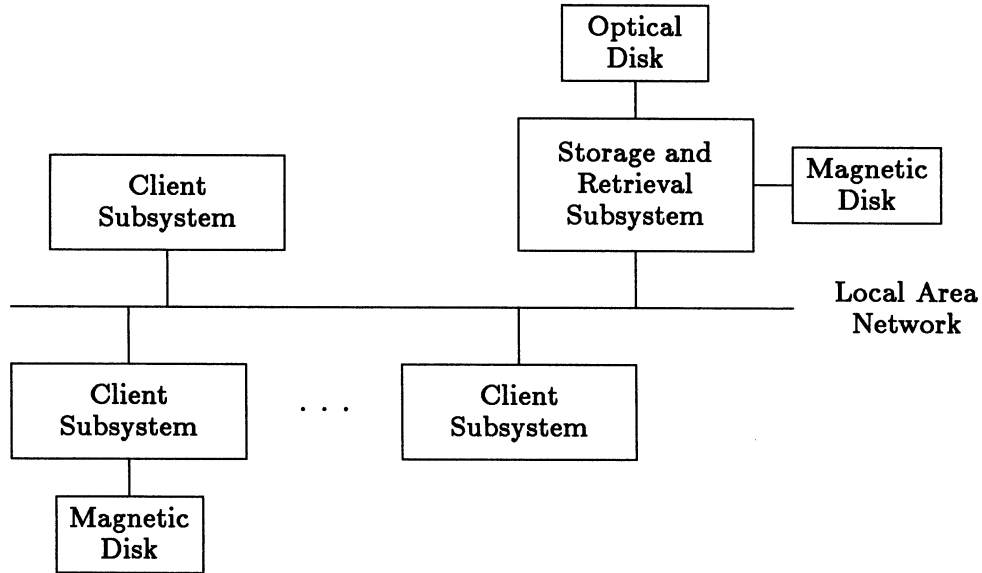


Figure 2: The Distributed Testbed System Architecture

Communication between subsystems may be divided into control messages and data messages. Messages containing control information such as user requests and high level acknowledgements are communicated between subsystems using the reliable interprocess communication primitives. All control information sent between subsystems passes through a set of transporter processes. Each subsystem has a receive transporter and a reply transporter. A receive transporter receives a message sent to the subsystem and relays it to the administrator. The reply transporter receives a reply from the administrator and relays it to a specified destination subsystem. Through transporters, an administrator may communicate with other subsystems without blocking.

Multimedia data messages move between subsystems using an end-to-end window based communication protocol[14]. This approach eliminates the overhead associated with the interprocess communication primitives and also bypasses the intermediary transporter processes. This optimized communication path is critical to system performance due to the large amounts of data that can be associated with multimedia documents. This path is made possible by setting up an end-to-end connection between subsystems, analogous to a high-bandwidth bus.

Although interaction is required between the server and its clients, a significant effort has been made to abstract the data structures and algorithms of the testbed. This so-called testbed approach

supports experimentation within the system. An abstraction may be altered or replaced without affecting other areas of the testbed.

A major area of experimentation is the scheduling problems associated with implementing an optical disk based multimedia document server. To this end, optical disk requests flow through a hierarchy of three schedulers. The long term scheduler receives client requests and enters them into the intermediate term scheduler. The intermediate term scheduler maps a client request to a set of individual page requests on a particular device to be entered into a short term schedule. The short term schedule consists of the current set of optical disk operations to be performed.

The testbed approach also supports the embedding of instrumentation for the purpose of performance monitoring. Through performance monitoring, scheduling policies and other design decisions may be evaluated. The current implementation provides a response time monitoring facility providing a metric for overall system performance.

An enhancement, is the development of a query processing facility based upon signature techniques. This requires the introduction of a *Query Evaluation Subsystem* that searches signature files retrieved from the server. The user will be returned identifiers for documents meeting conditions specified in the query. Document miniatures may also be returned to enable the user to quickly browse through a set of documents and select a set of documents according to some criteria. Query processing further implies the support of user interrupts to terminate queries.

The first version of the *Distributed Testbed* has been completed at the University of Waterloo. The system involves a set of SUN 3 workstations communicating over an Ethernet local area network. The document server runs on a workstation equipped with an Alcatel-Thomson WORM optical disk drive and Fujitsu magnetic disks.

The multimedia document management subsystem of MINOS is going to be ported as an application on top of the system. The functionality of the testbed is expected to grow over time as greater performance and user requirements are imposed.

4 Optical Disk Simulators and Emulators

Presently, both WORM and CD ROM optical disks are not rewritable storage media, and in the case of CD ROMs, the mastering process is lengthy and expensive.

To ensure optimal retrieval performance in a certain application environment, careful planning is required during the application design stages, regarding the placement of data on the optical disk and thereafter, the scheduling of queries.

There are two methods to aid the process of examining performance implications of alternative data placement and query scheduling algorithms. The first, employs analytical tools to produce exact results and at times can be difficult and/or prove expensive. In such cases, approximations may provide more efficient but sometimes, due to the necessary simplifications involved of the model under examination, unreliable results. In these situations, the only means for reliable experimentation is simulation.

The **Optical Disk Simulator (ODS)** package [7] runs on SUN workstations, and uses the storage capacity of one gigabyte of magnetic disk to provide full scale simulation of various WORM and CD ROM disk architectures, as well as simulation of expensive storage devices such as single drive and multidrive jukeboxes of WORMs and CD ROMs.

Simulated optical disk drive architectures include one of two types of reading servomechanisms. Those having an objective lens only, at the end of the laser beam path, and those that also have an adjustable reflecting mirror.

During sector retrievals, as dictated by document retrieval queries under various query scheduling algorithms, estimated delays include jukebox operation costs, long and short seek costs, track address identification costs, rotational delays and transfer costs. Furthermore, simulation of multidrive jukeboxes allows parallelism during reading, resulting in reduced transfer times.

The ODS system has been extensively tested and is presently validated against the Hitachi CDR-1503S CD ROM drive and the Alcatel-Thomson GIGADISC WORM drive.

To date, the ODS system has been used for experimentation related to the retrieval performance of multimedia documents with different storage requirements, for experimentation with various

query scheduling algorithms, for investigation of the performance characteristics of each of the two reading servomechanisms, for investigation of the various implications the span size of the access mechanism has on performance and for deriving important implications based on the property differences of WORM and CD ROM optical disks. The ODS system has also been used to validate exact and approximate analytical results in the above areas.

Optical Disk Emulator System (ODES) emulates the placement and collation of actual or randomly generated multimedia documents on simulated WORM and CD ROM optical disk and drive architectures, and enforces the time delays involved (jukebox operation costs, seek costs, address identification costs, rotational delay costs, transfer costs) when retrieval queries from different clients in a distributed environment are posed on the data.

The ODES system incorporates a number of software facilities, developed by members of the Minos group. The client management and data prefetching facilities of the Minos Distributed Testbed, the various placement and retrieval facilities of the WORM Optical Disk Based Document Filing System (ODAS) and the query scheduling, CD ROM and WORM drive, disk and jukebox simulation and retrieval cost prediction facilities of the Optical Disk Simulator package (ODS).

5 Analysis

Optical disk technology takes two basic forms, that of *Constant Angular Velocity* (CAV) disks of which WORM type disks are an example, and *Constant Linear Velocity* (CLV) disks of which CD ROM is an example. CAV disks rotate at a constant rate but have a variable storage density, while CLV disks rotate at a variable rate and have a constant storage density. The two forms embody trade-offs of speed and storage. CAV disks have faster access times but typically lower storage capacities than their counterpart CLV disks.

Some of the work performed at the University of Waterloo has been to develop basic models and analytic tools for studying the retrieval performance of both types of optical disks [6]. These tools are expected to prove useful for data base design decisions as well as for performance predictions of various file organizations. They may also be of value in the development of query optimization

schemes.

The retrieval performance of optical disks is generally determined by the time required to move the access mechanism (the seek time). A typical delay for a CAV WORM drive ranges from 100msec to 300msec, compared to that of a magnetic disk of 30msec. The issue is complicated slightly however, by the fact that some optical disk drives are capable of reading from more than one track without moving the access mechanism (and incurring the seek delay). The access mechanisms of such optical disk drives are equipped with an adjustable mirror that allows slight deflections of the beam of laser light used to read the data over a small set of tracks immediately beneath the access mechanism. This set of tracks is called a *span* and the number in the set is called the *span size*. Typical values of the span size are 10, 20 and 40 tracks. A span size of 40 tracks corresponds roughly to about 1 megabyte of storage on a CAV WORM disk. This capability can be likened to cylinders on magnetic disk packs, except that spans on an optical disk drive can be overlapping.

Our analysis concentrates upon determining the expected number of *spans* required to access a given number of objects for drives with or without a deflecting mirror in the access mechanism.

5.1 Models

The models used in the analysis of both types of optical disks are similar. The distinction between the two is made when considering the impact of track capacities on retrieval performance. The track capacity of a CAV type disk is constant. Because of the spiral recording scheme used, the capacity of tracks on a CLV type disk varies depending upon their position on the disk surface, tracks at the outer edge hold more data than those at the inner edge.

An optical disk is a device composed of T ordered tracks, an access mechanism and a viewing mechanism. A track is represented by its sequence number i within the tracks of the device, $i = 1, 2, \dots, T$ (to avoid discussion about boundary conditions we assume that track numbers are extended above T and below 1). Each track is composed of a number of sectors (or blocks). The access mechanism can be positioned at any track. When the access mechanism is positioned at a certain track i , the device can read data which completely exist within Q consecutive tracks (track i is one of them). In order to do that, the viewing mechanism *focuses* to a particular track with

qualifying data (within the Q tracks). We call the Q consecutive tracks a *span* and this capability of optical disks, *span access capability*. An anchor point a of a span, is the smallest track number within a span. The largest track number within this span is $a + Q - 1$. The anchor point of a span completely defines the tracks of the span. A span can therefore be described by its anchor point.

A number N of objects (or records) may qualify in a query. N is called the object selectivity (or record selectivity) of the query. In the case of optical disks the number of times that the access mechanism has to be moved for accessing the data which qualifies in a query is called the *span selectivity* of the query. Therefore, a first approximation of the cost of evaluating a query is given by the span selectivity of the query.

When the access mechanism is moved a seek cost and a rotation delay cost are incurred as in the case of magnetic disks. The seek cost depends on the distance travelled by the access mechanism. Transferring a track of data from the device involves a track transfer cost as is also the case in magnetic disks. When more than one track within a given span has to be transferred to main memory, the access mechanism does not have to be moved because of the deflecting mirror. However, there is a small additional delay (of the order of one millisecond) involved for focusing the viewing mechanism on each additional track (within a span) that has to be read. We call this delay a viewing cost and will ignore it for the remaining part of this paper. This model reduces to a model of magnetic disks when the number of tracks in a span is one and the track size is equal to the cylinder size.

A spiral scheme is generally used on CLV optical disks for the storage of data rather than the concentric rings found on CAV disks. The spiral consists of one long physical track of data recordings that begins at a distance from the centre of the disk called the *principal radius*, and continues until close to the outer edge of the disk where it terminates.

A track for our purposes starts from the intersection of the spiral with the radial line that begins at the centre of the disk and passes through the very beginning of the physical spiral. The track ends at the next intersection with this radial line. Since sectors are of fixed length, the radial line may intersect the body of a sector. By convention, the first sector of a track is the first whole

sector encountered after the intersection of the radial line and the track.

5.2 Analytic Results

We have succeeded in deriving results for predicting the retrieval performance of both types of disks. We have produced both exact and approximate analyses for a variety of object (record) sizes and configurations.

For CAV type optical disks, we derived estimates of the number of spans (movements of the access mechanism) required to retrieve small and large objects. For small objects, the two cases of records being allowed and not allowed to cross sector boundaries were examined. For large objects such as bit maps and digitized audio, which could require several hundred kilobytes or even megabytes of storage, we have produced results for their retrieval performance for the case where only large objects exist in a file and also for the case where small objects may exist.

For CLV type optical disks, similar exact and approximate results were derived. It was shown that for the case of small objects, that the corresponding CAV results were an excellent approximation. Such was not the case for larger objects however, but we were able to derive two results, one based upon the expected distribution of objects among the variable capacity tracks and another based upon multiple CAV solutions [6].

5.3 File Placement

We also examined the implications of the variable track capacities found on CLV disks, on the placement of files on the disk surface (i.e. into sets of different capacity tracks). We found that files which are expected to experience the greatest frequency of random access should be placed closer to the outer edge (larger capacity tracks) of the disk than those files that are expected to experience fewer random accesses.

References

- [1] *GC 1001 OEM Manual*. ATG Gigadisc, April 1987.
- [2] A. Bell. Critical issues in high density magnetic and optical storage. In *Proc. of the SPIE*, January 1983.

- [3] S. Christodoulakis. Issues in the architecture of a document archiver using optical disk technology. In *Proc. of ACM SIGMOD '85*, pages 34–50, May 1985.
- [4] S. Christodoulakis. Query processing in optical disk based multimedia information systems. *IEEE Database Engineering*, October 1987.
- [5] S. Christodoulakis and C. Faloutsos. Design considerations for an optical-disk based multimedia object server. *IEEE Computer*, December 1986.
- [6] S. Christodoulakis and D. Ford. Performance analysis and fundamental performance trade-offs for CLV optical disks. December 1988. Technical Report.
- [7] S. Christodoulakis and K. Hatzilemonias. Optical Disk Simulation package (ODS). *Technical Report, University of Waterloo*, November 1987.
- [8] S. Christodoulakis, F. Ho, and M. Theodoridou. The multimedia object presentation manager of MINOS: a symmetric approach. In *Proc. of ACM SIGMOD '86*, pages 295–310, May 1986.
- [9] S. Christodoulakis, E. Ledoux, and R. Ng. *A WORM Optical Disk Based Document Filing System*. Technical Report, University of Waterloo, December 1987.
- [10] S. Christodoulakis and T. Velissaropoulos. Architectural aspects of MINOS: a distributed multimedia information system. *IEEE Office Knowledge Engineering*, 1(1):18–27, February 1987.
- [11] G. Copeland. What if mass storage were free? *IEEE Computer*, 27–35, July 1982.
- [12] L. Fujitani. Laser optical disk: the coming revolution in on-line storage. *CACM*, 27(6):546–554, June 1984.
- [13] W.M. Gentleman. Message passing between sequential processes: the reply primitive and the administrator concept. *Software - Practice and Experience*, 11, 1981.
- [14] J.H. Saltzer, D.P. Reed, and D.D. Clark. End-to-end arguments in system design. *ACM Trans. on Computer Systems*, November 1984.

FEB 15 1989

88-05
out of
stock
gave Stavros a
note Nov. 20/89

Mr H. K. Lim
Fallener-Eggington Courts
Loughborough University
of Technology
Leics LE11 3TU
7th Feb 1989

University of Waterloo
Department of Computer Science
Waterloo, Ontario
Canada N2L 3G1

Dear Sir / Madam

Re: Optical mass storage systems and their performance
(Technical Report CS-88-05)
AND

Performance analysis for CLV optical discs
(Technical Report CS-88-06)

I will be grateful if you could send me copies of
the above two reports as soon as possible.

Thank you.

Yours faithfully

H.K. Lim

sent
with
compliments
Feb. 16/89



Southern Alberta
Institute of Technology
1301 - 16 Avenue N.W.
Calgary, AB. Canada T2M 0L4

July 14, 1988

Department of Computer Science
University of Waterloo
Waterloo, Ontario N2L 3G1

Dear Sir/Madam:

Please find enclosed a money order in the amount of \$4.00 which will cover the cost of Research Reports 1988 #CS-88-05 and CS-88-21.

Please forward to Southern Alberta Institute of Technology
Information Systems Division
G200, 1301 - 16th Avenue N.W.,
Calgary, Alberta
T2M 0L4

Attention: Terry Smith, Director

Thank you,

Meg Fox

Meg Fox
Administrative Assistant

*recd & sent reports
personal money order
#1732962
Credit Union Central
July 21*

Printing Requisition / Graphic Services

15053

1. Please complete unshaded areas on form as applicable.

2. Distribute copies as follows: White and Yellow to Graphic Services. Retain Pink Copies for your records.

3. On completion of order the Yellow copy will be returned with the printed material.

4. Please direct enquiries, quoting requisition number and account number, to extension 3451.

TITLE OR DESCRIPTION

CS-88-05

DATE REQUISITIONED

July 13

DATE REQUIRED

July 14

ACCOUNT NO.

1126894141

REQUISITIONER - PRINT

S. Christodoulakis

PHONE

4452

SIGNING AUTHORITY

BLDG. & ROOM NO.

☐ DELIVER

☒ PICK-UP

MAILING INFO -

NAME

DEPT.

Copyright: I hereby agree to assume all responsibility and liability for any infringement of copyrights and/or patent rights which may arise from the processing of, and reproduction of, any of the materials herein requested. I further agree to indemnify and hold blameless the University of Waterloo from any liability which may arise from said processing or reproducing. I also acknowledge that materials processed as a result of this requisition are for educational use only.

NUMBER OF PAGES 14 NUMBER OF COPIES 25

TYPE OF PAPER STOCK

☒ BOND ☐ NCR ☐ PT. ☐ COVER ☐ BRISTOL ☐ SUPPLIED ☐

PAPER SIZE

☒ 8 1/2 x 11 ☐ 8 1/2 x 14 ☐ 11 x 17 ☐

PAPER COLOUR

☒ WHITE ☐ ☒ BLACK ☐

PRINTING

☐ 1 SIDE PGS. ☒ 2 SIDES PGS.

NUMBERING

FROM TO

BINDING/FINISHING

☒ COLLATING ☐ STAPLING ☐ HOLE PUNCHED ☐ PLASTIC RING

FOLDING/PADDING

CUTTING SIZE

Special Instructions

COPY CENTRE

OPER. NO. BLDG. NO. MACH. NO.

DESIGN & PASTE-UP

OPER. NO. TIME LABOUR CODE

TYPESETTING

QUANTITY

P A P 0 0 0 0 0 0 T 0 1

P A P 0 0 0 0 0 0 T 0 1

P A P 0 0 0 0 0 0 T 0 1

PROOF

P R F

P R F

P R F

NEGATIVES

QUANTITY

OPER. NO.

TIME

LABOUR CODE

F L M C 0 1

F L M C 0 1

F L M C 0 1

F L M C 0 1

F L M C 0 1

F L M C 0 1

F L M C 0 1

F L M C 0 1

P M T

P M T C 0 1

P M T C 0 1

P M T C 0 1

P M T C 0 1

P M T C 0 1

PLATES

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

P L T P 0 1

TAXES - PROVINCIAL ☐ FEDERAL ☐ GRAPHIC SERV. OCT 85 482-2

To *Stavros*

From *Sue DeAngelis*

Date *June 21*

memo

University of Waterloo

*Could you please
advise as to what
reports you would
like me to send.*

Thanks.

*sent
88-05*

Sue

JUL 5 1988

*R: Only the technical
report on optical
mass storage system
& their performance*

5/25/88

Dear Sir,

Could you please send me a copy of any
available technical reports concerning your
MINOS research prototype? Thanks.

Sincerely,
Joe C. Roberts
5700 Coventry Lane
Austin, TX 78723
U. S. A.

Printing Requisition / Graphic Services

10359

1. Please complete unshaded areas on form as applicable.
2. Distribute copies as follows: White and Yellow to Graphic Services. Retain Pink Copies for your records.
3. On completion of order the Yellow copy will be returned with the printed material.
4. Please direct enquiries, quoting requisition number and account number, to extension 3451.

TITLE OR DESCRIPTION

Optical Mass Storage Systems and Their Performance

CS-88-05

DATE REQUISITIONED

DATE REQUIRED

ACCOUNT NO.

Feb. 22/88

Feb. 26/88

1 2 6 6 0 7 4 4 1

REQUISITIONER - PRINT

PHONE

SIGNING AUTHORITY

S. Christodoulakis

x3688

BLDG & ROOM NO.

☒ DELIVER

MAILING
INFO -

NAME

DEPT.

Sue De Angelis

Computer Science

MC 6081E

☐ PICK-UP

Copyright: I hereby agree to assume all responsibility and liability for any infringement of copyrights and/or patent rights which may arise from the processing of, and reproduction of, any of the materials herein requested. I further agree to indemnify and hold blameless the University of Waterloo from any liability which may arise from said processing or reproducing. I also acknowledge that materials processed as a result of this requisition are for educational use only.

NUMBER OF PAGES **12** NUMBER OF COPIES **60**

TYPE OF PAPER STOCK

☒ BOND ☐ NCR ☐ PT. ☐ COVER ☐ BRISTOL ☐ SUPPLIED ☐

PAPER SIZE

☒ 8 1/2 x 11 ☐ 8 1/2 x 14 ☐ 11 x 17 ☐

PAPER COLOUR

☒ WHITE ☐ ☒ BLACK ☐

PRINTING

☐ 1 SIDE ☐ PGS. ☒ 2 SIDES ☐ PGS. FROM TO

BINDING/FINISHING

3 down left side
☒ COLLATING ☒ STAPLING ☐ PUNCHED ☐ PLASTIC RING

FOLDING/
PADDING

CUTTING
SIZE

Special Instructions

Math fronts and backs enclosed.

COPY CENTRE

OPER. NO. BLDG. NO. MACH. NO.

DESIGN & PASTE-UP

OPER. NO. TIME LABOUR CODE

TYPESETTING

QUANTITY

P A P 0 0 0 0 0 0 T 0 1

P A P 0 0 0 0 0 0 T 0 1

P A P 0 0 0 0 0 0 T 0 1

PROOF

P R F

P R F

P R F

NEGATIVES

QUANTITY

OPER. NO.

TIME

LABOUR CODE

F L M

F L M

F L M

F L M

F L M

PMT

P M T

P M T

P M T

PLATES

P L T

P L T

P L T

STOCK

BINDERY

R N G

R N G

R N G

M I S 0 0 0 0 0 0 B 0 1

OUTSIDE SERVICES

\$ COST

TAXES - PROVINCIAL ☐ FEDERAL ☐ GRAPHIC SERV. OCT. 85 482-2