

LA-UR -84-3537

CONF 841114--2

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

LA-UR--84-3537

DE85 003748

TITLE INTEL. IDIS EVALUATION

AUTHOR(S) Jesse M. Cheadle  
Marcus B. Glassmire

SUBMITTED TO ASTUTE Conference to be held November 11-15, 1984,  
at Hilton Head Island, South Carolina

#### DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

**MASTER**

By acceptance of this article the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

*714p* DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

**Los Alamos** Los Alamos National Laboratory  
Los Alamos, New Mexico 87545

# INTEL IDIS EVALUATION

*Jeet Cheadle and Marc Glaesemire*

Los Alamos National Laboratory  
Los Alamos, New Mexico

## ABSTRACT

This paper describes an evaluation of a computer system called the Intel Database Information System (iDIS), which was recently performed at the Los Alamos National Laboratory. The evaluation consisted of the development of a possible application for the Compensation and Benefits Section of the Personnel Administration Division at Los Alamos.

## 1. INTRODUCTION

This paper describes an evaluation of the Intel Database Information System (iDIS) computer system, which was recently performed at the Los Alamos National Laboratory. It includes the evaluation of a possible application for the Compensation and Benefits Section of the Personnel Administration Division. The application was a part of the annual Salary Authorization Increase work the Compensation and Benefits Section performs. This test application lends itself well to database and spreadsheet applications. The Compensation System resides on a Control Data Corporation (CDC) 825 computer and uses Intel's SYSTEM 2000 software. Since the compensation data resides in two SYSTEM 2000 databases, the iDIS Extract facility was used. An additional goal was to evaluate the concept of downloading mainframe data to a frontend computer where the data could be evaluated by multiple users. This is opposed to the concept of hooking personal computers (PCs) to the mainframe and having users download data and work individually.

With the above in mind, we performed our evaluation of the Intel iDIS. As a national laboratory we are not permitted to categorically recommend or advocate the use of one vendor or hardware type versus another. However, we do feel the concept of frontend data processing as supported by the iDIS is a concept worthy of serious consideration by anyone proposing an application requiring downloaded distributed processing. In the following sections the reader can identify the pros and cons of the Intel iDIS computer.

## 2. DESCRIPTION OF THE HARDWARE AND SOFTWARE USED

### 2.1. Intel iDIS

The iDIS system had 768K of memory, 35MB of disk space, two Intel terminals, one Intel printer, and Version 1.5 iDIS software. This software consisted of the XENIX (a trademark of Microsoft(R) Corporation) operating system, MISTRESS (a trademark of Rhodnius, Inc.) relational database system, HORIZON (a trademark of Horizon Software Systems, Inc.) word processing system, MULTIPLAN (a registered trademark of Microsoft(R) Corporation) spreadsheet system, and Intel's SYSTEM 2000 Extract facility.

## **2.2. Compensation Databases**

Two Version 2.80 SYSTEM 2000 personnel databases residing on a CDC 825 running the NOS 2.1 operating system were the sources of the compensation data. The databases consisted of one level zero repeating group containing employee information. One database contained staff members and the other graded employees. Both databases were identical with regard to component names and numbers, data types, and key versus non-key elements.

## **2.3. Data Communications**

Communications between the iDIS and the CDC 825 were provided by a Paradyne uP 4800-baud modem and the Remote Batch Facility (RBF) of the NOS operating system.

NOTE: A HASP interface is required for iDIS communications.

## **3. DESCRIPTION OF THE EVALUATION**

The Compensation and Benefits Section of the Personnel Administration Division allowed us to use their data in our evaluation of the functionality of the iDIS. Our goal was to perform an extract from a SYSTEM 2000 database, populate a MISTRESS database, and use the MULTI-PLAN spreadsheet system with the MISTRESS database system. In this section the results of these efforts are described.

### **3.1. SYSTEM 2000 Extract**

With help from the iXTRACT manual, we were able to download selected information from the SYSTEM 2000 databases. The first step was to copy a NOS CCL file to the host job file cabinet. The file then had to be edited to contain the correct user, password, charge code, and SYSTEM 2000 calling sequence for a CDC 825 SUBMIT job.

The procedures to download the SYSTEM 2000 definition were easy to follow and took only a few minutes. It took ten minutes for the result of the batch job to come back to the iDIS.

The describe output was cleaned up, and the data selection for downloading was then completed. This took about forty minutes. The data extract from the SYSTEM 2000 database to a file on the iDIS took one hour and forty minutes. There were 84812 components (attributes) transferred from 3029 entries (tuples) from the staff member database. There were 4597 entries in the graded employees' database with 128716 components transferred.

To format the SYSTEM 2000 print output into MISTRESS database format took another forty minutes. To load (populate) the MISTRESS database took fifteen minutes.

It took an additional two hours to create the indexes. This took a long time because we created indexes on every (twenty-eight) attribute.

### **3.2. Spreadsheet Operations**

We developed several spreadsheet models on both the graded employees' and staff member databases. A model contains database and spreadsheet directives to create a usable spreadsheet. Once a model is developed, it is run to generate a spreadsheet.

One model created a spreadsheet containing counts of staff members, average salaries, and average evaluation dates for Laboratory staff members. Staff members were separated by race, sex, and degree. It took about one hour to enter the necessary database directives and format the model. It took thirty-eight minutes to run the model and generate the spreadsheet.

Another application was to divide staff members with bachelor's degrees by sex and supervisory or non-supervisory position. This involved the development of four models. Three of the models were able to have all the database calls, functions, and formatting done in them. These models took about thirty minutes each to develop and from five to fifteen minutes to run. These models were paralleled for the graded employees' database. The graded models were based on job series order.

Because of the 255-row limit in MULTIPLAN, we encountered problems in breaking out all possibilities of data as required by this application. There were 479 entries for male non-supervisory staff members with bachelor's degrees to be put into a spreadsheet. Each entry had three items, so three entries were put in each row, taking up nine columns. Only two database directives were put in to the model, one to put in staff members' information and one to retrieve the 1985 median salaries.

When the spreadsheet was created, the necessary work to get the data in the proper format was done with MULTIPLAN commands on the actual spreadsheet. To do the functions and formatting took approximately one hour and fifteen minutes.

The table lookup feature of MULTIPLAN was put to good use in these spreadsheets. By deriving the years of experience from the evaluation date, it was possible to look up and display what the new median salaries for staff members should be for the next fiscal year—a feature the Compensation and Benefits Section may find useful.

### **3.3. Word Processing Operations**

Because of a lack of time, we rarely used the capabilities of the word processing system. We would have liked to have generated a report or memo using the "interface" to MULTIPLAN.

## **4. EVALUATION RESULTS**

The following section presents the results of our evaluation of the Intel iDIS. From the above, it can be seen we were able to do a majority of the work we outlined. This section will be divided into different topics of interest.

### **4.1. HORIZON Word Processing Critique**

We were not able to do everything we wanted with the word processing system. However, with what we did accomplish, we found it to be cumbersome because of the numerous operations to perform some of the different functions. This was also the conclusion of a person familiar with word processing systems.

### **4.2. MULTIPLAN Spreadsheet Critique**

The spreadsheet system was easy to use, and we were able to accomplish our objectives. The spreadsheet system contains useful features to arrange and perform functions in the model and actual spreadsheets. One drawback to the spreadsheet system is the 255-row limit. There are ways to work around this, but it does create extra work for the user (see Section 3.2).

In the options menu of MULTIPLAN, one choice available to the user is "options." Within "options" is a recalculate option that is defaulted to "yes" or "on." More than once, a single cell function would be entered causing a recalculation of all cells. This consumed a lot of time if there were a lot of cells to recalculate. A default of "no" or "off" would be more appropriate. Recalculation is useful if you change one thing that affects everything else.

Another possible problem with the recalculate option is the question of whether or not you can turn off the auto-recalculate when the user runs a model. In the example above where it took thirty-eight minutes to generate a spreadsheet, if the run module automatically did a recalculate, significant time could be saved by turning off this option.

There are two other areas of concern with the model system. One concerns space usage. If the user performs a lot of database selects, space is used up very quickly. With 220 database selects and 100 alpha cells (320 out of 16005 cells), 88% of the space was used. After the model was run, the actual spreadsheet used only 13% of the available space. Part of this is because database selects use a lot of characters. This chews up the buffer the model uses. The other concern was about editing a model. If you change, say, a 1 to a 2, for some reason this uses up free space. We don't know why this happens, but we don't believe it should since no additional physical space is used.

### **4.3. MISTRESS Database System**

This section will be divided into two parts. The first part will concern itself with the use of the MISTRESS database system on the compensation application. The second will be a general evaluation of the overall capabilities of the MISTRESS database system.

#### **4.3.1. Compensation Database Application**

The Extract facility was used to download data from a CDC 825 SYSTEM 2000 database into an iDIS MISTRESS database. This was previously discussed in Section 3.1. By following the instructions in the iXTRACT manual, we were able to create the database easily. Very little interaction is required with the database when you are using it with the spreadsheet system.

There are two areas of concern with the database system. The first is the difference between upper and lower case. The database system recognizes the difference. Unless the user changes the attribute names to lower case in the user view section of the Extract, shifting between lower and upper case will have to be done frequently. The second area of concern is the length of attribute names. In the iDIS, the attribute name is the only reference to the attribute. This is unlike SYSTEM 2000, which has component numbers as short names. Short attribute names are also recommended for use with the spreadsheet system. A database select in the spreadsheet can be no longer than 150 characters including row column specifications. With long attribute names, it does not take much to use 150 characters.

One feature of the MISTRESS database system is the capability to select data (all or part of a relation) into another relation. This can have many useful applications, one of which is for use with the spreadsheet system. If you "select" is too long, you could go to the database and create a new relation. This relation would contain only the information your model selects from the database, and a "where" clause would not be necessary.

#### **4.3.2. MISTRESS Database Critique**

We found the database system easy to use. We were able to create and input data readily. The syntax for retrievals was easy to learn.

Manual data input is accomplished with the "insert" command. This asks for one attribute at a time. When all attributes of a tuple are entered, the database responds with a "Ready" message. The user can enter the relation (.e) or other options. Sometimes this process can be very cumbersome and time-consuming.

It would be nice if automatic menus or screens were an option of the database system. Currently, the user has to perform application development to create menus or screens. This involves creating a shell script to give the users a menu to interact with their database. Then macro or C programs have to be written to interface the screens with the database. Our understanding is that it is faster if the programs are written in C.

The capability of more than one function (count, average, sum etc.) in a "select" statement would also be a nice feature to have. In a spreadsheet model we had a count and two averages with the same "where" clause. We had to put in three "selects." One "select" would have made the work easier and probably faster.

After running a spreadsheet model that took five hours, we found caution must be used when creating relations from another relation. When the new relation is created, the previous indexes are removed and not re-created. After we put indexes into the new relation, the spreadsheet model took thirty-eight minutes, a significant improvement. The model had ninety "selects" in it. Each "select" computed a function. There were three conditions in each "select" "where" clause. As was mentioned earlier, the thirty-eight minutes would be a lot less if the automatic recalculate could be turned off.

#### 4.4. EXTRACT Facility

It is easy to extract data from a SYSTEM 2000 database by following the instructions in the iXTRACT manual. As seen in Section 3.1, it takes a considerable amount of real time to go from a SYSTEM 2000 database to a MISTRESS database.

One item a user has to be aware of during the Extract procedure is whether a database exists in which to put relations. If the database does not exist, the user has to create one before trying to put a relation in it. This possible error condition is not well documented.

#### 4.5. General iDIS Comments

The iDIS is slow going from screen to screen, and Intel is responding to this problem by implementing the 80286 technology with new iDIS machines. The iDIS we evaluated used the 8086 technology. We entered this paper into the HORIZON word processing system, and the system could not keep up with the typing speed. We have heard this does not happen with the 80286 iDIS.

Some work by Intel must be done to standardize the terminal keyboard. More than once we were confused by entering the wrong backspace sequence. For XENIX it is CTRL-H; for most of the iDIS it was the backspace; and for the spreadsheet and word processor it was a shift DEL. When viewing files, the iDIS says to stop viewing and hit the RUBOUT key, which is a shift DEL (this is not documented). When editing in the spreadsheet, the documentation said to use the F1 through F4 keys to move around, which did not work. Instead we had to use the F5 through F8 keys.

Intel has done a lot of work on interfacing IBM PCs to the iDIS. We have found the interface works well. The user can easily transfer files back and forth, which adds greater flexibility to the uses of the iDIS.

The iDIS comes with a set of documentation. Our set includes eight binders. Four are for the XENIX system, and the other four are for the iDIS system. For most of the work we did, we could easily find the documentation to help us. We found the documentation for the word processing system to be extremely hard to use. An overall index for the iDIS manuals would be a great help.

We found the iDIS menu system to be very complicated when we first started. After doing our first trial application on our own, we found ourselves becoming familiar with the menus. Now we feel comfortable with the menus and find ourselves having no trouble navigating through them.

If an iDIS is being considered for a standalone system, careful consideration should be given to the impact of learning a new system, XENIX; a new programming language, "C"; and how to do application development on the iDIS.

#### 4.6. Intel Support

We originally acquired the iDIS for a 90-day evaluation. This period started in March 1984. We have had the iDIS for a total of seven months. The first drawback we encountered was communications with the CDC 825. We decided to dial-in because Intel had indicated that dial-up would work. It took four weeks before we had a modem that would work.

Next we waited for a new release of the iDIS system. It was at this time we learned the CDC Extract was not on the system we had and was only available with the new release. It took six weeks to get the new system. We were informed the new release was so easy to install that "a three year old could do it." When we unpacked the new system, we found several new chips for the iDIS, and hardware changes too! We closed the box and called Intel in Albuquerque to install the new system for us. It was another two weeks before the new system was installed. It took two days to do the actual upgrade.

At last, we were ready to go. Now we found that the communications would not work. After four weeks, we, with help from CDC and Intel personnel, were able to get the communications to

work. It was also during this time that we found out we were the first CDC site to use the iDIS. Shortly after the Albuquerque Intel service representative went on vacation, the processor board on the iDIS started going bad. Before the iDIS was back up and running, a week and a half had gone by. A field service representative from Denver had to come down and look at the iDIS. He replaced the processor board on a Thursday, and the iDIS was down again on Monday. The service representative came back and replaced the board again, and it has worked ever since. When the service representative came back the second time, he brought two processor boards with him. This was a good thing because one of the processor boards was dead on arrival.

The above scenario is included to describe the sometimes laborious paths necessary to be a test-bed for new equipment.

## **5. SUMMARY**

The results of our evaluation demonstrated the capability to download host data and perform spreadsheet operations on that data. We were happy with the results we obtained. It does take a significant amount of time to perform certain operations, and we believe this should be a topic of concern at Intel.

We believe Intel learned a lot from the interactions with us. The iDIS performs well, and future enhancements should make it even better.

## **6. ACKNOWLEDGMENTS**

This research could not have been accomplished without the support and cooperation of numerous people at the Los Alamos National Laboratory and Intel. We thank Robert Ewald and Ann Hayes for letting us pursue this evaluation; Ron Hults, Dan Magnes, Rick Jones, Steve Jaynes, Mark Thompson, Chuck Haas, John Barrett, Mike Verba, and numerous other Intel personnel; Terry Senay and Tom Smith from CDC and Tom Stup from Los Alamos for providing support for the interface and hook up to the CDC 825. Finally, we would like to thank Fabian Martinez and the Compensation and Benefits Section at Los Alamos for providing us with an application.