TREETIPS

Issue #15 - May 1993

A Publication of TREEHOUSE SOFTWARE

Consultants' Symposium: A Smash!

As announced in the last issue of TREETIPS, Treehouse Software held its **First Annual International Consultants' Symposium** on April 22 and 23, 1993. Although we were optimistic from the beginning, the response was even better than we expected!

An International Event

The Symposium was truly an international event. The list of attendees included the founders of **Advanced Relational Technology** of Germany, our Brazilian affiliate **FYT Computer Software**, a variety of **consultants**, and a few of our **customers** from around the country. It was quite an impressive guest list! We were even joined by Charles Woratschek and two DPMA students (Lee Ann Burd and Elaine Sluka) from **Robert Morris College**. The entire **Treehouse Software staff** also attended, except those currently out of town on assignment.

Making a Difference, Together

George Szakach, the president of Treehouse Software, opened the event. George mentioned that the Symposium was an idea dating back to 1984, when several consultants discussed meeting annually to share their knowledge and experiences.

George reminded consultants that they are important, not only to the customer, but to vendors like Treehouse Software. "Together," he said, "we can make a difference in the user community. We can help the users, ourselves, and **Software AG**."

Franco Harris at the Consultants' Symposium

Richard Jacobson, vice president of Treehouse Software, presided over the proceedings. Rich welcomed everyone, reiterated the importance of the consultant community, and asked everyone to stand up and introduce themselves to the group.

Table of Contents

The First Day's Presentations

Scott Ferrell of TSI presented **SECURITRE**, fielding many interesting questions. Consultants wanted to know if SECURITRE created or eliminated jobs within sites (it doesn't, but it does help make existing people more effective). Consultants learned about how SECURITRE may secure the NATURAL Command Processor, how quickly SECURITRE reflects SSF rule changes, and that SECURITRE can protect against updates to mainframe data made from workstations.

Joe Gilley of TSI introduced ***adDB2***, the ADABAS to DB2 consolidation package. He discussed reasons why sites consolidate into DB2, how they often go about it manually, problems they encounter, and how *adDB2* is a cost effective alternative.

Mike Coane of Advanced Relational Technology (A/R/T) discussed the philosophy and implementation of *adDB2*. His exuberance for *adDB2* and his knowledge of ADABAS and DB2 immediately grabbed the audience's attention.

The *adDB2* presentation was followed by a **PROFILER for NATURAL** presentation by **Joyce Maguire**, the PROFILER Project Leader. Interest in PROFILER was extremely high, and every consultant appeared to have a story about where they could have used PROFILER, or where they would like to be using it.

(continued on next page)

Consultant's Symposium...

(continued from previous page)

Consultants were very pleased to find that PROFILER doesn't put "hooks" into the NATURAL source code and corrupt it, and that it monitors continuously, rather than taking "snapshots" of activity. They were not surprised to find out that over 120 sites are lined up to take a look at PROFILER.

Candlelight Banquet

On Thursday evening TSI provided a candlelight dinner for the 65 guests and companions. At the end of the dinner, a special guest, **William Westberg**, the Managing Director of the **Pittsburgh High Technology Council**, spoke about the emergence of high technology in the Western Pennsylvania area, and the importance of companies like TSI in this growth.

Franco Harris

Following Mr. Westberg's speech, George Szakach introduced **Franco Harris**, the former **Pittsburgh Steeler** and **Pro Football Hall of Famer**. Franco spoke a little about his career, what he has done since football, and opened the floor for questions. We found out that he is promoting Super Donuts, a vitamin enriched donut. He is also promoting Motorola Personal Pagers and Score sports cards. The attendees really enjoyed Franco's down-to-earth personality. Franco expressed a budding interest in computers, so George invited him to drop by the Treehouse when he is in Sewickley. Franco accepted the offer.

The Second Day's Presentations

Friday began with a very well-received **DynaDoc** presentation by **Ralph Partlow** of TSI. Ralph discussed the importance of documentation, and how documentation is a low-cost, high-return investment. Consultants could see that if management at a client site saw one application documented with DynaDoc, they would want all their applications to be documented the same way. One consultant said to us, "With DynaDoc, documenting might be fun for a change!"

Michael Laskowski of A/R/T made another ***adDB2*** presentation. He discussed the differences between NATURAL for ADABAS and NATURAL for DB2 and the specifics of the *adDB2* solution. Mr. Laskowski pointed out how DB2 performance can often be as good as, or better than, ADABAS performance. It was clear to everyone why Mr. Laskowski is one of the foremost DB2 experts in Europe.

Jim Lumpp, the N_2O Project Leader, presented **N_2O** and **N_2O/3GL**. Jim mentioned during his presentation that although N_2O fully supports the use of PREDICT, he has found that PREDICT is being underutilized by many U.S. sites. The consultants were happy to hear about the continued extensive N_2O development, including the new N_2O/3GL interface to ENDEVOR. The consultants saw that Jim is extremely knowledgeable about change management and N_2O. They quickly realized that with six developers working on it, N_2O will continue to stay far ahead of the competition.

Joe Gilley discussed the specifics of using **Natural ReSource** to convert Report Mode programs to Structured Mode, and shared some entertaining program conversion stories with the attendees. They agreed with Joe that there were several instances in the past where it would have been nice to have had Natural ReSource to recover accidentally deleted NATURAL source code from object code.

The last presentation was **TRIM**, made by **Paul Everson** of the TRIM Project Team. Consultants asked about being able to do screen prints directly from the TRIM RTM, and the possibility of logging Trace information to disk. Paul indicated that we would be investigating these possibilities at some later date. It was pointed out that TRIM development continues, with many enhancements planned for future versions.

Rich Jacobson closed the proceedings, telling consultants that if they would like to participate next year, presenting their own products and services, they should let us know.

The TWIG Bigwig

A consultant, **Dick McCann**, suggested a name for the group of attendees. Mr. McCann suggested calling it the "Treehouse Worldwide Interest Group" or "**TWIG**" for short. We've asked Mr. McCann to be the spokesperson for this group, and he has accepted.

(continued on page 15)

TREETIPS

Editing, Writing, and Design
Michael Salsbury
Nick Viola

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

Back issues available upon request.
Documentation Sets for all products also available.

Circulation: 10,200

First Annual International Consultants' Symposium

TREEHOUSE SOFTWARE

1993

Clockwise from foreground: Doug Wise, Cheryl Stamp, Mike Coane, Fernando & Sandra Boccanera, Sharon Carpenter, Nick Cvetkovic, and Debbie Lape

Clockwise from foreground: Jim Lumpp, Terry West, Carrie Love, Karen Hirst, Don Boring, Nick & Christine Viola, and Roslyn Lumpp

Clockwise from foreground: Harriet Johnson, Steve Robinson, William Anderson, Eric Kruger, Tom Milcic, Dale Farland, and Kim & Jim Rock

Clockwise from foreground: Terry & Rich Jacobson, Kim Canavan, Adam Hammer, Michelle Amato, Joyce Maguire, Rose Baloga, and Karen Dunlap

Clockwise from foreground: George & Emilie Szakach, Jose Freitas, Wagner Martins, Mike Salsbury, Keith Newsom, and William Westberg

Clockwise from foreground: Ralph Partlow, Bill Hill, Michael Laskowski, Joe Gilley, Larry Hemphill, Dick McCann, Dennis Gandy, and Elpedia Partlow

Editor's Sproutings

By Michael Salsbury

Our Products Have Worldwide Appeal

Interest in products continues to be very high around the world. About half of our staff has visited customer sites already this year, covering about a third of the states in the United States and reaching four countries. Our affiliates from France, England, and Brazil have been here, as well as the German developers of *adDB2*. We've enjoyed meeting all of these people, and sharing our experiences with them. We expect to receive visitors from Hong Kong, South Africa, Spain, and Australia very soon, and we have a few overseas trips planned for ourselves.

French Success Stories

Didier Dupeux of Fairware, our affiliate in France, visited the Treehouse headquarters in Sewickley. Fairware has been very successful in selling our products in his country, having made over 30 sales in just two years. For Fairware, N_2O has been a top seller (over 10 sales). According to Didier, the instructor at a recent NATURAL class in France asked four of his students if they had problems managing changes to programs at their sites. In each case, the students answered, "No. We have N_2O, so we don't have any change management problems!"

Brazilian Visitors

J.C. Freitas and Wagner Martins of FYT Computer Software, our Brazilian affiliate, joined us for our First Annual International Consultants' Symposium. They brought us a sample of the Brazilian soft drink named Guaraná, which we all enjoyed very much. After seeing how successful our symposium was, FYT plans to host a similar event in Brazil in September. We will be there.

UNIX Stories

We've been hearing a lot about UNIX lately. Software AG mailed out a flyer showing that ADABAS can perform more transactions per second than Oracle7 under UNIX.

The City of San Antonio, Texas, is putting its critical fire, police, etc., systems on UNIX because they believe it will allow them to keep these systems operational 24 hours a day, 7 days a week, unlike mainframes.

DATAMATION recently reported in January that "The message is clear: alternative platforms are beginning to approach and may soon surpass mainframes for pure throughput." With this in mind, you can be sure that TSI is investigating UNIX as an alternate platform for its products.

TRIM Benefits from Farmer's Advice

One of our salespeople was recently discussing TRIM with a potential customer who was evaluating performance monitors. The customer felt that only a third party's performance monitor could accurately monitor ADABAS and NATURAL. When our salesperson asked why, the customer said that a farmer told him you should "Never hire the fox to guard the hen house!" The customer felt that the same logic applied to choosing performance monitors.

TSI Getting Good Coverage

If you regularly read the new products sections of most data processing magazines, you've probably seen announcements for Treehouse Software products. For example, *adDB2* was announced in the April 1993 issue of Data Management Review. But we are getting other "good coverage", too. Software Magazine called for more information about PROFILER, to be included in an article on application performance tuning. Kerr-McGee Corporation requested permission to speak about SECURITRE at an upcoming Computer Associates security conference. You might also watch upcoming issues of S.L. Robinson and Associates' "Inside Natural" newsletter for NATURAL statement performance information gathered by PROFILER. We also have heard that a very prominent educator in the software development and testing fields tells his students that PROFILER is the only tool available for helping to test and quality assure NATURAL programs. Keep your eyes open. You never know where you will be hearing about Treehouse Software products next!

"It's in There!"

One of the TRIM customers recently called in to share a list of enhancement ideas with us. Three out of her first four suggestions had already been implemented in the current release, which she had not yet installed. This is one of the reasons we urge our customers to "keep current" and always run the latest versions of our products. Nearly every product release includes an enhancement, maybe the very one you're looking for!

(continued on next page)

Editor's Sproutings

(continued from previous page)

What are You Running?

One customer discussed a very specific need with one of our salespeople. The customer wanted to be able to see which NATURAL programs in a given application were actually being executed in the Production environment. They are considering contracting with TSI to develop this as an extension or offshoot of the PROFILER for NATURAL product. If you have an interest in such a feature, please let us know and we will step up the development efforts.

Change Management Evaluator Kits Available

We recently completed development of the Change Management and N_2O Evaluator Kit. The 50-page document contains good general information about evaluating change management products, and suggests some important considerations to keep in mind when evaluating various change management products for NATURAL and 3GLs. The kit highlights important N_2O features a site should be sure to evaluate, based on the recommendations of the N_2O development team.

METASTORE

METASTORE is TSI's new affiliate for Benelux (Belgium, Netherlands, and Luxembourg). METASTORE focuses on mainframe products which enhance functionality, performance, tuning, and interoperability. Their product line includes SQL Statement Prepare, SQL Catalog List, SQL Command Analysis, and the Treehouse product line. Benelux sites may contact:

Mr. Leo Van Dongen
METASTORE
Toekomstlaan 13
B-2200 Herentals
Belgium

Phone: 014-23.29.61
Fax: 014-22.47.70

PROFILER Released!

PROFILER Version 2.0 is officially released! Please call Treehouse to arrange for a free trial or to receive a free PROFILER manual. See the PROFILER article on page 9.

Treehouse Announces User Mini-Conferences

The success of our Consultants' Symposium has led TSI to consider hosting "mini-conferences" in regions of the United States with a relatively high number of ADABAS sites. We are planning to hold our first conference in Washington, D.C., in the fall of 1993. Next year, conferences may be held in Chicago, Dallas, Atlanta, or Orlando. The date for the Washington D.C. event has not yet been determined.

The mini-conferences may include Treehouse Software product presentations and discussions, technical presentations by users and consultants, and presentations by other vendors with products complementary to the ADABAS/NATURAL product line. There will be no charge to the users who wish to attend the mini-conferences.

If you would like to attend our Washington D.C. mini-conference this fall, or if you would like to make a presentation, please contact Michael Salsbury at Treehouse Software.

Keep that Material Coming!

We've recently received some very good articles from our affiliates, customers, consultants, and others. We plan to publish many of them. If you have some material that you think would make an interesting or useful TREETIPS article, please send it to us today!

TREEHOUSE
SOFTWARE

Current Releases

Product	Version
AUDITRE	1.2.0
AUTOLOADER	1.2.0
DynaDoc	3.3.1
N_2O	3.1.0
N_2O/3GL	3.1.0
PROFILER	2.0.0
ReSource	1.2.1
SECURITRE	2.2.1
TRIM	5.1.1/5.2.1

This list shows the versions of our products that you should be running at your site. If you are not running these versions please call Treehouse today.

Can We All Talk - Synchronously?

Gerd Diederichs, Advanced Relational Technology (A/R/T)

In our last issue, we announced adDB2, an ADABAS to DB2 Consolidation project that enables sites to consolidate their data from ADABAS into DB2 without requiring application changes. We mentioned in the article that the performance of applications using the adDB2 Data Bridge is often as good as ADABAS. Several of you have asked us how this can be the case. We asked Gerd Diederichs, an employee of Advanced Relational Technology (A/R/T), the developers of adDB2, to explain the role that Interregion Communication plays in this performance difference.

We would like to take you on a tour through the world of Interregion Communication. In the process, we hope to explain to you a few of those things many people wonder about, like, why is it such a problem for ADABAS to charge its CPU time to the users, while DB2 for instance just adds it to the bill that the user is running up in his own address space? It's got absolutely nothing to do with black magic, as you will see in a moment.

What is Interregion Communication?

Interregion Communication is that set of functions that allows your programs to exchange information with other programs, even though the makers of your operating system - with lots of assistance from the hardware gurus - wanted you to exist in "splendid isolation", so you wouldn't accidentally (or on purpose) step on anything that wasn't yours (and thus hurt someone).

As users of Software AG's products, you have grown quite accustomed to the concept of server systems - those programs that reside in the back of your machine and just wait for you to send them a request of some sort that they can handle for you. Database systems are like that, for instance, as are many others.

The ADABAS Router

And you probably also know that the interregion communication in the Software AG world is handled (transparently to you) by a piece of software called the "ADABAS ROUTER", usually implemented as a user SVC routine for your particular operating system. All the interesting stuff is handled inside that routine and you don't normally get to see it (or need to, for that matter).

In fact, this approach has some very nice properties. In particular, it isolates "client" software (the applications) and "server" software (the databases, etc.) from all the aspects of communication. It even allows you to replace your database nucleus with, say, a NET-WORK nucleus that sends all your database requests to another system. Of course, this approach also makes it extremely easy for us to replace your ADABAS nucleus with something entirely different, such as the *adDB2* Data Bridge nucleus which sends all your requests on to DB2.

Now, we would like to tell you how this interregion communication actually works inside the router.

The MVS manual covering the subject matter calls it "cross memory communications" and distinguishes between "synchronous" and "asynchronous" versions.

Asynchronous Communication

Asynchronous communication is "old hat". If your program runs in some address space and it would like to do something to the data in another address space, it can't just reach over and do it. (It could if you kept all of your interesting data in common storage, but that would be dangerous because anyone could access it or change it, without your knowledge.)

Instead, your program would "schedule an SRB". That is, it tells the operating system to quickly start a little high priority subtask in the other address space running this tiny program here that you loaded into common storage and that does exactly the one thing you needed done over there.

Interestingly enough, a huge portion of SRBs just deal with the posting of ECBs, telling the program in the other address space that you are done with something and that you now expect it to do something else. Usually, there is some protocol to pass additional information, like what exactly you expect done on the other side.

Synchronous Communication

Synchronous communication, in contrast, is something fairly new. It uses a new set of machine instructions that didn't even exist in the early 370 architecture. "Cross Memory Services", "Dual Address Space Feature", or in ESA architecture, "Access Register Mode" are some of the buzzwords to look up if you need to know the whole story. For the purposes of this article, it is sufficient to know that there are instructions that allow you to branch off to a piece of code that lives in another address space, to access data both in your address space and the other address space, and to eventually return to the next sequential instruction in your code.

(continued on next page)

Can We All Talk...

(continued from previous page)

The two main instructions here are appropriately named the "Program Call" (PC) and "Program Return" (PR). With these mechanisms, you can enter another address space's code at those entry points that they authorize for public use, manipulate data in exactly the way that these authorized entry points allow (so security is maintained), and return to your own address space.

DB2 Request Handling

DB2, for instance, uses exactly this method to handle your database requests. Your application calls the provided interface routine, and it performs a PC into one of the DB2 address spaces, gets your data processing done, and returns to you with the results.

ADABAS Request Handling

In fact, ADABAS also uses this mechanism. However, there is a catch. Owing to the internal architecture of the ADABAS nucleus, it wants to pick commands from a queue (the "Command Queue") one at a time.

When you call ADALINK, ADALINK branches off into some code inside ADASVC (without calling it as an SVC routine), where a PC is executed into the ADABAS address space. There, a Command Queue element is filled with your command's data, the nucleus is posted awake if necessary, and control is passed back to your address space. There, you wait for ADABAS to tell you that your command completed and that you can come pick up your data. While you waited, you were still inside ADALINK. Now, another PC is executed into another routine that gets data from the nucleus to your program. (The first PC is the 04-call, the second the 16-call, for those of you who have heard these terms before.)

So, what is the difference here?

ADABAS and DB2 as Public Libraries

Imagine you go into a library to get some information out of a book. Some libraries have a desk at the front where you talk to a librarian who will then go back to the shelves to find your book for you while you wait. (Hopefully, they don't ask you to come back tomorrow.)

In other libraries, they might just give you directions and send you to get the book yourself. (You will still have to show your library card at the desk to check out your book once you find it.) As long as all the shelves are clearly labeled, don't you think that you would get your book faster in the second library, especially if there are a lot of other people in line waiting to get a book (assuming that the first library can only afford one librarian)?

In other words, in ADABAS you still spend a considerable amount of time waiting in line for your information. In DB2, you can just run through the code and get it yourself. Along the way, you might still run into some situations where something is held for update, but that is also true in ADABAS.

Another aspect of this situation is how the CPU cycles are charged back. With asynchronous communications, the server address space spends all the CPU time needed to find your data or perform your update, etc., and therefore needs a way to measure it and charge it back to you.

Servers that use the synchronous method of communications, on the other hand, do hardly any work themselves. In fact, while you are running through the library your clock is ticking, while the librarian just sits there reading a magazine.

Another difference is that you don't get interrupted while you do the processing. In the ADABAS world, you call the nucleus, and you wait. In MVS, once you wait, you eventually find yourself awakened by a POST, but also placed in some queue of many tasks willing to work. It usually takes you longer to get back on a CPU after that than if you hadn't given it up in the first place. In addition, it takes MVS some work to switch between tasks all the time, so it is better for everyone if you don't get interrupted too often.

Add to that the overhead of the instruction path through MVS's WAIT and POST logic (which is surprisingly long), and you find that synchronous communication (the self service way) is much more efficient than the asynchronous alternative. Of course, this does not say anything about the performance of the servers behind the communication.

A Self-Service Database?

So, we borrow the bottom line from a Chinese Deli advertisement, which read:

"We recommend our efficient and friendly self-service."

Gerd Diederichs has spent 12 years of his professional life working for Software AG as a systems developer on COMPLETE, ADABAS, and NET-WORK. He is currently working as a systems developer and consultant on the adDB2 project for A/R/T.

N_2O Version 3.2 News

The N_2O Development Team continues to enhance the product. N_2O is becoming more and more efficient. In the last release of N_2O, the volume of audit records was reduced by 50%, dramatically lowering DASD usage. The efficiency of creating migration requests has been improved by over 30%. But the N_2O Team didn't stop with such "behind the scenes" improvements. N_2O Version 3.2 also includes some great new "visible" features.

Consolidated Migration Requests

N_2O now allows users to migrate all types of objects (i.e., NATURAL, PREDICT, SYSERR, and 3GL) with a single migration request. When users create a migration request, they indicate to N_2O the types of objects they wish to migrate as part of the request. N_2O then goes through the appropriate selection process for each object type and submits the request for approval and processing. This capability will be especially useful when migrating applications into Production, as it allows the entire application to move at once.

SECURITRE Interface

N_2O now allows sites to control the use of N_2O through Treehouse Software's SECURITRE product. Through SECURITRE, a site can use RACF, ACF2, or TOP SECRET to control migration activities and access to N_2O and its functions. Using this interface, sites can completely eliminate the need to use N_2O internal security mechanisms, if desired.

ENDEVOR Endeavor

The previous release of N_2O allowed sites to migrate 3GL objects stored in PANVALET, LIBRARIAN and Partitioned Data Sets (PDSs). N_2O Version 3.2 also allows sites to migrate 3GL objects stored in ENDEVOR.

Migration To/From Partitioned Data Sets

N_2O Version 3.2 allows sites to migrate objects between partitioned datasets (PDSs) and PANVALET, LIBRARIAN, or ENDEVOR. Since this allows a site to migrate 3GL objects from, for example, LIBRARIAN to a PDS, and from the PDS to ENDEVOR, N_2O/3GL can be considered as the "conversion" or "migration" utility. Sites can switch their 3GL repositories easily using N_2O/3GL, providing added value to the site. N_2O/3GL costs less than other tools designed to perform only this conversion/migration function.

Network Distribution Capability

N_2O will automatically distribute migrated objects to networked machines, using a package such as Systems Center's Network Data Mover (NDM). Using this capability, N_2O sites can send objects to remote CPUs via networks. Organizations which do all development on a local CPU, and run production applications on multiple CPUs distributed globally, can use this feature to distribute changed applications to all appropriate production environments.

Six New Reports

N_2O provides six new reports. Three of these reports provide directory listings of NATURAL, PREDICT, and 3GL objects in local and remote environments. These reports allow users to determine what objects are available for migration from a given environment without having to create a migration request.

The remaining three reports provide timestamp comparison between environments. For example, it will be possible for a site to determine if a remote CPU is running the correct version of a production application by comparing the timestamp of the code at the remote site with the timestamp of the code at the local development site.

Development Continues

N_2O development continues, and many more exciting and useful features are in the works. Contact the N_2O Development Team for more details.

TREEHOUSE
SOFTWARE

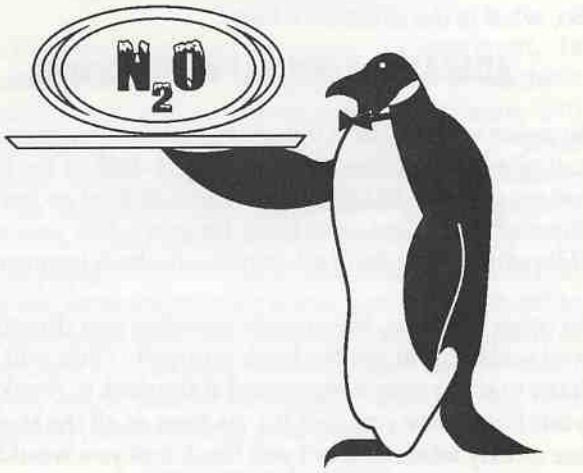

PROFILER for NATURAL

Treehouse Software is pleased to announce the release of Version 2.0 of PROFILER for NATURAL. PROFILER monitors the execution of a set of NATURAL programs and produces statistics for:

- Entire NATURAL applications
- Individual NATURAL programs
- Individual statements within programs

These statistics are available through a wide array of on-line and batch reports. PROFILER assists a site in performing many tasks, including Quality Assurance, Performance Analysis, Application Testing, and Education and Evaluation.

Quality Assurance

PROFILER is an excellent Quality Assurance (QA) tool. The PROFILER Application QA Report indicates the percentage of an application which has been tested, identifies untested program statements, and enables a site to more thoroughly test its code. The Application QA Report page below shows that 74.69% of the programs in an application library have been tested:

```
93-04-19                    *** PROFILER for NATURAL ***                    USER24
10:35:01                      Application QA Report                         PROLIB
Session:  3  TESTPROF                                                       Page 1

   Statistics for Library:  PAYTEST


            Total Objects Reported:                 162
            Total Objects Executed:                 121
            Percentage of Application Tested:     74.69


Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
     HELP         END                          +                      EXIT
```

In addition, the following report helps Quality Assurance personnel to determine which programs require additional testing:

```
93-04-19                    *** PROFILER for NATURAL ***                    USER24
11:09:01                      Application QA Report                         PROLIB
Session:  3  TESTPROF                                                       Page 2

        Statistics for Library:  NATLIB1

                      Run     Executbl   Executbl    %Executbl
Program  Type        Count      Stmts    Stmts Exec  Stmts Exec    Inconsistencies
-------  ----       -------    -------   ----------  ----------    ---------------
PROG001    P            3         54          45        83.3
PROG002    P            2         66          43        65.2     stats/object dates
PROG003    N            2        114          50        43.9
PROG004    M            4          2           2       100.0
PROG005    P            1         76          52        68.4
PROG006    S                     110
PROG007    P            2         23          15        65.2

-------------------------------------------------------------------------------

Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
     HELP         END                    -     +                      EXIT
```

The above report shows that only one program, the map PROG004, has been completely tested.

Performance Analysis

PROFILER helps analyze the performance of NATURAL applications by identifying problem programs, inefficient statements, poor application structure, poor design and expensive ADABAS access methods.

This report shows how the CPU usage was distributed across three programs:

```
93-04-19                    *** PROFILER for NATURAL ***                    USER24
10:43:01                      CPU Time Summary Report                       PROLIB
Session:  2  PAYROLL                                                        Page 1

  Starting Library/Program:  PAYTEST   PAYPROG1     Sort Order: PGM

  S
  *                          Run          CPU
  1  Library   Program      Count    Time (sec)          %CPU    CPU Graph
  -  -------   -------      -----    ----------          ----    ---------
  _  PAYTEST   PAYPROG1         1          0.05         16.67    **
  _  PAYTEST   PAYPROG2         1          0.04         13.33    **
  _  PAYTEST   PAYPROG3         1          0.21         70.00    *******

-------------------------------------------------------------------------------

  Total CPU Time for Session:  0.30
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
     HELP         END                                +     STMT  CPU   ADA   EXIT
```

The program PAYPROG3 consumed 70% of the CPU time used by the three programs profiled. The programmer might therefore want to concentrate tuning efforts on PAYPROG3.

If the programmer wishes to evaluate the ADABAS resources used by an application, the ADABAS Elapsed Time Summary Report may be consulted:

```
93-04-19                    *** PROFILER for NATURAL ***                    USER24
10:43:01                  ADABAS Elapsed Time Summary Report                PROLIB
Session:  2  PAYROLL                                                        Page 1

  Starting Library/Program:  PAYTEST  PAYPROG1     Sort Order: PGM

  S
  *                          Run     ADABAS Elapsed
  1  Library   Program      Count    Time (sec)          %ADA    ADA Graph
  -  -------   -------      -----    ----------          ----    ---------
  _  PAYTEST   PAYPROG1         1          0.12          3.75    *
  _  PAYTEST   PAYPROG2         1          0.20          6.25    *
  _  PAYTEST   PAYPROG3         1          2.88         90.00    *********


-------------------------------------------------------------------------------

  Total ADABAS Elapsed Time for Session:  3.20
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
     HELP         END                                +     STMT  CPU   ADA   EXIT
```

This report shows that 90% of the total ADABAS Elapsed Time consumed by PAYPROG1, PAYPROG2, and PAYPROG3 is consumed by the program PAYPROG3.

Application Testing

The NATURAL Statement Type Reports aid in application testing by pointing out untested programs, identifying code not executed, highlighting weaknesses in test data and procedures, and assessing the impact of new functions, additional users, data base changes, etc. There are four types of NATURAL Statement Type Reports:

- Data Base/Work File Loops
- Internal Subroutines/Non-Procedural Blocks
- FOR/REPEAT Loops
- IF/DECIDE Conditions

(continued on next page)

PROFILER for NATURAL

(continued from previous page)

Education and Evaluation

The Program Listing Report is a useful tool for education and evaluation. It provides insight to NATURAL internals for performance optimization and evaluation of NATURAL statement efficiency.

For example, PROFILER can illustrate the performance difference between two programs which perform the same function but use different algorithms. One program uses the "EXAMINE" statement. The other program uses the "FOR" statement.

When the programmer profiles the program containing the EXAMINE statement, PROFILER reports the following statistics:

```
93-04-19                  *** PROFILER for NATURAL ***                  USER24
10:52:01                   Program Listing Statistics                   PROLIB
Session:  2  PAYROLL                                                    Page 1

        Statistics for Library:  PAYTEST          Program:  PAYPROG8

           Program run count:                                      1
           CPU Time, milliseconds (msec):                          1.79
           ADABAS Elapsed Time, milliseconds (msec):
           Total Statement Executions:                             5
```

The above summary is followed by a listing of the source code for the program, showing the number of times each statement executed, and the total and average CPU time used by each statement. Unexecuted statements are indicated by an asterisk (*) in the "N/E" column:

```
93-04-19                  *** PROFILER for NATURAL ***                  USER24
10:53:01                   Program Listing Source Code                  PROLIB
Session:  2  PAYROLL                                                    Page 2

        Source Listing for Library:  PAYTEST          Program:  PAYPROG8

 Statement  Total      Average              Stmt
 Count      CPU (msec) CPU (msec)  N/E      Nbr    NATURAL Statement
 ---------  ---------- ----------  ---      ----   ----------------------------------
                                             0010   DEFINE DATA LOCAL
                                             0020   1 #ARRAY (A5/50)
                                             0030   1 #NUMBER(P2)
                                             0040   END-DEFINE
         1        .099       .099             0050   EXAMINE #ARRAY(*) FOR 'MONEY'
         1        .002       .002             0060   IF #NUMBER = 0
         1        .069       .069             0070     WRITE 'NO MONEY FOUND'
         1        .000       .000             0080   ELSE
                                           *   0090     WRITE 'YOU ARE RICH!'
                                             0100   END-IF
```

When the programmer profiles the version of the program using the FOR statement, the following statistics are reported by PROFILER:

```
93-04-19                  *** PROFILER for NATURAL ***                  USER24
10:54:01                   Program Listing Statistics                   PROLIB
Session:  2  PAYROLL                                                    Page 1

        Statistics for Library:  PAYTEST          Program:  PAYPROG8

           Program run count:                                      1
           CPU Time, milliseconds (msec):                          2.30
           ADABAS Elapsed Time, milliseconds (msec):
           Total Statement Executions:                             154
```

It is easy to see that it is more appropriate for the programmer to code this particular program using the EXAMINE statement rather than the FOR statement. The 'EXAMINE' statement is more efficient, using only 1.79 milliseconds of CPU Time compared to the 2.30 milliseconds used by the 'FOR' statement.

If the programmer thought that perhaps some other code was responsible for the difference in performance, the detailed listing could be consulted:

```
93-04-19                  *** PROFILER for NATURAL ***                  USER24
10:55:01                   Program Listing Source Code                  PROLIB
Session:  2  PAYROLL                                                    Page 2

        Source Listing for Library:  PAYTEST          Program:  PAYPROG8

 Statement  Total      Average              Stmt
 Count      CPU (msec) CPU (msec)  N/E      Nbr    NATURAL Statement
 ---------  ---------- ----------  ---      ----   ----------------------------------
                                             0010   DEFINE DATA LOCAL
                                             0020   1 #ARRAY (A5/50)
                                             0030   1 #NUMBER(P2)
                                             0040   1 #I(P2)
                                             0050   END-DEFINE
        51        .285       .006             0060   FOR #I 1 50
        50        .337       .007             0070     IF #ARRAY(#I) = 'MONEY'
                                           *   0080     WRITE 'YOU ARE RICH!'
                                           *   0090     ADD 1 TO #NUMBER
                                           *   0100     ESCAPE BOTTOM
```

This detailed listing shows the programmer that it was in fact the FOR statement which caused the significant performance difference.

The Right Way to Monitor Your Code

As you have seen, PROFILER provides useful reports that can help you tune and test your NATURAL applications. PROFILER is also the "right way" to monitor your NATURAL programs.

The PROFILER user interface is written in NATURAL, so its menus and reports are familiar looking and easy to use. However, all of the collection of statistics is performed by an efficient Data Accumulator written in Assembler.

PROFILER does not require any changes to be made to your application programs. No "probes" or "hooks" need to be inserted into the programs, and no special compilers, etc., are used or required.

PROFILER does not rely on figures extrapolated from a few "snapshots" of performance data. When activated by the user, PROFILER will monitor every program and statement execution, continuously collecting performance and testing information. Thus, PROFILER's statistics are very accurate and reliable.

Available Now for Trial

PROFILER is available for trial and operates in many MVS environments, including TSO, Batch, CICS, NATURAL 2.2 and 2.1, ADABAS, VSAM, and DB2.

Call Treehouse Software for more information.

TREEHOUSE
SOFTWARE

Treehouse Growth Report

TSI continues to branch out with new products, such as PROFILER for NATURAL, in new directions, such as DB2 with *adDB2*, and with new faces, such as Guil Hastings, Harriet Johnson, Debbie Lape, and Jim Rock.

For some time, **Guil Hastings** (pronounced "Gil") contracted with TSI through his company Sinolinc to help us locate affiliates in Asia. Recently, Guil decided to join TSI as a full-time employee. See the article by Guil below.

Harriet Johnson recently left Software AG in Reston, where she worked on ADAESI (security) and other areas. Her background is in Assembler and systems software. If you received the March 1993 issue of Software AG's *Enterprise Systems Today* newsletter, you probably read Harriet's article about ADAESI. At TSI, Harriet is involved in the development of the Assembler portion of PROFILER, and will provide technical assistance for SECURITRE, TRIM, and other products.

Debbie Lape joined TSI in January. Her initial duties include N_2O development. Prior to joining TSI, Debbie worked as a consultant for the Westinghouse Distribution and Control Business Unit, where she had the opportunity to use N_2O extensively. She has three years of experience in NATURAL and 8 years of experience in Customer Support.

Jim Rock joined TSI from US Steel. Jim will be overseeing TSI's Quality Assurance and Beta Testing programs, ensuring that our products are of the highest possible quality and that our beta sites put them through their paces. Initially, Jim has focused on PROFILER for NATURAL. Later, he will branch out to the other products.

Guil Hastings' Adventures in China

While I was living and working as a consultant for US companies in China, TSI asked if I could help find them an Asian affiliate. We discovered a very qualified company, System Networks Ltd. in Hong Kong. This gave TSI an ideal location from which to service much of Southeast Asia. And since English is a primary business language there, communication would not be a problem for TSI staff. So far the relationship has worked quite well.

Language was a key factor in deciding where to locate an affiliate. I knew from my experience with Chinese that Asian languages can be extremely difficult to learn. Speaking Chinese is actually not so difficult, it is reading and writing that pose the true challenge. Each Chinese character (there are 10,000) represents a sound and a meaning. But as it is paired with other characters its meaning may change. Like any language (computer or otherwise), Chinese requires tons of memorization and a bit more patience.

I finally decided to join TSI in March of this year. The decision was difficult because of the perks of living abroad that I would miss, not to mention leaving close friends behind. But the timing was right and I figured that the biggest perk, travel, was something I could count on in my new position with TSI.

My traveling in Asia has included places as infamous as Beijing and Bangkok, as remote as Siberia, and as strange as Harbin. Most travel focused around trains. Included in this was a trip I took on the Transiberian railway from Beijing to Moscow. Six days and nights crammed in a rail car with hundreds of Chinese, Russians and Mongolians could make anyone yearn for the wide open spaces of even New York City! Each place, and the trip to get there, was wonderfully unique, and full of odd adventures, such as the time the water buffalo tipped over my dining table in a restaurant in China, or when I was chased for two miles by wild dogs. Actually, it makes me think that traveling for TSI might actually be a lot better!

After three long years in Shanghai, it did feel like it was time to leave. The crush of people and the pollution were both tiring. It takes more than the architecture and mystique of its history to make up for the difficulties of daily life, and eventually it takes its toll. Nevertheless, it is a bittersweet farewell.

I do look forward to returning to China and seeing all the changes that will inevitably take place. Hopefully it will be to open the new People's Republic of China affiliate for TSI...but that is quite a few years away.

Guil will be handling our international marketing operations, joint ventures, communicating with our current affiliates, helping us find new affiliates, and improving communication between TSI and the international affiliates.

NATURAL Programming Tips

K.W. Patrick Ho, Manager, Data Management, Tax and Revenue Administration, Edmonton, Alberta, Canada

Treehouse Software makes no claims as to the accuracy or completeness of the material contained in this article, and the views expressed in this article are not necessarily those of Treehouse Software, its employees, or its affiliates.

Given a PE group defined as:

```
01 PE-GROUP (n)
   02 FIELD-A
   02 FIELD-B
```

When you DISPLAY all the elements of the PE group using the DISPLAY PE-GROUP(*) statement, the output will look like this:

```
  FIELD-A          FIELD-B
-----------      -----------
FIELD-A (1)      FIELD-B (1)
FIELD-A (2)      FIELD-B (2)
...              ...
FIELD-A (n)      FIELD-B (n)
```

If you use WRITE WORK FILE 1 PE-GROUP(*), NATURAL will write to the work file in the order of the individual fields, i.e.:

```
FIELD-A(1), FIELD-A(2),...FIELD-A(n),
FIELD-B(1), FIELD-B(2),...FIELD-B(n)
```

If you were exporting data out to COMPRESS and build another ADABAS file, or exporting data for PC or COBOL applications which prefer to view one occurrence of all the fields at a time, you would want to have the data written in the order of group occurrences, e.g.:

```
FIELD-A(1), FIELD-B(1),
FIELD-A(2), FIELD-B(2),
... FIELD-A(n), FIELD-B(n)
```

In order to do that with the original view, you would have to code:

```
WRITE WORK FILE 1 PE-GROUP(1)
                  PE-GROUP(2)
                  ...
                  PE-GROUP(n)
```

This will be a nuisance when you are writing out 99 occurrences, or when you want to change the occurrence count for different applications.

To make the process flexible, we can make use of a temporary array structure to rearrange the output grouping.

Take the MAJOR-CREDIT PE group in Software AG's demo FINANCE file as an example. Writing out MAJOR-CREDIT(*) to a work file will put the data in the order of CREDIT-CARD(*), followed by CREDIT-LIMIT(*), then CURRENT-BALANCE(*).

The following example will demonstrate how we can write to work files by group occurrences in a flexible manner:

```
0010 *
0020 DEFINE DATA LOCAL
0030 01 #CC-CNT (P3) CONST <3> /* Cr Card Count
0040 *
0050 01 FINANCE VIEW OF FINANCE      /*************
0060   02 MAJOR-CREDIT (1:#CC-CNT)   /* Original
0070     03 CREDIT-CARD               /* Structure
0080     03 CREDIT-LIMIT              /* in the DDM
0090     03 CREDIT-BALANCE            /*************
0100 *
0110 01 #WORK-CREDIT (1:#CC-CNT)     /*************
0120   02 #WORK-GRP (A26)            /* Different
0130   02 REDEFINE #WORK-GRP         /* Structure
0140     03 CREDIT-CARD (A18)       /* to be
0150     03 CREDIT-LIMIT (N4)       /* written to
0160     03 CURRENT-BALANCE (N4)    /* work file
0170 END-DEFINE                     /*************
0180 *
0190 ****************************************
0200 *
0210 * THE FOLLOWING LOOP WILL WRITE TO WORK FILE:
0220 *
0230 * CARD 1, CARD 2, CARD 3, LIMIT 1, LIMIT 2,
0240 * LIMIT 3, BAL 1, BAL 2, BAL 3
0250 ****************************************
0260 READ FINANCE
0270   WRITE WORK FILE 1 MAJOR-CREDIT(*)
0280 LOOP
0290 *
0300 ****************************************
0310 *
0320 * THE FOLLOWING LOOP WILL WRITE TO WORK FILE:
0330 *
0340 * CARD 1, LIMIT 1, BAL 1, CARD 2, LIMIT 2,
0350 * BAL 2, CARD 3, LIMIT 3, BAL 3
0360 ****************************************
0370 READ FINANCE
0380   MOVE BY NAME MAJOR-CREDIT(*) TO
0390     #WORK-CREDIT(*)
0400   WRITE WORK FILE 2 #WORK-CREDIT(*)
0410 LOOP
0420 END
```

This article was provided to TREETIPS by:

K. W. Patrick Ho
Manager, Data Management
Tax & Revenue Administration
Sir Frederick W. Haultain Building
9811-109 Street
Edmonton, Alberta
CANADA T5K 2L5

TREEHOUSE SOFTWARE

AIR Out Your Databases

Tom Visher, A.N.P. Solutions, Inc.

Treehouse Software makes no claims as to the accuracy or completeness of the material contained in this article, and the views expressed in this article are not necessarily those of Treehouse Software, its employees, or its affiliates.

The DASD residing at most sites represents a significant investment by the Data Processing Department. Many sites are concerned about the return they are getting on this investment. If much of the DASD is being used to store records which are infrequently (or never) being accessed, it makes much more sense to archive those records to tape. Tape is less expensive to purchase, requires less resources to maintain, and is a reliable storage medium for a site's less critical data.

It is this philosophy that led to the development of my Archive, Index, and Restore (AIR) System.

What is AIR?

AIR is a Tape ARCHIVE, INDEX & RESTORE System for ADABAS Records. AIR identifies records of Logical ADABAS Files that would be stored on Tape Media at 1/10 of the DASD storage costs.

AIR features an Index for determining the existence of records on Tape. Operators can quickly Query the AIR System through an On-Line Screen. A Set of Records from related files will be Restored to the database at Near-Line Speed (minutes) on 1 Restore Command.

AIR supports simple file structures including multiple logical files on 1 physical file. There are periodic groups and individualized Archive and Restore Criteria for each Logical File. Fully relational business rules and referential integrity rules are supported.

File modifications to ADABAS Files are supported in the Archive. Records from In-House Archive Systems can be incorporated into AIR. AIR's Index has capacity for: Hundreds of Millions of Records and Thousands of Logical Files.

AIR is an automated, efficient System with a data-driven, modular structure that is designed to have its functions customized.

Benefits of Downsizing Files

The advantages of Downsizing Files include:

- Reduced Operating Costs
- Improved Performance
- Easier Database Administration

One of the realities of data processing is: "Larger Is Expensive". This has driven the development of a comprehensive Archive System. As files get larger, database utilities take a longer time-window for running. Sometimes this requires the target file to be Off-Line, or even takes the whole database Off-Line. When you are dealing with a file corruption emergency and the files must be put back On-Line right away, the database administrator may have his repair options limited.

When a database experiences "performance problems", this can be an uncomfortable time for the DBA. Sometimes it seems "all eyes are on you" while the tuning process takes place. At one time or another, an effort is made to strategically place the files on each DASD Volume. The Largest Files are usually treated with different criteria than the smaller files.

Small files are usually easier to deal with when solving the "performance problem". Many times, a downsized file experiences improved performance simply because there is a higher concentration of active records in the buffers.

Another benefit of downsizing files is that you don't have to buy DASD as frequently as in the past. Buying DASD usually translates into weekend work for the DBA as he goes through the procedure for "Increasing" the database.

Benefits of Archiving Records

The advantages of Archiving Records Include:

- Putting Inactive Records on Inexpensive Tape
- Keeping Only Active Records on Premium DASD
- Removing the 16 Million Record Limit on Files
- Reducing the Physical I/O Bottleneck

For an archive system to work well, one must identify inactive records. Inactive records would be classified as having a low probability of being read in the short-term. Moving them to a tape archive conserves expensive DASD. There is minimal impact of not having the inactive records at your fingertips at sub-second speed. This kind of archiving can reduce operational costs considerably. The addition of DASD to accommodate file growth is expensive and labor-intensive.

(continued on next page)

AIR Out Your Database

(continued from previous page)

An ADABAS Physical File has the capacity for about 16 million records. In the past, when a file began to reach this limit, the least-important records were deleted. Now there is an alternative to deletion - Archive! Retention justification criteria only has to exceed the cost of a tape instead of the cost of DASD, which is 10X more expensive.

This is the "Information Age". Valuable Data stored cost-effectively gives a competitive edge to your organization.

As the proportion of active records in ADABAS increases, ADABAS Buffers will tend to hold a higher percentage of records that satisfy a Read-I/O Command. These are called "Logical I/O's" because no action is necessary to do the following "Physical I/O" Tasks. These tasks are some of the slowest operations of a computer and referred to as the "Physical I/O Bottleneck."

1) Take your turn going down the communication channel to the disk volume being addressed.

2) Mechanically position the Disk's Read-Head over the designated cylinder.

3) Wait for the Platter to mechanically rotate the designated track under the Read-Head.

4) Read a RABN of approximately 5,000 bytes of compressed records at the rate of the mechanically spinning Platter.

5) Send the RABN of data up the channel to the ADABAS Buffer.

6) Wait for the application program to be awakened to accept the 1 record.

On-line response time is very sensitive to I/O speeds. Successful ADABAS tuning involves improving the probability that when records are requested the disk's Read-Head will respond with a minimum of physical movement. Even better, is to have the records sitting in the buffers already.

Signs of Needing AIR

You know you need AIR when:

- Legal and/or accounting departments require that data be maintained for several years, even when that data is not frequently used.
- Your existing archive system takes hours to locate and restore sets of records.
- The number of records which have not been recently accessed is increasingly rapidly.
- On-line access is slower, and batch jobs are taking longer, because files are very large.
- Database and application efficiency is dropping because I/O and CPU usage for the database is increasing.
- You are adding DASD to the database every year.
- File tune-ups require a large time window.
- Record counts are reaching 1 million and files are just getting tighter and tighter - the database is running out of space.

Features of a Practical Archive System

AIR is a practical archive system which offers the following features:

- Easy to use, so people will want to use it.
- Offers on-line speeds for inquiry to determine if there are records on Archive.
- Provides near-line speeds for restores, so that you have the records available to you in a fairly short period of time.
- Handles hundreds of millions of records, and thousands of files.
- Minimizes human involvement by being as automatic as possible.
- Offers efficient movement of records from the database to the Archive, and from the Archive to the database.
- Stays in sync with database modifications and has the intelligence to deal with related sets of records.
- Pays for itself in reduced DASD requirements, increased database efficiency, etc.
- Is robust and rugged, doesn't easily break, and handles heavy use
- Provides simple recovery procedures in the event of failure.
- Incorporates the Archive and the Restore Business Rules for each file and contains the referential integrity and the relational parameters that apply.

Best of all, AIR is available today!

For More Information, Contact...

Tom Visher
A.N.P. SOLUTIONS, Inc.
Denver, Colorado
1-800/ 359-6209

TREEHOUSE
SOFTWARE

Consultants' Symposium: A Smash!

(continued from page 2)

Highly Rated

Each attendee had the opportunity to rate the speakers in a variety of categories. The general reaction was that the presentations were very professional, interesting, and informative, and that the products were excellent for TSI/consultant "partnering".

For example, **Steve Robinson** of **S. L. Robinson and Associates,** was very excited about the possibilities PROFILER offered, and has decided to take PROFILER into selected sites for the purpose of measuring the performance of NATURAL statements for his own educational pursuits (such as his *"Inside NATURAL"* Newsletter).

Kim Canavan approached Joyce Maguire to suggest another possible use for PROFILER. Kim thought that PROFILER would be a great tool to use for proving that her users tested the work she had done for them. PROFILER would help Kim to influence users to take the responsibility for testing her work. We think this is a great idea, and something more consultants and developers should consider.

Prizes for Some, but Something for All

Attendees received a Treehouse tote bag containing handouts, Treehouse literature, and a Treehouse coffee mug. On Thursday night, there were two drawings: one for TSI employees and one for the attendees. Each winner received an official NFL football, personalized and autographed by **Franco Harris. Ralph Partlow** won the employee drawing and **Cheryl Stamp** won as an attendee. Since Cheryl's so lucky, we decided to hire her. Everyone who attended got their picture taken with Franco, and also received a personalized autographed photo of him, courtesy of Score sports cards.

Same Time Next Year?

Consultants interested in attending or presenting at the Second Annual International Consultants' Symposium should contact **Michael Salsbury** at Treehouse Software.

TREEHOUSE SOFTWARE

Treehouse Software Affiliate List

Area Served	Affiliate	Phone Number
Australia	*Computer Consultants Australia*	61-3-416-3377
Benelux	*METASTORE*	32-14-23-29-61
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
Southeast Asia	*System Networks (Asia) Limited*	852-523-6949
United Kingdom & Eire	*Blenheim Software Ltd.*	44-75-357-1770
Venezuela & Caribbean	*Mega Soft Computacion C.A.*	58-2-483-53-70

Treehouse Software products and services include:

adDB2 - ADABAS to DB2 consolidation product/service
PROFILER - NATURAL program execution analyzer
ReSource - NATURAL source code conversion and recovery
DynaDoc - NATURAL application documentator
N_2O - NATURAL application change management system
N_2O/3GL - 3GL support within N_2O
SECURITRE - Centralized security administration for ADABAS/NATURAL
TRIM® - ADABAS/NATURAL performance monitor
AUDITRE - ADABAS auditing tool
AUTOLOADER - ADABAS file automatic unload/reload/dump utility

Consulting and Customized on-site classes designed to meet your unique needs.

Phone:
(412) 741-1677

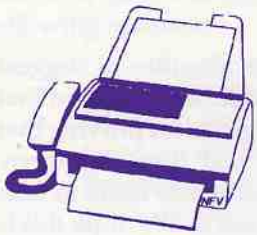

Fax:
(412) 741-7245

Treehouse Software, Inc.
400 Broad St., Suite 206
Sewickley, PA 15143

FIRST-CLASS MAIL
U.S. POSTAGE
PAID
Sewickley, PA 15143
Permit No. 89