N94-33813

Simulation of a Data Archival and Distribution System at GSFC

Jean-Jacques Bedet, Lee Bodden, Al Dwyer, P C Hariharan*
Hughes STX Corporation
7701 Greenbelt Road, Suite 400
Greenbelt MD 20770
jbedet@nssdca.gsfc.nasa.gov

John Berbert, Ben Kobler, Phil Pease
NASA/GSFC
Greenbelt MD 20771
berbert@nssdca.gsfc.nasa.gov
kobler@nssdca.gsfc.nasa.gov

Abstract

A version-0 of a Data Archive and Distribution System (DADS) is being developed at the Goddard Space Flight Center (GSFC) to support existing and pre-EOS Earth science datasets and test Earth Observing System Data and Information System (EOSDIS) concepts. The performance of the DADS is predicted using a discrete event simulation model. The goals of the simulation were to estimate the amount of disk space needed and the time required to fulfill the DADS requirements for ingestion (14 GB/day) and distribution (48 GB/day). The model has demonstrated that 4 mm and 8 mm stackers can play a critical role in improving the performance of the DADS, since it takes, on average, 3 minutes to manually mount/dismount tapes compared to less than a minute with stackers. With two 4 mm stackers and two 8 mm stackers, and a single operator per shift, the DADS requirements can be met within 16 hours using a total of 9 GB of disk space. When the DADS has no stacker, and the DADS depends entirely on operators to handle the distribution tapes, the simulation has shown that the DADS requirements can still be met within 16 hours, but a minimum of 4 operators per shift were required. The compression/decompression of data sets is very CPU intensive, and relatively slow when performed in software, thereby contributing to an increase in the amount of disk space needed.

Introduction

The Goddard Space Flight Center (GSFC) is building a Version 0 Distributed Active Archive Center (V0 DAAC) to support pre-EOS projects and test Earth Observing System Data and Information System (EOSDIS) concepts. This system will consolidate management and provide access, archiving, and distribution functions for Goddard's Earth Science data. This paper describes a study of the performance of one of the elements of the DAAC; the Data Archive and Distribution System (DADS). The DADS is responsible for the ingestion, archiving and distribution of pre-EOS data. To assess the storage needs and performance capability of the DADS, a discrete event simulation model has been developed using the NASA Data Systems Dynamic Simulator (DSDS) package. This study has identified potential bottlenecks in the utilization of the selected ingest, archival, and distribution devices (on-line disks, automated tape libraries, jukeboxes, and magnetic tape drives), and has identified the performance benefits to be gained by adding one or more stackers to the 4 mm and 8 mm tape drives.

* Present address: Systems Engineering and Security, Inc.
7474 Greenway Center Drive, Suite 720
Greenbelt, MD 20770

The GSFC DADS is expected to ingest 14 GB/day of data and distribute an estimated 48 GB/day of data over various media (4 mm, 8 mm, and 9 track tapes) and over the network. With these large volumes of data to be ingested and distributed, the GSFC DADS wanted to assess the amount of staging disk space and the number of tape drives required to meet the estimated DADS workload. To address these issues, a discrete event simulation model of the DADS has been developed using the NASA DSDS package. The model simulates the ingestion of regular and reprocessed data, and the fulfillment of standing orders and user requests for data distribution.

First, the GSFC hardware configuration and the main DADS activities that are simulated are described. A high level view of the DADS model is presented and the results obtained from the model are discussed. The contention for the robots of the Metrum RSS-600 Automated Tape Library (ATL) and the Cygnet optical disk jukebox, and the various tape and disk drives is explained. In particular, we looked into the effect of having human operators in the distribution process and quantify how 4 mm and 8 mm stackers could improve the performance. The impact of using compression and decompression techniques has also been studied. Finally, the lessons learned and future work are summarized in the last paragraph.

V0 GSFC DADS Configuration

First we examine the storage devices used to ingest, archive, and distribute data.The current hardware configuration of the V0 GSFC DADS, as of August 1993, is illustrated in Fig. 1.

Ingestion

Most of the data to be ingested at the GSFC DADS is received over an FDDI network (100 Mbits /s) and copied to Unix staging disks (2.7 MB/s). The ingestion operation is performed overnight to minimize the impact on the network. A small amount of data is received on 3480 cartridges.

Archival

To automate the archival and retrieval process, the GSFC DADS has acquired a Cygnet 1803 jukebox with 2 ATG WORM drives and an RSS-600 Metrum Automated Tape Library (ATL) with 4 RSP 2150 VHS drives. Based on the data type, a data set is either stored on the Cygnet jukebox, which can hold up to 131 12" WORM platters (9 GB per platter), or on the Metrum ATL which can accommodate 600 magnetic T120 VHS cassettes (14.5 GB per cassette). The file management is controlled by Unitree 1.7, which is running on an SGI 4D/440 workstation. Files are automatically migrated from the Unitree magnetic disk cache, which holds 13.8 GB, to either the jukebox or the ATL. Similarly requests for data already residing on the jukebox or the Metrum ATL are handled by Unitree, which retrieves the data and puts them in its cache. Table 1 provides the specifications of the two archive devices selected for the DADS.

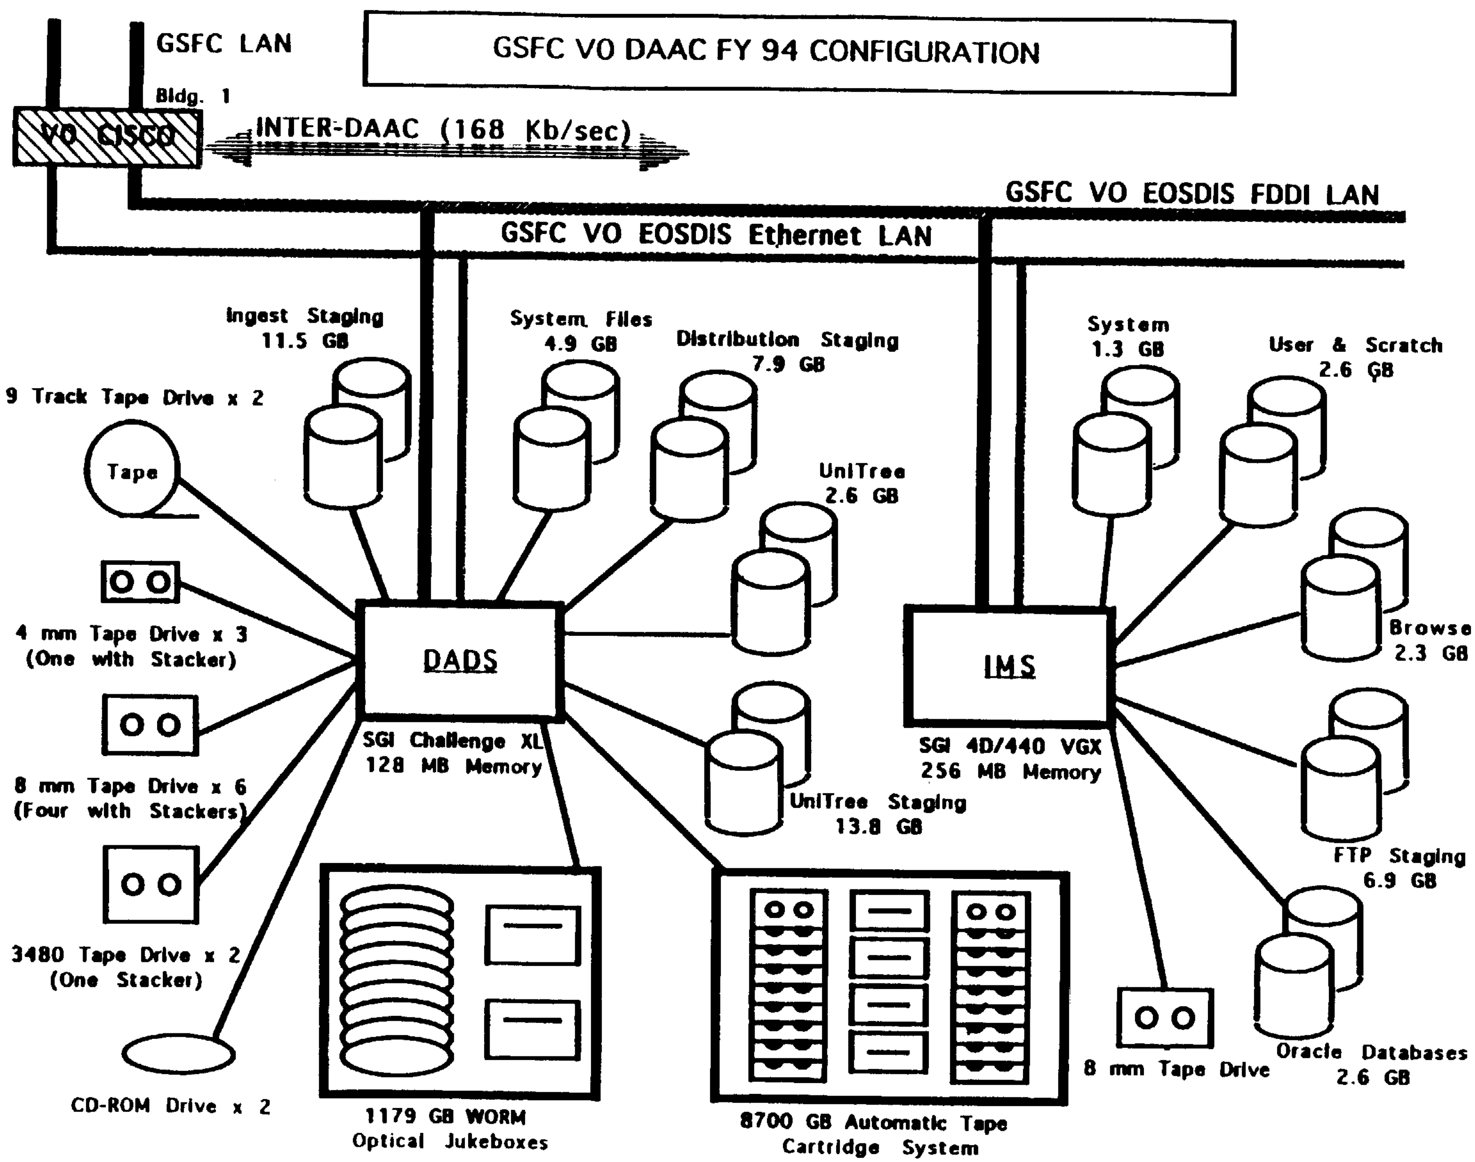

Fig 1. GSFC V0 DAAC Configuration

	1803 Cygnet jukebox	RSS-600 Metrum ATL
Media used	12" WORM platters	T120 VHS tape cassettes
Drive type	ATG WORM	Metrum RSP 2150 VHS
# of drives available	2	4
drive read/write rate (MB/s)	0.5	1
Media capacity (GB)	9 (4.5 GB/side)	14.5 (T120), 16 (T160)
Number of media	up to 131	600
System capacity (GB)	1179	8700 (T120), 9600 (T160)
Number of robot arms	1	1
Avg robot access time(s)	8	8

Table 1. Specification of DADS archive devices

Distribution

It is expected that most data will be requested on 8 mm and 4 mm cassettes. To automate the distribution process, the DADS has an 8 mm stacker and is investigating the possibility of purchasing additional 4 mm and 8 mm stackers. For users who still need their data on 6250 bpi tapes, the DADS has two 9 track drives. For quick delivery and for small files, the data may also be sent over the network. The characteristics of the distribution devices are summarized in Table 2.

	4 mm DAT	8 mm Exabyte (8500)	9 track drive
Number of drives	3	4	2
Manual fetch time (min)	1 or 3	1 or 3	1 or 3
Stacker fetch time (s)	60	60	N/A
Load time (s)	14	42	60
Unload time (s)	10	21	20
Manual return time (min)	1 or 3	1 or 3	1 or 3
Stacker return time (s)	60	60	N/A
Search rate (MB/s)	13	22.6	0.15
Rewind rate (MB/s)	25	28	1
Read transfer rate (MB/s)	0.17	0.40	0.17
Write transfer rate (MB/s)	0.17	0.43	0.17

Table 2. DADS distribution parameters

DADS Activities Simulated

The two main activities simulated in the model are the ingestion/archival and the distribution. The ingested data are subdivided into two categories: regular processing data and reprocessing data. For both categories the data are first copied to disks (Unix disks), compressed (optional), and transferred to the Unitree cache (referred to as Unitree disks) and then migrated automatically, under the control of Unitree, to the Cygnet jukebox or the Metrum ATL. In the case of ingested data, the metadata containing information about the data, are first extracted before being sent to the Unitree cache. In addition, some of the new regular ingested data are known in advance to be requested for distribution. The data used to satisfy these advance requests (called "standing orders") are kept on-line on the Unix disks until all the standing orders have been fulfilled. For the distribution requests that are not standing orders, the data are retrieved from one of the robotic devices (Metrum ATL or Cygnet jukebox), copied to the Unitree cache, decompressed (optional), staged to the Unix disks, and finally copied to one of the distribution media.

The sequence of actions, for each activity performed at the DADS, is as follows:

Ingestion/Archival

- Write incoming data to Unix disks
- Compress(optional) and copy data to Unitree cache
- If data are used in standing orders
  First complete all standing orders and then delete data from Unix disks
- If data are not used in standing orders
  Delete data from Unix disks
- Migrate data from Unitree cache to robotic devices archive
- Mark file as purgeable from Unitree cache

Distribution (non-standing orders)

- Retrieve data from robotic devices
- Copy data to Unitree cache
- Decompress (optional) and copy data to Unix disks
- Mark file as purgeable from Unitree cache
- Copy data to distribution media
- Delete data from Unix disks

Distribution (standing orders)

- Read staged data from Unix disks
- Write data to distribution media
- Remove staged data from Unix disks

Simulation model

Using DSDS, a model has been developed to simulate the various activities and devices at the DAAC. The block diagram illustrated in Fig 2, has four main components. The first one contains the elements that generate the files to be ingested or distributed. File sizes and inter-arrival times are both randomly computed by the use of appropriate distributions (e.g. uniform). This first component models the expected data volume to be ingested and distributed by the DADS. The second component (initialization), identifies the source and the destination of each file as well as the disk to which the file is temporarily stored. The third component acts as a switch, directing the file to the right device. The fourth component (devices) models the various storage devices and the resource allocation. After leaving the devices component, the step is incremented by the counter and the file is once again directed to the appropriate device by the switch component. This process is repeated until the file reaches the end component.

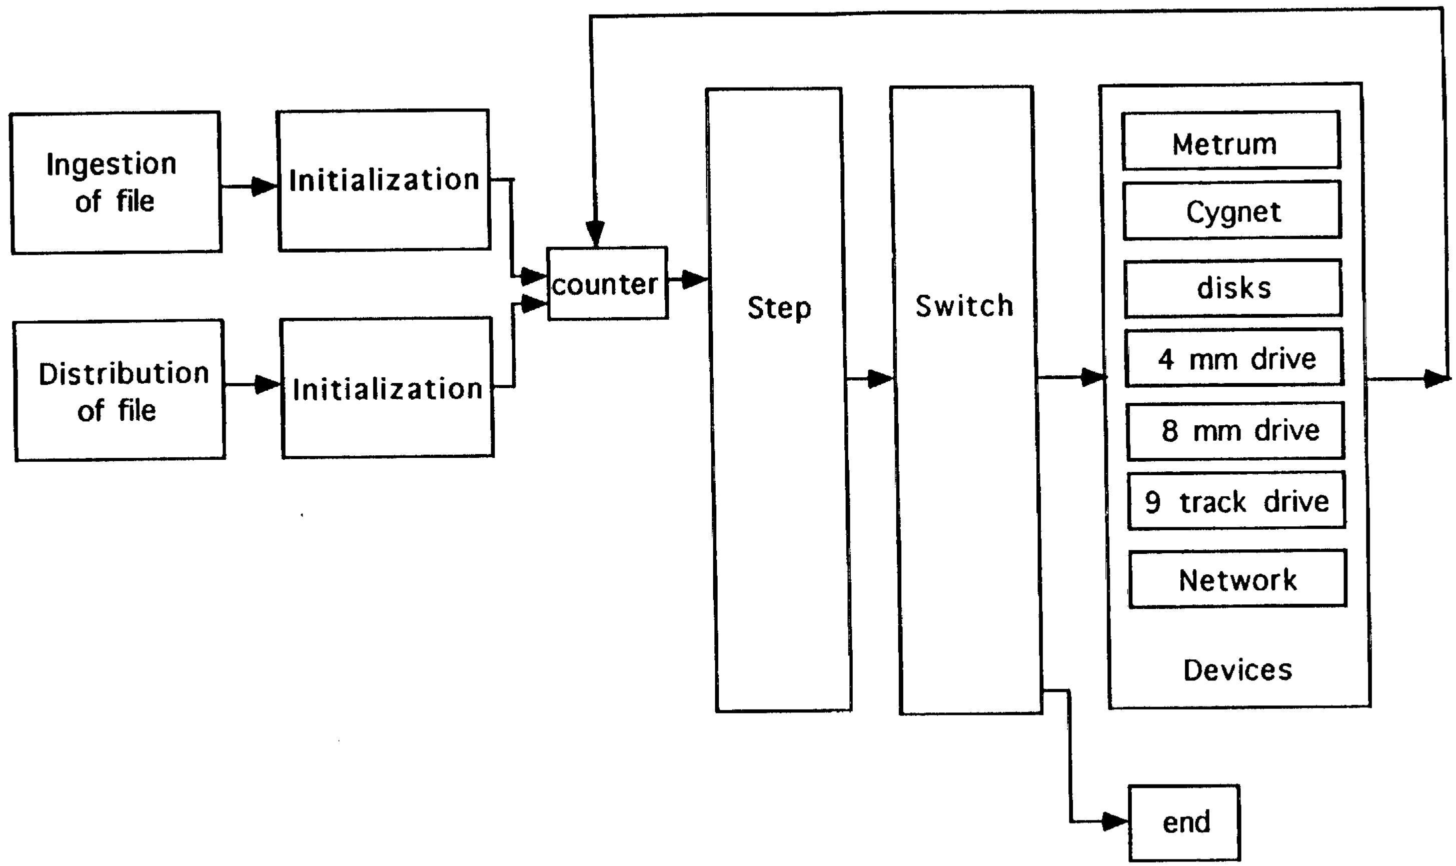

Fig. 2. Schematic Diagram of the DADS Model

DADS Simulation Model Assumptions :

- Any file ingested is first copied to the Unix staging disks and then to the Unitree cache disks for migration into the archive storage devices.
- Any file retrieved from the archive devices for distribution, is first copied to the Unitree cache disks and then to the Unix staging disks for tape copy.
- The simulation allows all the various DADS disk and tape storage devices to have different data transfer rates in read and write modes.
- UNITREE 1.7 supports multiple simultaneous read operations but can only support a single write operation at a time. This Unitree restriction has been implemented in the current version of the model.
- Each file read from or written to the jukebox requires a load and unload of a platter.
- Each file read from or written to the Metrum storage module, requires a load, and unload operation of a cassette.
- When not using a stacker, each file copied to a 4 mm drive, 8 mm drive, or 9 track drive requires a manual fetch of the blank tape to the drive load mechanism and a return of the copied tape.
- In the first two scenarios examined the requests are assumed to be distributed with an equal probability on each of the three types of media (8 mm, 4 mm, and 6250 bpi) in the proportion of 33%, 33%, and 33%. For later scenarios this was changed to 50%, 33%, and 17% after a survey of potential users was made.
- Distribution request files (non-standing orders) are uniformly distributed over 12 hours.
- Ingestion files for the SeaWiFS regular processing are uniformly distributed over 2 hours (except in scenario 2, when the 2 hours is changed to 16 minutes).
- Ingestion files for the SeaWiFS reprocessing are uniformly distributed over 16 hours.
- Ingestion files for the non-SeaWiFS regular processing are uniformly distributed over 6 hours.

DADS WORKLOAD

The largest volume data set to be ingested, archived and distributed by the Goddard DAAC is that from the SeaWiFS project (see Tables 3 and 4). The SeaWiFS project regular processing operation will send 1.59 GB/day to the GSFC DAAC over the network. In addition, periodically, the SeaWiFS project will reprocess all the data and redeliver replacement data at a rate of 8.9 GB/day. The total estimated distribution data volume for SeaWiFS (including the standing orders) is 40 GB/day (see Table 4).

In addition to SeaWiFS data, the GSFC DAAC will also service a number of other projects. These non-SeaWiFS data add 4.23 GB/day of ingest and 7.97 GB/day of distribution. In this report the SeaWiFS regular ingest and non-SeaWiFS ingest have been referred to as ingestion (regular). The workload was modeled to represent these separate categories so as to facilitate model validation with actual measurements of the DAAC operation with SeaWiFS test data.

In the simulation, using the Workload Model for Archive and Distribution of SeaWIFS Data (November 16, 1992), the daily volume of ingested data of each data type, has been estimated and is tabulated in Table 3. This table indicates also the percentage of this volume from each of the two sources to each of the two archive destinations . For instance, SeaWIFS L1A product is expected to have a volume of 694 MB per day. All SeaWIFS L1A data will be received over the network, and will be stored on the Cygnet jukebox. For simplicity, the simulation model assumes that all data ingested is transmitted over the network. This will have the effect of adding 1.23 GB/day to network ingestion which are currently assumed to be ingested by reading 3480 cartridges. Similarly Table 4 represents the workload for the distribution.

Table 3. Ingestion workload

		Source %		Destination %	
Data Type	**Volume (GB/day)**	**% From network**	**% from 3480**	**% to Jukebox**	**% to Metrum**
SeaWifs (regular)					
L1A	0.694	100		100	0
L2	0.461	100			100
L3	0.43	100			100
Total	1.585	100		43.79	56.21
SeaWifs (reprocessing)					
L2	4.61	100			100
L3	4.3	100			100
Total	8.91	100			100
Non-SeaWIFS (regular)					
AVHRR	1	100		0	100
TOVS	0.233		100	100	0
UARS	1	100		0	100
DAAC Climate data	1		100	100	0
CZCS	1	100		100	0
Total	4.233	70.87	29.13	52.75	47.25
Grand Total	**14.728**	**91.63**	**8.37**	**19.87**	**80.13**

Table 4. Distribution workload

Data Type	Volume (GB/day)	Source % % from Disk	% from Jukebox	% from Metrum	Destination % % to 4 mm	% to 8 mm	% to 9 track
SeaWifs							
Global data order	16	100			33	33	33
large chunks	14	50	20.91	29.09	33	33	33
small chunks	8	25	49.80	25.20	33	33	33
level 3	2			100	33	33	33
Total	40	62.5	17.28	20.22	33	33	33
Non-SeaWIFS							
AVHRR	5	80		20	33	33	33
TOVS	0.466	100			33	33	33
UARS	1	50		50	33	33	33
DAAC Climate data	1		100		33	33	33
CZCS	0.5		100		33	33	33
Total	7.966	62.34	18.83	18.83	33	33	33
Grand Total	**47.966**	**62.47**	**17.54**	**19.99**	**33**	**33**	**33**

DADS Performance

In order to estimate the amount of disk space necessary to ingest/archive and distribute data, and to determine the time required to satisfy the daily activities at the DADS, the discrete events model has been run for scenarios with varying assumptions. These are summarized in Table 5:

Table 5. Summary of assumptions by scenarios

	Scenarios 1-A	1-B	2	3	4	5	6	7-A	7-B
Assumptions									
Regular SeaWiFS ingestion (hours)	2	2	0.26	2	2	2	2	2	2
Reprocessing SeaWiFS ingestion (hours)	16	16	16	16	16	16	16	16	16
Non-SeaWiFS ingestion (hours)	6	6	6	6	6	6	6	6	6
Distribution SeaWiFS (hours)	12	12	12	12	12	12	12	12	12
% Distribution on 8 mm tapes	33.3	33.3	33.3	50	50	50	50	50	50
% Distribution on 4 mm tapes	33.3	33.3	33.3	33	33	33	33	33	33
% Distribution on 9 track tapes	33.3	33.3	33.3	17	17	17	17	17	17
Ingestion (GB/day)	14.7	14.7	14.7	14.7	29.4	14.7	14.7	14.7	14.7
Distribution SeaWiFS (GB/day)	25	40	40	40	80	40	40	40	40
Distribution non-SeaWiFS (GB/day)	0	0	0	0	0	0	0	8	8
Number of operators	∞	∞	∞	∞	∞	1-∞	1	1	1
Number of 8 mm stackers	0	0	0	0	0	0	0-2	2	2

Number of 4 mm stackers	0	0	0	0	0	0	0-2	2	2
Compress/Decompress (Y/N)	N	N	N	N	N	N	N	N	Y
Operator Avg. response/fetch time (min)	1	1	1	1	1	3	3	3	3

- Scenario 1-A: Disk space requirement for the ingestion and SeaWiFS standing order distribution.

First, the DADS has been examined ingesting all data over the network and processing the SeaWiFS standing orders. The disk space used on the Unix disks and the Unitree cache is illustrated in Fig 3. During the first two hours, the DADS receives regular SeaWiFS data, migrates them to the archive, and retains a copy of the data (1.6 GB) on Unix disks in order to fulfil the standing orders. All the other ingested data are rapidly migrated to the archive and do not accumulate on the Unix disks. Due to the large volume of standing orders (25 GB) to be copied to slow devices such as 8 mm and 4 mm tape drives, the standing order distribution operation continues up to 10 hours. At that time, the standing orders are completed and regular SeaWiFS data are deleted from disk, creating a big drop in the Unix disk space.

The Unitree disk space is also illustrated in Fig 3 and shows a peak of approximately 400 MB. During the ingestion process, data are migrated to the robotic devices archive as soon as possible. In this scenario the total ingestion rate approximately matches the archival rate (including robotic access times as well as jukebox and Metrum ATL write rates), so that only a small amount of data is retained in the Unitree cache. After migrating the files to the robotic archive devices, they are marked as purgeable in the Unitree cache. Only the non-purgeable files are plotted in the Unitree disk space in Fig 3.

- Scenario 1-B: Disk space requirement adding SeaWiFS non-standing order distribution.

Figure 4 represents the disk space used as a function of time when the SeaWiFS non-standing orders are added to the previous workload. With a non-standing order, the data are first retrieved from the Cygnet jukebox or the Metrum ATL, copied to the Unitree disk cache, and then copied quickly to the Unix disks. After writing the data set to the Unix disks, the space used in the cache is marked as purgeable. The Unitree cache used with non-purgeable files remains small (~400 MB) over time and is similar to the previous case (Fig 3). The daily volume of non-standing SeaWiFS orders to be distributed is quite large (15 GB) and the distribution tape device write rates are rather slow (see Table 2). This creates a bottleneck and the files are staged in the Unix disks for several hours, waiting to be copied to tapes. The Unix disk space used is illustrated in Fig 4 and it shows a peak of 5.5 GB and a sudden drop at 11 hours, when the standing orders are completed. The backlog of requests staged on Unix disks disappears at about 15.5 hours.

- Scenario 2: Effect of ingesting regular SeaWiFS data over a shorter time interval.

In the previous scenarios, the daily SeaWiFS ingestion was assumed to occur over a 2 hour period. A question of interest is to examine the DADS system when the ingestion is performed during a shorter interval of time. Fig 5 illustrates the case of an ingestion over 16 minutes. As expected, the completion of the standing orders, indicated by a sudden drop of 1.6 GB, occurs earlier (9 hours instead of 11 hours). The non-standing orders backlog disappears at 14.5 hours, or about 1 hour sooner than before. The Unitree cache during the first 2 hours is also much larger (1.4 GB). The data are ingested at a rate which exceeds the archival rate to the Metrum ATL and the Cygnet jukebox. This causes the data to be delayed in the Unitree cache. By controlling the ingestion schedule of the data (i.e., spreading it out), it is possible to keep the Unitree cache used at a minimum, but this increases the time required to eliminate the backlog in the distribution operations.

Unix disk space used as a function of time

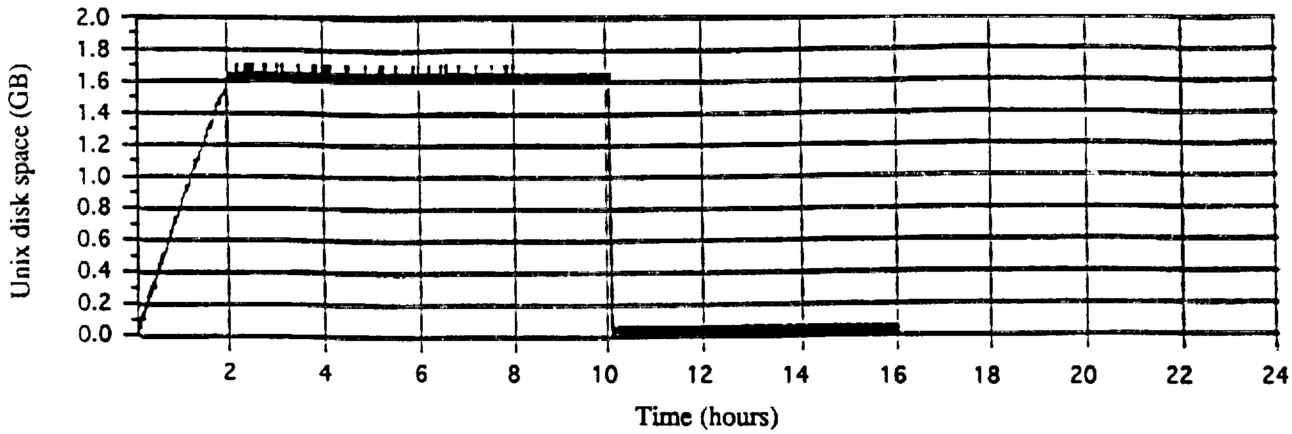

Unitree disk space used as a function of time

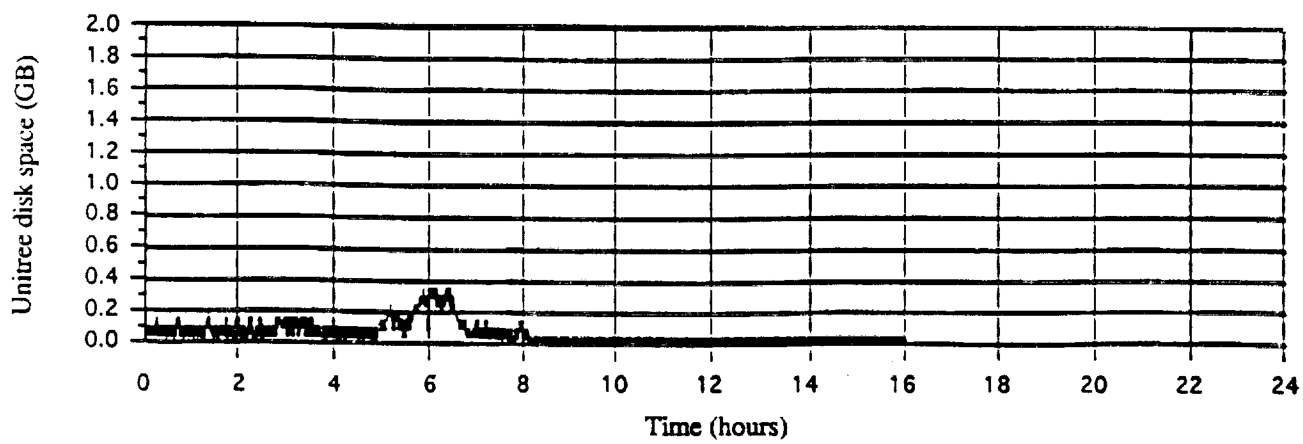

Fig 3. Disk space used in scenario 1-A

Unix disk space used as a function of time

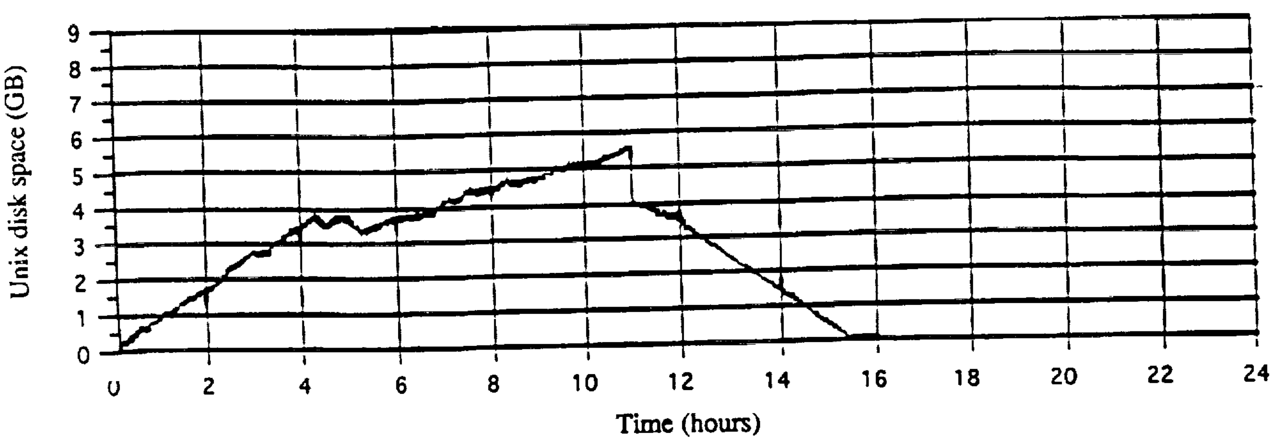

Unitree disk space used as a function of time

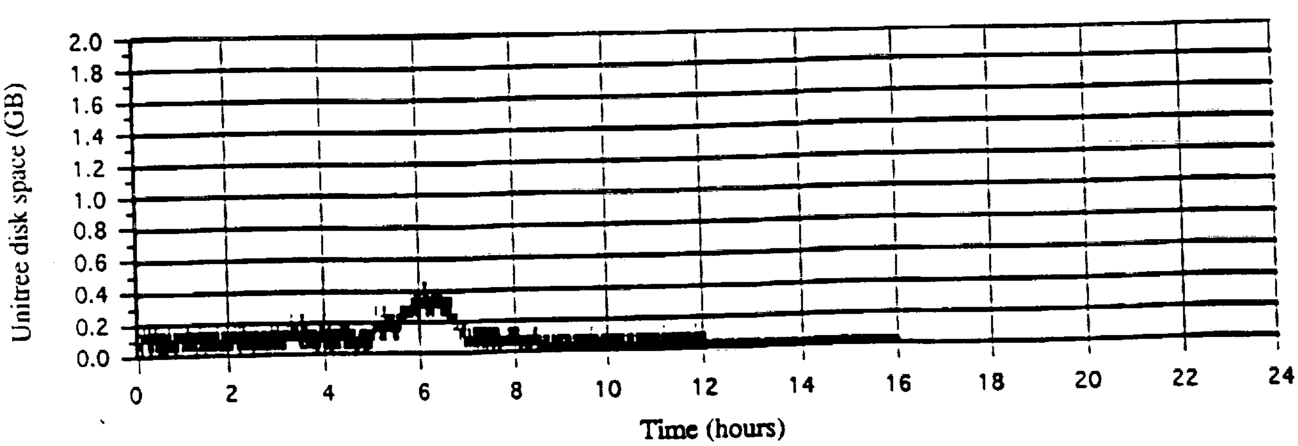

Fig 4. Disk space used in scenario 1-B

Unix disk space used as a function of time

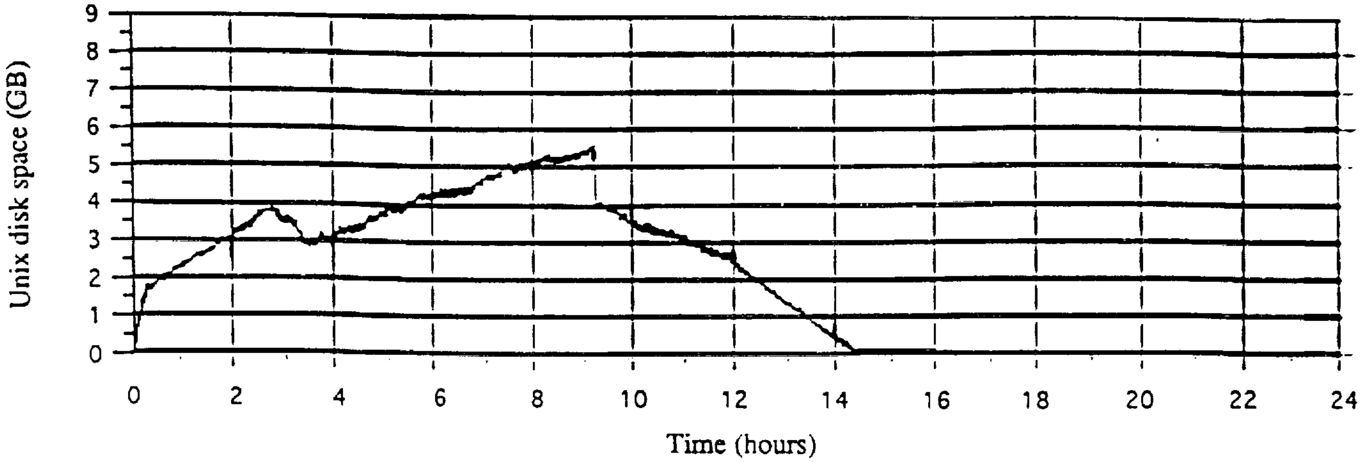

Unitree disk space used as a function of time

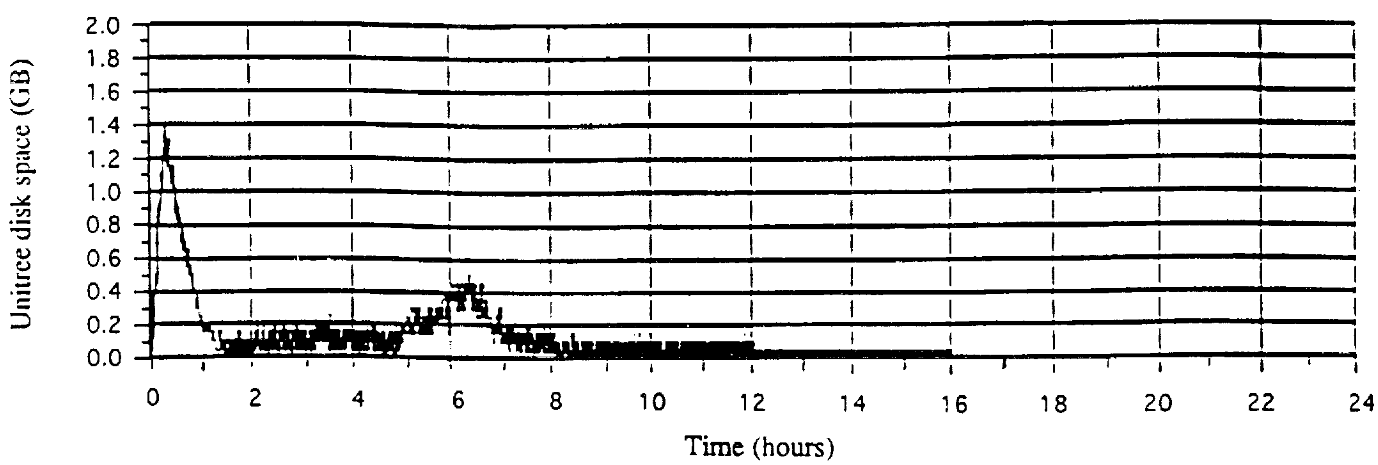

Fig 5. Disk space used in scenario 2

- Scenario 3: Effect of varying the proportion of distribution media requested.

When the model was originally developed, it was assumed that 1/3 of the requests were copied on 8 mm tapes, 1/3 on 4 mm tapes, and 1/3 on 9 track tapes. A small survey of scientists indicated that a more realistic proportion of media requested may be 50% for 8 mm, 33% for 4 mm, and 17% for 9 track tapes. The DADS model has been simulated for the same volume of data ingested and distributed as before (in scenario 1-B) but with the new proportion of distribution media (Fig 6). The change affected only the distribution process and the Unitree disk space remained the same. When compared to scenario 1-B, the Unix disk space needed is smaller (5 GB instead of 5.5 GB) and the distribution requests backlog was eliminated sooner (13 hours instead of 15.5 hours). By changing the proportion of media requests, the volume of data copied to 8 mm drives increased. The 8 mm drive write rates are about 2 to 3 times faster than the 4 mm drives. Consequently it took less time to fulfil the requests and the data were staged in the Unix disks for a shorter period of time.

- Scenario 4: Effect of processing 2 days worth of data in one day.

In this scenario we analyzed the ability of the DADS to be unavailable for 24 hours and to recover in the following 24 hours. In order to examine this case, the model has been fed with two days worth of data. The doubled amount of data to be ingested and distributed, results in a substantial increase in the disk space required (Fig 7) for both the Unix disks (18 GB) and the Unitree disks (7.5 GB). It requires almost 26 hours to complete the requests, thus slightly exceeding the 24 hours planned recovery period. The fallover could be easily accommodated the next day. If the disk space available at the DADS is less than the amount specified as used in the simulation (18 GB for Unix disks and 7.5 GB for Unitree disks), the ingestion and distribution functions would require additional time.

- Scenario 5: Effect of the number of human operators at the DADS

Human operators play a critical role in the performance of the DADS, since it may take them several minutes, on an average, to respond to a request and fetch distribution cassettes. In the previous scenarios, the simulation assumed that there was no restriction on the number of operators and each of them took 1 minute, on an average, to mount or dismount a tape. In the simulation, this 1 minute average was represented by a uniform distribution from 0 to 2 minutes. After discussion with the DAAC operation staff, this 1 minute delay to fetch and mount was found to be too optimistic, based on their experience, and has been replaced by an average of 3 minutes (uniform distribution from 1 to 5 minutes). The proportion of media requested are assumed to be respectively 50% for 8 mm, 33% for 4 mm, and 17% for 9 track tape, and the number of operators has been varied from unrestricted to 1.

Table 6 summarizes the results of these tests. Table 6, case # 1, differs from the scenario 3 assumptions only in that the operator response/fetch time was increased from 1 to 3 minutes. This resulted in an increase in total disk space (Unitree disks and Unix disks combined) from 5.4 GB to 7 GB, and an increase in the time to eliminate the backlog of requests from 13 hours to 16 hours. In cases 2 and 3, restricting the number of operators to 8 or 4 has little effect on the results. In case 4, with only 2 operators, the total disk space required, and the time to complete the requests, both begin to increase noticeably. In case 5, with a single operator, the total disk space is large (14.5 GB) and, even after 30 hours, the requests were still not completed. Thus, the DAAC needs more than one operator to keep up with the daily workload.

In summary, with 2 operators instead of 1, there is a substantial decrease in the disk used (11 GB) and a significant improvement in the time required to fulfill the distribution requests (17.5 hours). Having more than 4 operators does not change considerably the disk space requirement or the request completion time.

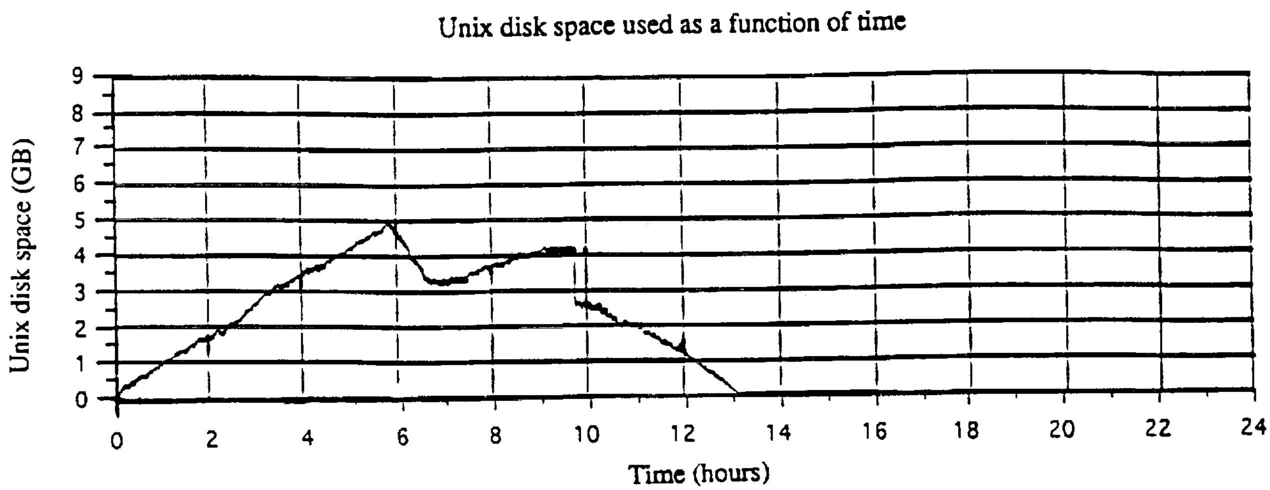

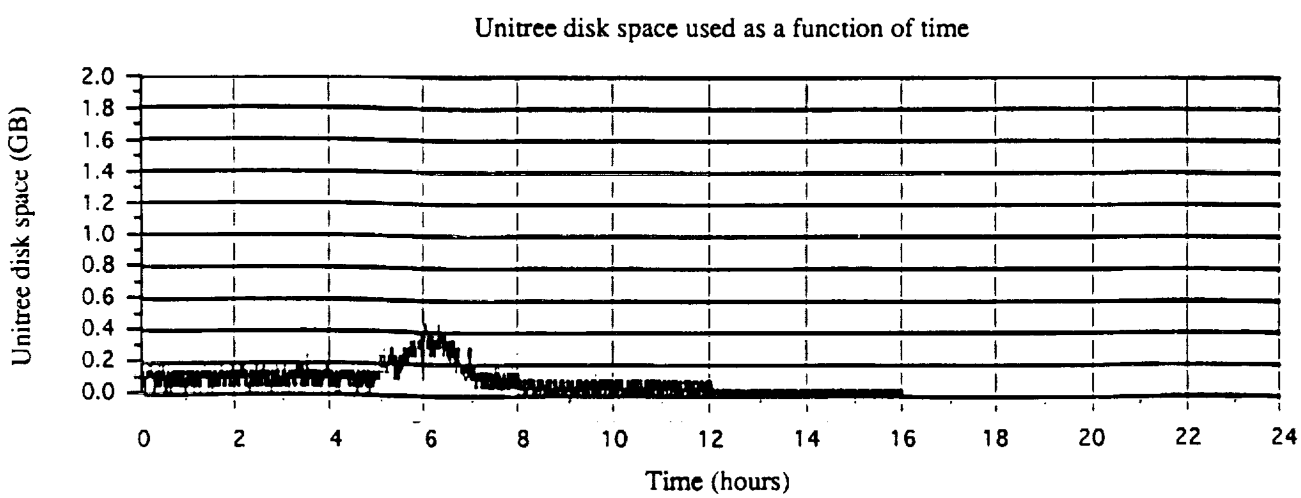

Fig 6. Disk space used in scenario 3

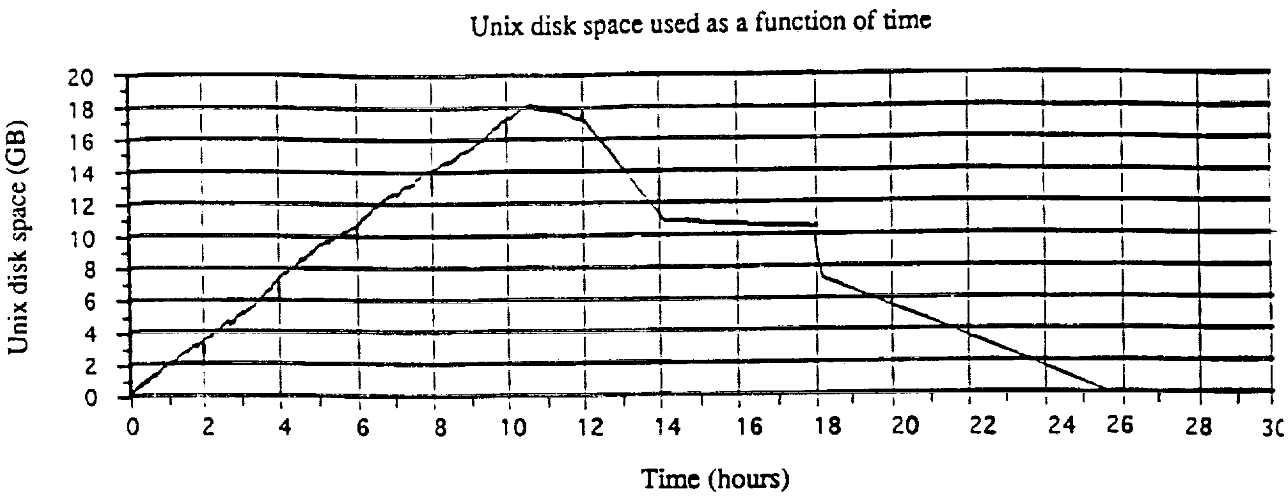

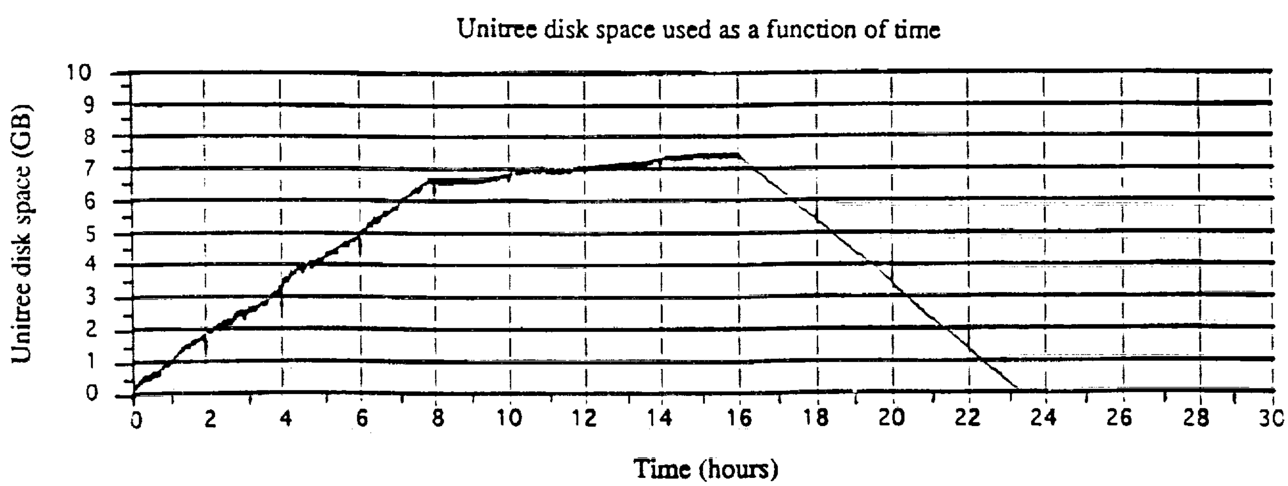

Fig 7. Disk space used in scenario 4

Table 6. DADS performance as a function of the number of operators

case #	1	2	3	4	5
Number of operators	unrestricted	8	4	2	1
Max disk space used (GB)	7	7	7.5	11	14.5
completion of standing orders (h)	8.5	9.5	9.5	11.7	21.5
completion of all requests (h)	16	16	16	17.5	>30

- Scenario 6: Effect of having stackers (4 mm and 8 mm) at the DADS

In order to automate the system and provide faster tape handling, the DADS model has been examined with one or several 4 mm and 8 mm stackers. It is assumed that the total number of stand-alone and stacker drives, for each type of drive remains constant (four 8 mm, three 4 mm, and two 9 track). There is a single operator to mount/dismount tapes from the stackers or the stand-alone tape drives.

The results obtained from the different cases examined are summarized in Table 7. Case 1 in Table 7 is the same as case 5 in Table 6 (i.e. one operator and no stackers). Adding a single 8 mm stacker and a single 4 mm stacker (case 4) has a significant impact on the performance of the DADS. The combined disk space required is reduced (13 GB instead of 14.5 GB), and the requests are completed much sooner (19 hours instead of > 30 hours) than when there is no stacker. As more stackers are installed at the DADS, the disk space used is decreased and the requests are finished over a shorter period of time. For cases 7,8, and 9, with 3 or more stackers, the results do not differ much from each other. In these 3 cases, the amount of disk space used is 9-10 GB and all requests are completed within 16-17.5 hours.

Table 7. DADS performance with and without stackers

Case #	1	2	3	4	5	6	7	8	9
Number of operators	1	1	1	1	1	1	1	1	1
# of 4 mm stackers	0	0	1	1	2	0	1	2	2
# of 8 mm stackers	0	1	0	1	0	2	2	1	2
Max disk space (GB)	14.5	13	14	13	11.5	9.5	9.5	10	9
Completion of standing orders (h)	21.5	19	18.5	13.7	16.5	15.5	12.25	12.25	10.7
Completion of all requests (h)	>30	24.5	26.2	19	23.7	21.5	16.5	17.5	16

- Scenario 7-A: Effect of ingesting all data and distributing all data without compression and decompression.

In the previous scenarios, the model had been executed when all data were ingested and when all the SeaWiFS data were distributed. For this scenario, the estimated distribution workload of AVHRR, TOVS, UARS, DAAC climate, and CZCS have also been included (8 GB/day). Disk space used when all data are ingested and distributed is illustrated in Fig 8. Comparing Fig 8 with Table 7, case 9, indicates that the additional 8 GB/day distribution workload results in an increase of total disk space required (from 9 GB to 16 GB) and takes longer to complete the SeaWiFS standing orders (from 10.7 to 13 hours). However, the time required to complete all requests (16 hours) is not changed.

Unix disk space used as a function of time

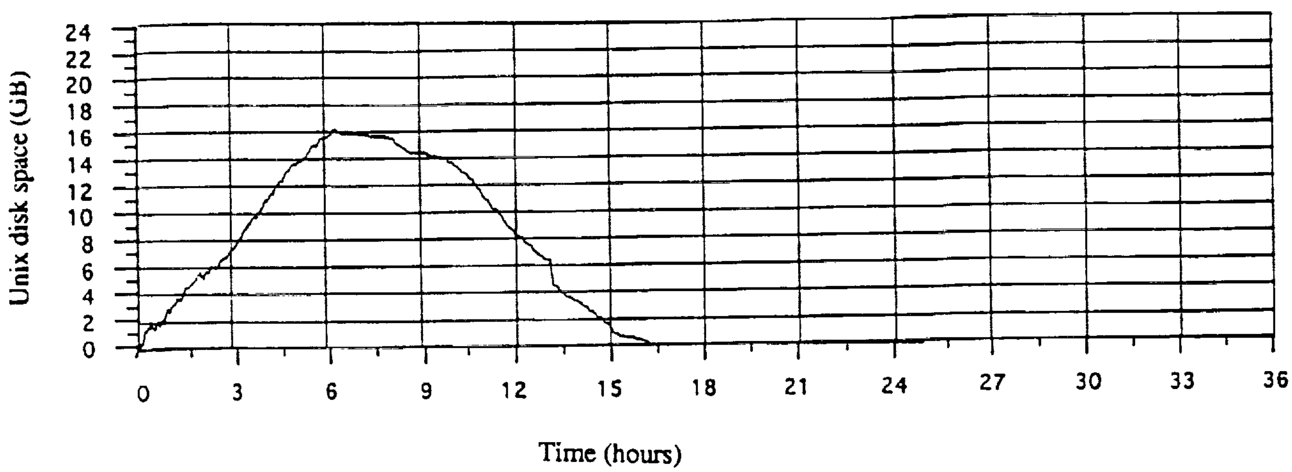

Unitree disk space used as a function of time

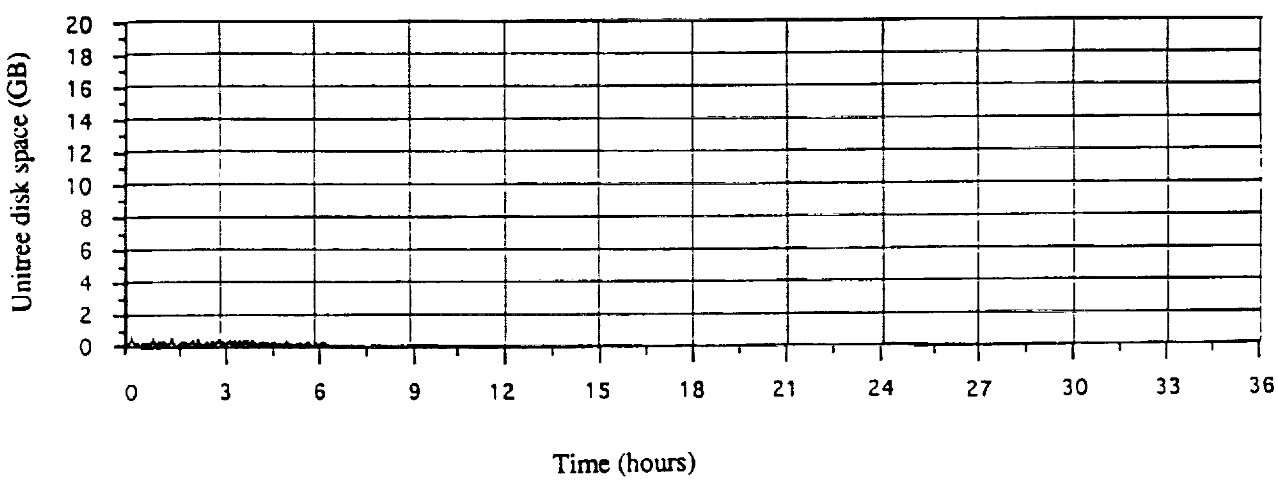

Fig 8. Disk space used in scenario 7-A

- Scenario 7-B: Effect of ingesting all data and distributing all data with compression and decompression.

The GSFC DADS is investigating the prospect of using data compression techniques to save storage space (1). Depending on the data set and the compression algorithm used, the original file can often substantially be reduced, thereby contributing to the mass storage solutions. However compression is a CPU intensive operation and can be rather slow if the compression is performed with software rather than hardware. The goal of this simulation is to estimate the impact on the DADS performance when using compression/decompression. In the DADS model the compression algorithm selected is the Unix compression (LZC), which does not have the best compression ratio, but is faster than the other algorithms evaluated and is a quasi-standard. The compression rate varies from data set to data set and, based on the results of the compression investigation, a compression rate of 200 KB/s was chosen for this simulation. It is assumed that all files ingested are archived in compressed form. The standing orders and the other distribution requests are sent to the users in an uncompressed form. With a slow compression/decompression rate it is expected that this may cause a bottleneck in the system.

The penalty for performing compression/decompression is indicated in Fig 9. The Unix disk space has increased from 16 GB to 19 GB, and the Unitree cache, which was under 1 GB, has now a peak of 10 GB, so that the total disk space is now 29 GB. The time to fulfill the distribution requests has increased slightly from 16 hours to about 17 hours. The large increase in disk space required is due to the slow compression/decompression rate assumed, which delays the ingestion and distribution processes, thereby causing data to build up on the disks. If the total amount of disk space had been constrained to less than 29 GB, the time required to fulfill the distribution requests would have been increased.

Unix disk space used as a function of time

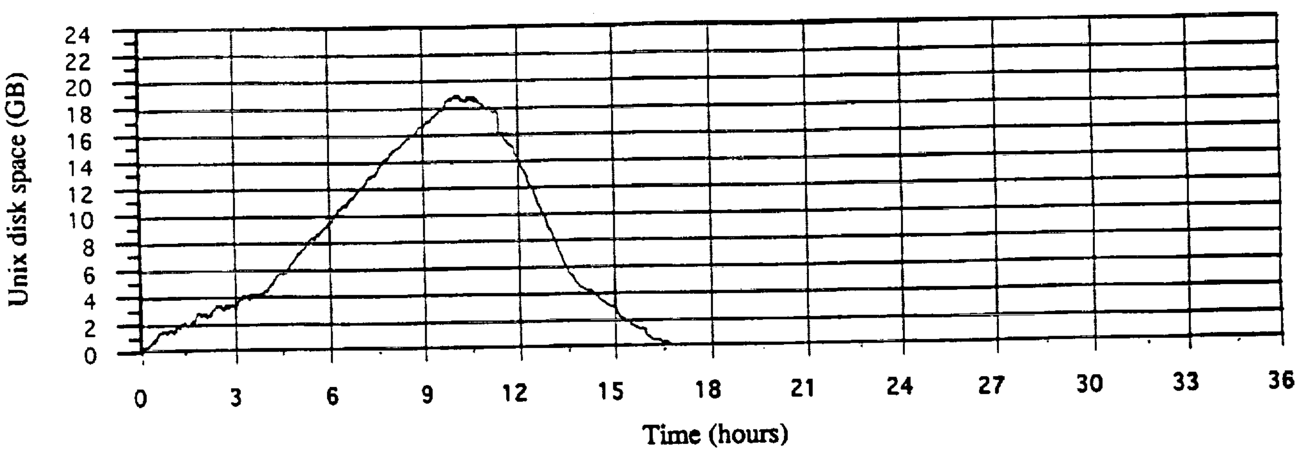

Unitree disk space used as a function of time

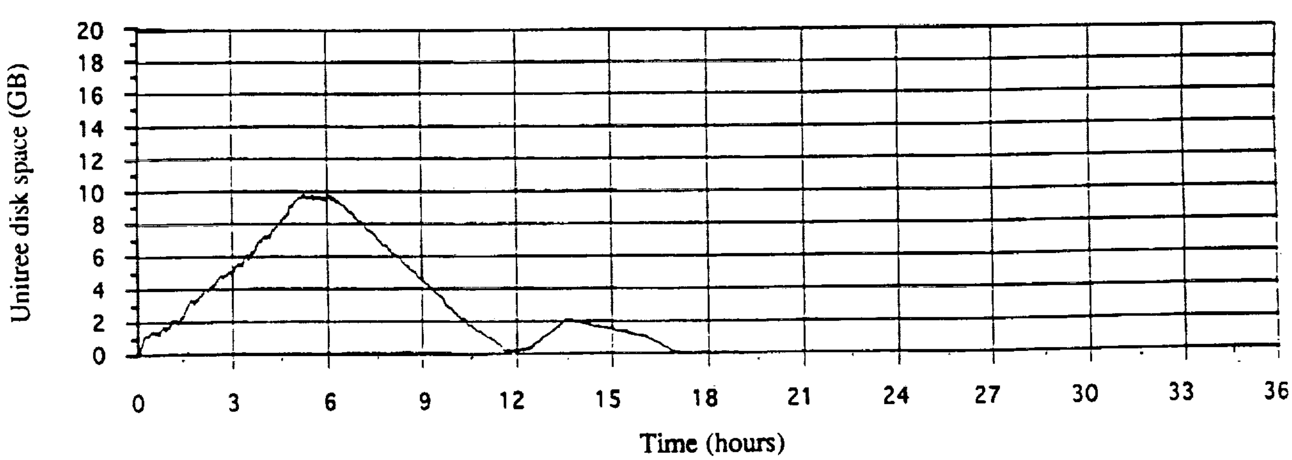

Fig 9. Disk space used in scenario 7-B

Conclusion

The discrete event simulation model has proved to be a very useful tool to evaluate the performance of the V0 DADS. The amount of disk space necessary to fulfill the daily ingestion and distribution requests at the DADS has been estimated under various conditions. The model has demonstrated the importance of matching the ingestion rate to the archival rate to prevent data build-up in the Unitree cache, thus minimizing the amount of cache disk space required (scenario 2). Human operators play a critical role at the DADS since the time of about 3 minutes required to manually mount/dismount tapes is a limiting factor in the DADS performance. However, having too many operators (4 or more) does not improve the performance of the DADS (scenario 5). Stackers (4 mm and 8 mm) can substantially help in automating and processing the DADS requests more quickly (scenario 6). With a single operator, under the assumptions of scenario 6, a single 8 mm stacker reduces request completion time from > 30 hours to 24.5 hours. The combination of two 8 mm stackers and two 4 mm stackers further reduces the request completion time to 16 hours. The compression and decompression operations are very CPU intensive and, if performed at slow software rates, will require substantial additional disk space (scenario 7-B).