




I. P. Sharp

newsletter

JULY/AUGUST 1983
Volume 11 Number 4



in this issue

Global Limits Control System	1
CAPERS	3
New Release of SUPERPLOT.....	4
DEXUS: A System for Telecommunications Circuit Design	5
New Singapore Stock Exchange Data Base	7
Lundberg Survey's Share of Market Now On Line	9
STATSBUND Data Base	10
Join the APL Users' Group: SIGAPL ...	11
APL84 Call for Papers	11
New Local Dial-Up Numbers	12
New PROMIS Offices	13
Singapore, Toronto	13
<i>A Working Introduction to APL Now in German.....</i>	13
Bookends.....	14
Technical Supplement 45	
CPU Limits in SHARP APL.....	T1
Making Things Fit, Part 1	T2

Global Limits Control System

I. P. Sharp is a leader in bringing together a reliable worldwide communications network, application software products, and one of the world's largest collections of online numeric data bases. By integrating these products and services, I. P. Sharp has developed software packages that are changing the way businesses operate.

To help banks stay competitive in today's dynamic financial environment, I. P. Sharp's Global Limits Control System can give accurate and timely information on their worldwide lending and trading operations so they can make effective decisions. Global Limits enables banks to monitor their exposure to different countries, institutions, and currencies minute by minute. By monitoring their risk exposure, banks can diversify their lending, and thus minimize their exposure in any one area.

Banks customarily set lending limits or ceilings which all their trading offices must adhere to. The head office usually sets the limits for each institution, country, currency, and financial market that the bank deals in.

... Global Limits Control System

Banks that operate internationally make thousands of financial deals each day (many with different branches or subsidiaries of the same client) from various locations. A typical transaction might involve a bank's London trading office's wanting to make a \$10-million money market deal with a Mexican bank. Without the Global Limits Control System, the London office cannot find out *immediately* if the transaction will exceed overall trading limits either for the Mexican bank or for Mexico as a whole. Telexing the head office introduces long delays, because the head office paperwork is likely to be several days to several weeks behind. *The Wall Street Journal*, in a front-page article on the Global Limits Control System, quoted a spokesman from a British clearing bank as saying that "we're drawing on information that's 75 days old—and we're the rule rather than the exception."

The solution for many banks has been to assign limits to each trading office. Local limits, however, can restrict the use of available funds. For example, there may be capacity unused in Hong Kong that could be used in Paris. To overcome underutilization of funds, local limits may be increased, but this can introduce other problems: if all trading offices were to reach their dealing limits, the bank would be exposed to unacceptably high risk.

A better solution

Now banks can choose a more practical and reliable way to enforce these limits. A bank can use the Global Limits Control System to record its trading transactions, update its files, and produce a variety of reports—in seconds with a minimum of keystrokes and computer expertise.

A bank can run Global Limits on the I. P. Sharp timesharing service, or on its own computer. In either case, the I. P. Sharp Network makes the information accessible to all of the bank's trading offices 24 hours a day.

Besides providing current and complete data on a bank's limits and exposures, the system can communicate other information of importance to all traders. For example, Global Limits can keep them informed of any political or financial situations that may affect the bank's trading and lending activities.

Since banks have different routines and established procedures, each bank requires slightly different control of global lending. The Global Limits Control System has been designed by I. P. Sharp as an 80-per-cent solution. I. P. Sharp will then customize each bank installation for the remaining 20 per cent. One advantage of writing the Global Limits Control System in SHARP APL is that the system can be modified easily to suit the needs of each user.

To support the banks, I. P. Sharp has assembled a team of software specialists who know and understand the needs of the banking industry. They work with each bank to determine the changes and enhancements that are required, and then tailor the system to the bank's exact specifications. Experienced staff in I. P. Sharp's 60 offices around the world complement this team by providing local participation in the customization process. Because I. P. Sharp offices are in most of the major financial centres, its staff can also provide local support and training for a bank's dispersed trading offices.

I. P. Sharp representatives, who have been demonstrating and discussing the Global Limits Control System with senior bank executives, have received extremely favourable responses. It is apparent that this is indeed the system bankers have long been waiting for.

If you would like to learn more about the Global Limits Control System, contact your local I. P. Sharp office. We keep a test system up and running that we invite you to try out. ■

Wendy Sexsmith, Toronto

CAPERS

Decisions are the life-blood of a healthy enterprise. These decisions often involve capital expenditure projects which are characterized by an initial outlay of large sums of money, with the returns, be they revenue or cost savings, spread over a number of years. For example, most organizations in the private and public sectors typically make decisions concerning:

- Plant expansion
- Replacing manual operations with a computerized system
- Introducing new product lines
- Acquiring additional equipment

In each of these situations, an enterprise must first answer the basic question: Is it economically feasible to proceed with a particular course of action? That is, can a capital expenditure be justified now by future returns?

The approach to solving such problems goes by a number of different names, such as cost/benefit analysis, capital budgeting, or project evaluation, but the techniques involved are basically the same.

Now I. P. Sharp offers CAPERS, a Computer Assisted Project Evaluation and Reporting System, to help you evaluate capital expenditure projects. For each project under consideration, you first supply the relevant information:

- Annual capital expenditures
- Future revenues or cost savings
- Operating expenses
- Income tax information
- Financing information

Using this data, CAPERS performs the various calculations to arrive at the annual net cash flow after tax. The cash flows are then discounted to provide net present value and other profitability indicators, such as the internal rate of return. You can request reports that range from detailed cash flow statements to summary profitability statistics. The cash flow summary report accompanying this article is for a ski resort considering the expansion of its facilities.

Using CAPERS

The calculation routines and standard reports are already in place; you supply only the relevant data associated with the specific project you want to evaluate. In general, you can use CAPERS in any situation where you need to perform a cost/benefit analysis. Some of the features of CAPERS are:

- Unlimited number of revenue and cost streams
- Flexible reporting capabilities
- Case handling facilities for managing data on different scenarios
- "What if" analysis
- Ranking of cases using any profitability indicator
- Ability to consolidate results from a number of projects

CAPERS is in workspace 57 *CAPERS*. To start using the system, type:

```
)LOAD 57 CAPERS
```

For more information, ask for the *CAPERS Users' Guide* from your I. P. Sharp representative. If you are enrolled in I. P. Sharp's MAILBOX, send your queries to the group CAPQ.

CAPERS is intended primarily as a tool for analysis at a fairly general level. Where a more in-depth analysis of a particular project is required, ask your I. P. Sharp representative for assistance in tailoring the software to suit your requirements. ■

Frank Arthur, Sydney

... CAPERS

SKY-HIGH EXPANSION

CASH FLOW SUMMARY

YEAR	GROSS REVENUE	TOTAL EXPENSES	NET PROFIT	CAPITAL EXPENDITURE	CASH FLOW BEFORE TAX	INCOME TAX	NET CASH FLOW	CUM. NET CASH FLOW
1983	360	240	120	1,350	(1,230)	(37)	(1,192)	(1,192)
1984	416	242	174		174	(78)	252	(940)
1985	480	45	436		436	92	343	(597)
1986	555	54	501		501	156	344	(252)
1987	641	58	583		583	218	365	112
1988	740	61	678		678	280	399	511
1989	855	74	780		780	339	441	952
1990	987	80	908		908	407	500	1,453
1991	1,140	85	1,055		1,055	482	573	2,025
1992	1,317	103	1,214		1,214	561	653	2,678
1993	1,521	111	1,410		1,410	656	754	3,432
1994	1,757	133	1,623		1,623	758	865	4,298
1995	2,029	144	1,885		1,885	882	1,002	5,300
1996	2,344	173	2,171		2,171	1,018	1,153	6,453
1997	2,707	205	2,502		2,502	1,174	1,328	7,781
TOT	17,847	1,808	16,039	1,350	14,689	6,908	7,781	
PRESENT VALUE 10.0%/o					5,661		2,856	
PRESENT VALUE 15.0%/o					3,620		1,723	
PRESENT VALUE 20.0%/o					2,313		987	
PRESENT VALUE 25.0%/o					1,442		487	
PRESENT VALUE 30.0%/o					839		134	
PRESENT VALUE 40.0%/o					82		(325)	

*FIGURES ARE EXPRESSED IN THOUSANDS.

New Release of SUPERPLOT

A new version of SUPERPLOT, I. P. Sharp's general-purpose business graphics system, is now available. With SUPERPLOT, you can plot your own data from any one of I. P. Sharp's public data bases or from any private data base on a wide variety of plotting devices and terminals.

Version 5 offers many new options and enhancements, and is fully compatible with previous SUPERPLOT releases. In all cases tested, Version 5 proved to be less expensive to use than its predecessor. Using the new version of SUPERPLOT, you can reduce your CPU usage by 15 to 30 per cent depending on the type of plot you run. In addition to that saving, you pay lower character charges with several of the more popular terminal types. If you switch to Version 5, you will also find that SUPERPLOT is more powerful and flexible than ever before.

All versions of SUPERPLOT are available in the workspace 3 SUPERPLOT and in all MAGIC workspaces. Version 5 is now the default in all the workspaces. Full online documentation for all the new SUPERPLOT features is available through SUPERPLOT's NEWS option. A brief summary of the new options and enhancements follows.

New options

- **ANNOTATION** allows the placing of user-specified text anywhere on the plot. This option was used to place text in the top left corner of the Lundberg Share of Market data base plot in this issue.
- **BOX** specifies the colour, style, and number of passes for boxes which can be drawn around a number of areas of the plot. The STATISBUND data base plot in this issue uses this option to draw a single-pass box around the outside of the plot, and a multiline box to accentuate the legend.

... SUPERPLOT

- *TEXTSPEC* specifies the characteristics of any of the many areas of text on a plot. You can now control the size, colour, font, and justification of the characters; and the direction, slant, and rotation of the text. In this issue, both the Singapore Stock Exchange data base plot and the STATISBUND data base plot use this option. The SINGSTOCK plot uses *TEXTSPEC* to print the *YLABELS* horizontally instead of vertically. The STATISBUND plot uses *TEXTSPEC* to rotate the *XTEXT* to make the plot more readable.
- *PENORDER* facilitates the production of transparencies, foils, and camera-ready copy.
- *OVERLAP* specifies the arrangement of *MBR* or *MBS* line types. In the March/April issue of the *I. P. Sharp Newsletter*, the two plots on page 4 used this option to give a three-dimensional appearance.

Enhancements

- *XLABEL* and *YLABEL* options now allow multiple-line labels. The Singapore Stock Exchange data base plot uses multiline *YLABELS*.
- *SEPARATION* option has been extended to help you produce transparencies, foils, and camera-ready copy. For example, all the plots in this issue were prepared for printing using this feature.
- *PLOT* now accepts an empty right argument, or a vector or matrix of text. In both cases, SUPERPLOT suppresses the printing of everything except the title, boxes, timestamp, and footnote of a plot, as well as any text included in the right argument or specified through the *ANNOTATION* option. The STATISBUND plot uses this feature to print an overall title, footnote, and highlight box around the two plots.
- Version 5 supports four new video display terminals: the Ramtek 6211, the Tektronix 4027A, the Microtel VTX202, and the Human Designed Systems HDSGVT.

All of these new SUPERPLOT features came about as a direct result of your comments. If you have a request or a suggestion for SUPERPLOT, consult your I. P. Sharp representative, or send a message to the MAILBOX group *SUPQ*. ■

Phillip King, Toronto

customer application

DEXUS: A System for Telecommunications Circuit Design

The ARIES Group is a consulting organization that specializes in designing both voice and data networks in Canada or the United States. They offer a number of programs for analysis of voice and data networks under various tariffs, for pricing circuits, and for designing telecommunications networks. They also maintain a data base of telephone rate centres, local calling areas, and telephone tariffs for the United States and Canada. Both the programs and the data base are available on the SHARP APL Service.

I. P. Sharp was selected by the ARIES Group to supply its computer services because of I. P. Sharp's technical expertise, worldwide network, and quality of service.

The ARIES System for Network Design and Analysis is a completely interactive, comprehensive set of programs written in SHARP APL to help telecommunications managers and analysts design and plan cost-effective networks easily.

The issues and options involved in designing a telecommunications network are so diverse that evaluating a particular option or a vendor's proposal can be difficult and time-consuming.

For example, in the design of a voice communications network, one of the issues facing a telecommunications manager is whether to use a tie line or direct dial to service two points in the network; and if a tie line is selected, what the expected busy rate would be.

...DEXUS

The tie-line feature of the DEXUS Circuit Optimizing System is one of ARIES' simpler programs. In its breakeven analysis of direct distance dialing (DDD) versus tie lines, it determines:

- Tie-line circuit cost under pertinent tariff (e.g. MPL, California Series 2008)
- Direct dial cost (DDD)
- Number of calls needed to justify various numbers of tie lines
- Busy rate which can be expected for various numbers of tie lines

In the following example, a breakeven analysis for a tie line between Washington, D.C., and New York City for calls lasting five min-

utes is shown. In this example, the cities are specified by name—you may also use an area code/exchange combination. DEXUS determines that both cities are Category A Rate Centers, what their V and H coordinates are, and that the tie-line tariff is at MPL Schedule 1 rates.

If you would like to try the tie-line analysis program, the ARIES Group will privilege your user number for a one-month trial. For further information on DEXUS or accessing the system, contact:

Ed Garner
 The ARIES Group, Inc.
 Suite 380, 1395 Piccard Drive
 Rockville, Maryland 20850
 (301) 963-7555 ■

```

)LOAD 4237 DEXUS
SAVED 20.17.57 05/27/83
** THE ARIES GROUP -- THE DEXUS CIRCUIT OPTIMIZING SYSTEM **
RATE CENTER FILE AS OF: MAY 21,1983

ENTER COMMAND: TIELINE
NO TARIFF ENVIRONMENT SET

ENTER GEOGRAPHIC ENVIRONMENT: US
CONTINENTAL UNITED STATES -- CORRECT? Y
USE CURRENTLY EFFECTIVE TARIFF? Y
TARIFF EFFECTIVE AS OF APRIL 2,1982 SELECTED -- OK? Y

CONTINENTAL UNITED STATES TARIFF ENVIRONMENT
AMERICAN TELEPHONE AND TELEGRAPH
** TYPE 2000 MPL VOICE GRADE SERVICE **
ENTER AVERAGE CALL LENGTH IN SECONDS: 300
300 SECONDS PER CALL -- CORRECT? Y
DEFAULT CALL PARAMETERS

STATION-TO-STATION DAY RATES
8 HOURS PER BUSINESS DAY
22 BUSINESS DAYS PER MONTH

CHANGE DEFAULTS? N
USE THIS PROGRAM ONLY FOR POINT-TO-POINT CIRCUITS
ENTER ONLY TWO CITIES:
CITY,ST:NUMBER OF DROPS -- OR END
WASHINGTON,DC
NEWYORK,NY
END

CONTINENTAL UNITED STATES TARIFF ENVIRONMENT
AMERICAN TELEPHONE AND TELEGRAPH
** TYPE 2000 MPL VOICE GRADE SERVICE **

FROM          TO          DISTANCE  NODES  LINE CHG  SVC TRM  TOT CHG
+WASHINGTON DC +NEWYORK NY    206      2      $347.22   $71.00   $418.22

TOTALS:                206      2      $347.22   $71.00   $418.22
    
```

(cont)

... DEXUS

THE ARIES GROUP
COMMUNICATIONS ANALYSIS

CONTINENTAL UNITED STATES TARIFF ENVIRONMENT
AMERICAN TELEPHONE AND TELEGRAPH
** INTERSTATE MESSAGE TELECOMMUNICATIONS SERVICE **

	FROM: DC WASHINGTON			
TO	CODE	SECONDS	MILES	COST
NY NEWYORK	A	300	206	\$2.14

THE ARIES GROUP
COMMUNICATIONS ANALYSIS
TIE LINE TRADEOFF

CALL PARAMETERS
STATION-TO-STATION DAY RATES
8 BUSINESS HOURS PER DAY
22 BUSINESS DAYS PER MONTH

NUMBER OF LINES	BREAKEVEN CALLS	P(BUSY) (PERCENT)
1	196	8.49
2	391	1.43
3	587	0.27
4	782	0.05
5	978	0.01
6	1173	0.00

ANOTHER TIE-LINE ANALYSIS? *N*

**data
bases**

**New Singapore Stock Exchange
Data Base**

You can now get online descriptive and daily trading information through I. P. Sharp for over 350 stocks traded on the Singapore Stock Exchange. The SINGSTOCK data base also contains information for over 20 indices covering the Singapore, Kuala Lumpur, and other Asian stock exchanges.

For each stock listed on the Singapore exchange, the daily facts available are:

- Volume
- Low
- High
- Close-Buy
- Close-Sell
- Last Price
- Price/Earnings Ratio

Descriptive information on the stocks includes:

- Security code
- Security number
- Full name
- Date stock added
- Date of last quote
- Delist indicator

The industrial classification of each stock is indicated by the initial letter of the security code:

- A Industrial and Commercial
- B Finance
- C Hotels
- D Properties
- E Debentures, Loans, and Bonds
- F Plantations
- G Mining

... Singapore Stock Exchange

Indices information

The various indices can also be identified by the initial letters of the stock code:

- SES* Singapore Stock Exchange Indices
- ST* Straits Times Indices
- KL* Kuala Lumpur Indices
- OI* Other Asian Indices

The *SES*, *ST*, and *KL* classifications include indices for Industrial and Commercial, Finance, Hotels, Properties, and Mining. The *OI* classification includes composite indices for the Taiwan, Bangkok, Hong Kong, South Korea, Manila, Tokyo, and New Zealand exchanges.

The Singapore, Kuala Lumpur, and Tokyo Stock Exchanges and Hong Kong indices are updated daily, the morning after the day's trading at 10 a.m. Singapore time. Data is available from January 1982. *The Singapore Stock Exchange Daily Journal* supplies the information.

The other Asian indices are updated each week with the daily figures for the preceding week. The data is from the *Far Eastern Economic Review* and starts in January 1983.

The data is adjusted for stock splits by I. P. Sharp's Singapore office.

Accessing the data

You can access the SINGSTOCK data base through the workspaces 39 *MAGIC* or 55 *RETRIEVE*. You can get online documentation through these workspaces or through the workspace 1 *DATABASES*. For a two-page description of SINGSTOCK, type:

SINGSTOCKD 'DESCRIBE' in 39 *MAGIC* or
ΔDESCRIBE 'SINGSTOCK' in 55 *RETRIEVE*

To obtain a directory of all the stock codes in the data base, and relevant details such as timeframe and full name, type:

SINGSTOCKD 'DIRECTORY' in 39 *MAGIC* or
ΔDIRECTORY 'SINGSTOCK' in 55 *RETRIEVE*

By using the screening facilities in the data base, you have great flexibility in selecting which items you want displayed. The following example displays the values of all the Kuala Lumpur Stock Exchange indices.

The enclosed plot shows some of the daily facts available in the SINGSTOCK data base for the stock of the Malaysia Mining Corporation.

If you would like more information on SINGSTOCK, contact your I.P. Sharp representative or the I.P. Sharp MAILBOX group, MONQ. ■

Mark Seltzer and Loon Wai Sum, Singapore

KUALA LUMPUR STOCK EXCHANGE INDICES

	11APR83	12APR83	13APR83	14APR83	15APR83
<i>LAST</i>					
<i>KL FINANCES</i>	1725.1	1728.8	1720.0	1720.8	1721.8
<i>KL HOTELS</i>	831.1	829.8	828.2	828.2	826.9
<i>KL INDUSTRIALS</i>	574.9	578.3	573.3	574.4	574.9
<i>KL MINING</i>	252.9	253.3	252.9	252.2	251.4
<i>KL PLANTATIONS</i>	950.0	957.9	955.4	959.9	946.1
<i>KL PROPERTIES</i>	1244.5	1254.5	1240.6	1242.5	1247.2



I. P. Sharp

newsletter

JULY/AUGUST 1983
Volume 11 Number 4

technical supplement 45

CPU Limits in SHARPA PL

During the development of software, programming mistakes frequently occur. For example, some errors in logic can cause the computer to repeat a sequence of steps with no possible escape from the routine. Such an infinite loop can expend vast amounts of computer resources (CPU units) when run on today's powerful computers. Since the use of SHARPA PL is charged by the number of CPU units consumed, an infinite loop can be a costly mistake.

Thus protective mechanisms to limit the consumption of CPU units are available to all SHARPA PL programmers. You can set a default CPU limit which is activated every session. Or during a session you may opt to define a CPU limit through the use of the system command, `)CPULIM`. Both alternatives are discussed in this article.

Default CPU limits

You may have a default CPU limit associated with your user number. This default limit may be set either when the account is created or at any time thereafter. Ask your I. P. Sharp representative to have a default limit set. Once the limit has been set to the user-defined value, it is in effect during every subsequent session, unless you override the default CPU limit through the use of the `)CPULIM` command (which is explained later in this article).

The default CPU limit is really a *computation* limit. A computation limit is reset each time a new keyboard transaction begins. In other words, each time you hit the carriage return key while entering data, the default computation limit is decreased by the

number of CPU units required to execute the expression. If this required number of CPU units exceeds the default limit, an *INTERRUPT* error message is displayed along with the line of the function or expression in which the interrupt occurred (see *Programmer notes* at the end of this article). Otherwise, the default limit is reset and you are ready to enter more data. If the execution of the expression is terminated by an *INTERRUPT*, the default limit is immediately reset, and you are free to continue the session either by resuming execution of the expression or by entering any other input into the terminal. The maximum default computation CPU limit which can be set is 3200; however, in practice, a much lower value would normally be used.

CPULIM system command

You can control the number of CPU units consumed in two different ways through the use of `)CPULIM`. Until this command is used in a session, only the default CPU limit is in effect. Once a CPU limit is set through the use of `)CPULIM`, this limit is in effect throughout the rest of the session or until the command is reissued. It should be noted that CPU limits which are set through the use of `)CPULIM` override the default CPU limit. The following example sets a *computation CPU limit* of 30 CPU units and a *session CPU limit* of 100 CPU units:

```
)CPULIM 30 100
```

Once you have hit the RETURN key after typing in the required `)CPULIM` command, the computer returns the value of the previous limit settings which were in effect. The computation limit is reset each time a new key-

... CPU Limits

board transaction begins. If the computer requires more than 30 CPU units to execute an APL expression or function after you have hit the carriage return key, an *INTERRUPT* would be displayed after 30 CPU units have been expended. Then the computational CPU limit would be automatically reset to 30 CPU units.

The session CPU limit simply counts down the consumption of CPU units until it expires. Unlike the computation CPU limit, the session CPU limit is independent of typed input transactions. After setting a session limit of 100 CPU units as in the above example, an *INTERRUPT* diagnostic is displayed as soon as 100 CPU units have been expended since the setting of the session limit. The session limit would then automatically be reset to 100 CPU units. Note that the default session CPU limit is always zero, that is, no session limit at all.

In general, if you are contemplating specifying a CPU limit through the use of *)CPULIM*, you should consider the nature of the work to be performed. If there is to be much input with moderate intermixed processing, a session limit would be preferable. If functions are to be developed and you wish to guard against an expensive infinite loop, a computation limit is more effective. If you don't expect to enter much input, either limit can be set.

Further examples of the use of the *CPULIM* command are:

```
)CPULIM 50 makes the computation CPU  
limit 50 CPU units but does not alter the  
session CPU limit
```

```
)CPULIM 0 200 removes the computation  
CPU limit and makes the session CPU limit  
200 CPU units
```

```
)CPULIM 0 0 removes both the computa-  
tion CPU limit and the session CPU limit
```

Programmer notes

1. The rules for default CPU limits and the use of *)CPULIM* as described above also apply to S-tasks.
2. In an N-task, the session CPU limit is explicitly set in the *□RUN* parameters and overrides any limits set in the calling task.

Similarly, the limit set for a B-task overrides any limits set in the submitting task.

3. The default computation CPU limit of an N-task or B-task is always zero (that is, no limit). The session CPU limit controls the consumption of CPU, since input in an N-task or B-task cannot occur.
4. Expiry of a CPU limit is a trappable event (number 1003) in a T-task or S-task. In an N-task or B-task, expiry is not trappable, and instead the task ends.
5. Once the default CPU limit has been overridden in a session through the use of *)CPULIM*, the original default value can be reinstated in the same session by a subsequent use of *)CPULIM* with an appropriate argument.
6. *)CPULIM* cannot be invoked as a line of an APL function.
7. The user must provide a computation CPU limit when setting a session CPU limit, since the computation limit appears first in the argument to *)CPULIM*. ■

Jeannie Crabtree, Ottawa

Making Things Fit, Part 1

The demand for computing power seems to be infinite. Faster processors and larger memories are announced regularly, and just as regularly they are consumed by new users and new applications. As long as unlimited resources aren't instantaneously available to every user, a substantial part of every programmer's job will continue to be making things fit. This is the process of taking sets of information and algorithms, and arranging them to fit within the constraints imposed by hardware, software, and the real world.

This process is analogous to putting the proverbial square peg into a round hole. There are at least six ways to do this:

- 1) Reduce the volume of the peg
- 2) Expand the volume of the hole
- 3) Change the shape of the peg

... Making Things Fit

- 4) Change the shape of the hole
- 5) Cut the peg into pieces
- 6) Make it *look as if* the peg is in the hole

Each of these strategies of volume alteration, dimension distortion, segmentation, and prestidigitation will be used in this series of articles.

We will consider the problems associated with making things fit in the following:

- 1) Rectangular arrays
- 2) Finite workspaces
- 3) External files
- 4) Printed pages
- 5) Display screens

We will see that many of the problems APL programmers face every day fall under the rubric of "making things fit".

1) Rectangular arrays

APL is an *array processing language*. Until recently, it was limited to processing simple arrays. The elements of simple arrays are indivisible scalar values. The structure of these arrays was, and continues to be, rectangular. That is, the length of each subarray along a given dimension is the same. Not all real world problems fit conveniently into rectangular arrays however. The typical corporate structure immediately comes to mind. There aren't many companies that require every manager in the organization to supervise the same number of people!

But now SHARP APL and other APL implementations provide for the enclosure of arrays. If an array is enclosed, it is treated as a scalar, and thus can be an element of a rectangular array. Subarrays of uneven length can be enclosed, and catenated together. The uneven, hierarchical corporate structure will fit quite nicely in this scheme. Why then should we worry about making ragged data fit in rectangular arrays? There are at least two reasons.

- 1) Some APL users haven't moved up into the world of APL implementations which support enclosed arrays.
- 2) Even when data is inherently ragged, it is sometimes more efficient, in terms of space or time, to use one of the methods

developed for handling ragged data before the introduction of enclosed arrays.

Let's consider a real-world problem. Suppose you were developing a budgeting and cost accounting application for a large corporation. There would be lots of categories (accounts) you would associate data with. But in many cases, you might want to make finer discriminations between transactions, and so you would associate subaccounts with each account. An actual organization which structured its accounting information this way had approximately 250 accounts. The number of subaccounts associated with a single account ranged from 0 to 200.

There are at least five reasonable solutions to the problem of how to store the data.

- 1) Create a matrix with as many rows as accounts and as many columns as the maximum number of subaccounts. Fill in the empty locations, where there are fewer than the maximum number of subaccounts, with a special **fill element**. The subaccounts are stored in the matrix, and the account numbers are stored in a corresponding vector.
- 2) Create a vector which contains all subaccounts for all accounts catenated together. Each group of subaccounts begins with a special **delimiter value**. The accounts are stored in a separate vector. The number of delimiters in the subaccount vector equals the number of accounts.
- 3) Create a file in which the first component contains the vector of accounts, and each succeeding component contains the subaccounts for one account.
- 4) Create a vector which contains all the subaccount numbers catenated together. The groups are distinguished by a parallel **partition vector**, in which a 1 marks the beginning of each new subaccount group. The account numbers are stored in a separate vector. The number of 1's in the partition vector equals the number of accounts.
- 5) Enclose each group of subaccounts, and catenate them together. Store the account numbers in separate vector.

... Making Things Fit

Storing data

What are the disadvantages of each method? The first will create a numeric matrix 250x200, which has 50 000 elements, or 200 000 bytes. This will be much too big on some APL systems. The storage requirements of the second will be more modest. However, this method assumes that there is some delimiter value which will never be part of the data. This is reasonable in this case, since the data are identifiers, which will all be integers greater than 0. It would not be reasonable to assume this if the numbers were arbitrary values. The third method may incur a heavy CPU penalty for doing file I/O as opposed to calculations in the workspace. The fourth has no way to represent account numbers which have no accounts, since the partition vector always requires a 1 to mark the start of the subaccount group. The fifth method isn't available on some APL systems.

What operations will we want to perform on each of these data structures? We can get along fine if we can do the following:

- 1) Add an account
- 2) Delete an account
- 3) Add a subaccount to an account
- 4) Delete a subaccount from an account
- 5) Get all the subaccounts belonging to one account
- 6) Find the account associated with each of several subaccounts

Relative speeds

The first four activities are maintenance work, and will be done relatively infrequently. The fifth activity will be done during production work, like data entry and reporting. Also it must be done by the other operations. So we would like to know the relative speed of performing this operation on each of the five data structures.

Programs which do the retrieval activity for each of the data structures follow. In each case, the argument is the account number whose subaccounts we wish to retrieve.

```

▽ S+GETSUBS1 A
[1] S←,SUBACCTS[ACCTS\A;]
[2] S←(Φ\ΦFILL#S)/S ▽

```

```

▽ S+GETSUBS2 A
;N
[1] N←ACCTS\A
[2] S←((SUBACCTS=1ρSUBACCTS)/
1ρSUBACCTS),1+ρSUBACCTS
[3] S←S[N]+(S[N+1]-1)+SUBACCTS ▽

```

```

▽ S+GETSUBS3 A
[1] S←READ FILE,1+ACCTS\A
[2] A ASSUMES ACCTS HAS BEEN READ
FROM FILE ▽

```

```

▽ S+GETSUBS4 A
[1] S←(ρACCTS)ρ0
[2] S[ACCTS\A]+1
[3] S←(≠PART\S≠1+0,S)/SUBACCTS ▽

```

```

▽ S+GETSUBS5 A
[1] S←,SUBACCTS[ACCTS\A] ▽

```

The following table lists the relative speed of each of these five methods in three situations:

- A) Get the subaccounts for a single account which has a single subaccount
- B) Get the subaccounts for a single account which has the maximum number of subaccounts
- C) Get the subaccounts for all accounts, one account at a time

Rather than list the absolute amount of CPU time used by each method, each time was divided by the smallest time in that category to produce a relative timing.

	A	B	C
1	2.22	2.22	2.18
2	55.67	56.97	54.08
3	6.86	2.5	2.35
4	16.94	17.03	16.34
5	1	1	1

The enclosed array version is clearly the fastest. The padded array version and the file version come next. The partitioned vector and delimiter vector versions bring up the rear. The speeds vary because of the amount of work each function must do to identify which items to return.

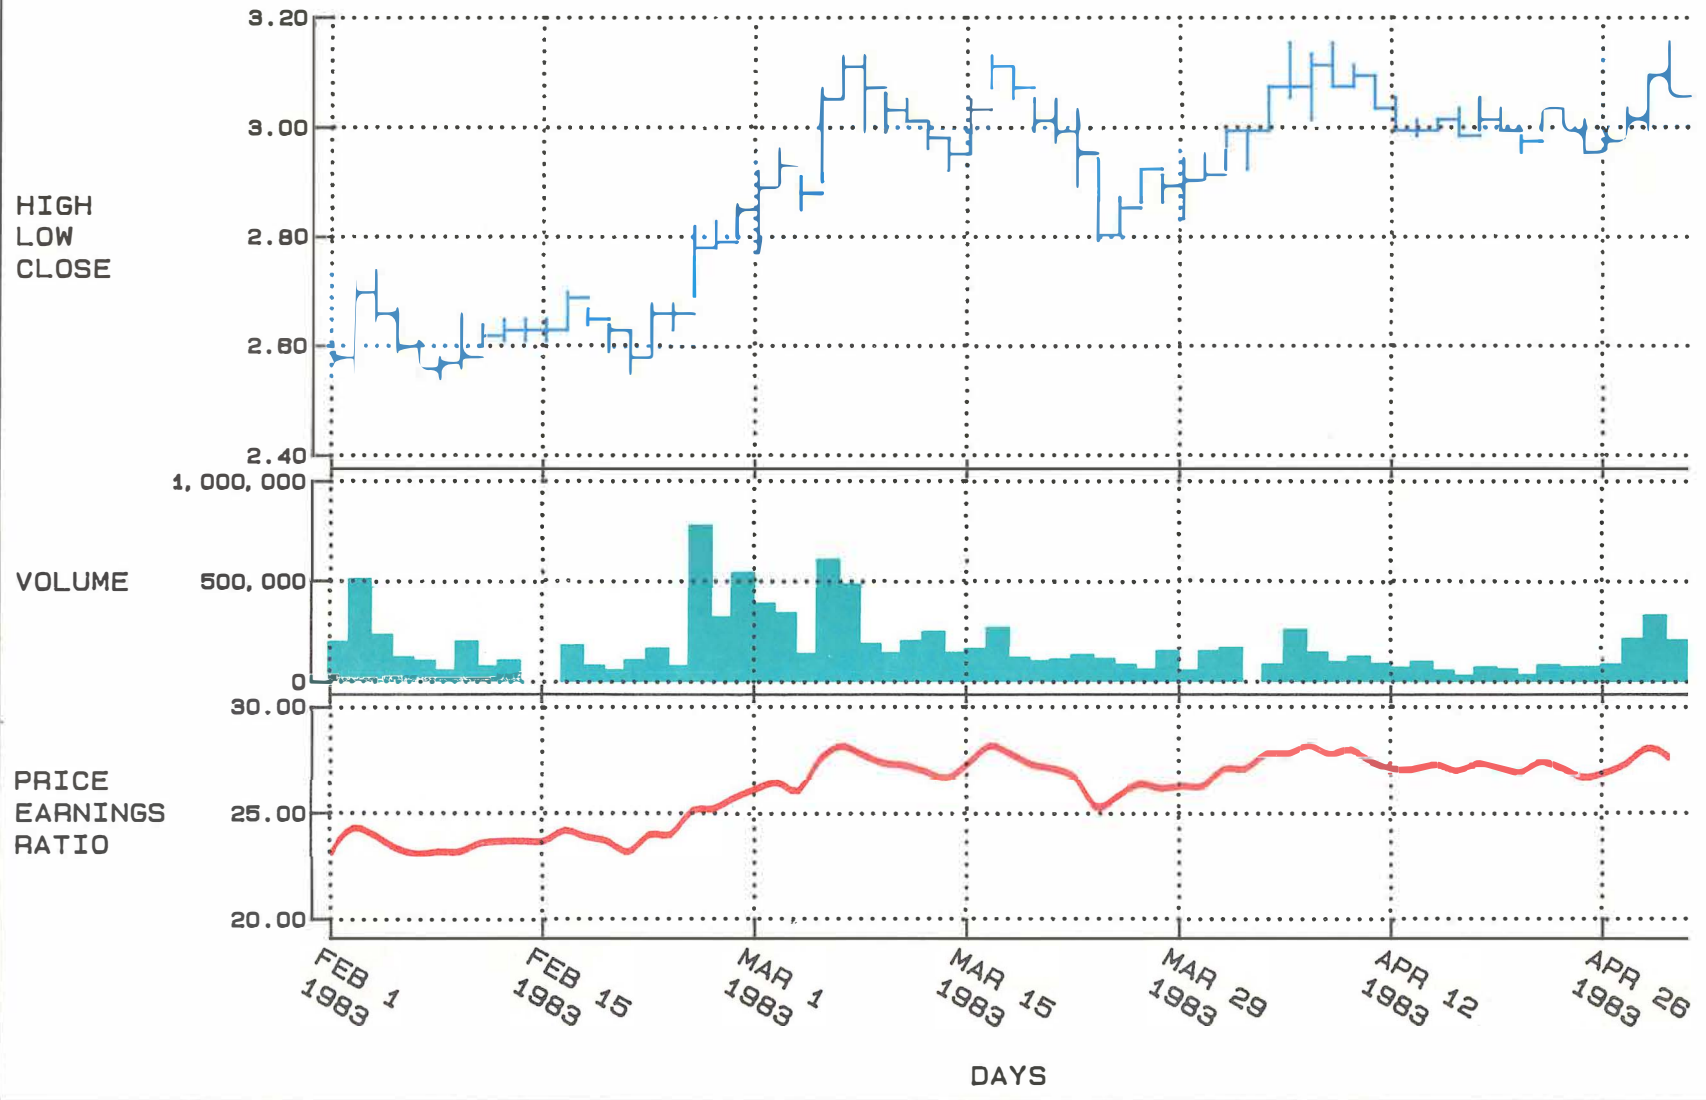
We can see that several of the strategies identified at the beginning of the article have been employed. The first and fifth methods use the strategy of "dimension distortion". The first extends the length of the groups, and the fifth reduces each group to a scalar. The second and fourth methods use the strategy of "volume alteration". The second method adds delimiter values to the data and the fourth a separate partition vector. The third method uses the "prestidigitation" strategy, and doesn't even store the data in a rectangular array at all.

In the next article in this series, we will consider ways to make things fit in finite workspaces. ■

Robert Metzger, Rochester

I. P. SHARP NEWSLETTER


MALAYSIA MINING CORPORATION




SOURCE: SINGAPORE STOCK EXCHANGE DATA BASE

AUSSENHANDEL 1982

EINFUHR/AUSFUHR
DIE WICHTIGSTEN
HERSTELLUNGLAENDER

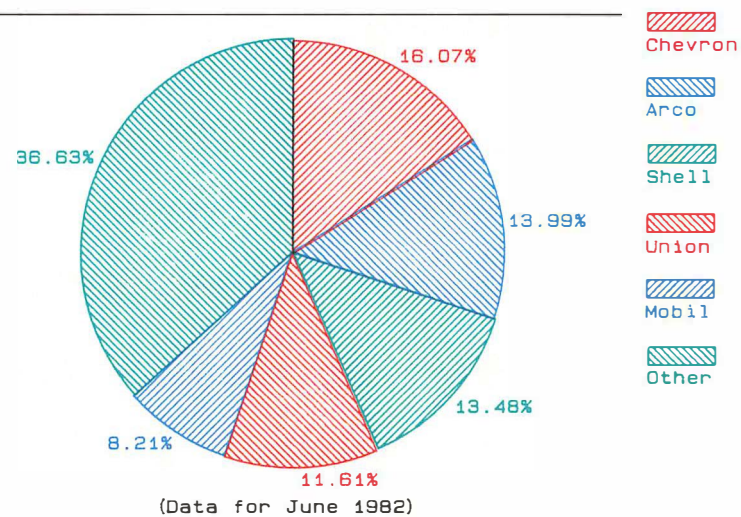
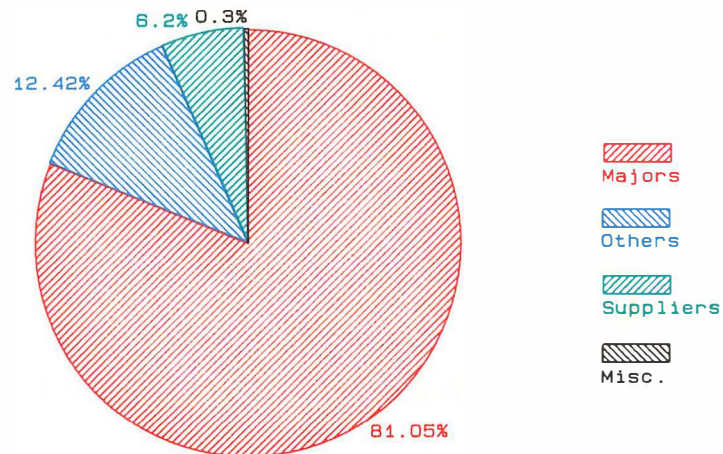


ANTEIL DER EINFUHR/AUSFUHR
NACH LAENDERGRUPPEN



QUELLE: STATISTISCHES BUNDESAMT WIESBADEN

Share of Market for California



Source: Lundberg Survey - Share of Market Data Base.

Lundberg Survey's Share of Market Now On Line

If you are involved in marketing oil products, you will be interested to learn that Lundberg Survey's Share of Market is now available on line through the SHARP APL Service. Share of Market contains gasoline sales volumes for the wholesale market, and market shares by brand and federal state. Using this data on line, market analysts can interpret time series quickly and easily, with the sophisticated searching and cross-referencing facilities of SHARP APL.

Lundberg Survey, one of the U.S.'s leading suppliers of statistics relevant to oil marketing, calculates volumes and market shares from a number of sources, including state government reports of the taxes paid by each company. The data is collected on a monthly, statewide basis, and is updated as new data becomes available (approximately weekly). The Share of Market (SOM) data base contains the following facts for the past five years.

- Number of gallons sold
- Share of market (percentage of gallons sold in a given state by a given company)
- Report date (date data was last updated)
- State
- Brand
- Functional classification (e.g. major integrated oil company, cooperative)
- Start date (of period over which figures were tallied)
- End date (of period over which figures were tallied)
- Description (state and brand)
- Symbol (abbreviated description)
- Data base number (a number representing a particular brand in a particular state)
- Discontinued indicator (1 if current data is available for that brand and state, 0 if it is not)

For a directory of fact codes, load the workspace 121 *PETROSERIES*, and type:

SOM 'DIRECTORY'

You can search for specific data from any combination of fact codes. For example, using

MAGIC you can generate a report that selects data only from those major integrated oil companies that sold over one million gallons during March 1983. Or you could select data for those states in which a particular brand held a market share of at least 10 per cent during the previous month.

For more information on accessing Share of Market through *MAGIC*, load 121 *PETROSERIES*, and type:

SOM 'DESCRIBE'

Custom or standard reports

In addition to creating your own custom reports with *MAGIC*, you can use standard reports similar to those provided in hard copy by Lundberg Survey. Two standard reports are currently available through the interactive program *INFOMAGIC* in 121 *PETROSERIES*. The first shows (in descending order of sales volume): the name of the company, the number of gallons sold, the change in volume from the previous period, the change in volume from the same period during the previous year, and the change in year-to-date sales for both volume and market share. You may select the states for which you want the report printed. The second report presents similar data, but organized by functional classification rather than sales volume.

Share of Market is the second Lundberg data base available through the SHARP APL Service; Dealer Buying Price has been on line since January this year. I. P. Sharp will be adding other Lundberg data bases soon. These include Energy Détente (petroleum retail price series for over 60 countries), and Retail Price Survey (U.S. retail prices for gasoline and diesel by market, state, and Petroleum Administration for Defense districts..

The Lundberg Survey data bases are available by subscription only. For subscription information, please contact:

Lundberg Survey Inc.
12041 Strathern
North Hollywood, California 91605

...Share of Market

Learn how you can use Share of Market and other online Lundberg data bases at a free half-day seminar. To register, please contact your I. P. Sharp representative. Courses are scheduled in the following cities:

Houston	August 16, 1983
Dallas	August 18, 1983
Chicago	August 23, 1983
Cleveland	August 25, 1983
Washington	September 13, 1983
New York	September 15, 1983

Or send a message to the MAILBOX group, LUNDQ.

The enclosed plots show the share of the market for various suppliers by functional classification in the upper one and by sales volume in the lower one. ■

Deborah Bodinger, Newport Beach

STATISBUND Data Base

The STATISBUND data base contains a wide range of statistics relating to the economy of West Germany. STATISBUND gives a very detailed picture of the economic, financial, political, and social position of West Germany. Researchers, econometricians, and planners will find the information useful for a variety of studies, including forecasting, strategic planning, and growth and trend analysis.

The Statistisches Bundesamt, West Germany's Statistics Office, collects the data from both government and private sources, and updates it monthly. The data base contains monthly, quarterly, and yearly time series, with some starting in 1950. Almost 50 000 time series are currently available, and the number is likely to be increased in the near future. The subject categories available now are:

- Demographics
- Labour
- Business
- Agriculture, forestry and fishing
- Manufacturing

- Construction and housing
- Tourism and domestic trade
- Foreign trade
- Transport and communications
- Money and credit
- Finance and taxes
- Social benefits
- Personal consumption expenditure
- Wages and salaries
- Prices
- National accounts

You can access and display STATISBUND data easily using MAGIC. You can also use MAGIC for forecasting and for other types of analysis. To retrieve data, load the workspace 39 MAGIC and type:

STATISBUND series codes

For example, *STATISBUND 2022,2023* returns some unemployment figures. For a comprehensive directory of series codes, in German, load the workspace 767 *STATISBUND*, and type *DESCRIBE* or *BESCHREIBUNG*. The directory also contains other useful information, such as reporting units, source, frequency, and references to further information.

In addition, you can combine STATISBUND data with data from other I.P. Sharp public data bases. For example, you can compare the economy of West Germany with the economies of other countries.

For more information on the STATISBUND data base, load the workspace 39 *MAGIC* and type:

STATISBUND 'DESCRIBE'
—for English text

STATISBUND 'BESCHREIBUNG'
—for German text

There is no subscription fee to access the STATISBUND data base.

The enclosed plot shows the imports and exports of principal trading countries and by groups of countries, using data drawn from the STATISBUND data base. ■

Sally Drew, Düsseldorf

... STATISBUND Data Base

An article on the STATISBUND data base in German appears in the *Financial and Economic Newsletter*, July 1983, Volume 3 Number 3. Ask your I. P. Sharp representative for a copy.

organizations

Join the APL Users' Group: SIGAPL

The Special Interest Group on APL (SIGAPL) is an affiliate of the Association for Computing Machinery (ACM). Both ACM and SIGAPL are international organizations—over 20 per cent of SIGAPL's members are outside the United States.

SIGAPL is dedicated to promoting the development and application of APL. The two chief means of achieving this are sponsoring annual international APL conferences and publishing *APL Quote Quad*.

Over 800 people attended APL83 in Washington, D.C., the most recent conference sponsored by SIGAPL and ACM. FinnAPL, the Finnish APL Association, in cooperation with SIGAPL and ACM, is sponsoring APL84 in Helsinki, from 11 to 15 June 1984. The next conference in North America will be APL85 in Seattle, Washington.

APL Quote Quad is a quarterly publication containing news about APL organizations, correspondence, contributed articles, APL algorithms, and a recreational APL column. Back issues of APL conference proceedings and *APL Quote Quad* are available from:

ACM Order Department
P.O. Box 64145
Baltimore, Maryland 21264

With a SIGAPL membership, you receive:

- *APL Quote Quad* subscription
- SIGAPL conference proceedings
- Discounts on purchase of back issues of *APL Quote Quad* and conference proceedings
- Discount on SIGAPL conference registration fees

Annual membership dues are \$10 U.S. for ACM members, and \$25 U.S. for non-ACM members. If you would like an application form, write:

Association for Computing Machinery
11 West 42nd Street
New York, New York 10036

Or send a message to Robert Metzger, *RCM*, in the I. P. Sharp MAILBOX, or contact your I. P. Sharp representative. ■

conferences

**Call for Papers**

APL84, the international conference of the APL community, will be held in Helsinki, from 11 to 15 June 1984. This conference is organized by the Finnish APL Association (FinnAPL) in cooperation with the Association for Computing Machinery (ACM) and its Special Interest Group on APL (SIGAPL).

The conference is devoted to both the APL language and its applications, the latter especially from an end user's point of view. The program will include contributed and invited papers, and a commercial exhibition and forum. A parallel program of social and cultural events for attendees and their families will be held before, during and after the conference.

APL84 opens on Monday 11 June with a full day of intensive tutorials, at various levels of expertise, by leading teachers of APL. The

... APL84 Call for Papers

formal proceedings, from Tuesday to Friday, will be divided into streams covering different interest areas, consisting of presentations about a half-hour in length. There will also be special interest sessions, an exhibition hall, and a reading room for materials of specialized interest.

The program committee invites you to submit papers on any APL-related subject, in particular:

Application papers

- Finance and budgeting
- Planning and control
- APL in management
- Data bases
- Engineering
- Applied mathematics
- APL in various professions
- Personal computing
- Teaching with APL

Technical papers

- Professional methods of APL program development
- Extensions to APL and their differences
- APL and other languages
- APL on microcomputers
- Implementation of APL systems
- User interface: ease of learning, full-screens, and graphics
- Teaching of APL

Paper submission requirements

Drafts of full papers should be submitted by 5 September 1983 to:

FinnAPL
APL84 Program Committee
P.O. Box 1005
SF-00101 Helsinki 10, Finland

I. P. Sharp MAILBOX: *PC84*
IBM VNET address: APL84 at HEKVM

Further information is available in the workspace 1 *APL84* on the I. P. Sharp time-sharing service. Questions regarding the conference can be directed to:

Program Chairman:

Sten Kallin
IBM Svenska AB
S-16392 Stockholm, Sweden

Telephone: (46) 8-7524000
IBM VNET address: SE44930 at KISTAVM
I. P. Sharp MAILBOX: *SKAL*

General Chairman:

Timo Seppala
TMT-Team Oy
P.O. Box 452
SF-00101 Helsinki 10, Finland

Telephone: (358) 0-4521155
I. P. Sharp MAILBOX: *TIMO*
IBM VNET address: APL84 at HEKVM ■

communication news

New Local Dial-Up Numbers

Birmingham

1200 cps (021) 359-4315

Coventry

1200 cps (0203) 27327

Dublin

1200 cps (01) 764905

Frankfurt

300 cps (0611) 747935

New PROMIS Offices

I. P. Sharp is pleased to announce the recent appointment of NICOTEC as the agent representing PROMIS in the following countries: Austria, Belgium, Finland, France, Germany, Holland, Italy, Norway, Spain, Sweden and Switzerland. PROMIS is I. P. Sharp's computer-aided manufacturing system for microelectronics.

Representatives of European companies who would like to obtain information about PROMIS are invited to contact Nic Albertsen, president of NICOTEC, at the following address:

NICOTEC
Ragetenstrasse 6
CH-6318 Walchwil
Switzerland

Telephone: 042-316637
Telex: 862523NICO

In addition to the PROMIS Group in the Toronto headquarters of I. P. Sharp, there is a PROMIS sales and support office in Palo Alto, California:

PROMIS Group
I. P. Sharp Associates Inc.
Suite 201, 220 California Avenue
Palo Alto, California 94306

(415) 327-1700 ■

Singapore

Christina Chang has been appointed general manager of Singapore International Software Services (SISS), I. P. Sharp's representative in the Republic of Singapore.


After receiving a Bachelor of Commerce degree from the University of New South Wales, she stayed in Sydney for five more years working as a programmer/analyst for Allied Manufacturing and Trading Industries Limited.



In 1977, Christina joined Singapore Airlines, and was senior system analyst and project leader before her secondment to SISS in December 1981. Until her recent appointment, she was responsible for customer services at SISS. ■

Toronto

Brian Traquair has been appointed branch manager of the Toronto office. Brian obtained his Bachelor of Mathematics (Computer Science) degree in 1979, from the University of Waterloo. He joined the Toronto branch of I. P. Sharp that same year. He worked in consulting and supervisory roles until 1981, when he was involved in the creation of a branch marketing group. Brian has been marketing manager in Toronto since that time. ■



publications

A Working Introduction to APL Now in German

A German translation of Ken Iverson's introductory APL course is now available as *Eine Arbeitseinführung in die APL-Sprache*.

Over a hundred people have used the German version at the I. P. Sharp education centres in Austria, Germany and Switzerland. Students are enthusiastic about this course because it offers them the opportunity to work independently and at their own speed.

Eine Arbeitseinführung in die APL-Sprache is a valuable contribution to APL education in the German-speaking region. You can install it on your own in-house system running SHARP APL or other versions of APL, such as STSC's APL*PLUS, IBM's VSAPL, or DEC's APLSF.

...A Working Introduction to APL

For more information on the German version of *A Working Introduction to APL*, contact the I. P. Sharp office in Düsseldorf, Vienna or Zurich. ■

Ken Iverson's *APL-Einführungskurs jetzt auch auf Deutsch*

Wir freuen uns, eine sorgfältige deutsche Übersetzung der APL-Einführung *A Working Introduction to APL* vorstellen zu können.

Die deutsche Version ist in Österreich, Deutschland und der Schweiz bereits von ca. 100 Kursteilnehmern in den Schulungszentren von I. P. Sharp erprobt und mit Begeisterung aufgenommen worden. Geschätzt wird vor allem die Möglichkeit, selbständig und dem eigenen Lerntempo entsprechend aktiv am Terminal zu arbeiten.

Der Kurs heisst *Eine Arbeitseinführung in die APL-Sprache* und ist ein wertvoller Beitrag für die APL-Ausbildung im deutschsprachigen Raum. Der Kurs wird in den Schulungszentren von I. P. Sharp angeboten, kann aber auch gekauft und an der eigenen Anlage unter SHARP APL, APL*PLUS, VSAPL, APLSF und anderen APL-Versionen installiert werden.

Nähere Auskünfte erhalten Sie von I. P. Sharp Düsseldorf, Wien oder Zürich. ■

Margarete Buch, Vienna

Bookends

Rationalized APL (new)

Rationalized APL, by K. E. Iverson, is the first in a new series of I. P. Sharp Research Reports. The series will address technical questions in the design and use of APL. This first report presents a broad program of extensions to the language that are in various stages of design and implementation. Three aims underlie this program:

- Retain complete compatibility so that existing programs continue to work

- Rationalize and simplify the language and its syntax

- Develop further functions and operators that fit into the simplified scheme

Some of the facilities discussed in the report—such as the rank operator—have already been incorporated into SHARP APL. The report concludes with some general comparisons of extensions to APL that have appeared in other versions of APL.

April 1983, 52 pp., \$4.00

MICROMAGIC Users' Manual (new)

MICROMAGIC is a facility which uses the MICROCOMM diskette and allows IBM Personal Computer users to communicate with the I. P. Sharp timesharing service. (See the article in the *I. P. Sharp Newsletter*, March/April 1983). You can use MICROMAGIC to access all the I. P. Sharp public data bases. You can also display retrieved data in a variety of formats, or store it as an IBM PC file.

March 1983, 47 pp., \$4.00

QUICKDRAW Users' Guide (new)

QUICKDRAW is intended for APL programmers who require graphic output for custom applications. (See the article in the *I. P. Sharp Newsletter*, May/June 1983.) QUICKDRAW is designed particularly for production business graphic applications, using the Hewlett-Packard 7221 series of four- and eight-pen plotters, to produce large numbers of specialized graphs on a regular basis. Many different options and types of plots are possible.

February 1983, 118 pp., \$8.00

MAGICSTORE Users' Guide (new)

MAGICSTORE is a multidimensional file management system you can use easily to create, maintain, list, and calculate time series, that is, numeric data associated with regular periods, such as days or months. You can use MAGICSTORE alone or combine it with other I. P. Sharp time series manipulation languages such as MAGIC. There is also a four-page MAGICSTORE brochure available at no charge.

publications

... Bookends

(See the article on the front page of the May/June issue of the *I. P. Sharp Newsletter*.)
March 1983, 114 pp., \$7.00

Energy Data Bases Reference Guide (revised)

The *Energy Data Bases Reference Guide* is in the same handy, pocket-sized format as the *Financial and Economic Data Bases Reference Guide* and the *Aviation Data Bases Reference Guide*. The *Energy Data Bases Reference Guide* is designed to be a quick reference source for data base accesses and data definition. This completely revised guide replaces the May 1980 edition.
June 1983, 93 pp., \$3.00

ACTPAK: The SHARP APL Actuarial Package (new)

The *ACTPAK Reference Manual* has now been superseded by this completely revised publication, which fully describes the new version of ACTPAK (*I. P. Sharp Newsletter*, May/June 1983). Since the new ACTPAK features a simplified calling syntax for the actuarial functions, this manual is essential reading for proper use of the package. Also included in the manual is an up-to-date directory for the SHARP APL Actuarial data base.
June 1983, 114 pp., \$8.00

Grant Clarke, Toronto



mailing request

- Please amend my mailing address as indicated.
- Please add the following name(s) to your Newsletter mailing list.
- Please send me a Publications Order Form.
- Please add my name to the Aviation Newsletter mailing list.
- Please add my name to the Energy Newsletter mailing list.
- Please add my name to the Financial and Economic Newsletter mailing list.
- Please add my name to the PROMIS Newsletter mailing list. PROMIS is a computer-aided manufacturing system for microelectronics.

Name: _____

Title: _____

Company: _____

Address: _____

Please send me information about your courses in:

City: _____ Tel.: _____

The *I.P. Sharp Newsletter* is published by I.P. Sharp Associates, Suite 1900, 2 First Canadian Place, Toronto, Canada M5X 1E3.
Your comments and contributions are welcome.

Editor: Irene Shimoda
Circulation: Mary Kopfensteiner
Printed in Canada
July 1983
ISSN 0226 854X

international offices

Aberdeen

I.P. Sharp Associates Limited
5 Bon Accord Crescent
Aberdeen AB1 2DH
Scotland
(0224) 25298

Amsterdam

InterSystems BV
Kabelweg 47
1014 BA Amsterdam
The Netherlands
(020) 86 80 11
Telex: 18795 ITS NL

Atlanta

I.P. Sharp Associates, Inc.
1210 S. Omni International
Atlanta, Georgia 30335
(404) 586-9600

Boston

I.P. Sharp Associates, Inc.
1 Liberty Square
Boston, Massachusetts 02109
(617) 542-2313

Brisbane

I.P. Sharp Associates Pty. Ltd.
6th Floor, Dalgely House
79 Eagle Street
Brisbane, Queensland 4000
Australia
(07) 229-8330

Brussels

I.P. Sharp Europe SA
Boulevard de la Cambre 36, bte 5
1050 Bruxelles
Belgium
(02) 649 99 77

Calgary

I.P. Sharp Associates Limited
Suite 550, Bow Valley Square 4
250-6th Avenue S.W.
Calgary, Alberta T2P 3H7
(403) 265-7730

Canberra

I.P. Sharp Associates Pty. Ltd.
16 National Circuit
Barton, A.C.T. 2600
Australia
(062) 73-3700

Chicago

I.P. Sharp Associates, Inc.
Suite 3860
55 West Monroe Street
Chicago, Illinois 60603
(312) 782-3177

Copenhagen

I.P. Sharp ApS
Østergade 24B
1100 Copenhagen K
Denmark
(01) 11 24 34

Coventry

I.P. Sharp Associates Limited
7th Floor B Block
Coventry Point, Market Way
Coventry CV1 1EA England
(0203) 56562

Dallas

I.P. Sharp Associates, Inc.
Suite 1148, Campbell Center
8350 North Central Expressway
Dallas, Texas 75206
(214) 369-1131

Denver

I.P. Sharp Associates, Inc.
Suite 416
5680 South Syracuse Circle
Englewood, Colorado 80111
(303) 741-4404

Dublin

I.P. Sharp Associates Limited
Segrave House
Earlsfort Terrace
Dublin 2, Ireland
(01) 763605

Düsseldorf

I.P. Sharp GmbH
Kaiserswertherstrasse 115
4000 Dusseldorf 30
West Germany
(0211) 45 20 52

Edmonton

I.P. Sharp Associates Limited
2358 Principal Plaza
10303 Jasper Avenue
Edmonton, Alberta T5J 3N6
(403) 428-6744

Halifax

I.P. Sharp Associates Limited
Suite 706, Cogswell Tower
2000 Barrington Street
Halifax, Nova Scotia
B3J 3K1
(902) 423-6251

Helsinki

TMT-Team Oy (Agent)
Westendintie 1
SF-02160
Espoo 16, Finland
(0) 452 1155

Hong Kong

I.P. Sharp Associates (HK) Limited
Suite 606, Tower 1
Admiralty Centre, Hong Kong
5-294341

Houston

I.P. Sharp Associates, Inc.
Suite 375, One Corporate Square
2600 Southwest Freeway
Houston, Texas 77098
(713) 526-5275

London, Canada

I.P. Sharp Associates Limited
Suite 510, 220 Dundas Street
London, Ontario N6A 1H3
(519) 673-4060

London, England

(European Headquarters)
I.P. Sharp Associates Limited
132 Buckingham Palace Road
London SW1W 9SA, England
(01) 730-4567
Telex: 8954178 SHARP G

Los Angeles

I.P. Sharp Associates, Inc.
Suite 1230
1801 Century Park East
Los Angeles, Ca. 90067
(213) 277-3878

Madrid

I.P. Sharp Associates Limited
Serrano 23, Piso 8
Madrid 1, Spain
(91) 276 70 54

Melbourne

I.P. Sharp Associates Pty. Ltd.
520 Collins St., 7th Floor
Melbourne, Victoria
3000, Australia
(03) 614-1766

Mexico City

Teleinformatica de Mexico SA
(Agent) Mail to:
Arenal N 40, Chimalistac
Mexico 20 D.F., Mexico
(905) 550-8033

Miami

I.P. Sharp Associates, Inc.
Suite 240
15327 N.W. 60th Avenue
Miami Lakes, Florida 33014
(305) 556-0577

Milan

Informatical Society Italia Srl
(Agent)
Via Eustachi 11
20129 Milan, Italy
(02) 221 612

Montreal

I.P. Sharp Associates Limited
Suite 1610
555 Dorchester Boulevard W.
Montreal, Quebec H2Z 1B1
(514) 866-4981

New York

I.P. Sharp Associates, Inc.
Suite 210
230 Park Avenue
New York, N.Y. 10169
(212) 557-7900

Newport Beach

I.P. Sharp Associates, Inc.
Suite 1135
610 Newport Center Drive
Newport Beach, Ca. 92660
(714) 644-5112

Oslo

I.P. Sharp A/S
Postboks 486 Sentrum
Dronningens gate 34
Oslo 1, Norway
(02) 41 17 04

Ottawa

I.P. Sharp Associates Limited
Suite 600, 265 Carling Ave.
Ottawa, Ontario K1S 2E1
(613) 236-9942

Palo Alto

I.P. Sharp Associates, Inc.
Suite 201, 220 California Ave.
Palo Alto, Ca. 94306-1683
(415) 327-1700

Paris

I.P. Sharp SARL
9 Rue du Cirque
75008 Paris
France
(1) 225 98 20

Philadelphia

I.P. Sharp Associates, Inc.
Suite 604, 437 Chestnut St.
Philadelphia, Pa. 19106
(215) 925-8010

Phoenix

I.P. Sharp Associates, Inc.
Suite 503
3033 N. Central Avenue
Phoenix, Arizona 85012
(602) 264-6819

Rochester

(United States Headquarters)
I.P. Sharp Associates, Inc.
1200 First Federal Plaza
Rochester, N.Y. 14614
(716) 546-7270

Rome

Informatical Society Italia Srl
(Agent)
Piazza Della Rotonda 2
00100 Rome, Italy
(06) 656-5925

San Francisco

I.P. Sharp Associates, Inc.
Suite C-415, 900 North Point St.
San Francisco, Ca. 94109
(415) 673-4930

Saskatoon

I.P. Sharp Associates Limited
Suite 303, Financial Bldg.
230-22nd St. E.
Saskatoon, Saskatchewan
S7K 0E9
(306) 664-4480

Seattle

I.P. Sharp Associates, Inc.
Suite 223, Executive Plaza East
12835 Bellevue Redmond Road
Bellevue, Washington 98005
(206) 453-1661

Seoul

Daewoo Corporation (Agent)
541, 5-Ga, Namdaemoon-Ro
Jung-Gu (CPO Box 2810)
8269 Seoul, Korea.
771-91/2
Telex: 23341-5/24295 DAEWOO K

Singapore (Far East H.Q.)

I.P. Sharp Associates(S) Pte. Ltd.
Suite 1601, CPF Building
79 Robinson Road
Singapore 0106
Republic of Singapore
2230211
Telex: 20597 IPSAS RS

Singapore

Singapore International
Software Services Pte. Ltd.
(Agent)
Suite 1601, CPF Building
79 Robinson Road
Singapore 0106
Republic of Singapore
2230211

Stockholm

I.P. Sharp AB
Kungsgatan 65
S111 22 Stockholm, Sweden
(08) 21 10 19

Sydney (Australian H.Q.)

I.P. Sharp Associates Pty. Ltd.
8th Floor, Carlton Centre
55 Elizabeth St.
Sydney, New South Wales 2000
Australia
(02) 232-6366

Tokyo

INTEC, Inc. (Agent)
37-18, 3-chome, Hatagaya
Shibuya-ku, Tokyo 151
Japan
(03) 320-2020
Telex: 2322008 INTECA J

Toronto

(International Headquarters)
I.P. Sharp Associates Limited
Suite 1900
2 First Canadian Place
Toronto, Ontario
M5X 1E3
(416) 364-5361
Telex: 0622259 I P SHARP TOR

Vancouver

I.P. Sharp Associates Limited
Suite 902, 700 West Pender St.
Vancouver, B.C. V6C 1G8
(604) 687-8991

Victoria

I.P. Sharp Associates Limited
Chancery Court
1218 Langley Street
Victoria, B.C. V8W 1W2
(604) 388-6365

Vienna

I.P. Sharp Ges.mbh
Renngasse 4
A-1010 Wien, Austria
(0222) 66 42 48

Warrington

I.P. Sharp Associates Limited
1-3 Dolmans Lane
Warrington, Cheshire
WA1 2ED England
(0925) 50413/4

Washington

I.P. Sharp Associates, Inc.
Suite 305, 2033 K Street N.W.
Washington, D.C. 20006
(202) 293-2915

Wayne

I.P. Sharp Associates, Inc.
Suite 303
155 Willowbrook Blvd.
Wayne, New Jersey 07470
(201) 785-8050

White Plains

I.P. Sharp Associates, Inc.
Suite 312 West
701 Westchester Avenue
White Plains, New York 10604
(914) 328-8520

Winnipeg

I.P. Sharp Associates Limited
Suite 208
213 Notre Dame Avenue
Winnipeg, Manitoba R3B 1N3
(204) 947-1241

Zurich

I.P. Sharp AG
Fortunagasse 15
CH-8001 Zurich, Switzerland
(01) 211 84 24

I.P. Sharp communications network

APL OPERATOR VOICE (416) 363-2051 COMMUNICATIONS (416) 363-1832

Our private, packet-switched network connects with the Value Added Networks—Datapac, Datex-P, PSS, Telenet, Telepac, Transpac, and Tymnet—to provide access from the 46 countries listed below:

- Argentina • Australia • Austria • Bahamas • Bahrain • Barbados • Belgium • Bermuda • Brazil • Canada • Chile • Denmark
- Dominican Republic • Finland • France • French Antilles • Germany • Great Britain • Hong Kong • Hungary • Ireland
- Israel • Italy • Japan • Korea • Kuwait • Luxembourg • Mexico • The Netherlands • New Zealand • Norway • The Philippines
- Portugal • Puerto Rico • Qatar • Republic of Singapore • Republic of South Africa • Saudi Arabia • Scotland • Spain
- Sweden • Switzerland • Taiwan • United Arab Emirates • U.S.A. • Virgin Islands

SHARP APL is accessible from over 600 places via a local telephone call. Please ask your nearest I.P. Sharp office or representative for a complete list of access points and access procedures. Our private network also connects with the worldwide Telex network via the Amsterdam and Rochester, New York nodes.