

Vol. 7, No. 4, May 1982

Software AG International Users' Group Newsletter

Published by the Software AG
International Users' Group and
Software AG of North America,
Inc.

Contents

	page
Users' Group News	
International Users' Group to Meet at Las Vegas	1
South American Users' Hold Fifth Annual Conference	1
Software AG of Far East Creates New Users' Group Region	2
German-Speaking Region for Triumvirate	2
Regional Reports	
Southeast Asia Region	3
Australian Region	3
Eastern Canada Region	4
Far East Region	4
Far East Region—Kansai Regional ASSIGN	5
Far East COM-LETE Meeting	5
Eastern Region	6
Sierra Pacific Region	7
Northwest Region	8
Pacific Southwest Region	8
ADABAS News	
Product Representative's Report	9
ADABAS Considerations in an Online Update Environment	9
Version 4 MPM Dual Disk Logging	10
ADABAS Special Interest Group Chairpersons	12
NATURAL News	
NATURAL Product Representative's Report	13
NATURAL Design Considerations	13
NATURAL Special Interest Group Chairperson	22
COM-LETE News	
COM-LETE DOS Available	23
COM-LETE Special Interest Group Chairpersons	23
ADABAS-M News	
New Features	24
Software AG News	
International Presidents Meet	25
SAGNA Expands Regional Facilities	25
Software AG Class Schedule	26
Release Levels & SAGUTILS	29
Early Warning	29
Software AG Contacts	30
Regional Representatives	30
Dates to Remember	31
International Software AG Users' Group Executive Committee	32

International Users' Group to Meet at Las Vegas

The International Software AG Users' Group [SAGGROUP] will hold its Executive Committee meetings during the 10th International Conference at Las Vegas. At the meetings the Committee will discuss the past year's activities and prepare plans for the coming year.

Each year at the annual Conference both the retiring and newly-elected SAGGROUP officers, representatives and Special Interest Group [SIG] chairpeople get a unique opportunity to meet together with users from around the world to discuss product status and futures, to listen to presentations by other users on their applications and ideas, and to participate in workshops and tutorials.

Those who have participated in SAGGROUP deserve special thanks for performing a vital function for the users of Software AG products and for accepting responsibilities outside of their own professional activities. The contributions made by these people help to provide international communications among users and to ensure that the products progress to meet the users' increasing needs.

Since the Innisbrook Conference of 1981 new regions have been formed in Iberia and Southeast Asia. The German-speaking countries formed a unique triumvirate, and South American users held a highly successful conference in Brazil.

SAGGROUP is one of the most active and unique users' groups in the industry and in the world. Your participation and support of the people involved can only help it continue to prosper and grow.



South American Users' Hold Fifth Annual Conference

The Fifth Annual South American Users' Conference was held from Monday, March 29 through Friday, April 2.

Attendees representing most of the Brazilian user community and some from Argentina attended courses on ADABAS internals, ADABAS design and heuristic development during the first two days of the Conference.

For the remainder of the Conference, some 100 users and guests listened to 14 Users' presentations covering a wide range of topics from Performance and Tuning to Pre-processors and Programming Techniques. The presentations were judged by two expert judges and by a consensus of the participants. The top four papers received prizes that included a round trip from South America to Las Vegas to present their papers at the 10th International Users' Conference.

The two judges were: Mr. Paulo C. Busnardo, President, *Brazilian National Association of Computer Users*; Mr. Waldemar W. Setzer, Ph.D., *Head*

of Applied Mathematics Department, University of Sao Paulo.

The winners of the prizes were:

First Prize:

Nelson Luiz Correa Robello
Prodemge-Data Processing for the State of Minas Gerais
"Data Dictionary—A Proposal"

Second Prize:

Antonio Rubens Anciães
Amaral Filho
Centrais Eletricas Brasileiras S/A—ELETROBAS
"Factors of Influence in the Performance of Data Base Systems"

Third Prize:

Paulo Roberto Suzuki
Abril S/A—Cultural and Industrial
"Abril Experience with NATURAL 1.1"

A special prize was awarded by CONSIST to the ARUANDA Project Team from SERPRO—Data Processing Services for the Brazilian Federal Government entitled: "ARUANDA—An Information System of Generic Use"

We look forward to hearing these presentations and others from the South American user community at Las Vegas.

SOFTWARE AG of the Far East Creates New Users' Group Region

Software AG of Far East, Inc. (SAGFE), one of the leading software package vendors in Japan, announced that the Southeast Asia users will meet in Singapore for the first time and form a Southeast Asia Regional Group—a first among software vendors in Japan. This indicates that Southeast Asia has grown to be a profitable and attractive marketing spot next to the Japanese market, where the software industry has recently started to burgeon.

According to SAGFE's President, Mr. Kazuro Fujimoto, there are now six users in Southeast Asia: two in Singapore, two in the Philippines and two in Hong Kong. Also, two organizations from Singapore and one from Hong Kong will purchase ADABAS soon.

SAGFE has a representative office in Singapore and expects 15 to 16 participants at the first Users' Group meeting. The main purpose of the meeting is to exchange information among users on the use of SAGFE's products such as ADABAS and NATURAL and to allow SAGFE to realize the users' needs and requirements regarding the products.

SAGFE established the representative office in Singapore in mid-1978. The regional meeting is a result of that office's efforts since that time. *(Based upon a translation from "Nikkei Sangyo Shinbun" February 9, 1982)*



German-Speaking Countries Form Triumvirate

The German-speaking countries formed a unique triumvirate at their Tenth Users' Group Conference held in the fall. The three members at the Triumvirate represent all the German-speaking countries: West Germany, Austria and German-speaking Switzerland.

The members of the Triumvirate are:

Mr. Bernard Möller
Westdeutsche Landesbank
D-4400 Muenster

Mr. Manfred Knöpfle [speaker]
Stadtverwaltung Augsburg
D-8900 Augsburg

Mr. Hans Ulrich Räetz
Bundesamt Für Statistik
CH-3003 Bern

At the meeting the triumvirate formed work groups for NATURAL, COM-LETE, Utility DBA-Aids and for Siemens users.

According to Mr. Möller, member of the triumvirate, the general disposition among the users was "sheer sunshine and bliss."

Southeast Asia Region

Christian Perez
Asian Development Bank

A new Software AG Users Region has been established. Users from the Southeast Asia Region held their first meeting in Singapore on March 5, 1982. It was attended by representatives from six user organizations from Singapore, Hong Kong and the Philippines; by prospective users from Singapore; and by Software AG, which included Mr. Fujimoto, President of Software AG of the Far East (Tokyo). Mr. Christian Perez was elected regional representative. The Group is expected to grow substantially in the near future.

After the opening remarks, all users presented their configurations and applications, and arrangements were agreed upon for future meetings. It was agreed that a general meeting would take place annually; however, two subgroups—users from Singapore and users from Hong Kong and the Philippines—would meet more frequently and that papers would be exchanged between the two subgroups.

Technical discussions which took place in the afternoon session mostly focused on NATURAL. Software AG representatives said that NATURAL 1.2 was not yet released and that no decision has been announced on the pricing policy for NATURAL Security Features. Users objected to a possible separate pricing of NATURAL security. In that connection, Raymond Wong, Cathay Pacific Airways, made a presentation on a system they are implementing in order to control NATURAL program libraries, ensure consistency of source and object modules, and provide a trail of program changes.

Some users noted that they had to send their staff to the United States for training sessions in English as training in Japan was conducted in Ja-

panese. They suggested that Software AG organize a training program in English for Southeast Asia Users. The need was expressed for ADABAS and NATURAL internals. The lack of NATURAL internals documentation was emphasized by several users.

Australian Region

John Palmer
Australian Bureau of Statistics

Two meetings have been held in the Australian Region since the last November report on activities in that Region. The first was held on 12 November 1981 and the second on 4 March 1982.

November Meeting

The November meeting was held in Sydney at the Reserve Bank of Australia. The 25 representatives of user organizations in attendance were pleased to welcome Gary Bowers (SAGNA) and Lewis Folb (SPL South Africa). Users appreciated hearing firsthand the latest news from Software AG. Lewis Folb was able to outline the steps being taken to increase the strength of the SPL support team. Users at the meeting commented that problems with support, due to experienced staff leaving SPL, were starting to impact their operations. The plans outlined by Folb to increase Sydney and Canberra-based support were, therefore, welcomed by attendees. SPL also reported that:

- two new users in the region are Electric Power Transmission and Department of Social Security;
- the Data Base Machine is about to be installed at the Western Australian TAB; Frank Zaknic (SPL) is visiting U.S. and Germany for training in the DBM.

In the next segment of the program Steve Maxwell provided a brief history of the Reserve Bank and its evolution from a batch installation to an increasingly significant user of

DB/DC facilities. He predicted a rapid growth in the use of ADABAS in future development projects.

Bob Hayes provided an outline of the development methodology adopted at the Bank. The approach combines both data and event analysis and the synthesis of the events into the overall conceptual schema. Although the methodology described has not yet been the basis for an implemented project, considerable interest was expressed in the approach.

Delegates were guided through one of the major production areas of the Bank utilizing ADABAS and NATURAL.

After lunch, Gary Bowers provided a presentation on the new Software AG Data Dictionary outlining the background of the current ADABAS data dictionary, the basic objectives and philosophy of the new dictionary and the current status of development work. This led to numerous questions from users relating to the continued support of the existing dictionary and the user extensibility capability in the new dictionary.

The meeting concluded with the regular "Problems and Solutions" session.

Eastern Canadian Region

*Gene Miklovich
Systemhouse*

The Eastern Canada Users met on January 12 in Ottawa. The meeting was hosted by Industry, Trade and Commerce and 25 people attended.

The meeting began with a presentation by Eric Valentine of Systemhouse on pseudo-online updating with NATURAL. This presentation described in the PEMD system at ITC, which runs as a single-user, read only version under TSO. The advantages of this system were to allow many users access and to generate updates (which were applied later by an overnight batch NATURAL run)

while at the same time not requiring MPM to be running. Natural was used online under TSO to take advantage of the screen editing as well as other capabilities. One problem area was the collection of the transactions on a work file. This was overcome through using a COBOL subroutine to write to a sequential file with DISP = MOD since each fetch causes an open. This led to some problems due to the COBOL routine being reloaded each time it is called. Eric said he prepared a two-page paper on guidelines and options to use at compile and link time to overcome this.

Discussions on the use of NATURAL by different users followed. Eric pointed out that ITC uses an enhanced version of the JOHNSON and JOHNSON PDS exit.

This version includes a re-numbering capability and also resolves the problem of losing a member on a SAVE when the PDS runs out of space.

Ted Hillis followed with a presentation of how Consumer's Gas runs with multiple MPM's. He cited the major reasons for this environment as:

- 1) Improved Recovery Capability (Time);
- 2) Aid Conversion;
- 3) Ability to adjust files;
- 4) Operational Fine Tuning.

Ted distributed notes and described in detail Consumer's Gas operation. He indicated that they used 1 SVC per version of ADABAS (i.e., 3 or 4) and a BDAM routing table. The route table was accessed through a modified version of ADAHOOK which then determined the appropriate data base and file number to access. In addition, Ted illustrated some of the DBA aids used to maintain the route table. The use of this approach gave Consumer's Gas much more control and flexibility and Ted indicated future enhancements are already planned.

The next discussion was led by Ernie Greco of Ultracom regarding Version 4 testing. He

indicated that there were not too many problems in testing Version 4. However, he did note the following:

- 1) Mass Update and Regenerate require SM03;
- 2) Recovery under Single User is different from under MPM and is not documented;
- 3) Version 4 resolved the problem in 3.2.1 of the Application aborting, without ADABAS knowing about it or doing anything, by backing out transactions after the time limit is exceeded. Ernie indicated there were three or possibly four types of users:
 - 1) ET/BT;
 - 2) File Cluster;
 - 3) Exclusive user;
 - 4) Exclusive with ET.

He was unsure, and no one was able to confirm, whether there was actually a Class 4 or if he reverted to a Class 1 user. Ernie suggested that for rerunnable systems, all online users be ET users and all batch jobs be exclusive use ET users. A discussion on the modes of recovery ended with the request that Cliff present a discussion on recovery at the next meeting based on questions he received before the meeting.

Cliff indicated that there is a new internal sort algorithm available for users at the SM04 level. In addition, based on feedback from the users, more SAG courses could be offered in Canada. Before adjourning, it was agreed the next meeting would be held March 23 in Toronto.

Far East Region

*Hiroki Kawamura
Nippon Kokan K.K.*

The ASSIGN (ADABAS Special Interest Group of Nippon) meeting held in Tokyo on September 21, 1981, had an attendance of 35 persons from 24 companies.

Mr. Terada of Kyodo Oil company gave a presentation on Kyodo Oil's accounting system using NATURAL. The user uses this system interactively

via CRT which allows Kanji characters and ruled lines on the screen. It has 70 kinds of screens and consists of 170 programs (50,000 steps). There is a dictionary of Kana/Kanji character conversion which is one of the ADABAS files.

Mr. Ishii of Software AG gave a presentation on a Database Machine and reported the findings from the tests that SAGFE performed.

At the November 24, 1981 ASSIGN meeting held in Tokyo, an attendance of 28 persons from 17 companies was reported.

Mr. Sakamoto and Mr. Hasegawa of Kinokuniya Bookstore Co., Ltd. gave a report on ADABAS troubles which occurred at their installation. They voiced out "It took two to three weeks to solve troubles because SAGFE had to send dump lists to SAG (Darmstadt). If SAGFE had the source lists, SAGFE could modify it tentatively. Therefore, we strongly request that SAGFE have the source lists in Japan."

Each company explained briefly about Change/Enhancement requests they had submitted. ASSIGN will send all requests from Japanese users to North America.

At the January 20, 1982 ASSIGN meeting held in Tokyo, an attendance of 31 persons from 23 companies was reported.

Mr. Hasegawa of Kinokuniya Bookstore reported the troubleshooting on S5 command and its solution. The solution is as reported by the Urgent Circular No. 3.

Mr. Ishizaka, Senior System Engineer of SAGFE, explained about ADASUM, an ADABAS command Simulation Program to simulate the day's applications at night, and ZAPCHK, a ZAP Check Program to check whether the zap is properly applied. Currently missing, ADASUM OS version and ZAPCHK DOS version are now under conversion and will be released this spring.

It was determined that the user status reports would be

distributed to those users who handed in the reports and consented to distribution.

Far East Kansai Region ASSIGN

Masayasu Kimura
Kansai Electric Power Co., Ltd.

Because the number of installations has increased in the western part of Japan recently, KANSAI REGIONAL ASSIGN was established in September, 1981. KANSAI REGIONAL users met twice in Osaka.

The KANSAI ASSIGN/COM-LETE meeting held on September 30 had an attendance of 24 persons from eight companies.

Mr. Ishii of Software AG gave a presentation on a Database Machine and reported the findings from the tests that SAGFE performed.

Mr. Suemasu of Software AG introduced NATURAL/KANJI based on Chinese Characters under COM-LETE and its examples of usage.

We discussed how to conduct the future meetings and decided the following:

- Masayasu Kimura from Kansai Electric Power Co., Ltd. was elected chairperson of the KANSAI REGIONAL ASSIGN;
- Each company will introduce their actual implementations and experiences with ADABAS.

On December 3, 1981, the second meeting was held and 19 persons from seven companies were in attendance.

Conversion of general purpose master to ADABAS was discussed. Nissho-Iwai explained the process of conversion of general purpose master to ADABAS and presented the comparative result of performance.

We discussed Change/Enhancement requests which users submitted.

Far East Region COM-LETE Meeting

Yasuhiko Katada
Tokyo Sangyo Credit Association

The COM-LETE meeting held in Tokyo on September 10, 1981, had an attendance of 13 persons from six companies.

Tokyo Sangyo Credit Association presented how COM-LETE is being used for program development, and how NATURAL is being used for their applications. In their company, Management Information System was implemented using NATURAL. The speaker emphasized, "COM-LETE has contributed a great deal because it attains high performance but does not devour much resources."

We discussed how to manage the Program Source File under COM-LETE. The users requested higher level management functions than LIBRARIAN or PANVALET.

Software AG reported some technical information and the status of the new version.

On November 19, 1981, the COM-LETE meeting held in Tokyo had an attendance of 13 persons from five companies.

Software AG gave a summary on the past Change/Enhancement requests and briefly explained this year's requests from all users. We discussed those requests.

In addition, Software AG announced additional utilities which users and SAGFE developed.

On February 19, 1982, the COM-LETE meeting held in Tokyo reported an attendance of 16 persons from nine companies.

Mr. Fukumoto, Manager of Isetan Data Center, introduced their organization, the content of existing applications and the history of system implementation for the past 15 years. Further, the attendees visited Isetan's EDP room with special arrangement and observed the



demonstration of 3270 KANJI terminal and the KANA terminal.

The questionnaire on utilities prepared by SAGFE was distributed by exchanging information among COM-LETE users. The JCL submit utility system which Isetan Data Center introduced is an example.

SAGFE encouraged the users to attend the upcoming International Users' Conference in Las Vegas.

The ballot procedure of Change/Enhancement requests was explained. The users were requested to return the ballots to SAGFE by March 9, and the results will be announced at the next meeting in May.

Eastern Region

*Dan Nolan
Computer Sciences
Corporation*

The Winter '82 meeting of the Eastern Region was attended by nearly one hundred representatives of over 60 companies in the area. We began with an introduction of new members and selection of officers for the next year. The new roster for 1981-82 reads as follows:

Programme Directors: DBA & Dictionary

Tom Bradbury
Hay-Huggins Data Svc's
215-875-2718

Host Languages

Julie Cottrell
National Education Ass'n
202-722-7515

Natural

Laura Sadowski
State of Delaware
302-571-3380

Adabas Internals

Marilou Sazon
The Boston Company
617-833-4372

TP, COM-LETE & Equip.

(unfilled)
(unfilled)
212-976-3838

Vice President

Bill Gordon
E.I. DuPont
302-366-3553

Regional Representative

Dan Nolan
Computer Science Corp.
202-237-2000

Past President

Rob Collins
Computer Science Corp.

To review these positions: Programme Directors field questions in their respective areas either to more experienced users or to Software AG and will arrange for presentations or work sessions at each regional meeting; Regional Representatives co-ordinate each meeting; Vice Presidents will succeed the Regional Representative for the next term and stand in the event of absence; Past Presidents stand by for the term following their service as Regional Representative and will stand in absence after the Vice President. Members, of course, have agreed to support these activities as eagerly as they avail themselves of the benefits.

Rob Collins began with a management view of how DBMS has changed the systems development process. He observes that the new tools have literally made the traditional methods obsolete and made feasible the once-theoretical areas of modeling new systems (designing online rather than on paper) and user commitment (dedication rather than mere involvement). Rob summarized with a statement that it would be preferable to withdraw rather than continue without total support of management and users, that to develop a system without a user to guide the way was to virtually guarantee user dissatisfaction and eventual failure.

Dan Keller of Chemical Abstracts introduced the incorporation of new positions, the "Para-Professionals" or "Data Base Co-ordinators" in a DBMS Environment. CA has been training generalists to assist the professionals in carrying out specific tasks such as file and field definition, dictionary and language maintenance, modifications and periodic review of the data bases. This releases the DBA members and development teams and allows them to concentrate on

their specialties. CA's success supports the concept of "Librarian" in the earlier structured development team descriptions and reminds us that the "Objective of Technology is Simplicity." (And encourages us that we may, in fact, be more sophisticated than we think.) Dan emphasized the needs for controls and standards, gave an extremely concise and encompassing account of the tools they employ, and passed along the standards they set up.

Jim Bagby of Software AG conducted a work-session on Tuning which followed lunch. We apologize for subjecting those who couldn't have cared less to such a technical discussion but, as Jim readily admitted, this is about the only way of nailing the super-SAG-ers down.

Whatever they've got on their desks *has* to be more interesting than writing documentation. Jim observed that the similarities of each site end installation and any resemblance to other sites, living or otherwise, is coincidental. In other words, the configuration and parameters which work best depend on *what* is being done on your machines and will probably not produce even marginally satisfactory results on someone else's. Jim began by describing the installation (priority of ADABAS in queue to other functions) and allocations of storage sets, and went on to review each of the "run-nuc" parameters in turn, describing their purposes, relationships to internals and other factors, under what conditions they might be changed and the situations which might be expected from altering (or ignoring) each key.

For those unable to attend, Jim's discussion is available in a single tape cassette as is a similar presentation given earlier by Paul Peterson in Denver. Only a single copy is available (from me), so a call is out for volunteers to copy or transcribe these rare documents.

Jim Forestell presented Graphics with NATURAL. It ap-



pears impressive, albeit hampered by equipment unavailability. Currently, the ability of NATURAL to present data to graphics packages in a form closest to that expected with a minimum amount of fuss, bother and processing remains the least onerous alternative for those not owning large blocks of IBM stock.

Jim Crawford, user group liaison, gave us a preview of the Las Vegas conference directed more toward the work ethic than to those recreation aspects which gave us so much trouble last year. Jim promises a prospectus and agenda geared to gird ourselves for the inevitable ordeal with our management and budget people and make our justifications less of a gamble.

The February meeting ended with general discussions on each problem area with members presenting their topics of concern to other users and representatives of Software AG. One of the more vital issues was that meetings are too far apart. It was proposed (finally) that local meetings be held at six- or eight-week intervals locally as a filler, to be coordinated by the hosting member and supported by the regional representative. Two sessions were tentatively set up (ComSat in D.C. and DuPont in Delaware). At the same time, Delaware Valley members questioned the formation of a group in their region to which they were not invited. Jim Crawford will investigate and clarify the situation for the members at the next meeting which will be, incidentally, on June 1, 1982, same time, same place. To avoid the time-consuming process of having each member introduce themselves, only new members will be announced at future meetings. "Established" members are asked to wear their own ID tags (from some other function) so that we may readily recognize and welcome newcomers.

A letter of thanks was sent to our host, Mrs. Blackman, for the quick reaction of her staff

upon discovery that no refreshment or feeding plans were made, and for graciously accepting our head-count in lieu of the actual attendance. We should also like to express our appreciation for your patience and acceptance of the unorthodox manner in which the tab was collected (by passing the hat). Please be assured this will not recur. Arrangements have been made for more formal dining and more parsimonious (\$5-6) prices for the next meeting.

Adjournment was at 4:00 p.m. for some, 5:30 p.m. for others and 9:00 p.m. for a hardy (thirsty) few. We will be contacting you regarding welcoming new members, filling the remaining position of Programme Director for Tele-Processing, COM-LETE and Equipment and with the agenda of the June meeting.

Sierra Pacific Region

*Colette Farabaugh
Castle & Cooke*

Our March meeting was hosted by Omex Corporation and held at the Ramada Inn in Santa Clara. The topic of this meeting was Version 4 Performance and Tuning.

The first speaker was Greg Lohman, from the Consulting and Education Center in Seattle. His presentation centered on tuning the ADABAS nucleus. Some of his points were:

- If the last thread is processing more than 10% of the commands (excluding updates), another thread should be added.
- The default for the ADARUN LP parameter should be sufficient unless more than 15,000 commands are processed in eight hours.
- The default for the LS parameter will allow only one-quarter of the records which were sorted under Version 3. LS=60K will equal V3.2.1.
- Throwbacks—if the number is

high, do some analysis of the command log.

- Take advantage of the user exit for ADALOG and trap only desired commands.
- NAB is applicable under MVS, TSO and COM-LETE only. There should be at least one NAB for every thread allocated. Each NAB takes 4K of real core.
- Hold queue processing—if a needed record is unavailable the command will be put at the end of the hold queue. The command will be tried 20 times before timing out.
- A4, N4, E4, L3 do not release the record at the end of the update. If they are not released by an ET or manually the hold queue may fill up.

At the conclusion of Greg's presentation, Mel Stertz, Systems Support in the Software AG Los Angeles office, spoke about the functions of his office. He stated that his primary responsibility is to coordinate on-site support.

Bob Taylor, Systems Support Representative, spoke on the design aspect of ADABAS systems. Some of his points were:

- Do not duplicate an old access method system.
- Minimizing I/O will increase performance and reduce response time.
- Avoid Algorithm 4—use super-descriptor.
- Keep ISN lists short.
- FIND SORTED has been improved under V4.
- Match the shorter ISN list against the longer list.
- READ physical—best command in terms of I/O.
- READ logical—reading in the order the file was loaded is better than reading in a purely logical order.
- Updates—performance is directly related to the number of descriptors being updated. Use as few descriptors as possible.
- Adds—also directly related to the number of descriptors.
- Deletes—try setting a delete flag instead of actually doing the delete. The actual delete could be done during a low-use period, probably using batch.

- File design—use the minimum amount of files in the system—more files mean more I/Os.
- Application design—the penalty is the greatest when every file is processed sequentially, doing something to every record.
- DBA should be involved in the system design, even before the walk-through phase.

During these presentations, there were questions from the attendees plus several experiences were related. Overall, the meeting was a very enjoyable and informative one. These tips should make converting to and running under V4.1 easier.

Northwest Region

*Dian Odell
Evans Products*

The topic of the morning was "Managing the Test and Production Environments in Database." The discussion concerned the ideals and assumptions of ideas, plans, and requirements in this area as well as the realities of standards, monitoring, and enforcement. The wide range of approach and attitudes shown in the discussion indicates that this is not an area of absolute truths, but is extremely dependent on the characteristics and mood of each installation.

The informal 'free-for-all' session prompted questions about the hazards of running more than one "parallel" utility at a time; the appropriate dispatching priority and non-swappable condition for MPM; Natural response time and performance, logical vs. physical design—how to proceed; use of dictionaries; DOS problems; dual and single logging, etc. There was also a request from the Canadian members that the USA members use a courier service when sending timely mail items to Canada.

Rich Hahn from Weyerhaeuser discussed some of his company's efforts and discoveries in the area of NATURAL—using

SAS for various analyses of the command log; an online interactive problem tracking system where users can enter their own problems and comments; investigation and solution of RC 145 (HOLD queue overflow); beware the seemingly innocent use of parentheses in the syntax of NATURAL statements—performance results can be drastic; and, broader use of the WHERE clause to improve performance.

In general, the informal discussions at breaks and lunch seemed livelier as members sought out persons with similar situations or problems to exchange ideas and solutions. The network is growing, and we are becoming familiar enough with each other to see the members as extensions of our own local resources.

Pacific Southwest Region

*David Berg
Hughes Aircraft*

The meeting was hosted by Nissan Motor Corporation on January 22, 1982.

The meeting opened with a self-introduction of attendees because there were so many new faces. This was followed by reading and accepting the minutes of the last meeting and incidental announcements.

Jim Crawford, users' group liaison, was introduced and discussed highlights of the 1982 Conference. He indicated that tutorials would be aimed at three levels of user sophistication: new, intermediate and advanced. He also said that 35 papers have been received from users, thus offering attendees a good variety of quality presentations and that workshops were being planned for each of the special interest groups. The next newsletter should contain the technical agenda for the Conference to allow users plenty of time to plan to attend.

Mel Stertz, Western Region's Marketing Support Rep., re-

ported that the local technical reps now report to the regional office to allow better utilization of manpower resources. Technical support will still originate from the Denver office; however, by following the trouble logs generated by local users, the regional office will be able to focus on serious problems and fill in the gaps left by Denver. We were encouraged to call the regional office any time there is a problem in getting satisfactory support.

John Weil added that the technical support survey taken by the regional office indicated that, with a couple of exceptions, users felt Denver was doing a good job.

George Kamida followed with a presentation on ADAMINT utilization at Nissan. He illustrated the hierarchical structure of One Entry Point modules, MULT modules and single views, emphasizing the naming conventions used to maintain consistency and integrity among all the modules. Nissan uses copy libraries for the source modules, tying the copy library names with the Data Dictionary file names and module names. The Data Dictionary is also used to cross-reference the ADAMINT modules and the calling programs.

Following a short break, Jeff Walker of Walker Interactive Products gave a technical presentation on his family of Integrated Interactive Financial Systems (IIFS) which are ADABAS based and whose ADABAS interface was written by David Del Rio of Software AG. He illustrated the modular construction of the system and described the details of the ADABAS interface, which provides an unusually high degree of flexibility in the use of the applications software.

Before adjourning, Walt Hopkins warned everyone that the LP parameter must have a consistent value across all utilities and the MPM or unpredictable AUTOBACKOUT results may occur.

The meeting was adjourned at noon.

ADABAS Product Rep's Report

*Jennifer Lee
Aerospace Corporation*

I would like to take this opportunity to thank the Regional Representatives, Special Interest Group Chairpersons, and those users, who have submitted articles to the Newsletter, for their participation and contributions to the user community. Paul Barbuto deserves special commendation for his work in coordinating the Special Interest Group workshops for the May Conference. I especially would like to thank Alan Shue for reviewing the Conference papers and for his assistance in preparing the Conference agenda. Also, to those users who offered to present topics at the Conference, I would like to extend appreciation on behalf of the Users' Group for their active interest and participation in the Conference.

I have enjoyed serving as the ADABAS Product Representative and offer my continued assistance to the newly elected Representative. I would encourage everyone to support the Special Interest Groups by attending the Conference workshops and expressing a willingness to share in the work.

Hopefully, with the cooperation of the user community, we can look forward to a very active and productive year for the Special Interest Groups.

I look forward to seeing you at the Conference and welcome your comments prior to the Conference with regards to ADABAS and Software AG.

ADABAS Technical Notes

ADABAS Considerations In An Online Update Environment

*Bob Becker
Foremost Insurance*

Recently, Foremost Insurance addressed a major project in the area of policy administration utilizing an ADABAS file in which new business policy transactions would be added, deleted or updated online interactively during normal business hours. Also, a batch job would run daily to pull off these policy transactions so that they could be entered into the evening batch applications. We decided to look at ADAM as the access method for these ADABAS files described above.

Using the unique descriptor, we were able to save the find command prior to adding a policy to this file. That is accomplished by performing the N1 command and then checking the response code for a 199. If one received a 199, the record for that policy is already on file. That is, the record has been added by someone else for the unique descriptor of the ADAM key. Therefore, the find command was virtually eliminated to determine whether or not to add the record.

Another important improvement in using the ADAM ap-

proach is by giving an example. Program or CRT A might do a find command for policy 123. At the same time CRT or program B might do a find command for policy 123. Both would find ISN quantity of zero, that is, no records found. In the next instant, program A would add the record and receive a zero response, program or CRT B would add the record and receive a zero response (providing the descriptor was a non-unique key as in previous releases to Version 4).

Therefore, two people could add the same policy at the same time and not know about each other. Using the unique key nomenclature under ADAM, the find command is no longer necessary and program or CRT B in the above example would receive the 199 response code knowing that it cannot add that record. Therefore, programs become more simple to code and error conditions more straightforward under Version 4 of ADABAS using the unique descriptor method.

However, a word of caution is needed. If a program receives a 199 response code and is an online ET/BT logic user, the transaction timer is started regardless of the fact that the programmer has not received a zero response code on the add. Therefore, the program should issue a BT or ET before starting the next logical transaction. That is because of the transaction timer being started by the N1 command regardless of the 199 response code.

We are quite pleased with the ADAM approach to online

update ADABAS files. However, another word of caution. ADAM files cannot be reorganized using the "reorganize data function" of ADAORD. ADAM files must be unloaded and reloaded to be moved within data storage and cannot enjoy the simple straightforward approach of "reorganized data." As special interest group chairman for performance of ADABAS, I will be addressing these and other topics of discussion in a presentation at the International Conference. I hope to see you there.

Version 4 MPM Dual Disk Logging

*Keith Newsom
Washington State Data
Processing Service Center*

In our implementation of Version 4, we decided to use dual disk logging for our command and protection logs. It was an excellent choice. Dual logging gives us two capabilities that we never had with Version 3: (1) the capacity for an unlimited number of commands per MPM run, and (2) the ability for users to access the command log immediately for debugging purposes.

Timestamps

Version 4 makes considerable use of the Time-of-Day (TOD) clock, a doubleword which contains the elapsed time since the beginning of the century! (See IBM's Principles of Operation manual for more information.) The high-order word of the clock, whose precision is exactly 1.048576 seconds, is used as a timestamp in each command log record and in the header records of dual logs.

When a dual log is formatted, its timestamp is binary zero. When MPM opens a dual log, it writes the current time in the log header record (the 7th record on a 3350). After ADARES copies a dual log, it resets the timestamp to binary zero (SAGUTILS update series 8202, TL1284).

The ADARES CLCOPY and PLCOPY functions copy the older log of a pair if there is anything there to copy (see non-valid timestamps: they are considered logically empty, and no copying is done). If one log has a valid timestamp, it is copied. If both logs have valid timestamps, the log with the earlier timestamp is copied. Thus, two executions of ADARES ensure that both logs of a pair are copied.

Creating dual logs

When creating dual logs, remember that they are similar to WORK datasets. They must be formatted by ADAFRM before they can be used (see Installation Manual, Appendix A, and Utilities Manual, p. 129.1). We currently use 50 cylinder logs on 3350s. Each dual command log for our production MPM (for which we log basic, control block, and I/O information—about 130 bytes per record) holds about 200,000 commands. Each command log for our test MPM, for which we log everything possible so the user can have maximum information, holds about 80,000 commands.

Using dual logs

The ADARUN parameters we use on our MPM step are: LOGGING = YES, DUALCLD = 3350, DUALPLD = 3350, DUALCLS = 8994, DUALPLS = 8994, LOGCB = YES, and LOGIO = YES for both test and production MPMs. In addition, for test MPM we add LOGFB = YES, LOGRB = YES, LOGSB = YES, LOGVB = YES, and LOGIB = YES. Why 8994?



DUAL LOGGING

(cartoon courtesy of Gary Duffield)

Remember that ADABAS never uses the first track of any direct dataset. Therefore, for a 50 cylinder dataset on a 3350: (50 cyls * 30 trks/cyl - 1 trk) * 6 blks/trk = 8994 blks.

Copying dual logs

As a general rule, keep your dual datasets empty, ready for MPM to use (see DBA Reference Manual, pp. 70.1-70A.1). To keep MPM happy, we copy our dual logs at four different times: (1) prior to running MPM; (2) as needed, during MPM; (3) after running MPM; and (4) during daily backup.

(1) Prior to MPM: At the beginning of an MPM run, we run two CLCOPY and two PLCOPY steps to ensure that all dual logs are empty prior to starting MPM. These steps are actually part of a pre-MPM jobstream, so that any possibility of dataset name contention disappears when the pre-MPM job ends. The logs are copied to

the +1 and +2 generations of tape generation groups for the command and protection logs. In its last step, this pre-MPM job submits the MPM job to the internal reader.

(2) During MPM: Whenever a dual log fills up, user exit 2 is called with register 1 indicating which of four possible situations has occurred: 1 full command log; 2 full command logs; 1 full protection log; or two full protection logs. Our exit 2 (which is on the DBA tools tape available from Reston) reads a CLCOPY or PLCOPY jobstream from a PDS member pointed to by a DD card within the MPM job JCL. It modifies the jobname in the JOB and /*SETUP cards so the names follow a pattern. (For example, command log copy jobs for the "Z" data base, our test data base, begin with CZ1, CZ2, . . . CZ9, CZA, . . . CZZ, and back to CZ1 if necessary.) The copy jobname and the MPM's jobname are logged with a WTO. In the rare case that both dual logs are full, the copy jobname and the MPM's

jobname are logged with a WTOR which tells the operator that MPM must wait until the copy job finishes, at which time he should reply OK to the outstanding reply.

Editors Note:

(R1 = -1) means 1 CLOF is full
 (R1 = -2) means 2 CLOG is full
 (R1 = non-zero) means 1 PLOC is full
 (R1 = 0) means 2 PLOC is full

This situation of both logs being full should never be allowed to happen. Unfortunately, it does. We discovered the hard way that copying a log directly to tape was too slow in our shop because of the procedures we require for jobs which create output tapes. Things were fine until the day a user went into a loop and filled



up our command logs 27 times. The operators got bogged in the WTOs and WTOs and replied prematurely to several of the WTOs, whereupon MPM started writing over uncopied information. Fortunately we didn't lose anything valuable, but it forced a quick redesign. Now our copy jobs do a CLCOPY or PLCOPY to generation +1 of a disk GDG and then submit a follow-on job to the internal reader. This follow-on job simply copies the newly-created generation to tape verbatim. The important part, clearing the dual log for MPM's benefit, is done immediately.

(3) After MPM: At the end of an MPM run, a follow-on job is submitted to the internal reader. This job copies all valid disk logs to tape. Of course, if the system should die under MPM, this job doesn't get submitted, but that's not important because when MPM comes up again after IPLing, the pre-MPM job described in (1) will copy the logs instead.

(4) During backup: We also copy any valid disk logs during the daily midnight backup run. Sometimes the backup runs before the post-MPM job described in (3), but it doesn't matter which runs first.

Command log statistics

We have a command log statistics program, written in SAS, which produces an exception report of abnormal commands (undesirable response codes, high elapsed times, batch and TSO open and close commands), and statistical summaries of elapsed and calculated CPU times for each command type and user type. We run this program after each CLCOPY of a valid dual log. We determine whether a log is valid by writing the CLCOPY step's DDRUCK messages to disk,

and then analyzing them with a small DYL-260 program. If the log was not empty, DYL-260 writes the statistics jobstream to the internal reader.

Browsing dual logs

The dual logs must have DISP = SHR in all references by MPM and copy jobs (see Operations Manual, pp. 48-49). This has a side benefit: As soon as MPM writes a log buffer, users can look at it. We have SPF under TSO, so our users can browse a command log within seconds to see where they erred. However, because the log is a bit difficult to interpret this way, we have written a DYL-260 program which selects and formats in-

formation for a desired job, TSO session, or terminal and/or transaction under CICS.

Editors Note:

With ZAP TEM50372 you are required to run MODE = MULTI to prevent copying the active file. This will not be required with SM5.

Summary

I hope this description of our experiences is useful to any of you who are considering dual logging. It is working very well for us and should do the same for you.

ADABAS Special Interest Group Chairpersons

Data Dictionary

Ms. Virginia Sullivan
INCO, Ltd.
1 New York Plaza
New York, New York 10004
(212) 742-4165

Education/ Documentation

Mr. John Kaye
Foremost Insurance Company
5800 Foremost Drive, S.E.
Grand Rapids, Michigan 49501
(616) 942-3364

Performance

Mr. Bob Becker
Foremost Insurance Company
5800 Foremost Drive, S.E.
Grand Rapids, Michigan 49501
(616) 942-3364

Restart/Recovery

Vacant

Retrievals/Interfaces

Mr. John Colthar
R.J. Reynolds Industries
World Headquarters-4-A
Winston-Salem, N.C. 27102
(919) 777-4285

Data Administration

Ms. Carole Schermer
Chemical Abstracts
P.O. Box 3012
Columbus, Ohio 43210
(614) 421-6940

Large Data Bases

Mr. Eric Siefert
CCSI
1 Huntington Quadrangle
Melville, N.Y. 11750
(516) 752-3208

DBA Tools

Mr. James Buckner
World Health Organization
Geneva, Switzerland

or

Mr. Jeff Schutzer
INCO, Ltd.
1 New York Plaza
New York, New York 10018
(212) 742-4584



NATURAL Product Rep's Report

Steve Froneberger
Duke Power Company

I would like to thank Jim Reiner for the excellent job he has done as Chairman of the NATURAL Features Special Interest Group. His articles demonstrate that he not only has a thorough understanding of NATURAL and how it works, but also is aware of the needs and desires of the user community. His comments (as well as comments from the users) indicate the need for NATURAL to be more "user-friendly". Version 1.2 should be available soon and promises to have new features which should aid in end-user understanding.

However, not all of our needs will be answered by this new version. There are two major problem areas which I feel will still exist. First, when processing data from more than one file, the end-user has to understand what happens or how the data will be retrieved. Failure to know this results in very inefficient programs and affects the response for all users. Second, if end-users are allowed to perform updates, they must understand HOLD logic and program accordingly. Failure to do this may result in the HOLD queue filling up and we all know by now what this means.

Until the end-users do not have to be concerned with how the data is processed, NATURAL cannot be given to them without imposing restrictions.

The NATURAL user community must keep Software AG informed of their needs for NATURAL to be an end-user program development tool.

As stated above, NATURAL 1.2 should be available soon. Prices for the NATURAL Security System and NATURAL Advanced Facilities are now available. This is a good indication that we can look for these products soon also.

I would like to encourage all NATURAL users to contribute articles to the newsletter. Tell us how you are implementing NATURAL in an end-user environment, what techniques you have developed for handling the NATURAL System File, good and/or bad experiences you have had with NATURAL, and what DBA tools you have developed to manage the NATURAL environment. Everyone welcomes new ideas about how they can improve their operations.

NATURAL Technical Notes

Natural Design Considerations

Sierra Pacific User's Group

Several users in many regions have expressed a need for Advanced NATURAL Programming training from Software AG. Some have also expressed a need for NATURAL internals. Until such a training project is launched, users continue to learn by trial and error. Regional meetings present an outstanding opportunity to share techniques learned. Here's a few from the last Sierra Pacific User's Group meeting:

1. HOLD LOGIC

NATURAL holds records at the module level when using transaction logic. The SET GLOBALS TR=ON will hold all records in a system hold queue. The hold queue, which can hold a limited number of records, is shared by all ADABAS users who use hold logic. In the event that a NATURAL program holds an excessive amount of records, ADABAS will flush the system hold queue to handle the abnormal volume the NATURAL program generates. This will affect *all* ADABAS users.

All records held in a NATURAL program are released when:

1. ET (end of transaction) is issued;
2. Another NATURAL program is fetched, or,
3. The NATURAL session is terminated.

Editors note: 'NATURAL will not release held records when an 'OPRB=' parameter is coded in the PARM module at generation time.'

Individual records are also released under V3.2.1 only when an A4 (update with hold), is issued; however, NATURAL does not generate A4 commands.

Editors note: 'A fix is now available to issue A4/L4 commands.'

Under ADABAS V4.1, only the ET will release records. Time parameters are also used for holding records; thus, the hold can be broken. Under V3.2.1 a program will terminate with NATO145 when it tries to get to a record already held. Under V4.1, this will not happen since the system hold

queue cannot be monopolized. However, a NATURAL program can still cause ISN queue overflows if it tries to hold too many records. NATURAL under ADABAS V4 requires the same design considerations as NATURAL under ADABAS V3.2.1, when holding records since NATURAL holds records at the module level.

One way to keep from overflowing the hold queue in the instance of a mass file update is to use two program modules to update a file; one to position

in the file and another to perform an update on a single record. This is because the update command generated in a READ or FIND is an A1 which does not release the record from the hold queue. This two-module approach can be used to avoid overflowing the ISN hold queue.

Now, the following program model can be used to pass a file and update selected records without using two separate program modules:

```
0010 *          HOLD LOGIC SUBSET
0020 *
0030 SET GLOBALS TR=ON
0040 *
0050 FIND PERSONNEL WITH AGE GT 22,
0060 IF JOB = 'DATA PROC ANALYST'
0070 DO
0080 GET IN FILE PERSONNEL *ISN (0050)
0090 UPDATE RECORD IN (0080)
0100 WITH SALARY = 40000
0110 DOEND
0120 LOOP
0130 *
0140 END TRANSACTION
0150 *
0160 END
```

Editors note: 'The example shown for hold logic can still blow the hold queue if many records are updated since the END TRANSACTION is outside the FIND loop.'

The preceding program will only hold those records that pass the IF test. It will issue an L4 for the GET in line 0090 and L1's for the records that do not meet the IF criteria. An ET is still needed to release the records held.

ADABAS COMMANDS GENERATED WITH THE ABOVE PROGRAM:

```
L1 (L1 for records not passing IF test)
L1
L4 (L4 & A1 for records passing IF test)
A1
```



HOLD LOGIC COMMAND SUMMARY

ADABAS V3.2.1

A1 - UPDATES RECORD
 A4 - UPDATES RECORD AND REMOVES IT FROM THE HOLD QUEUE
 ET - RELEASES ALL HELD RECORDS FOR USER

ADABAS V4.1.1

A1 - UPDATES RECORD
 A4 - UPDATES RECORD
 ET - RELEASES ALL HELD RECORDS FOR USER

Speaking of ADABAS commands for NATURAL statements, here's what I found:

NATURAL STATEMENT	ADABAS COMMAND	
	TR=OFF	TR=ON (HOLD LOGIC)
FIND NUMBER	S1	S1
FIND FIRST	S1,L1	S1,L1
FIND UNIQUE	S1,L1	S1,L1
FIND	S1,S1,L1,...L1	S1,S1,L1...L1
FIND SORTED	S1,S9,L1...L1	S1,S9,L1,...L1
FIND...SORT	S1,S1,L1...L1	S1,S1,L1,...L1
FIND...DELETE	S1,S1,L1,E1	S1,S1,L4,E1
FIND...UPDATE	S1,S1,L1,A1	S1,S1,L4,A1
STORE	N1	N1
GET	L1	L1
HISTOGRAM	L9	L9
READ	L3	L3
READ...SORT	L3	L3
READ...DELETE	L3,E1	L6,E1
READ. . .UPDATE	L3,A1	L6,A1
READ PHYSICAL	L2	L2
READ LOGICAL	L3	L3
READ BY ISN	L1	L1

BELOW TO BE CUT FOR PERFORATED RULE



2. SORT

Online sorts have been frowned on ever since NATURAL was made available for use in our environment. Online sorts involve a tremendous amount of overhead. For those not familiar with what ADABAS does with an online sort, the following is presented:

A NATURAL program will find all the records in the personnel file with the name between BAKER and DOUGLAS. It will print the names in alphabetical order.

The FIND retrieves the records in *ISN sequence, so without the SORT, the records are not in alphabetical order. The SORT will cause ADABAS to pass the entire normal index for the NAME key and order the found ISNs accordingly, then select the records for the range

criteria in the new sequences. It is important to note that the selection of records is performed after the pass of the normal index. This is the high overhead involved. When a sort is used, the normal index is passed, beginning to end, then the data selection is performed.

This is why sorting should be done in batch. Batch NATURAL will use the OS sort instead of the ADABAS sort. The 'SORT' statement must be used in batch mode, not the 'SORTED' statement. If 'SORTED' is used in batch, the program will function as if an online sort were being processed. So, perform sorts in batch and use the 'SORT' verb when possible.

```
0010 FIND PERSONNEL WITH NAME = 'BAKER' THRU 'DOUGLAS'  
0020 SORT BY NAME  
0030 DISPLAY NAME  
0040 END
```



HOW THE ADABAS SORT WORKS

Example:

Find personnel with name EQ 'BAKER' thru 'BROWN' sorted by age:

1. NORMAL INDEX

(PHYSICAL LAYOUT)

BAKER	21, 32
BANKS	12, 15
BENSON	72
BROOKS	1
BROWN	4, 9, 51, 52

2. ISN SEQUENCE - Returned by FIND

NAME	ISN
BROOKS	1
BANKS	12, 15
BAKER	21, 32
BROWN	49, 51, 52
BENSON	72

The range search returns the records in ISN sequence as in the above right table. In order to sort these by age, the entire associator list for AGE & ISN must be passed. This is the overhead that should be avoided.

Associator by AGE Sequence

Which must be passed

Names sorted by AGE-final output

AGE	ISN	NAME	AGE	ISN
10	14, 7, 9	BENSON	15	72
15	72, 12	BANKS	15	12
25	21, 27	BAKER	25	21
26	32, 3	BAKER	26	32
33	52	BROWN	33	52
47	1, 49	BROOKS	47	1
53	51, 15	BROWN	47	49
62	52, 5, 43	BANKS	53	15
		BROWN	53	51
		BROWN	62	52

Alternatives for online ADABAS SORTS include the following:

Instead of:

```
1570 FIND IN PRECINCT-FILE WITH
1580 ELECTION-PORTION EQ #START-VALUE THRU #END-VALUE
1590 SORTED BY ELECTION-PORTION
```

Use: (If Possible)

- 1570 READ PRECINCT-FILE BY ELECTION-PORTION
1580 STARTING FROM #START-VALUE
1590 ENDING AT #END-VALUE
- 1560* IF USING SUPERDESCRIPTOR
1570 READ PRECINCT-FILE BY ELECTION-PORTION
1580 STARTING FROM #START-VALUE
1590 IF ELECTION-PORTION GT #END-VALUE
1600 ESCAPE
- 1560 FIND NUMBER PRECINCT-FILE WITH ELECTION-PORTION
1570 EQ #START-VALUE THRU #END-VALUE
1580 MOVE *NUMBER (1560) TO #N (N8)
1590 READ (#N) PRECINCT-FILE BY ELECTION-PORTION
1600 STARTING FOR #START-VALUE



3. SEARCH METHODS AND PERFORMANCE

The following program demonstrates time measurements by performing the same search using six different search methods. Note the difference between reading the data records and just reading the keys in the associator from the inverted list.

```
ADHOC
0110*****
0120*
0130*          SEARCH METHODS AND PERFORMANCE
0140*
0150*      1. READ ALL RECORD IN PERSONNEL FILE (1000).
0160*      2. LIMIT RECORD SELECTION TO 100 RECORDS.
0170*      3. LIMIT RECORD SELECTION TO 25 RECORDS.
0180*
0190*****
0200*
0210 WRITE 'METHOD' 15T 'RECORDS FOUND' 30T 'BEGIN TIME'
0220      45T 'END TIME' 60T 'ELAPSED TIME'
0230 MOVE ALL '-' TO #UNDER (AB4)
0240 WRITE #UNDER
0250 SKIP 2
0260*
0270*****
0280 METHOD 1          READ LOGICAL *
0290*****
0300 SETTIME
0310 MOVE *TIME TO #INTIME (A10)
0320 READ PERSONNEL BY NAME
0330 LOOP
0340*
0350 MOVE *TIMD (0300) TO #HOLD-TIME (N7)
0360 REDEFINE #HOLD-TIME (2X #MINUTES (N2) #SECONDS (N2.1))
0370 WRITE 'READ LOGICAL' 15T *COUNTER (0320) 30T #INTIME
0380      45T *TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS' //
0390*****
0400 METHOD 2          READ PHYSICAL *
0410*****
0420 SETTIME
0430 MOVE *TIME TO #INTIME
0440 READ PERSONNEL IN PHYSICAL SEQUENCE
0450 LOOP
0460*
0470 MOVE *TIMD (0420) TO #HOLD-TIME
0480 WRITE 'READ PHYSICAL' 15T *COUNTER (0400) 30T #INTIME
0490      45T *TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS' //
0500*****
0510 METHOD 3          READ BY ISN *
0520*****
0530 SETTIME
0540 MOVE *TIME TO #INTIME
0550 READ PERSONNEL BY ISN
0560 LOOP
0570*
0580 MOVE *TIMD (0530) TO #HOLD-TIME
0590 WRITE 'READ BY ISN' 15T *COUNTER (0550) 30T #INTIME
0600      45T *TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS' //
0610*****
0620 METHOD 4          FIND NUMBER *
0630*****
0640 SETTIME
0650 MOVE *TIME TO #INTIME
0660 FIND NUMBER PERSONNEL WITH NAME GT '
```



```

0670 MOVE *TIMD (0640 to #HOLD-TIME
0680 WRITE 'FIND NUMBER' 15T *NUMBER (0660) 30T #INTIME
0690     45T (TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS' //
0700 *****
0710 METHOD 5             FIND *
0720 *****
0730 SETTIME
0740 FIND PERSONNEL WITH NAME GT ' '
0750 FIND PERSONNEL WITH NAME GT ' '
0760 LOOP
0770*
0780 MOVE *TIMED (0730) to #HOLD-TIME
0790 WRITE 'FIND' 15T *NUMBER (0750) 30T #INTIME
0800     45% *TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS' //
0810 *****
0820 METHOD 6             HISTOGRAM *
0830 *****
0840 SETTIME
0850 MOVE *TIME TO #INTIME
0860 HISTOGRAM PERSONNEL NAME
0870 LOOP
0880*
0890 MOVE (TIMD (0840) to #HOLD-TIME
0900 WRITE 'HISTOGRAM' 15T *COUNTER (0860) 30T #INTIME
0910     45T *TIME 60T #MINUTES 'MINUTES' #SECONDS 'SECONDS'
0920*
0930 END

```

*Richardson
won't compile*

This program was run three times standalone on a IBM 4341. There were no other batch jobs running and no TP activity. The important thing is not so much the elapsed time for an individual search but the relation between the various search methods and the time they use. Note that the HISTOGRAM is quite fast. This is an overlooked verb in most shops. Consider using it.

Method	Records Found	Begin Time	End Time	Elapsed Time
READ LOGICAL	976	22:11:16.8	22:11:37.4	0 MINUTES 19.6 SECONDS
READ PHYSICAL	1000	22:11:36.4	22:11:49.6	0 MINUTES 13.2 SECONDS
READ BY ISN	1000	22:11:49.6	22:12:03.4	0 MINUTES 13.8 SECONDS
FIND NUMBER	976	22:12:03.4	22:12:03.6	0 MINUTES 0.2 SECONDS
FIND	976	22:12:03.6	22:12:18.3	0 MINUTES 14.7 SECONDS
HISTOGRAM	269	22:12:18.3	22:12:21.9	0 MINUTES 3.6 SECONDS

All Searches limited to 100 records

READ LOGICAL	100	22:12:21.9	22:12:23.7	0 MINUTES 1.8 SECONDS
READ PHYSICAL	100	22:12:23.7	22:12:25.1	0 MINUTES 1.3 SECONDS
READ BY ISN	100	22:12:25.1	22:12:26.4	0 MINUTES 1.3 SECONDS
FIND NUMBER	976	22:12:26.4	22:12:26.5	0 MINUTES 0.1 SECONDS
FIND	976	22:12:26.5	22:12:28.1	0 MINUTES 1.6 SECONDS
HISTOGRAM	100	22:12:28.1	22:12:29.5	0 MINUTES 1.3 SECONDS

All Searches limited to 25 records

READ LOGICAL	25	22:12:29.5	22:12:30.0	0 MINUTES 0.5 SECONDS
READ PHYSICAL	25	22:12:30.0	22:12:30.3	0 MINUTES 0.3 SECONDS
READ BY ISN	25	22:12:30.4	22:12:30.7	0 MINUTES 0.3 SECONDS
FIND NUMBER	976	22:12:30.7	22:12:30.8	0 MINUTES 0.1 SECONDS
FIND	976	22:12:30.9	22:12:31.3	0 MINUTES 0.5 SECONDS
HISTOGRAM	25	22:12:31.3	22:12:31.7	0 MINUTES 0.4 SECONDS



4. GET SAME VS. OBTAIN

In the last Newsletter, Jack Reitz of Software AG presented information on referencing occurrence of a MU or PE group in NATURAL without using a GET SAME. Here are the details on why the GET SAME produces overhead to avoid:

GET SAME

```
0010 READ (10) FINANCE BY ISN
0020 FOR J 1 to 10
0030 GET SAME CREDIT-CARD (J)
0040 WRIT CREDIT-CARD (0030/J)
0050 LOOP
0060 END
```

ADABAS Commands Generated:

COMMAND	FNR	ISN
L1	3	1
L1	3	1
L1	3	1
L1	3	2
L1	3	2
L1	3	3
L1	3	3
L1	3	4
L1	3	4
L1	3	5
L1	3	5
L1	3	5
L1	3	6
L1	3	6
L1	3	6
L1	3	7
L1	3	7
L1	3	7
L1	3	8
L1	3	8
L1	3	9
L1	3	9
L1	3	9
L1	3	10
L1	3	10
L1	3	10

reads record with ISN=1
reads occurrence #1 of credit-card
reads occurrence #2 of credit-card

An extra L1 is issued for every GET SAME that retrieves data; the more L1's, the more time used.

OBTAIN

```
0010 READ (10) FINANCE BY ISN
0020 OBTAIN CREDIT-CARD (1-10)
0030 WRITE CREDIT-CARD (1-10)
0060 END
```

ADABAS Commands Generated:

COMMAND	FNR	ISN
L1	3	1
L1	3	2
L1	3	3
L1	3	4
L1	3	5
L1	3	6
L1	3	7
L1	3	8
L1	3	9
L1	3	10

The OBTAIN issues no data base commands. An L1 is issued for each read.

Use the OBTAIN and MOVE INDEXED if you know how many occurrences exist, or if only a few exist. Be aware that the OBTAIN sets aside enough core to hold all occurrences in the OBTAIN, even if only a few really exist. Use the GET SAME if the number of occurrences might vary unpredictably: say from 1 to 70 occurrences.

