

TREETIPS



Issue #13 - October 1992

A Publication of

TREEHOUSE
SOFTWARE

Treehouse Introduces PROFILER!

What is PROFILER?

PROFILER for NATURAL is a powerful tool for monitoring the execution of a set of NATURAL programs. It produces statistics which show:

- How often each program is executed
- How much CPU time each program consumes
- Within each program:
 - How often each statement is executed
 - How much CPU time each statement consumes

In addition to the above statistics, the Profiler for NATURAL also records the elapsed time taken to complete ADABAS calls issued from NATURAL programs.

Benefits of Using PROFILER

PROFILER is useful to any NATURAL site performing the following tasks:

- performance analysis
- debugging
- application sizing
- evaluating the effectiveness of testing procedures
- education

During **performance analysis**, Profiler identifies:

- problem programs and statements
- inefficient code
- inappropriate use of external subroutines

- poor application structure and design
- expensive ADABAS access methods

PROFILER enables you to pinpoint problem programs and statements while the system is being developed **and** once it is running in Production.

During **debugging**, PROFILER traces program flow and reveals statement execution counts. This helps programmers determine if programs are executing properly.

During **application sizing**, PROFILER accurately determines the resources required by an application before it is moved into Production. PROFILER also assesses the impact of new functions, increased number of users, or other changes to the application or workload.

During the **evaluation of the effectiveness of test data and procedures**, PROFILER determines which parts of an application were tested, and which were not. This

helps the programmer identify serious weaknesses in the test data and procedures used to evaluate NATURAL applications.

During **education**, PROFILER provides an insight for both experienced and novice programmers into NATURAL internals for purposes of optimization. PROFILER also makes it possible to assess the performance impact of the NATURAL Optimizing Compiler and other performance tools.

PROFILER is Very Friendly

PROFILER is easy to use, and is the perfect complement to TRIM and other performance monitoring tools. PROFILER is also a perfect complement to N₂O, as it helps the site to confirm that code has been fully tested before migration to the Production environment.

Special Introductory Offer

In order to help make PROFILER successful, we want to give it a happy home. We believe the best home for PROFILER is with sites who have purchased N₂O, or who **purchase N₂O by December 31, 1992**. These customers will be able to purchase PROFILER for a very special (and extremely attractive) price.

To discuss PROFILER, receive sample PROFILER reports or other literature, contact Treehouse Software or your TSI affiliate today!

Table of Contents

Treehouse Introduces PROFILER	1
Editor's Sproutings	2
ADABAS 5.2	3
John Maguire Leaves SAG	5
In Memoriam	5
Asking "Where?", Part II	6
N ₂ O Improvements	8
SECURITRE Improvements	9
DynaDoc	10
Where Are They Now?	11
Affiliate Profile: ITS	13
El Articulo En Espanol, Parte II	14

TREEHOUSE
SOFTWARE

Editor's Sproutings

by Michael Salisbury

Pittsburgh Gets a New Airport

Those of you who have taken trips on USAir have probably been through the Greater Pittsburgh International Airport. With the recent opening of a new airport terminal, Pittsburgh is now set to handle transportation needs of Treehouse Software well into the 21st century. It is the first major airport built in the U.S. since 1974. The new airport is a modern, state-of-the-art facility.

If you happen to be at the new airport someday and have a few hours between flights, give us a call. We're only a few minutes away. We'll come out and pick you up, give you a tour of the Treehouse, and get you back in time for your flight.

On-Site Service!

Although it's nothing as new as the airport, our technical, management, and sales staffs have made many, many trips. Destinations included Florida, Georgia, Delaware, Texas, Oregon, New York, New Jersey, Washington, Washington DC, Oklahoma, Kansas, Virginia, Kentucky, Illinois, and Alberta. Most of the trips were to present, demonstrate, and assist with products. We want to make sure our customers get off on the right foot and get the most from their purchases.

Treehouse representatives also visited **France** to meet customers there and discuss our products with them.

George Szakach, the president of Treehouse Software, even went to Australia to look at some potential products and visit CCA, our affiliate there. George also met some potential customers. **PROFILER**, the product announced on the cover of this issue, came back from Australia with George. We think he carried it in his pouch. Wonder what Customs would say about that?

Software AG in the Spotlight

Software AG continues to make headlines. For example, a recent Computerworld article entitled "ADABAS Stays One Step Ahead" related users' opinions that the strengths of ADABAS outweighed its limitations. This article referred to last year's survey in which ADABAS also ranked tops among the DBMS products. Last year, system security integration and SQL extensions were the biggest weaknesses for ADABAS. This year's biggest weaknesses no longer include security. We believe **SECURITRE**, **AUDITRE**, and **N₂O** are primarily responsible for eliminating the security weakness.

Playing Fair

Some salespeople don't play fair. We understand that a competitor's salesperson told a customer "**Our best customers are old N₂O users.**" To our knowledge, all the old N₂O users (and also the young ones) are still N₂O users and, in fact, are staunch proponents of N₂O. More often than not, if we lose an N₂O sale it is because another product is offered to the customer for little or no money, not because the other product is better. In these cases, management often elects to "shelve the other product until it is as good as N₂O." We've gotten many apologetic calls from customers who say, "We're sorry, but we were forced to buy another product even though we know N₂O is vastly superior."

But don't take everyone's word that N₂O is better. Decide for yourself. Request a trial of N₂O. Request a trial of any other products you are considering. If you like N₂O better, let us know how we can help you convince your management. If needed, we will come to your site and help you convince your management that a solution is needed immediately, and that the best solution is N₂O. We don't want any more of those apologetic calls.

The Elusive Line 4040 Bug

In our last issue, we reported a possible bug in NATURAL which sporadically occurs at line 4040 (hmmm). Bob Ide of UNISYS reported that he tested the information under NATURAL 2.1.7 and had no problems.

The very next day, one of our developers had the strangest experience with one NATURAL program in our software at a user site. A full day of investigation finally got it down to "NATURAL didn't like line 4040 in this one particular program." The developer inserted

(continued on page 12)

TREETIPS

Editing, Writing, and Design
Michael Salisbury

Layout, Production, and Distribution
Nick Viola
Terri Hammerschmitt
Rose Baloga

*Back issues available upon request.
Reference manuals for all products also available.*

Circulation: 8,000

ADABAS 5.2

by Greg Dzieweczynski, State of Minnesota

Introduction

Database Management System applications tend to be **I/O intensive**. In a worst case scenario in which a system is I/O bound, the system fails to function faster due to the limitation caused by this aspect of processing.

An important aspect of performance and tuning is managing your I/O subsystem. The technological advancements to increase CPU processor speed have surpassed those of the I/O subsystem. The wait time due to physical I/Os is becoming a greater factor in increasing system response time and delays. The **elimination of physical I/Os** can make a significant difference in overall system throughput.

DBMSs that adapt to hardware architectural enhancements will deliver superior performance or throughput. The ability of DBMS applications to exploit certain extensions of operating system features create dramatic benefits in preventing or eliminating this bottleneck.

Data / Hiperspaces

One of the greatest ADABAS performance gains can be achieved by reducing the number of physical I/Os through the use of the MVS/ESA architectural facility of **dataspaces** and **hiperspaces**. The MVS/ESA operating system utilizes this new architecture to provide these two new types of storage spaces. Dataspaces and hiperspaces were created to provide an efficient access method to large quantities of data.

Dataspaces are virtual storage or real memory areas. A dataspace, which can be viewed as an extension to an address space, is byte-addressable and backed by a combination of central, expanded and auxiliary storage.

A **hiperspace** is comprised of 4K blocks that are not byte-addressable and reside in expanded storage. Hiperspace data is addressable only by special instructions that move data 4K bytes at a time.

The primary benefit in using dataspaces and hiperspaces is the **elimination of the amount of conventional I/O to DASD**. Enabling data to reside in a virtual storage dataspace in real memory or a hiperspace in expanded memory avoids the conventional I/O to DASD and associated wait time.

"The primary benefit in using dataspaces and hiperspaces is the elimination of the amount of conventional I/O to DASD."

The ability to dynamically cache the Relative Address Block Numbers (RABNs) with the highest read activity into a data/hiperspace will improve the performance of ADABAS.

The optimization technique to identify and select the files with the heaviest read activity and largest amount of generated I/Os is known as **"fencing"**. The utilization of the **ADABAS Dynamic Caching** feature to augment the ADABAS buffer pool manager **reduces the number of physical read EXCPs**, improves the performance of ADABAS, reduces the number of CPU service units, increases the throughput of the database, and decreases response time.

Data In Memory

Using data/hiperspaces provides a performance boost to database and on-line transaction systems and employs the **Data In Memory (DIM)** concept. The (DIM) concept states, "put as much data as you can in processor storage and access the rest as fast as possible." Eliminating the I/O operations by keeping data in processor storage for synchronous and spontaneous access is the primary objective of the DIM concept. The ADABAS Dynamic Caching feature implements the principles of the DIM concept. By embracing the techniques of DIM, response time and throughput can be improved.

The relative access time for disk I/O for a random 4K page from a 3380/3390 disk drive is approximately 15 to 25 milliseconds. Accessing a page from expanded storage, a hiperspace, requires about .50 to .75 milliseconds; and .05 to .09 milliseconds for a dataspace in real memory. The DIM concept succeeds by avoiding the elapsed time delays and CPU overhead associated with executing thousands of instructions in the operating system to initiate and complete I/O operations.

Buffer Pool Management

Another significant enhancement to ADABAS Version 5.2 is the redesigned buffer management concepts and the ability to perform asynchronous I/O. Previous versions of ADABAS halted processing of all commands to the database during a buffer flush. The

(continued on next page)

ADABAS 5.2

(continued from previous page)

queuing caused by a buffer flush at high on-line transaction rates impacted the performance and created fluctuations in ADABAS throughput. In the buffer pool management arena, Upper/Upper headers (square root tables) have been added to the buffer pool to lower the number of loops needed to scan the Buffer Pool (LBP) for a block. To locate a specific RABN in the buffer pool, the new algorithm for the number of loops during the scan is approximately $1/2$ the square root of the number of headers contained in the LBP. If there are 10,000 headers to search, the number of loops is about 50. The old formula was the square root of the number of headers. Using the old method to scan 10,000 headers would require 100 loops. Also, an "importance header" has been added to reduce the scan time to determine if a block can be overlaid within the buffer pool. Applying this formula to 10,000 headers yields 5,000 iterations to locate a block to be overwritten. CPU time in buffer pool management for the ADABAS nucleus has been reduced considerably to deliver a more efficient version of ADABAS.

Asynchronous Buffer Pool

A new **asynchronous buffer pool flush** parameter (LFIOF) is included in ADABAS V5.2. With the asynchronous I/O feature, a database process does not have to wait for the completion of an I/O request before continuing to perform extra beneficial work. Asynchronous I/O means the I/O subsystem can retrieve data or write data while the processor is performing other functions.

Normally, a processor issues an I/O request and then waits until the I/O request is completed. Previous versions of the ADABAS database management system did not contain this feature. During a buffer flush, the processing of commands to the database stopped, causing a queuing condition within the database. Using this standard support of the asynchronous I/O mechanism in the database buffer flush process enables you to **take advantage of the inherent parallelism in your disk subsystem**.

Considerations

The dynamics of the increased real and expanded memory requirements used to accommodate the DIM architecture may influence the amount of paging in the system. A level of awareness should exist to avoid the size threshold at which paging could have a negative impact on response time. The size of the LFIOF parameter can affect the interval of time between buffer flushes, completion time of a buffer flush, IOSQ elongation time on the devices, and database response

time. Too large a LFIOF parameter can cause queuing in the disk subsystem and elongate the device service time. This can cause you to reach the threshold where the database suffers an impact and increases the response time of the application system.

When the database issues a request to a device in the form of an EXCP request, the I/O Supervisor program (IOS) will check to see if there is an active I/O already on the device. Only one I/O from an MVS system can be active on the device at a time. If there is an active I/O on the device, IOS will queue the new request on the UCB (Unit Control Block) and return. After the active I/O completes, IOS will issue a start subchannel command for the next pending I/O.

IOSQ time is the duration an I/O request sits in the UCB's pending queue waiting for service. The goal in DASD tuning is to provide zero IOSQ time, with an average of 1 or 2 milliseconds viewed acceptable. This guideline is now obsolete.

The ADABAS 5.2 version has the ability to continue to process database commands while its asynchronous buffer flush is occurring. The database is now sending multiple read and write I/O requests to a single file. This common phenomenon will appear as an elongated IOSQ time in the device response time. An accurate assessment of the device service time can be obtained by subtracting IOSQ time from average response time (AVG RESP TIME) of the device. (For example, AVG RESP TIME of 156 milliseconds and 136 milliseconds of IOSQ time results in a 20 millisecond device service time.)

Conclusion

The intent of these enhancements is the **removal of the I/O performance bottleneck** and a **delivery of higher performance than would be available** if the enhancements were not utilized. The larger the amount of memory, the more you can store in the data-in-memory (DIM) machine. The DIM technique will improve the overall system performance by using processor storage to decrease the amount and time in physical I/O processing within the system. The LFIOF database parameter allows continual processing of commands in the database during a buffer flush and the ability to leverage the inherent parallelism of the I/O subsystem. The use of these features will **dramatically increase performance and throughput** of the ADABAS database management system.

(continued on next page)

ADABAS 5.2

(continued from previous page)

COMPARISON

ADABAS 5.1.9 / ADABAS 5.2.2

IBM 9000/720 (117 MIPS)

384M Real, 768M Expanded

1 TOR, 3 AORs, CICS 2.1

Data at 15 minute intervals from EPILOG/RMF/OMEGAMON

DYNAMIC CACHING Dataspaces (135M)

	ADABAS 5.1.9	ADABAS 5.2.2	IMPROVEMENT
ADABAS Statistics			
TCB service units	5.84M	5.62M	4%
% CPU utilization	77	76	-
CPU service units/cmd	17.32	13.75	20%
Cmds/sec	375	454	21%
I/O service units/cmd	2.6	2.5	-
LBP (Meg)	4	4	-
LFIOF (Meg)	N/A	1	-
Internal Response Time			
Avg. Response time (sec)	.62	.39	38%

Greg Dzieweczynski is the Technical Manager for Database Services at the State of Minnesota. He has 17 years of IS experience and has worked with ADABAS for 5 years.

John Maguire Leaves SAG

This summer the *Washington Post* exposed a "disagreement" between John Maguire (SAGNA founder, former President, CEO, and Chairman of the Board) and Software AG. We're still sorting this one out, and so are the great legal minds of Reston. What is clear is that John has been "relieved" of his consultant position with SAG.

The official stance from a Software AG spokesperson is: "Software AG thinks it's unfortunate that John Maguire took this route. We would like to resolve this as reasonably as possible. It is not in our interest or John Maguire's to discuss the details because it is in litigation. This lawsuit does not affect the continuing success of our business."

George Szakach, President of Treehouse Software says, "I have known John and SAG for over 17 years. Under any circumstances, if you are SAG, you keep your founder and past President happy because he can do you a lot of good. The last thing SAG needs is another personal vendetta, aimed at their founder (of all people), and of such significance that it gets media attention."

In Memoriam - Barry Lipsker

The SAG community is saddened by the death of a good friend and associate, Barry Lipsker. Barry worked in various capacities for SAGNA, SAGD, and as a consultant over the past 17 years.

Those of us who had the pleasure of knowing Barry were often reminded of his love of rock music, especially 50's and 60's oldies. Barry was a member of a band called "The Thunderballs" in 1966 and played in England, Germany, and Vietnam.

Our sympathies go out to Barry's wife, family, and friends.

TREEHOUSE
SOFTWARE

Asking "Where?", Part II

by Ralph Partlow

Editor's Note: No claims as to the accuracy or completeness of the material contained in this article are made by Treehouse Software. Readers are encouraged to make their own evaluation of the topics discussed in this article, based on their own observations and analysis. The views expressed in this article are not necessarily those of Treehouse Software, its employees, or affiliates.

In my last article, I discussed the business need for Geoprocessing and the potential benefits an organization might expect to achieve. This article will start to explain how this is accomplished in a data processing environment.

First, we must begin to understand the data structures and accessing mechanisms required to support high-volume business transactions.

Traditional geographic data structures are built in such a way that the data placement is dictated by geography itself. If one is drawing maps, deriving routing information, or dealing with the surface of the earth in a geographic sense, where one event or data item is spatially related to the previous, then this structure works well. However, when it is necessary to deal with events which have almost no spatial relationship to each other, then such a structure tends to break down with increased transaction arrival rates.

It therefore becomes necessary to build a data structure which lends itself to high-volume, unrelated events. Database systems lend themselves very well to tabular data, but not well at all to spatial data. The problem then becomes one of representing spatial data in a tabular means.

One such implementation was developed by the Census Bureau in a structure called the DIME (Dual Incident Mapping Element) file.

Basically, this file is made up of three types of records:

- 0-Cells
- 1-Cells
- 2-Cells

Each of these records has a certain range of attributes associated with it.

0-Cells

These records are intended to represent a point in geography, typically the center of an intersection of two or more roads, though they may represent any point on the surface. 0-Cells then have Latitude and

Longitude as their main attribute. Thus, if it were possible to identify a specific 0-Cell out of millions, then it would be possible to identify exactly where that point existed in relation to any other Latitude/Longitude pair.

1-Cells

A 1-Cell is a line which connects two 0-Cells. Typically a road, a 1-Cell can represent any line between two 0-Cells. A 1-Cell is assigned an imaginary left side and right side. Attributes assigned to a 1-Cell include the following, for each side:

- 1-Cell (right or left side) identifier code - "GeoCode"
- From 0-Cell
- To 0-Cell
- Street Name
- From - beginning house number (address)
- To - ending house number
- 2-Cell (one for each "side" or GeoCode)
- Shape

2-Cells

A 2-Cell is a two-dimensional area which must be completely enclosed by multiple 1-Cells. 2-Cells have attributes such as:

- City Name
- Zip Code
- County
- State
- Country
- Flood Plain
- Voting Districts
- Delivery Agent

Given all of the above, it becomes relatively easy to determine the correct "Delivery Agent" when presented with a street address.

In a geographic data structure, these records are arranged in such a way that they will be in close physical proximity to each other to support dealing with the data in a true geographic sense; mapping, routing, etc.

In a mainframe database environment, the structure of these records is changed somewhat as we are not dealing with the data in a true geographic manner. The three record types are "folded" into one type, best described as a 1-Cell.

(continued on next page)

Asking "Where?"...

(continued from previous page)

Each 1-Cell will represent that place which is one side of a street or another linear divider between any two 0-Cells. These 1-Cell records are then referred to as "GeoCodes".

Each GeoCode record is assigned a unique number which designates that particular record and may or may not carry other meaning. In some situations, it might be desirable, if not advisable, to embed information within the GeoCode number, such as Country, State, County, etc. A GeoCode must identify a single side of a single street segment uniquely within the world. This, obviously, would require very large numbers.

When applied in this fashion, the 0-Cells and 2-Cells disappear, their data becoming a part of the individual 1-Cell records. This is a bit redundant and will require more data storage, but the efficiency of retrieving correct information is well worth the price in data storage overhead.

If the proper key fields are available, a simple Find of 1-Cells with a STREET-NAME = "MAPLE" and FROM-ADDR equal to or less than "437" and TO-ADDR equal to or greater than "437" will always find the correct 1-Cell and possibly numerous incorrect 1-Cells.

Selecting a single 1-Cell (GeoCode) is the goal. Incorrect GeoCodes cause the problems.

In virtually all cases, we know the Country and State. This limits the potential volume significantly. Normally, data for different states is maintained separately.

The more information supplied with the address, the more likely:

1. That the correct single GeoCode will be returned
2. That the correct GeoCode will not be returned at all

This second problem is simply due to the probability of invalid data either through incorrect entry or ignorance of the correct data.

If, for example, the real street name was "Maple Parkway North West", then what happens if someone enters "Maple Avenue" or "Maple Street"? Systems of this nature take care of these inconsistencies by breaking street names up into pieces. A single street name as above might reside in four indexes:

- Street Name = Maple
- Street Type = Parkway
- Direction1 = North
- Direction 2 = West

This alone does not solve the problem of finding the correct GeoCode. But, if the system uses only "Maple" to find the 1-Cell entries and then uses "Parkway", etc., to attempt to resolve which of the Maples is the correct entry, then this offers a better chance of getting the proper response.

Other mechanisms are required to effectively resolve street addresses into GeoCodes:

- Synonym Street Names:
US1 = Pacific Coast Highway, etc.
- Place Names:
City Hall
Times Square
- Phonetic Search
- Nearby Intersections

The most interesting of these is Nearby Intersections.

A single question "Could you tell me the name of a nearby street which crosses Maple?" will generally get a positive response. Given that information, it is generally possible to identify the 0-Cell associated with that intersection, thereby giving the exact Latitude/Longitude.

Each GeoCode has associated with it a pair of 0-Cells. Each 0-Cell has a Latitude/Longitude and a beginning/ending address range. For each Street Address and GeoCode "found" which could possibly contain that Street Address, it is possible to develop and approximate the Latitude/Longitude for that address. "Approximate" is the keyword.

There is a formula which would have to be dug out of books and does not belong in an article such as this which, given two Latitude/Longitude pairs and a percentage of distance between these points, will provide another Latitude/Longitude.

Figuring out the percentage along this path is simple:

```
Beginning number = 401
Ending number = 499
House number = 437
(499 minus 401) = Range
Range minus (499 minus 437) = Position
(Position / Range) * 100 = Percentage
```

(continued on page 14)

N₂O Improvements

The last issue of TREETIPS highlighted the release of N₂O Version 3.0, including N₂O/3GL, which supports the migration of 3GL objects (such as COBOL, JCL, or Assembler code) between environments. Even though that was only four short months ago, both N₂O and N₂O/3GL have grown substantially.

N₂O Version 3.1, to be released soon, offers these exciting new features:

- Static SQL Support for DB2
- Migration of MVS PDS Members by N₂O/3GL
- LIBRARIAN Interface for 3GL
- Enhanced PREDICT Support
- Optional Timestamp Comparison Before Migration
- N2OKILL Feature
- Support for Up to 10 Levels of Authorization
- Warnings for Concurrent Checkouts

Static SQL Support for DB2

N₂O has always been capable of migrating NATURAL for DB2 code stored in an ADABAS FUSER. However, until Version 3.1, N₂O did not provide Static SQL support.

As an extension of the Autocompile process, N₂O will now generate a DBRM containing the Static SQL used in the NATURAL for DB2 program. N₂O will then bind the DBRMs into a DB2 Plan.

Migration of MVS PDS Members

N₂O/3GL now supports the migration of 3GL programs and other objects from one MVS partitioned data set (PDS) to another.

The N₂O Development Team has put considerable effort into the design of this capability. They contacted many sites to discuss how 3GL code and partitioned data sets are used in their environment.

For example, many sites use a single PDS for private development areas. This PDS might contain COBOL, JCL, and Assembler code for a particular project. When the code is migrated to Production, the COBOL, JCL, and Assembler members may each go into a separate PDS, even though they came from the same Development PDS. N₂O/3GL will support this type of environment, as well as an environment where more than one PDS is used for both Development and Production.

N₂O/3GL offers complete archiving and recovery of 3GL objects. N₂O/3GL optionally compiles, assembles,

and links 3GL programs. All N₂O/3GL migration activity is recorded for future auditing and reporting needs.

LIBRARIAN Interface for 3GLs

N₂O/3GL also supports the migration of 3GL code stored in LIBRARIAN libraries, as well as 3GL code stored in PANVALET libraries. In these situations, N₂O is the "engine" driving the 3GL migrations, while PANVALET and LIBRARIAN are the "repositories" for the 3GL objects.

PDS, PANVALET, and LIBRARIAN Support for ONE PRICE!

Rather than charge users for three different versions of N₂O/3GL, Treehouse Software bundled support for PDS migrations, PANVALET, and LIBRARIAN into one package. This gives sites the flexibility to handle 3GL migrations in the manner they feel is best, and does not tie them to a particular 3GL migration method or product. In fact, the site may use all three methods if desired.

Enhanced PREDICT Support

N₂O PREDICT migration capability has been enhanced to include the following functionality:

- Option to move all or selected userviews for a file when migrating the PREDICT information for that file
- Selection and loading of PREDICT objects in the "preferred" order: Keyword, User, Database, File, Verification, Relationship, System, Program, Module, and Report
- Option to check for the existence of PREDICT documentation for a NATURAL program prior to the migration of that program

Additional functionality for PREDICT is being considered for future releases of N₂O.

Optional Timestamp Comparison Before Migration

N₂O will optionally verify that the timestamp of object code is greater than the timestamp of the source code, before migration takes place. This helps to ensure that changes made to the source code are reflected in the object code.

(continued on page 11)

SECURITRE Improvements

SECURITRE Version 2.2 offers many interesting new features. These features include:

- Field-level security
- NATURAL Security System (NSS) Conversion Facility
- Improved tracing features and flexibility
- Parameter updates on-line via the Real-time Monitor

Field-level Security

SECURITRE now allows sites to secure specific ADABAS fields against access or modification. The site specifies to SECURITRE which fields on which files require Field-level Security and how this security should function. After checking file security to determine if the user may access the file, SECURITRE will verify that the user has access to the fields requested.

NSS Conversion Facility

Many sites are currently using Software AG's NATURAL Security System (NSS). These sites need a way to provide a **smooth transition from NSS to SECURITRE-based security**. For this reason, SECURITRE now includes an NSS Conversion Facility.

Conversion involves two processes: converting NSS data into SECURITRE for NATURAL parameters, and converting NSS rules into rules to be entered into the site's System Security Facility (RACF, ACF2, or TOP SECRET).

This on-line NSS Conversion Facility generates the rules and parameters needed to accomplish Logon Security, NATURAL Session Initialization Security, Program Security, Run Security, and DDM Security. The output also specifies which users should be granted access to which pseudo dataset names (DSNs), based on current NSS rules.

Several sites have already made use of this facility. Many of them have been surprised to discover that their NSS database still contains access rules for users who left their site years ago! They have discovered one of the benefits of centralized, single rule base security: it is much easier to keep current!

Improved Trace Facility

The SECURITRE Trace Facility has been improved to better enable SECURITRE users and TSI support personnel to diagnose any problems which might arise. This leads to quick solutions to a site's possible security problems.

On-line Parameter Updates

A new SECURITRE Real-time Monitor facility allows updates to most SECURITRE for ADABAS parameters to be made on-line without bringing the database down and up. This feature makes SECURITRE's security more dynamic than in previous releases, and simplifies the security phasing-in process.

ADABAS 5.2 and NATURAL 2.2 Compatible

Of course, SECURITRE Version 2.2 is compatible with NATURAL 2.2 and ADABAS 5.2.

Pricing Change - NOT!

Even with all these new features, it has not been necessary to increase the price of SECURITRE. There is still time to purchase SECURITRE at the current price and keep your costs down. Contact your TSI representative for more information.

Futures

SECURITRE is poised to support the VM/CMS and DOS operating systems. However, we are currently looking for volunteers to assist us in beta testing the product. If you are

interested in helping us, let us know. We can certainly make an attractive proposition for you to thank you for your assistance.

Other enhancements planned for the future include security for ADABAS On-line Services (AOS), security for PREDICT, further integration with N2O security, and NATURAL Utility Security. We are also determining how to implement the 8-steplib capability familiar to NSS V2.2 users. Many other enhancements are in the discussion stages and will be announced when nearer to implementation.

TREEHOUSE
SOFTWARE

DynaDoc & CONSTRUCT: Partners in RAD

Due to the fact that the user community has a great need for the ability to document existing NATURAL applications, Treehouse Software presents DynaDoc primarily as a documentation generator.

The "Top Down" Approach

A little-known fact is that the original DynaDoc was developed to support front-end design and development. This is called the "Top Down" approach in today's DynaDoc.

DynaDoc provides numerous facilities to assist in the design/development effort:

- Functional decomposition diagrams
- Working documents which detail business application level data
- Detailed specifications
- Sign-off documents
- Technical Reference Manuals (which may be printed before programming starts but after skeleton programs have been generated and scanned by DynaDoc), to serve as the detailed specifications for programmers
- User Manuals which may be produced and reviewed by end users prior to development activities

Additionally, DynaDoc supports user-defined entities being entered into the DynaDoc Control File, thereby becoming a part of the application documentation both before and after system development.

Code Generation with DynaDoc

Another little-known fact is that DynaDoc can actually GENERATE SKELETON PROGRAMS!

Once the designer has gotten to the point of defining actual programs and has entered program names, titles, and overview descriptions into DynaDoc, DynaDoc can generate and SAVE skeleton programs into the library related to the application.

These programs have no logic or functionality, as would programs produced by CONSTRUCT or DynaGen, but they do contain:

- A standard Program Comment Block with the program name, title, and overview filled in
- DEFINE and END-DEFINE statements
- INPUT USING MAP statements for each map referenced by the program

- Any other standard references placed in the shell by the DynaDoc Administrator at the site (e.g., SETs, GDAs, LDAs, INCLUDEs, etc.)

DynaDoc supports any number of different shells, e.g., ONLINE, REPORT, UPDATE, etc., which may be tied to a program as the program is being specified through the DynaDoc Edit Program Specification function. Five shells are supplied with DynaDoc to serve as examples.

Each site has complete control over the content of these shells. Therefore, a shell can begin with a CONSTRUCT block. This block can be filled in with everything CONSTRUCT needs to generate a program, including the name of a CONSTRUCT model to be used.

The CONSTRUCT programmer in this environment always starts with an existing CONSTRUCT program. All the programmer needs to do is program the standard user-exit and specify screen and report layouts.

While Rapid Application Development (RAD) products such as CONSTRUCT and DynaGen offer significant benefit, even more can be achieved by having the design tool generate the first version of the program.

Of course, if DynaDoc helps to generate the first version of the program, you can be assured of **Rapid Documentation Development** as part of your Rapid Application Development effort!

Treehouse Software has available the following DynaDoc material, free of charge:

- Product Overview
- Color Demonstration Diskette (3.5" IBM)
- A NEW Sample DynaDoc Output set

You can also call Treehouse to arrange for a free trial, receive the new documentation set, or talk technical about the product with Ralph Partlow.

TREEHOUSE
SOFTWARE



Where Are They Now?

Some of you old-timers in the Software AG world may remember Ralph Partlow. Ralph, a native of Kansas City, Missouri, had been the Corporate Database Administrator for International Nickel (INCO), from 1973 to 1975 and decided in favor of ADABAS as the corporate DBMS. Ralph left INCO to join SAGNA in 1975 to serve as the Vice President of Technical Operations.

He remained with SAG for over six years, leaving in 1981 and going into various SAG-related consulting ventures until rejoining SAG in 1987 to head up development on NATURAL GEOGRAPHIC. He left again in 1988 when the company was completely restructured following the buy-out by the German company.

He once again went into consulting. Between assignments, Ralph got married during one of his many trips to the Philippines. Ralph and his new wife, Elpidia, live in a 38-foot motor home, as this makes it much easier to move from contract to contract.

Ralph's days of moving around may be over, though, as he contracted with Treehouse in April 1992 to direct the DynaMacs project. Since that time, he has devoted his efforts to molding DynaDoc and its related documentation into shape. The new and much improved DynaDoc 3.3.1 is being released this month. DynaGen and DynaMenu are next.

You may contact Ralph by calling Treehouse Software.

Editor's Note: Be sure to see Ralph's articles "DynaDoc and Construct: Partners in RAD" and "Asking Where?, Part II" in this issue.

TREEHOUSE
SOFTWARE

N₂O Improvements

(continued from page 8)

N2OKILL Feature

The N2OKILL feature allows a site to archive unneeded programs in an environment and SCRATCH them from that environment. All of the activity of the N2OKILL feature is recorded in the N₂O audit trail information and will appear on all appropriate reports.

Support for Up to 10 Levels of Authorization

N₂O supports up to 10 levels of authorization prior to migration. This capability allows N₂O to more easily accommodate a site's approval process.

Warnings for Concurrent Checkouts

Some sites allow multiple users to check out the same program for concurrent development activity. If a program is checked out by two or more developers, N₂O will warn all checkout users that other changes or enhancements may need to be assimilated into their program before it may be checked in to the target environment.

This capability helps to ensure that bug fixes applied during the maintenance cycle of a program are also applied to other versions being concurrently developed.

Great News!

The worldwide response to the previous version of N₂O was impressive. Sites in Germany, France, Venezuela, Australia, and South Africa are becoming

very excited about N₂O, and Version 3.1 will add to that excitement.

We are committed to providing the best possible product, service, and support. If you request a free trial of N₂O, we will come directly to your site to assist in the installation, setup, and training process. We have already developed a simple checklist and diagram to help you prepare for our visit. If you follow these steps correctly, we will usually have N₂O configured and **running on your machine in less than an hour.**

One of our users said, "Your competitor only sent us a demo diskette. **You sent an N₂O developer** to do a presentation. Your competitor gave us three references, but we got voice mail from all three people. We were able to **talk to all the Treehouse references** and got **favorable reviews.** The competition's product looks like it requires specification controls, but we don't want all that control. We like N₂O's **flexibility** to let us decide which controls we need."

Act Now!

Due to the high level of response N₂O is receiving, we ask that you call us as soon as possible to schedule your visit by our staff. Demand for site visits is high, and will increase as word about N₂O Version 3.1 spreads!

Be sure to read the cover story about PROFILER for NATURAL. PROFILER is an excellent complement to N₂O, and current N₂O users as of **December 31, 1992**, may purchase PROFILER at a special reduced price. Call for more information.

TREEHOUSE
SOFTWARE

Editor's Sproutings

(continued from page 2)

a comment on this line and the problem disappeared. If your NATURAL program has a problem with line 4040, you might want to see if this solves it.

SECURITRE a Hit at CA Security Conference

Dean Clothier of Kerr-McGee Corporation in Oklahoma City, Oklahoma, and Ken Shaurette of American Family Life Insurance in Madison, Wisconsin, made presentations at the 1992 Computer Associates Security and Audit Conference. Each of them described the security problems their site was experiencing and how these security problems were solved through SECURITRE and ACF2. Response to the sessions was quite positive, and all of the attendees felt that this type of session should be offered at the next conference. Ken volunteered to facilitate a "Birds of a Feather" (BOF) session and/or form a "SECURITRE User Group" for next year's conference. He is asking all SECURITRE users and potential users to get in touch with him to share American Family's experiences with SECURITRE. You may reach Ken at (608) 249-0100, extension 33845.

Customer Says N₂O Does it All!

We recently received a customer's review of N₂O. This customer evaluated N₂O based on the features they needed in a NATURAL Change Management product. They listed 23 critical capabilities, weighted their importance in the decision, and indicated the degree to which N₂O supported these capabilities. N₂O provided full support of 21 of the 23 needed functions, and at least partial support of the remaining two. The N₂O team is seriously evaluating the best way to fully implement these two functions and completely satisfy the customer. We will be happy to provide a copy of this customer's evaluation to you if you ask for an N₂O trial.

DBA to Management: TRIM Beats Competition!

Another customer, a large international manufacturing company, made the decision to purchase TRIM over the competition. According to the evaluation document the senior DBA provided to management, "The support from Treehouse Software has been outstanding, many of the Treehouse staff are ex-Software AG (SAG) employees and claim to drive to maintain an excellent working relationship with SAG. Given the diversified capabilities, ease of use, and stability of the product, I recommend the purchase of TRIM." If you would like to see what the customer had to say about the competition, give us a call. We will provide you with

a copy of their evaluation if you request a trial. However, at their request we have eliminated all references to their company name, location, or employees' names.

TRIM Deals with Stress

Our technical staff spent a considerable amount of time stress testing TRIM versions 5.1.1 and 5.2.1. They fooled the TRIM Real-time Monitor into thinking that some pretty serious ADABAS activity was going on. They uncovered a couple of problems which might be experienced by sites who are issuing about a billion commands a minute, or which had over 44,000 simultaneous users of a single database. TRIM users need not worry, however. The bugs were fixed and TRIM now handles "stress" admirably.

Only Name Brand Presentations From Now On

Treehouse Software will no longer provide user groups with "generic" presentations. We have received too many complaints from users that generic presentations are boring. They tell us that they want to see product specific presentations. Unfortunately, SAG policies prohibit us from making specific presentations at SAG sponsored meetings. We are, however, willing to hold Treehouse sponsored meetings at which we will talk about our products and answer your questions about them. We'll even let SAG drop by to tell you what they have to offer. After all, they're our partners.

New Demo Diskettes

A copy of the full-color, totally reworked N₂O demo diskette was mailed to our primary contact at each known ADABAS site. We also developed an impressive DynaDoc demo diskette. If you did not receive a copy of either demo and would like one, feel free to call us and request it.

Of Course You Can Have a Manual!

One user commented to us that they were pleased that Treehouse Software sends out product reference manuals to potential customers upon request. They commented, "other vendors don't do that." Don't hesitate to ask us for a reference manual if you want to know more about how a Treehouse product is installed and used. We'll be happy to send you one.

Don't forget, with Treehouse, operators are not standing by to take your call. Developers are!

TREEHOUSE
(SOFTWARE)

Affiliate Profile: ITS

Information Technology Services (ITS) recently became Treehouse Software's affiliate for South Africa.

ITS was founded seven years ago, to focus on providing information processing services for distributed systems and mainframes. ITS has grown significantly since then, and is now comprised of four business units: Commercial Systems, Application Management Systems, Special Projects, and Local Authorities. ITS currently has about 75 full-time staff members.

ITS Application Products

ITS develops application systems covering many business functions, i.e., budgeting, general ledger, accounts payable, and payroll. The ITS Local Authorities Division is rapidly emerging as the dominant supplier of municipal systems in South Africa. Their Galaxy application software line consists of four products: Mercury, Saturn, Jupiter, and Venus. Each of these is applicable to the needs of a specific size of government organization. Mercury, for example, is an entry-level system designed for municipalities. Venus is a large-scale system for use at large government sites.

Multiple Platform Support

ITS uses only Software AG, TSI, and their own products to develop its application software. ITS has successfully ported its ADABAS/NATURAL applications across all hardware platforms for which ADABAS and NATURAL are available locally. These platforms include MVS, VM, DOS/VSE, VAX VMS, Siemens BS2000, and Wang VS. In addition, Software AG has recently selected ITS to beta test the Software AG UNIX product line on Hewlett-Packard and Siemens-Nixdorf machines.

DynaMacs

With its emphasis on ADABAS/NATURAL projects, it was logical that in order to establish a more competitive edge in the marketplace, ITS should seek quicker and more efficient methods of system development. A powerful product called **DynaMacs** was developed. DynaMacs grew to include such features as automated documentation (**DynaDoc**), code generation, dynamic menus, prototyping, and context-sensitive help facilities.

Treehouse Software Affiliation

ITS has been marketing DynaMacs commercially for a number of years. ITS has introduced the product to the

international market through **Treehouse Software**. Treehouse Software is currently focusing its efforts on DynaDoc, and the response has been excellent.

Treehouse products are popular in South Africa. Several **N₂O** sales are imminent. **AUDITRE** has helped ITS and a government client uncover evidence in a major investigation of fraud. **AUDITRE** reports are expected to be entered into evidence in the government's case! The other products are also finding good homes with ITS clients.

If your organization is located in South Africa and is interested in Treehouse Software products, contact Paul Maré at ITS.



INFORMATION TECHNOLOGY SERVICES

Phone: (011) 886-7690

Fax: (011) 886-0815

TREQUEST

When we travel to customer sites and far away places, we sometimes bring back items which can be displayed in the office - mostly pictures and posters which we have framed. We want to continue with the decorating scheme, and are requesting reasonably priced pictures or posters of your city, company, office, etc. Logos, flags, banners, and models are good.

We have been sent items before, like chewing gum, coffee mugs, and champagne - usually attached to a change/enhancement request. In some cases, your management has sent us a letter instructing us not to send you gifts. But, they have never said you cannot send us gifts. And, our policy is that we certainly can receive gifts.

Visitors often comment on the unique memorabilia in our office. Favorites include Scott's inflatable penguin and a photo of a truck in Alaska covered (inside) with snow. We are hoping that many of you would like to be remembered in the halls of the Treehouse.

- Emilie Szakach, V.P. & Chief Decorating Officer

El Artículo En Español, Parte II

Mega Soft Computacion CA

La Importancia de una política de "Control de Cambio"

Una mañana muy activa, el departamento de Base de Datos y Soporte Técnico en plena acción tratan desesperadamente de recuperar la versión anterior del programa PRM4000P que determina la ruta de distribución de los productos de la empresa. El Presidente ha llamado varias veces al Vicepresidente de Informática haciéndole saber la cantidad de dinero que se pierde por hora al no poder despachar la mercancía. El PROBLEMA... un cambio en un programa que no fue adecuado y la falta de una POLITICA de Control de Cambio que pudiera detectar y evitar el problema a tiempo.

Se identifica Ud. con esta situación? En la medida que las funciones de negocios son más dependientes de "Sistemas Estratégicos" se hace necesario un control estricto y adecuado de su funcionalidad. Se convierte entonces en una necesidad la implantación de una POLITICA de Control de Cambio que permita establecer para los programas que conforman las aplicaciones un esquema de Cambio controlado, Recuperación, Control de Actualización Concurrente y Operacionalización entre otros.

La mayoría de las empresas en Latinoamérica se han dado cuenta de la importancia de estos esquemas y han creado Gerencias, Departamentos o Jefaturas con la finalidad de establecer una adecuada POLITICA de Control de Cambio. Este esquema por lo general es representado por dos componentes: Normas de Control de Cambio y Herramientas que permitan implantar estas Normas.

En los casos conocidos de éxito, se ha establecido un proyecto piloto que implanta las Normas a través del uso de una Herramienta a lo largo de toda la organización. Familiarizando a todos los usuarios y organizaciones afectadas por posibles "cambios" con las Normas y el uso de la herramienta elegida. De esta manera, se permite la ejecución descentralizada del "cambio" manteniendo al mismo tiempo un control centralizado dentro del Ciclo de Vida de Desarrollo de Sistemas.

Por razones NO necesariamente técnicas nos vemos obligados a responder con velocidad cada vez mayor a los cambios en las reglas de negocios que a su vez afectan las aplicaciones de una Organización.

Principalmente por esta razón vemos como el mundo de negocios obliga a las organizaciones de Informática a seguir un enfoque cada vez más predecible y controlado tanto bajo el punto de vista gerencial como técnico.

"Control de Cambio" a través de políticas y herramientas propias se convierte entonces en uno de los aspectos fundamentales para transformar el desarrollo de aplicaciones basado en un enfoque de artesanía a un enfoque de Ingeniería.

Mega Soft
Computación C.A.

Mega Soft is the Treehouse Software affiliate serving Venezuela and the Caribbean.

Asking "Where?"...

(continued from page 7)

Now we have an approximate Latitude/Longitude for each of the Street Address/GeoCode combinations, and we have a Latitude/Longitude for a nearby intersection. Figuring out which GeoCode is closest to the nearby intersection is very easy.

There are many other intricacies in resolving GeoCodes but if the above is understood, then the concept is clear.

Next time - Coverages, Overlays, and Granularity

Ralph Partlow is an independent consultant with 30 years of data processing experience. He has been involved with Software AG products for most of his career. Ralph is currently leading the DynaMacs project team at Treehouse Software. To reach Ralph, contact Treehouse Software.

TREEHOUSE
SOFTWARE



Treehouse Data Processing Glossary

For the last two issues of TREETIPS, we have included the "Treehouse Data Processing Glossary" and asked for your help in adding to it. We received some funny ones recently:

Abend:	When a "flexible" program stretches just a little too far. (<i>Allen Lamb, Alton Ochsner Hospital</i>)
Assumed Decimal Point:	Located two positions to the right of a programmer's salary in estimation of his own worth. (<i>Modern Data Magazine, per Keith Newsom, State of Washington</i>)
Batch:	A unit of chocolate chip cookie measurement.
Buffer:	One who waxes a car or polishes a floor.
Downtime:	The time between when your program brings the system down and the time that technical support discovers it was you. (<i>Steve D'Aprile, American Fidelity</i>)


Hard Drive:	Duluth to Buffalo in February.
Memory Cycle:	Recollections of your first two-wheeler. (<i>Roger Cole, University of Delaware</i>)
Microfiche:	Guppies, minnows, and the like. (<i>Roger Cole, University of Delaware</i>)
MVS:	How you feel when someone has a better computer than you. (<i>Kevin Ménager</i>)
RAM:	Alamadingdong.
Windows:	One of the perks of a promotion to management. (<i>Steve D'Aprile, American Fidelity</i>)

The Treehouse coffee mugs we spoke about in the last issue have arrived. Those of you whose definitions appeared in TREETIPS, and all of those whose articles were published in the past, should receive one in the mail soon. If you'd like a Treehouse mug, it's time to get your material in the mail to us!

Treehouse Software's Affiliate List

<u>Area Served</u>	<u>Affiliate</u>	<u>Phone Number</u>
Australia	Computer Consultants Australia	61-3-416-3377
Brazil	FYT	55-11-288-1094
France	Fairware	331-39-65-0688
Germany	MaK DATA SYSTEM	49-431-3995-134
South Africa	Information Technology Services	27-11-886-7690
United Kingdom & Eire	OCS Software Limited	44-81-207-5434
Venezuela & Caribbean	Mega Soft Computación C.A.	58-2-483-53-70

To expedite the distribution of new releases of TSI products, please return old Treehouse Software tapes that you may have in your possession. Thank you.



ADABAS, ADABAS Dynamic Caching, ADABAS On-Line Services (AOS), NATURAL, NATURAL Security System (NSS), NATURAL Geographic, NATURAL Optimizing Compiler, CONSTRUCT, and PREDICT are all products of Software AG. UNIX is a product of AT&T. DB2, RMF, MVS, DOS, RACF, and CICS are products of IBM. CA-ACF2, CA-TOP SECRET, PANVALET, and LIBRARIAN are products of Computer Associates. Omegamon is a product of Candle Corporation. The penguin is the property of Scott. Any other product names mentioned are trademarks of their respective holders. The mention of any product name in TREETIPS should not be considered to represent support or endorsement by Treehouse Software Inc., its employees, or affiliates.

Treehouse Software products and services include:

PROFILER - a NATURAL program execution analyzer

DynaMacs - a family of productivity tools for NATURAL 2

DynaDoc - a NATURAL program documentation tool

N₂O - a NATURAL application Change Management system

SECURITRE - a centralized Security Administration package for ADABAS/NATURAL

TRIM® - an ADABAS/NATURAL performance monitor

AUDITRE - an ADABAS auditing tool

AUTOLOADER - an ADABAS file automatic unload/reload/dump utility

Consulting and Customized on-site classes designed to meet unique needs, including:

ADABAS,

NATURAL 2 (Beginning, Intermediate, and Advanced),

PREDICT, and

the Treehouse Software products listed above



Phone:

(412) 741-1677



Fax:

(412) 741-7245

Treehouse Software, Inc.

400 Broad St., Suite 206

Sewickley, PA 15143

FIRST-CLASS MAIL

U.S. POSTAGE

PAID

Sewickley, PA 15143

Permit No. 89