

I. P. Sharp



NEWSLETTER

EASY An Econometric Analysis System

Caroline Haythornthwaite, Toronto

EASY, I.P. Sharp's Econometric Analysis System, is an interactive system designed primarily for economists and social scientists. EASY provides a flexible and convenient environment for the analysis and forecasting of time series data.

In econometric analysis, you start by specifying a relationship between factors in the economy or in your business. For example, suppose you believe that beer sales over the summer depend on two temperatures: the atmospheric temperature and the degree of baseball fever (measured by the attendance records for the local team). You have, in effect,

specified a single-equation model that can be stated as: daily values of the variable *beer sales* (called the **dependent** or **endogenous** variable) are a function of the variables *temperature* and *attendance* (called the **independent** or **exogenous** variables).

The next step is to quantify this relationship, that is, to measure the effect on *beer sales* of a change in *temperature* or *attendance*. You estimate the equation, using observations from previous time periods. Then you can use various statistics to assess how well the model fits the observed data. If the model is a

"good fit", you can then predict future values of *beer sales*, based on estimates of future temperatures and attendance.

EASY allows you to perform all these steps—specification, estimation, statistical analysis, and simulation—with meaningful English commands such as *SPECIFY*, *REGRESS* and *SIMULATE*.

Once you have estimated the basic equation, it is a simple matter in EASY to vary the specifications and try different formulations of your model. You can apply EASY to your own corporate data or draw from I.P. Sharp's many public data bases.

In This Issue

Applications Software

EASY, An Econometric Analysis System 1

Conferences

Register Now for the 1982 APL Users Meeting 6
 APL83 Call for Papers 18
 Australian APL Conference 18

Special Systems

Computer-Aided Manufacturing of Semiconductors 7

Data Bases

Aviation Data Bases 10
 The UKCSO Macro-Economic Data Base 11
 Fund Monitor Unit Trusts and Insurance Bonds Data Base 13
 Commodity History 15

Network News

Changes in the Communications Network 15

Bulletin Board

Toronto, Dallas, Houston, Montreal 17

Corporate News

Boston Office on the Move 19
 New Dial Access Numbers in Boston and Hartford 19

Technical Supplement 38

Wheels within Wheels T1
 Searching, Part II T4
 Shuffle off to Buffalo T7
 System Reliability 1981 T8

These data bases contain a wide range of statistics from housing starts to airline load factors, from currency exchange rates to fat lamb futures. And if you want to see your results illustrated in colour, just use SUPERPLOT.

Using EASY

EASY is available in the workspace 39 EASY and any MAGIC workspace. In either case, simply type EASY to begin.

As an example of a single-equation model, let's consider a simple two-variable case that postulates personal consumer expenditures (CONS) to be dependent on personal disposable income (INCOME). Data for this example has been stored previously in a private data base, and is retrieved using MAGIC's VARDEF function. The data timeframe is set, and the variables CONS and INCOME are defined. Then we enter EASY's interactive mode: (see figure 1)

The command REGRESS initiates the most commonly used estimation technique, the method of ordinary least squares. This regression estimates the relationship between CONS (the dependent variable) and INCOME (the independent variable). The standard output is shown, giving the values of the regression coefficients and various descriptive statistics. You also have available additional regression results, such as an analysis of variance table or correlation matrix.

Now you can save the estimated equation, under the chosen name CONSUMPTION, by typing:

* REGRESS COMMAND: SAVE CONSUMPTION

Later you might want to examine a table or plot of the estimated residuals, or use these residuals in a subsequent specification.

```

)LOAD 39 MAGIC
SAVED 13.20.21 03/26/82
QUARTERLY, DATED 1 72 TO 4 81
TIE '39 EASYEXAMPLE'
VARDEF EASYEXAMPLE 'CONS,INCOME'
EASY

* REGRESS COMMAND: REGRESS CONS,INCOME

CONS = C0 + (C1 x INCOME)

TIMEFRAME: QUARTERLY, DATED 1 72 TO 4 81

MEAN OF DEPENDENT VARIABLE CONS          648.652

      VARIABLE          MEAN          COEFFICIENT          STD. ERROR          T-VALUE
CONSTANT TERM          26.186          13.314          1.967
INCOME          837.475          0.743          0.016          46.962

COEFFICIENT OF DETERMINATION R-SQUARED.....          0.9831
CORRECTED R-SQUARED (RBAR-SQUARED).....          0.9826
F-STATISTIC FOR SIGNIFICANCE OF REGRESSION( 1, 38)          2205.4146
STANDARD ERROR OF THE ESTIMATE.....          7.9158
DURBIN-WATSON STATISTIC.....          0.5635
COEFFICIENT OF VARIATION (AT THE MEAN OF Y)..(o/o)          1.2203
    
```

Figure 1

More complex equation estimation

The preceding example illustrates the most basic regression facility in EASY. You can also adjust for such common econometric pitfalls as autocorrelation, multicollinearity and heteroscedasticity, by using time-dependent error correction and various distributed lag schemes.

Simulation

Now you are ready to forecast future values of the dependent variable. This is accomplished in

EASY through the command *SIMULATE*.

To produce the forecast you will need estimates of the future values of the independent variables. These might be obtained from your sales department, from government figures or from astrological tables—whichever you consider more accurate. Alternatively, there are several methods available in EASY to generate these forecasts. One such method is the *PREDICT* function used in the example below.

(*PREDICT* forecasts your data by taking into account both the series trend and seasonal variation.)

Let's recall the consumption function saved earlier. Then we can simulate it over the next two years to produce predictions for the dependent variable *CONS*. (see figure 2)

You can also use *SIMULATE* to see how well values calculated from the equation match the observed values. All you have to do is specify a simulation timeframe that includes some or all of the data timeframe.

A colourful summary of the analysis is given in the *SUPERPLOT* on the following page.

* REGRESS COMMAND: **RECALL CONSUMPTION**

EQUATION RECALLED -- CONSUMPTION

* REGRESS COMMAND: **SIMULATE**

* SIMULATION TIMEFRAME? **1 82 TO 4 83**

* EXOGENOUS VALUES ARE REQUIRED FROM 1 82 TO 4 83 (8 PERIODS).

* EXTEND INCOME VIA? **PREDICT 3**

PREDICTION

PERIOD	95% LOWER LIMIT	FORECAST	95% UPPER LIMIT
1982 1	756.272	770.477	784.682
1982 2	768.194	782.538	796.882
1982 3	779.799	794.290	808.781
1982 4	791.020	805.663	820.307
1983 1	802.578	817.388	832.199
1983 2	817.330	832.369	847.407
1983 3	831.824	847.102	862.380
1983 4	845.976	861.502	877.028

Multiple-equation models

Often an exogenous variable, such as *INCOME* in the simulation above, is itself dependent on one or more other variables. In our initial example the degree of baseball fever is measured by attendance records. Instead it might be more accurately represented as a function of two variables: local sales of baseball paraphernalia and advertising expenditures on game promotion.

Figure 2

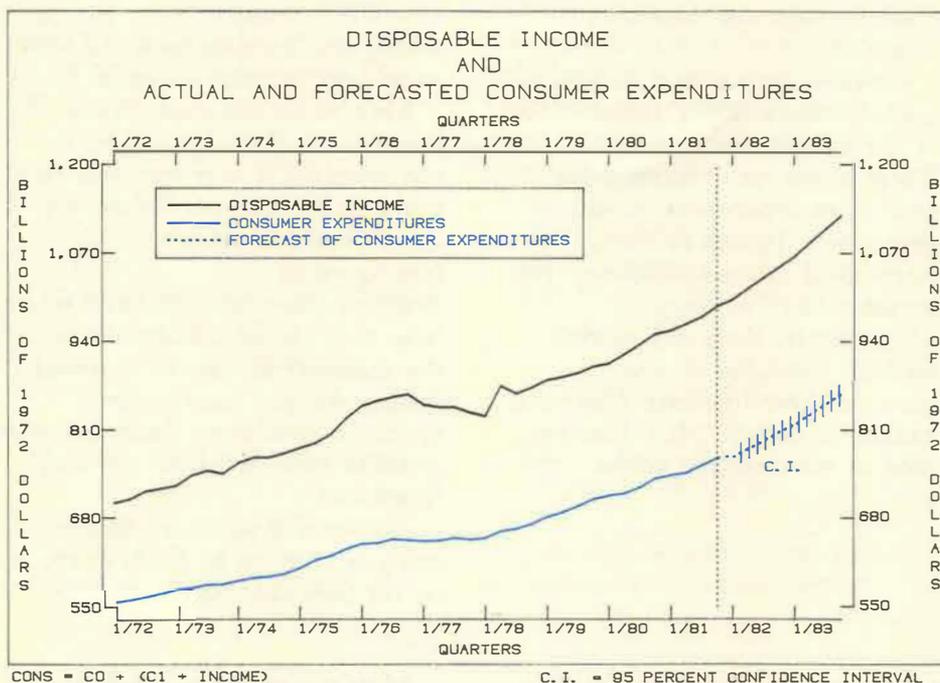


Figure 2 SUPERPLOT

In order to integrate this new equation with the original specification of *beer sales*, you need to specify a multiple-equation model.

As an example, let's postulate a slightly different model of consumption that could have been estimated and saved in EASY by the following commands:

* REGRESS COMMAND: `REGRESS PERCONS,GNP-TAXES, 1 LAG PERCONS`

* REGRESS COMMAND: `SAVE CONSLAGGED`

CONSLAGGED represents personal consumer expenditures as a function of gross national product less taxes and personal consumer expenditures lagged one period (i.e. how much a consumer spends this quarter is related to how much he spent last quarter).

TAXES is an exogenous variable, but *GNP* is not. *GNP* is instead defined, and saved, as the sum of personal consumer expenditures, investment and government expenditures by the following commands:

* REGRESS COMMAND: `DEFINITION GNP,PERCONS, INVEST,GOVT`

* REGRESS COMMAND: `SAVE OUTPUT`

Note that the definition of *PERCONS* depends on the definition of *GNP*, and vice versa. EASY can then resolve these two equations simultaneously. The sequence below follows through the stages to the solution. Assume that data already exists up to the last quarter of 1981.

A two-equation model is specified:

* REGRESS COMMAND: `SPECIFY CONSLAGGED,OUTPUT`

The solution interval is chosen as the last two quarters of 1981 and the first two quarters of 1982.

* REGRESS COMMAND: `INTERVAL 3 81 TO 2 82`

The next step is to obtain data values for the solution of the model. These can be generated by EASY, using a command such as:

* REGRESS COMMAND: `CREATE ALL FROM ANYEQUATION`

Or the analyst can supply the values by entering the following command:

* REGRESS COMMAND: **CREATE ALL FROM PROMPT**

Once all the data values have been supplied, the model is ready to be solved.

* REGRESS COMMAND: **SOLVE**

You can also save models for future reference or further analysis. The following command saves the model, with its dataset, under the name *BALANCEDBUDGET*.

* REGRESS COMMAND: **SAVE BALANCEDBUDGET WITH DATASET**

Now you can try variations of this model, such as assessing the effects of a change in tax revenues. For example, the variable *TAXES* can be reassigned the value 391 over the whole time interval, by the command:

* REGRESS COMMAND: **ASSIGN TAXES:391**

Then the model can be solved with these new values and saved under the name *TAXCUT*.

The *COMPARE* command contrasts the solution values for the new *TAXCUT* model with those of the *BALANCEDBUDGET* model so you can see the effects of the change in tax revenues.

(see figure 3)

* REGRESS COMMAND: **COMPARE ALL WITH BALANCEDBUDGET**

FIRST ROW IS TAXCUT. SECOND ROW IS BALANCEDBUDGET.

	8103	8104	8201	8202
<i>PERCONS:</i>	822.082 815.633	834.167 815.745	857.761 826.766	880.545 837.409
<i>GNP:</i>	1439.082 1432.633	1463.467 1445.045	1517.761 1486.766	1540.545 1497.409
<i>TAXES:</i>	391.000 407.603	391.000 422.396	391.000 425.000	391.000 425.000
<i>INVEST:</i>	214.300 214.300	226.700 226.700	235.000 235.000	235.000 235.000
<i>GOVT:</i>	402.700 402.700	402.600 402.600	425.000 425.000	425.000 425.000

* REGRESS COMMAND: **STOP**

Figure 3

Further documentation

A new *EASY* manual will be available soon. Contact your local I.P. Sharp representative to obtain this manual or information on I.P. Sharp's data bases and other statistical functions useful in econometrics.

**Register Now
for the 1982
APL Users Meeting**

Register now for the third APL Users Meeting, to be held at the Westin Hotel, in Toronto, from 4 October to 6 October 1982. There will be something for everyone: whether you are just thinking of introducing APL into your organization or you have been using APL since its innovation; whether you are a manager, an end user or an application programmer.

In the past, our Users Meetings have focused on APL applications. This year, applications sessions will be complemented by special interest groups, special technical topics, tutorials and workshops.

The applications sessions are still, of course, a very important aspect of the meeting. This year's program features sessions on:

- International banking
- Programming tools and techniques
- Budgeting systems
- Management science applications
- Management information systems
- Financial applications
- Production planning and control systems
- Multinational financial applications
- Exploiting public data bases
- Actuarial and insurance applications

Special interest groups, held each evening, will provide an informal atmosphere for the exchange of ideas in subject areas

such as arrays, in-house APL, or graphics, or gather together people with a mutual interest in particular applications. We welcome suggestions for groups you would find especially interesting.

Three half-day sessions will focus on particular aspects of APL systems design. The special technical topics are: designing user friendly systems, designing maintainable systems, and designing efficient systems. The material covered in these sessions will be published as a separate volume to the conference proceedings and promises to be a handy reference tool.

An addition to the main conference this year will be the Sunday tutorials. These full-day courses, to be held on October 3rd, will introduce those unfamiliar with APL to the language. The Sunday tutorials will help the novice user get the most out of the following three days. The sessions are tailored to four groups: financial planners, statisticians and economists, managers, and actuaries.

Another new feature this year are workshops on the following topics: *Introducing APL*, *Managing APL*, and *APL Training*. These workshops will be of interest to a broad segment of

the APL community and will provide an opportunity for the audience to participate in the discussion. The floor will be opened for questions, comments and debate after brief presentations from a panel of people who have dealt with the various problems posed in these areas.

Of course, there will be plenty of time to meet other APL users throughout the day or during the evening in the hospitality suites. We expect over 800 users at the meeting who are involved in every conceivable APL application, from all continents. We hope to see you there.

The registration fee for the 1982 APL Users Meeting is \$250 U.S./\$295 Canadian for the three days. If you wish to attend the Sunday tutorials as well, the all-inclusive fee is \$300 U.S./\$355 Canadian. Registration forms and further information are available from:

1982 APL Users Meeting
I.P. Sharp Associates Limited
Suite 1900, Exchange Tower
2 First Canadian Place
Toronto, Ontario
M5X 1E3

Or Call Rosanne Wild or David
Manson at:

(416) 364-5361

Computer-Aided Manufacturing of Semiconductors

Douglas Scott, Toronto

The improvements in productivity achieved by semiconductor manufacturers in the past decade have not been matched by any other industry. These improvements have resulted from the increased density of integrated circuits (i.e. reduced feature sizes), the successful use of larger silicon wafers, yield improvements, and the ability to process wafers more quickly.

In the 1980s, semiconductor manufacturers are using automation more as a tool to enhance productivity even further. One reason for this move towards automation is the large number of variables involved in semiconductor manufacturing. Integrated circuits are produced on silicon wafers, usually three to four inches in diameter. Depending on the complexity of the circuit and the wafer size, there may be 100 to 600 integrated circuits on a single wafer. An average wafer fabrication line today produces about 12 000 wafers per month.

A typical manufacturing process consists of 150 process steps, many of which have 10 to 20 control parameters. Most wafer fabrication lines produce 30 to 40 different products, and a single process area may use 60 different types of equipment. A typical fabrication line can easily generate 400 000 items of information per day. Efficient analysis and correlation of this data is impossible without the aid of computers.

Economic justification

Construction of a wafer fabrication area today costs between \$20-

million and \$50-million. A computer aided manufacturing system for such a facility, including hardware and software, may cost from \$350-thousand to \$800-thousand.

A productivity improvement of only two percent is sufficient to justify the system cost. Dataquest Inc., independent consultants to the semiconductor industry, state that "while it is difficult to quantify, an effective computer management system in wafer fabrication can increase productivity from 10 percent to 30 percent".

PROMIS overview

PROMIS is the **Process Management and Information System** which arose out of a joint development program between I.P. Sharp's Special Systems Division and General Electric's Solid State Applications Operation in Syracuse, New York. It is believed by many semiconductor industry experts to be the most advanced system of its kind available.

Work on the system began in 1978, and 30 man-years of software development effort have been spent on the system so far.

The prime objective of PROMIS is to improve the productivity of semiconductor engineering and manufacturing by providing more accurate and timely information about

production operations. Major functions performed by PROMIS include:

- Process management
- Automatic routing and tracking of work in process
- Up to the minute production reporting
- Accurate inventory records
- Engineering analysis
- Equipment status recording and reporting
- Facilities monitoring and control, with automatic distribution of alarm messages to appropriate work centres

PROMIS can be used to analyse trends in yield by product, by process, by recipe, by equipment used, by type of wafer, by dozens of other variables, or by any combination of these, over any time period.

Production reports include forecast and actual versus planned completions, work in process inventory, production activity, performance to schedule, production throughput performance, yield variance, reworks, rejects, slow moving lots, equipment availability, and equipment usage.

In addition to the above, features being released this year include production planning and control, production scheduling, equipment capacity planning and control, materials planning, and standard costing.

PROMIS runs on a Digital Equipment VAX computer, with facility sensors connected to front-end microprocessors.

Programmable manufacturing

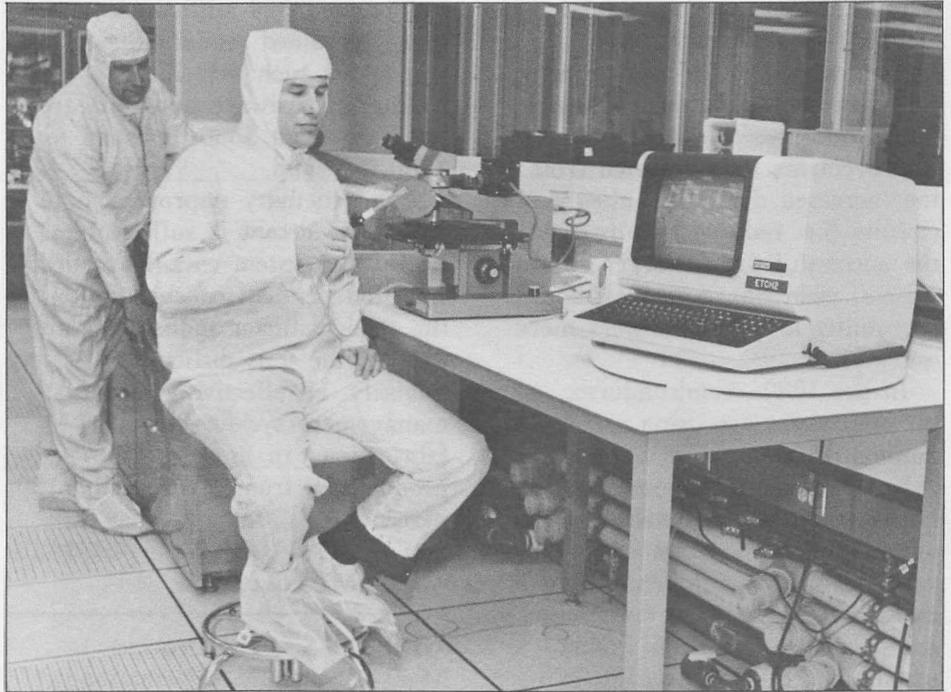
PROMIS provides development and production engineers with the ability to create, modify, and freeze entire manufacturing processes on line. Users require no knowledge of any programming language. They are provided with a series of simple, English-language prompts for all information that they need to enter.

All wafer processing is controlled by recipes maintained on the central computer, and PROMIS automatically distributes the right recipe to the right work centre at the right time, for all lots in the fabrication area. Depending on whether the fabrication area is an engineering lab or a production line, there may be 300 to 2 000 lots active at one time.

Paperless factory

One of the difficulties associated with semiconductor manufacturing is that process yields are severely affected by the intrusion of stray particles into the clean rooms. For example, it has been found that yields can be significantly affected by particles of cosmetics, and even by carbon particles exhaled from the lungs of a smoker up to six hours after having a cigarette.

Since ordinary paper can be a significant source of stray particles, some semiconductor manufacturers are now banning all paper documents from their clean rooms. PROMIS provides the mechanism by which all information required



This PROMIS terminal in General Electric's Semiconductor Products Dept. is used to enter defect codes during visual inspection of wafers.

by operating personnel is available to them via video displays, and all feedback of data to the management system is via video display keyboards, or direct interfaces to facility sensors and equipment.

APL prototype

The procedure usually followed in developing a new computer system is to submit the requirements specification to users for their approval. The problem with this approach is that users often find the specification dull to read and difficult to understand. They cannot easily visualize from a paper description how a highly interactive paperless system will

actually function in their own environment. Furthermore, it is difficult to determine whether such a specification is complete, correct, consistent, and unambiguous.

This problem of communication between user and developer can be greatly alleviated if a prototype of the proposed system is evolved. In the case of PROMIS, I.P. Sharp's Special Systems developed a complete prototype of the system in APL before the specifications were finalized. Users were able to exercise the prototype in the same way they would use the production system, and hence they provided valuable feedback to the developers early in the creation phase.

The major benefit that resulted from this modelling exercise was that, after three to four revisions of the prototype, the users knew and understood exactly what they would be getting. They liked the user interface and, when the production software was completed, they got what they wanted first time.

A panel discussion at the Fourth International Conference on Software Engineering dismissed software prototypes on the grounds of cost. In the case of PROMIS, however, the total cost of the prototype, which includes all manpower and computer usage costs, both for developing and for exercising the model, amounted to less than \$50-thousand; i.e. less than 10 percent of the Phase I PROMIS software cost.

Other engineering disciplines do not embark on major production jobs without first building a scale model, a prototype, a pilot plant, or whatever. We, as software engineers, are complacent if we imagine we can readily dispense with the prototype phase. APL is the only tool that makes it feasible economically.

User-friendly design

The consensus of the semiconductor industry is that PROMIS has given new meaning to the term "user friendly" system design. The tree-structured command menu provides quick and easy access to all PROMIS commands, usually via a single keystroke plus carriage return.

The *HELP* function provides access to more than 200 pages of information on topics such as:

- Sign on and sign off
- How to use the keyboard
- User introduction to PROMIS
- Glossary of terms
- Facility and work centre layouts
- Production and engineering document control system
- Useful phone numbers

In addition, *HELP* files are provided for every branch of the command tree and for every command within each branch; i.e. the entire user manual is available on line. There are presently about 150 PROMIS commands, most of which have multiple options.

The assistance provided at any point depends on where the user is in the command tree, and two levels of detail are provided for every command: one oriented towards the experienced user who wants a quick reminder; the other towards the novice user learning the system for the first time. In this and other ways, PROMIS has been made "user friendly" for both experienced users and novices, whose needs often conflict.

Other examples of the friendly user interface are:

- Use of English-language prompting and messages throughout.
- When a user is prompted for a limited choice response, hitting the RETURN key provides a list of the available options.
- Video highlighting (reverse video, dual intensity, underline, flashing).
- Use of split screen techniques.
- Automatic paging.
- Electronic mail, both for individuals and for work centres.
- Ability to request a printed copy of anything from a screen to an entire terminal session, on a hardcopy printer located outside the clean area.
- Ability to save user input sequences ("scripts" or "command procedures") to facilitate re-execution of periodic reports or complex multiple-search-criteria requests.
- Screen as well as hard copy graphics.

The use of automation in semiconductor manufacturing facilities is increasing rapidly as more and more companies discover the benefits of improved information systems, including the potential for productivity enhancement.

Aviation Data Bases

Keith Wright, Toronto

I.P. Sharp's involvement in the field of numeric data bases began with aviation data. At present, we are one of the world's major suppliers of aviation data.

For the aviation analyst, we offer a wide variety of data bases ranging from the Official Airline Guide (OAG), which contains information on all scheduled flights in the world, to the Commuter Online Origin-Destination (COAND) data base, which contains information about passengers and cargo carried by U.S. commuter airlines, or air taxi operators.

The following summary of aviation data bases maintained by I.P. Sharp is divided into data bases covering international airlines and those with U.S. data only.

INTERNATIONAL AVIATION DATA

Official Airline Guide (OAG2)

In the field of international aviation data, the Official Airline Guide data base contains over 20-million unique flight lines. The information in this data base is supplied to us in machine-readable form by Official Airline Guides, Inc. of Chicago.

For all of the scheduled flights of approximately 660 international airlines, the OAG data base contains more than 20 pieces of information, including origin, destination, equipment type, service class, departure and arrival times, and weekly frequency. The

SHARP APL system allows easy and effective analysis of this information.

International Civil Aviation Organization (ICAO)

The ICAO data base contains the traffic statistics of more than 400 airlines and 300 airports worldwide. The International Civil Aviation Organization is a United Nations organization which collects such statistics as aircraft movements, passengers embarked and disembarked, number of departures, hours in flight, and freight and mail tonne-kilometers performed.

The particular strength of this data base is its international data, with a common set of definitions and reporting rules for all reporting airlines.

Aircraft Accidents (AISL)

The AISL data base contains information on aircraft accidents as supplied by the Aviation Information Services Limited, in London, England. The data goes back to 1952.

Association of European Airlines (AEA)

For analysts who are interested strictly in data on European airlines, there is the AEA data base. This data base is maintained by the Association of European Airlines, which summarizes and forecasts the traffic and financial data of its member airlines. In this

data base there are over 175 000 time series, which contain statistics on operations, intra-European point-to-point traffic, operating revenues and costs, and forecasts of scheduled passenger traffic.

U.S. International Air Travel Statistics (INS)

The INS data base, compiled by the U.S. Department of Transportation in cooperation with the Department of Justice, Immigration and Naturalization Services, contains information supplied by both U.S. and non-U.S. carriers. Passenger counts between U.S. and foreign ports are broken down according to passenger citizenship (U.S. or other), flight type (scheduled or chartered), and the carrier's nationality.

U.S. AVIATION DATA

Form 41 (FORM41)

The name for the FORM41 data base is derived from the series of forms filed to the U.S. Civil Aeronautics Board by all U.S. certificated commercial air carriers. Form 41 contains balance sheet, income statement, and traffic statistics down to the aircraft type and airport level for over 60 carrier entities. I.P. Sharp obtains copies of the Form 41 schedules as they become available, edits and corrects them, then releases them within two or three days of receiving them.

I. P. Sharp



NEWSLETTER

technical supplement 38

Wheels within Wheels

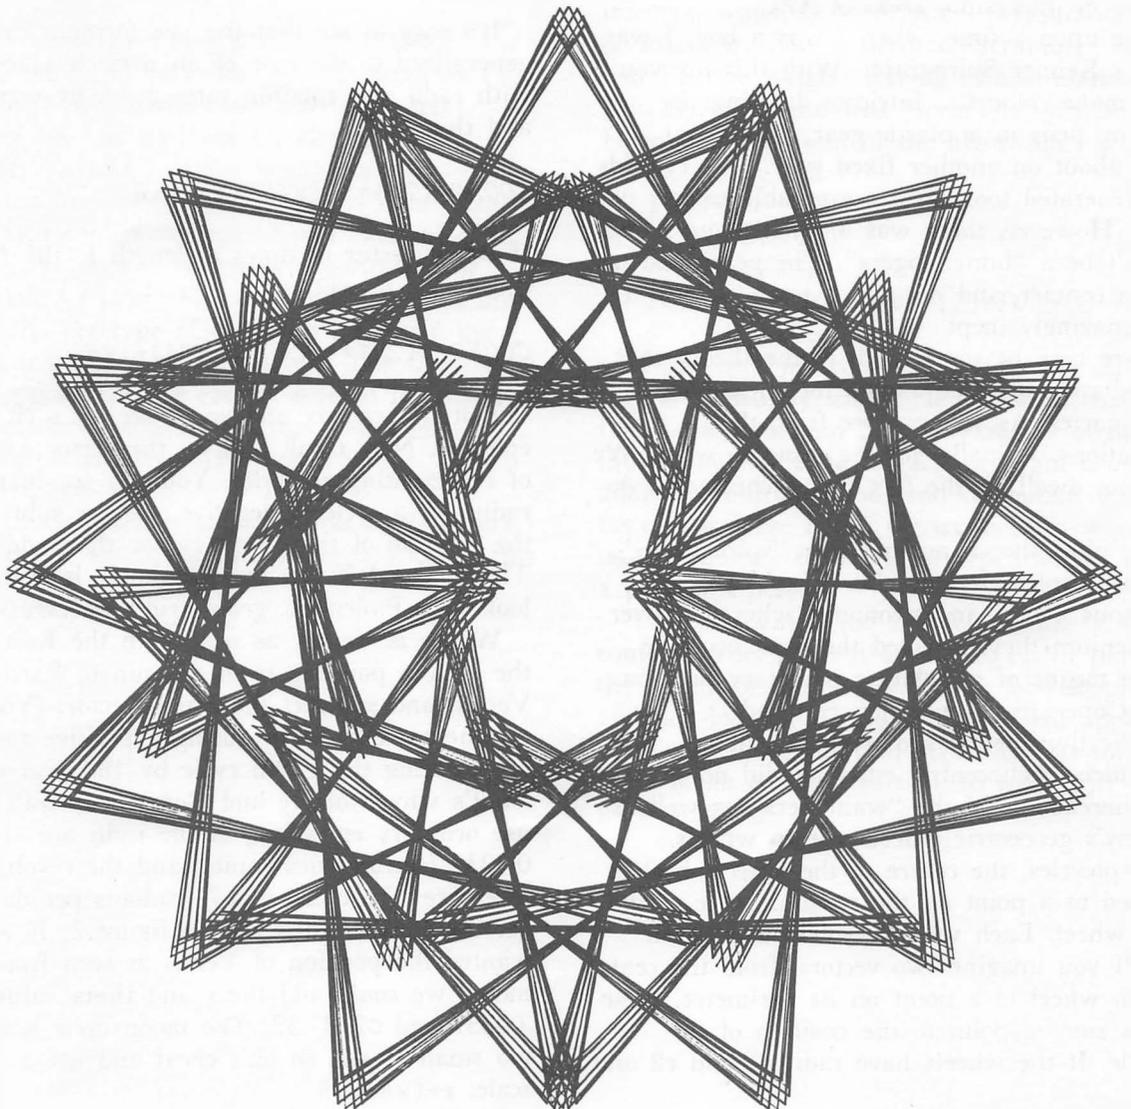


Figure 1 is produced using the following values: $r=1, (1\div 15), 1.3$ and $\text{Period}=(7\div 3), (11\div 6), 9$ and $\text{Repeat Time}=\wedge/7 \ 11 \ 9=693$. Each small circle has radius $(1\div 15)$ and is composed of 11 points. There are $9\times 7=63$ of them.

Or the plot thickens

Clement Kent, Toronto

It has often been said that the computer is a wonderful tool, but I think that a computer plus APL makes a great toy. No, I'm not joking. One of the greatest advantages of APL is that you can learn it by playing with it. For instance, the topic of this article started out as an attempt to emulate a toy, which instead led me to a number of interesting areas of APL.

Once upon a time, when I was a boy, I was given a Kenner Spirograph. With this toy you could make colourful, intricate drawings by inserting pens in a plastic gear, which you rolled about on another fixed gear. The cycloids thus generated took on a remarkable variety of forms. However, there was a catch—you couldn't be a "butterfingers". The gears had to stay in contact, and not slip a tooth or two. I was amazingly inept.

There may be some truth to the theory that as adults we overcompensate for childhood inadequacies. As you can see from the illustrations, I finally got the thing to work. We need not dwell on the fact that I cheated to do it.

The plots shown are not true "spirographs", they are Ptolemaic epicycles. They have an illustrious history in astronomy, where for over a millenium they provided the most accurate known means of calculating planetary positions. Even Copernicus' revolution revolved epicyclically, but the sad truth is that Copernicus' heliocentric epicycles did not predict the whereabouts of the "wanderers" as well as Ptolemy's geocentric wheels within wheels.

In epicycles, the centre of the outer wheel is attached to a point on the circumference of the inner wheel. Each wheel revolves at a given rate. If you imagine two vectors, from the centre of each wheel to a point on its perimeter, these vectors sum to point to the position of the epicycle. If the wheels have radii **r1** and **r2** and

revolve at rates **theta1** and **theta2** (in radians per unit time) the x,y coordinate vector at time **t** is:

$$COORD \leftarrow (R1 \times 2 \ 1 \circ T \times THETA1) + R2 \times 2 \ 1 \circ T \times THETA2$$

or,

$$COORD \leftarrow (2 \ 1 \circ . \circ T \times THETA1, THETA2) + . \times R1, R2$$

It's easy to see that the last formula can be generalized to the case of an **n** circle epicycle, with radii and rotation rates given by vectors **r** and **theta**.

$$COORD \leftarrow (2 \ 1 \circ . \circ T \times THETA) + . \times R$$

If **t** is a vector of times of length **k**, the (**k**, 2) coordinate matrix is:

$$COORD \leftarrow \circ (2 \ 1 \circ . \circ T \circ . \times THETA) + . \times R$$

Plot this matrix and you'll see an **n**-circle epicycle. Now think back to the vector sum way of representing epicycles. You can see that if the radius of a cycle is negative, we are subtracting the position of that circle rather than adding it. That's useful if you want to know how things look in a Ptolemaic, geocentric planetary system.

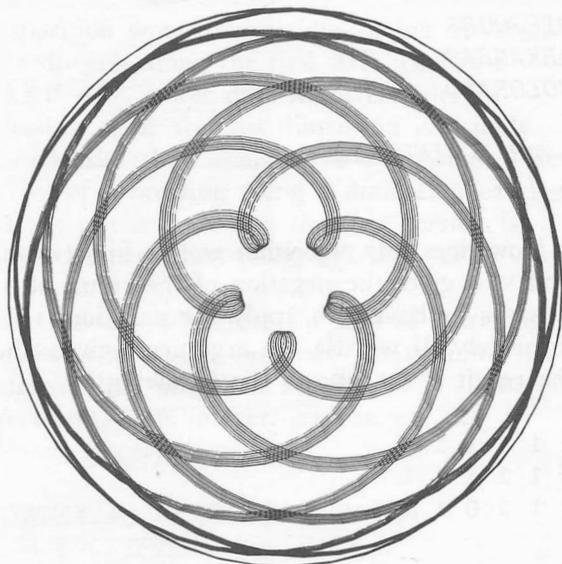
Where is Venus, as seen from the Earth? Get the vectors pointing from the sun to Earth and Venus, and subtract the Earth vector. (You can get the same result by using a positive radius but rotating the Earth cycle by 180 degrees. That's why Ptolemy and Copernicus both could use ordinary epicycles.) If the radii are -1. and 0.7233 (astronomical units) and the revolution rates are $0.2 \times 365.26 \ 224.7$ (radians per day) we get the Earth-Venus plot in figure 2. If we wanted the position of Venus as seen from the moon, we could add the **r** and **theta** values of $-.00257$ and 0.2×27.32 . The moon cycle would be too small to see, so let's cheat and use a log scale: $r \leftarrow (\times R) \times \circ R$.

Unfortunately, the planets refused to fit a two-cycle model, so more cycles had to be added (with five parameters, you can fit an elephant).

It is, of course, the eccentricity of the planetary orbits which complicates things, and Ptolemy certainly had an eccentric solution for it. You might be interested to see how close you can get to an ellipse using a two-cycle model. *Hint:* the eccentricity e of an ellipse is 1 minus the ratio of the shortest radius (semiminor axis) to the longest radius (semimajor axis). If the outer circle has radius $e \div 2$, what should its period be? What should be the radius of the inner circle? For comparison, the eccentricities of the orbits of Mars and Mercury are .093 and .206. (Users interested in calculating true planetary positions should see workspace 13 *STARMAP*.)

Now, planetary motion does not produce very nice plots due to the obstinate refusal of the planets to have periods that are nice round numbers. With plastic gears on the heavenly spheres, we could guarantee that after a fixed number of revolutions the cycle would close and we would be back where we started. How can we do it in APL?

Figure 2



Let's restrict ourselves to theta values that are of the form $0.2 \div PERIOD$, where $PERIOD$ is an integer vector. How long does it take before we return to the starting point? Each of the circles must do an integral number of revolutions, so we need a period $PTOT$ which is a multiple of each individual period. The answer, of course, is:

$$PTOT \leftarrow \wedge / PERIOD$$

which is very neat indeed, and an unexpected use of $\wedge /$. (See *Technical Supplement 30*, January/February 1981.)

The final gimmick I use to produce interesting "spirographs" is ellipsis, that is, I leave things out. If instead of calculating points on the epicycle at short intervals one gets them spaced farther apart, and connects successive points with straight lines, rather striking effects can be obtained. Figure 1 is an example, using $r \leftarrow 1, (1 \div 15), 1.3$ and $\theta \leftarrow (7 \div 3), (11 \div 6), 9$. Producing interesting plots is an intuitive process, but experience will show you that relative primeness of the periods used is important. So $v / PERIOD$ is important too. A toy for astronomers and number theorists!

Next issue, I will discuss other cyclic games, and some more serious tools for analysis of cycles in data. Stay tuned for "Turning the lambda knob", or "Is your limit set ergodic?"

Figure 2 is a plot of the Geocentric Venustian Orbit. The ratio of Period Earth to Period Venus is almost 13 to 8. After 8 Earth-years (13 Venustian-years) the orbit almost repeats itself. The number of inner lobes (5) is 13-8. Coincidence??

Searching, Part II

Robert Metzger, Rochester

In the first part of this series, we explored different algorithms for searching vectors. In this article, we will investigate ways to search tables.

2) Searching tables

There are two different ways you might want to search in a matrix. You can search for individual elements, or you can search for sections of the matrix, e.g. rows or columns.

When searching for elements, we might want to know which row contains specified elements. The function below does that.

ELMNTINROW $\omega v . = \alpha$

```

PRIMES
2 3 5 7
11 13 17 2
3 5 7 11
    
```

```

7 ELMNTINROW PRIMES
1 0 1
    
```

Notice that we use $v . =$ instead of ϵ . This is because the rules governing argument shape for inner product allow us to search by rows or columns. You may recall that we said in the previous article that ϵ treats the right argument as a vector.

To determine the number of the row in which the element was found, we could change the function to the following:

ELMNTINROW $(\omega v . = \alpha) f_{11} \uparrow \rho \omega$

Now, we might want to find out both the row and column positions where the element was found. This function does that for us.

ELMNTLOC $\square IO + \square (\rho \alpha) \uparrow (-\square IO) + (\omega = , \alpha) / \uparrow \times / \rho \alpha$

```

PRIMES ELMNTLOC 7
1 4
3 3
    
```

Note that we must convert the indices to origin 0 for the encode to work properly. This is done by subtracting $\square IO$. To make the result origin independent, we add $\square IO$ back in again.

Searching for subarrays (rows or columns) is a little more complicated. There are four basic approaches. Since the problems are analogous, we will only discuss searching rows.

- A) Compare every row against every other row, using $\wedge . =$ or $v . \neq$.
- B) Encrypt each row into a single number and then use dyadic \uparrow .
- C) Sort-merge the rows using Δ .
- D) Enclose the rows and use dyadic \uparrow .

A) The inner product solution is the most commonly used one. The function below assumes that the number of columns in both arguments is the same.

ROWINDEX1A $\square IO + \uparrow / \wedge \backslash \omega v . \neq \square \alpha$

<i>R</i>	<i>S</i>
OHIO	COLORADO
TENNESSEE	DELAWARE
MISSISSIPPI	MISSOURI
MISSOURI	
ARKANSAS	
COLORADO	

```

R ROWINDEX1A S
6 7 4
    
```

How does this algorithm work? First we note that $v . \neq$ gives the negation of the result of $\wedge . =$. It saves us having to apply the negation separately. If we use the arguments given above, the result of $\omega v . \neq \square \alpha$ is the following:

```

1 1 1 1 1 0
1 1 1 1 1 1
1 1 1 0 1 1
    
```

We can replace with 0's, those 1's which are on the right of 0's by using **and scan** (^).

```
1 1 1 1 1 0
1 1 1 1 1 1
1 1 1 0 0 0
```

The row summation of this matrix (+/) gives the origin 0 position of the first occurrence in each row.

```
5 6 3
```

Then we add the index origin to make the calculation origin independent.

```
6 7 4
```

As an experiment, you might try comparing the CPU time used by this function to the following variation on the previous algorithm.

```
ROWINDEX1B< [IO++f^λv.≠ω
```

You might think that if the matrix being searched (α) is a great deal larger than the matrix being searched for (ω), the cost of the transposition would overshadow doing the scan and reduction along the first axis. (Because SHARP APL stores data in ravel order, processing along the last dimension is usually less expensive than along other dimensions.) But the cost of processing along a dimension other than the last is such that the first version is always cheaper when using SHARP APL.

B) One way to reduce the cost of searching matrices is to reduce each row (or column) to a single number, and then use dyadic \uparrow . If the matrix already is integer, we can use the functions given below.

```
ROWINDEX2A< (α ENCRYPTNUM B)⊥
ω ENCRYPTNUM B←α BASE ω
Δ [CT←0
```

```
ENCRYPTNUM< ω⊥ωα
```

```
BASE< 1+(↑fA)-L fA←α, [IO] ω
```

Now this encrypting algorithm is only effective up to a point. If the integers resulting from the decode operation exceed the largest integer distinctly representable (2×56 in SHARP APL), then the low order digits are lost and different rows may map into the same number. So, the longer the row, or the larger the numbers, the more likely that this algorithm will not work correctly.

The same concept can be applied to character data. Each character is mapped into its index in a set of characters. These numbers are then encrypted into a single number.

```
ROWINDEX2B< (α ENCRYPTCHAR SET)⊥
ω ENCRYPTCHAR SET
Δ [CT←0
```

```
ENCRYPTCHAR< (ρω)⊥(-[IO]+ω⊥(ωSIZEα)←α
```

```
SIZE< (1↑ρω), (L(ρα)⊗MAX)⊥1↑ρω
```

```
Δ< α
```

You can see that *ROWINDEX2B* parallels *ROWINDEX2A*, and *ENCRYPTCHAR* parallels *ENCRYPTNUM*. The variable *SET* which is global to *ROWINDEX2B* is a vector containing all the characters which will be found in the argument to be encrypted. The variable *MAX* which is global to *SIZE* can contain two possible values: $1 \times 2 \times 31$ or 2×56 . The first is the biggest integer representable in IBM/360-like computers in four bytes, the second in eight bytes. As we noted in the previous article, the speed of the index search will depend upon which data representation is used.

It is also possible to encrypt into complex numbers in SHARP APL. This can double the maximum number of columns or square the size of the alphabet compared to encrypting into

eight-byte real numbers. Since residue (and thus encode) is currently undefined for complex numbers, however, it will be more difficult to write an inverse function *DECRYPTCHAR*.

The function *SIZE* computes the maximum number of columns which can be encrypted. The table below shows the relationship between the size of the character set and the number of columns.

<u>MAX</u>	<u>ρSET</u>	columns
1+2*31	22-35	6
	36-73	5
	74+	4
2*56	26-34	11
	35-48	10
	49-74	9
	75+	8

If the left argument of *ROWINDEX2B* is constant, i.e. the same table is being searched over and over again, it can be encrypted once and passed in its encrypted form. This will save a good deal of CPU time.

```
ROWINDEX2C⊙
α1SET ENCRYPTCHAR ω Δ [CT←0
```

C) The third approach to searching for rows involves sorting the matrix. The algorithm for the function listed below was developed by J. Henri Schueler. It searches character matrices.

```
∇ INDEX←LEFT ROWINDEX3 RIGHT
;GRADE;ROWS;SELECT;IOTA
[1] LEFT←LEFT,[ρIO] ρAV[ρIO+255]
[2] GRADE←[AV]ρRIGHT,[ρIO] LEFT
[3] ROWS←1+ρRIGHT
[4] SELECT←GRADE<[ρIO+ROWS]
[5] INDEX←IOTA←ρROWS
[6] INDEX[SELECT/GRADE]←
((~SELECT)/GRADE)
[ρIO+(SELECT/ρSELECT)-IOTA]-ROWS
[7] INDEX[(∇/RIGHT≠LEFT[INDEX;])
/IOTA]←[ρIO+(1+ρLEFT)-1 ∇
```

Let's look at this algorithm in more detail. In line [1], the last element of ρAV is added as a row in order to insure that a row of the left argument will sort to the last position. In line [2], the SHARP APL dyadic grade up is used to grade the character matrix formed from joining the arguments. Prior to the implementation of dyadic grade up, we had to use a defined function in this situation. While it was efficient, the primitive function is much more so. In line [4] of *ROWINDEX3*, the rows of the left argument are identified. For each row of the right argument, line [6] computes the index of the first row in the left argument which equals it, or else exceeds it by the smallest amount. Line [7] marks those which are not found.

Since SHARP APL also allows numeric arguments of grade up to have rank greater than 1, we can make this algorithm work for numbers. We change two lines as shown below.

```
[1] LEFT←LEFT,[ρIO] ρRIGHT
[2] GRADE←ΔRIGHT,[ρIO] LEFT
```

D) The fourth approach to searching matrices requires enclosing each row as an element of a vector. This approach makes sense only if the data in the matrix is not fundamentally rectangular. Such a matrix has probably been padded with fill elements to justify short rows. An example is a matrix of names, one per row, in which blanks have been used to pad the shorter names.

The function listed below provides yet another means of searching the rows of a matrix.

```
ROWINDEX4⊙ (ENCLVECFROMMAT α)
∇ ENCLVECFROMMAT ω
∇ VECTOR←ENCLVECFROMMAT MATRIX
;SELECT;LENGTH
[1] SELECT←ϕ∇\ϕMATRIX≠1+0ρMATRIX
[2] MATRIX←(,SELECT)/,MATRIX
[3] LENGTH←+\+/SELECT
[4] VECTOR←(1+0,LENGTH)†">
LENGTH†"><MATRIX ∇
```

Like the second approach, significant reductions in cost can be gained if the left argument can be enclosed once, instead of every time it is used. The searching function would be simplified as follows:

```
ROWINDEX4B⊔ α1ENCLVECFROMMAT ω
```

Let's investigate how *ENCLVECFROMMAT* works. Line [1] identifies the fill elements (1 or 0) with which the rows have been left justified. In line [2] the fill elements are removed. In line [3] the lengths of the partitions are computed. In line [4] the data from each row is segregated and enclosed.

We have seen a variety of algorithms for searching matrices. The best choice will be one which fits the type of data you have and the system features available to you. The fact that four good algorithms can be defined for doing what amounts to table lookups shows how powerful a language APL really is.

Acknowledgments: Many thanks to J. Henri Schueler for assistance in the preparation of this article.

Shuffle off to Buffalo

Clement Kent, Toronto

In rewriting the fast Fourier transform, I found an interesting little piece of code. At first, it threw me for a loop, but then I figured out how to do it without iteration. Can you?

You are given that n , the length of the series to be transformed, is a power of two—let's say $n=2^*M$. The idea is to develop a vector *POSN* of integers, which is an origin 0 permutation vector of length n . This is used to "scramble" the data and Fourier weights during transformation. Initially, all it did was "scramble" my brains. Here's the old form of the algorithm:

```
POSN←,0
LP:→(N≤ρPOSN)ρ0
S←Y+Y←ρPOSN
POSN←(Sρ0,Y)+Sρ⊔(2,Y)ρPOSN ⊔ →LP
```

It can be replaced by a single APL statement that is pleasingly simple. It is also faster than the above algorithm for many values of n . But I found that a revised looping program using only addition, multiplication, and catenation was fastest for large n . As Gertrude Stein would tell us, we mustn't be afraid of repetition. A loop is a pool is a polo.

This problem is too simple to make a good contest, so let's call it a puzzle. I will publish the names of the authors of the best looping and non-looping solutions, if you will append your name, address, and algorithm (as three parts of an enclosed vector), to workspace 999 *CONTEST*.

SYSTEM RELIABILITY

System Reliability 1981

SHARP APL SYSTEM RELIABILITY

.....

	<i>OUTAGES</i>	<i>MINS-OUT</i>	<i>o/o UP</i>
<i>JAN 81 **</i>	5	88	99.8
<i>FEB 81 **</i>	11	296	99.1
<i>MAR 81 **</i>	5	460	98.8
<i>APR 81 **</i>	4	89	99.8
<i>MAY 81 **</i>	3	127	99.7
<i>JUN 81 **</i>	8	315	99.2
<i>JUL 81 **</i>	8	120	99.7
<i>AUG 81 **</i>	7	194	99.5
<i>SEP 81 **</i>	6	106	99.7
<i>OCT 81 **</i>	14	318	99.2
<i>NOV 81 **</i>	20	619	98.3
<i>DEC 81 **</i>	8	285	99.3

.....

The Civil Aeronautics Board currently produces reports, such as the quarterly *Air Carrier Financial Statistics*, using our FORM41 data base.

Service Segment (ER586)

The ER586 data base consists of detailed monthly summaries of all of the activity across each city-pair flown by U.S. certificated air carriers. This is a large and detailed data base that allows historical analysis of all flight activity down to the flight number and aircraft-type level. The data, filed with the U.S. Civil Aeronautics Board, consists of such items as number of flights, available pounds, passengers transported by class, and transported freight and mail. In all, there are approximately 60 facts available for each of the city-pairs flown by the reporting carriers.

Origin-Destination (OAND)

The OAND data base, another large compilation of reports submitted to the U.S. Civil Aeronautics Board, describes the actual manner in which the public utilizes the carrier schedules. The data is collected by a survey of the itineraries of 10 percent of all passengers travelling on U.S. certificated carriers, and consists of detailed itineraries, local and connecting passenger counts by

carrier, and total industry passenger counts by city-pair.

**Commuter Online
Origin-Destination
(COAND)**

For analysis of commuter airlines, there is the COAND data base, a compilation of information filed with the U.S. Civil Aeronautics Board by U.S. air taxi operators. These airlines have regularly scheduled passenger service using aircraft which seat fewer than 60 passengers. Air taxi operators report to the CAB, the number of passengers and the number of pounds of mail and cargo carried on a city-pair basis.

I.P. Sharp has offered aviation data bases to its customers since 1973, and continues to enhance existing data bases and develop new ones. These data bases, combined with the analytic and reporting systems available on the SHARP APL system, provide the aviation analyst with a powerful and efficient decision-making tool.

**The UKCSO Macro-
Economic Data Base**

*David Fowler
London, England*

The United Kingdom Central Statistical Office (UKCSO) data base consists of about two thousand time series containing data related to a variety of economic and statistical areas in the United Kingdom. The data is compiled by the CSO from a number of sources in the British Government, such as the Economics Group at the Treasury.

Updates are issued every month and a major revision is undertaken each October. Data is available from 1948 to the present, although not all series cover this entire period. In all, three frequencies are available: annual, quarterly, and monthly, with versions of the quarterly and monthly series seasonally adjusted.

The UKCSO data base is divided into six sections:

- Index of Industrial Production
- National Income and Expenditure
- Balance of Payments
- Prices
- Wages and Earnings
- Population and Manpower

Each group listed above is subdivided into tables of related series. The columns of these tables contain series titles and corresponding series codes. When accessing a series, its location in the data base is specified by entering both the series code and the number of the table in which the code appears.

Each table is identified by a five-digit number. The first two

digits are used to identify the group to which it belongs. The remaining three digits are used to identify a specific table within that group. For example, the first two tables in the section *National Income and Expenditure* are:

- 20001 Expenditure on the GDP at current prices—annual
- 20002 Expenditure on the GDP at current prices—quarterly

The columns in each table contain series codes. The series codes consist of eight alphanumeric characters containing the following information: base year, series identification, frequency, and an indication of whether or not a series has been seasonally adjusted. The title, series code, frequency, and range of all series available in the UKCSO data base can be obtained from the *CSO Macroeconomic Data Base Index*. The *Index* can be purchased from the London, England office of I.P. Sharp Associates Ltd.

Accessing the data base

Series contained in the UKCSO data base are accessed through MAGIC. Series can be accessed and displayed with the retrieval command *UKCSO*, followed by the table number and the series code of the series required. For example, to retrieve the series

containing the index of retail prices for transport and vehicles, type:

```
UKCSO '40003/56RPTV 9'
```

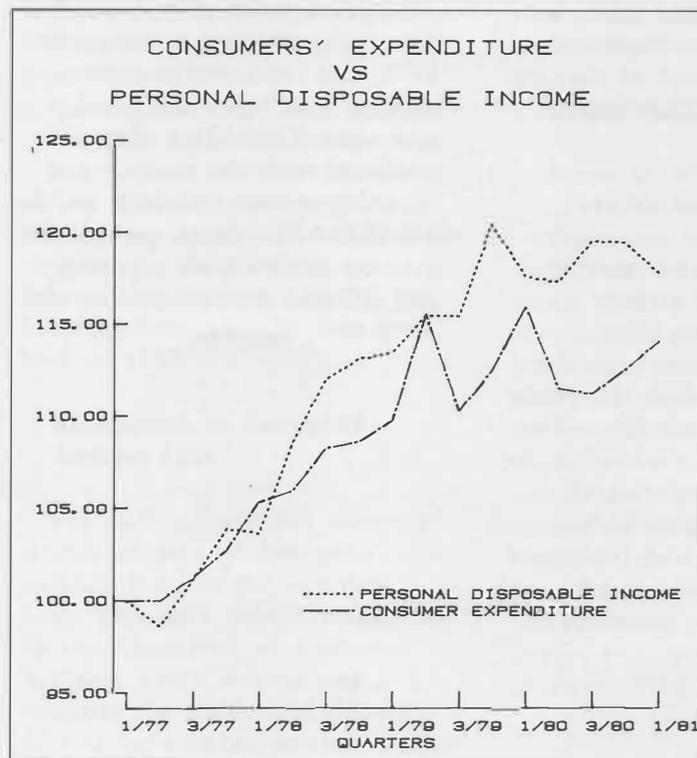
The table number and series codes are separated with a slash and enclosed in quotes. More than one series from the same table can be retrieved at once by separating the series codes with commas. For example, to retrieve the national income accounts for food, alcohol, and tobacco type:

```
UKCSO '20073/75CEFT 1,75CEADT1,75CETT 1'
```

To display a time series in tabular or graphical form, with headings, and labels within a certain timeframe, the MAGIC commands

DISPLAY, *TABLE*, and *PLOT* contained in workspace 39 *MAGIC* can be used.

The following plot, using data from the UKCSO data base, in conjunction with *MAGIC* and *SUPERPLOT*, illustrates how personal disposable income has outdistanced consumers' expenditure. The main reason for this growth was the rise in real wages, while inflation was moderated by the strength of sterling in '79 and '80.



1977 = 100 (1975 PRICES)

**Fund Monitor Unit
Trusts and Insurance
Bonds Data Base**

*Mona Albano and
Andy Neilson, Toronto*

The Fund Monitor Unit Trusts and Insurance Bonds data base (FUNDMONITOR) contains current and historical offer (asking) prices for units of two classes of United Kingdom investment funds. These two classes are all U.K. authorized unit trusts open to the public and most U.K. life insurance funds that are traded as units. Non-authorized unit trusts and unit trusts open only to special categories of investor (such as pension funds) are excluded. The unit trusts are open-ended mutual funds. They can be marketed to investors only if they are authorized by the United Kingdom Department of Trade, which enforces safeguards for the investors.

The data is monthly from December 1973, and weekly from January 1978. Its source is *The Financial Times* of London. The weekly data consists of each unit's asking price as of the end of the week, usually Friday, and the monthly data is the price as of the last Friday of the month. There are about 850 funds reported in the data base; the number increases as new funds are introduced.

The information in this data base is of interest to investment houses, financial institutions, brokers, trust managers, and investors. The data can be used for portfolio management, investing, forecasting, and for general economic research, including studies of portfolio performance and management.

Funds are divided into 16 groups, called **sectors**, according to the type of fund and its investment objectives. Each sector is represented by an index (see below). This division into sectors allows investors to compare the performance of a fund against its sector of the market, or one sector of the market against another. There are 4 sectors for life insurance funds and 12 sectors for unit trusts. Some funds are not classified, and have a sector number of 0. Descriptions of the sectors follow.

1. Income Unit Trusts
2. Growth Unit Trusts
3. General Unit Trusts
4. Small Company Unit Trusts
5. Special Situations and Recovery Unit Trusts
6. Financial Unit Trusts
7. Commodity Unit Trusts
8. International Unit Trusts
9. North American Unit Trusts
10. Far East Unit Trusts
11. Offshore Gilt Funds
12. Managed Life Funds
13. Property Life Funds
14. Equity Life Funds
15. Fixed Interest Life Funds
16. Fixed Interest Unit Trusts

As well as the data for each individual fund, the data base also contains 17 indices: one for each of the 16 sectors and a composite index for sectors 1 to 7 (i.e. trusts invested in the United Kingdom). Each index is a geometric mean of the prices of typical funds in that sector. The funds that comprise an index may change from time to time. The base date for all indices is January 1973 (1973=100), which is the start date for the monthly data.

DATA BASES

The following example, using workspace 39 *MAGIC*, retrieves weekly prices for seven funds during September and October 1981.

```

RESETOPTIONS ◇ AUTOLABEL
5 WEEKLY DATED 1 9 81 TO 31 10 81
TITLE 'SECTOR 1 TRUSTS, WEEKLY'
TITLE 'SEPTEMBER AND OCTOBER 1981'
ROWTITLE 'DATE'
COLWIDTH 9 12 11 8 8 11 9
'H' TABLE 317 297 121 265 282 411 419 FUNDWEEKLY 1
    
```

SECTOR 1 TRUSTS, WEEKLY
SEPTEMBER AND OCTOBER 1981

DATE	QUILTER QUADRANT INCOME	NATIONAL WESTMINSTER INCOME	CRAIGMOUNT HIGH INCOME	M AND G MIDLAND AND GENERAL	MIDLAND DRAYTON HIGH YIELD	HEDDERWICK WICKMOOR DIVIDEND	TYNDALL SCOTTISH INCOME
4SEP81	132.00	51.70	39.20	188.40	69.70	78.70	186.20
11SEP81	132.00	51.60	38.40	185.30	69.50	77.50	186.40
18SEP81	126.40	47.70	37.40	181.90	66.80	75.90	179.80
25SEP81	126.40	43.70	33.60	163.70	62.10	69.90	168.60
2OCT81	117.40	43.80	34.10	161.60	59.20	69.10	165.00
9OCT81	117.40	45.80	34.40	167.80	61.30	71.40	170.40
16OCT81	117.10	44.00	33.60	166.50	60.10	70.20	168.80
23OCT81	117.10	43.50	33.40	164.10	59.70	68.70	167.60
30OCT81	116.70	44.30	33.90	160.60	60.50	68.00	167.80

Commodity History

Andy Neilson, Toronto

The EUROCHARTS COMMODITIES data base contains over 10 000 price and volume statistics for all major commodities traded on London, Chicago, and New York futures markets. Data includes open, high, low, close, volume, and open interest for both metals and soft commodities. There is no subscription fee or surcharge for using the data base, above the usual timesharing charges. Data is available monthly from 1960 and daily from 1973, and is maintained by Eurocharts Limited of the U.K.

Recently, several Chicago commodities were added, including Treasury bills and Treasury bonds. A large amount of history was included in these new additions to facilitate historical analysis.

The following table shows the commodity, code, and dates for which historical data was added.

Commodity	Commodity Code	Available From
Treasury bills	TB	January 1976
Treasury bonds	US	August 1977
Lumber	LB	July 1977
Plywood	PW	July 1978
Feeder cattle	FC	July 1978
British pound	BP	October 1975
GNMA	GM	October 1975
IMM gold	AU	December 1974

Changes in the Communications Network

*David Chivers and
Stephen Crouch, Toronto
Geoff Oxer, Amsterdam*

Since November 1981, some changes in the I.P. Sharp Network support of asynchronous terminals have been taking place progressively. For most users, no changes to established procedures are necessary. The main purpose of the changes is to improve network reliability, and to extend the range of terminals supported.

Attention signal

Two changes are being made in the recognition of the attention signal.

When BREAK or ATTN is used, while APL is in input mode, the command always deletes the entire input line. You will get a new prompt for input. Previously, when you entered attention on a terminal directly connected to an Alpha port, a variable number of garbage characters was added to the input line, if APL was in input mode. This change brings the Alpha ports into harmony with the previously existing situation for users with terminals connected to an IBM 3705, or for those accessing the SHARP APL system through the public networks Telenet, Tymnet or Datapac.

When APL is not in input mode, in addition to signalling attention or interrupt by using the BREAK or ATTN key, you can enter the ASCII control signal ETX (CONTROL C) for this purpose. It has the normal effect of attention or interrupt. When APL is in input mode, ETX

continues to be recognized as input data. This change was introduced to accommodate those using word processors and personal computer terminals, which may not be equipped with a BREAK key.

Invalid host selection

When you attempt to access a network host destination on the I.P. Sharp Network, other than the default system (usually I.P. Sharp's production system), a destination command with an invalid destination was previously treated as a request to access the default system. In the new release, an invalid destination returns *LINES DOWN*. A carriage return immediately following the initial *O*) continues to be treated as a request to access the default system. A destination selection command now may consist of three or four characters between the *)* and the carriage return. Thus *)IPS* is a valid destination command, and *)ABCD* would also be valid if the destination *ABCD* existed.

New host selection

Immediately after signing off, you have the opportunity to make a network call to a new host by entering a new destination selection command. The connection to the terminal is maintained for 20 seconds, instead of being dropped as soon as the signoff message is printed. This capability, however, is not available when access is via Telex.

Also, if a network call is set up to a SHARP APL system which is not currently available for any reason, or if a successful sign on has not been completed, you can return to the destination selection state by typing *)OFF* followed by a carriage return. This might be appropriate, for example, after receiving *NUMBER NOT IN SHARP APL SYSTEM, INCORRECT SIGNON or APL PROBABLY DOWN*.

These changes to destination selection also allow an improved chance of a successful connection to SHARP APL systems served by more than one network IBM 3705.

Echoplex

For ASCII terminals only, characters typed at your terminal may optionally be echoed back to the terminal for verification that they have been correctly received by the network and not corrupted by the local telephone connection. To use this option, set the terminal to FULL DUPLEX, and type a letter *E* between the initial letter *O* and the right parenthesis of the destination selection command, that is, *OE*). Subsequent destination selection commands in the same call will then continue to assume echoplex operation. To revert to normal, sign off, then enter a destination selection command with a letter *L* before the right parenthesis, and set the terminal to HALF DUPLEX. With either option, the modem or acoustic coupler should always be set to FULL DUPLEX.

When using echoplex with full speed input, such as by use of the repeat key, there is a possibility that transmitted characters will be lost if the terminal sends characters faster than they can be echoed. The characters echoed to the terminal will be the same as those sent to APL, so visual fidelity is preserved. If the terminal is set to transmit two stop bits, there is no danger of characters being lost.

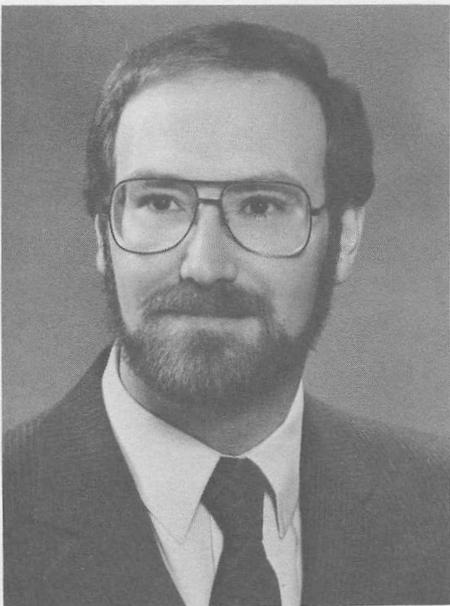
Other changes

In a future release, locally generated messages will be further distinguished from other responses by including a double linefeed in all local messages. Local messages include *LINES DOWN* and the prompt for a new destination selection command.

There has been a steady increase in the number of alternative paths through the I.P. Sharp Network, enhancing even further the network availability to users. Should you receive *LINES DOWN* in the middle of a session, it is worth trying immediately a new destination selection command, as there is a good chance that an alternative route to the destination exists.



Grant Clarke



Guy Bitton

Toronto

Grant Clarke is the new manager of the Publications Department in Corporate headquarters. Grant joined I.P. Sharp in May of last year as technical writer and editor. He has twelve years of experience in technical writing and publications. For the past few years, he was with Scotiabank at their Toronto Systems Centre. Grant also has several years systems programming experience, gained with the Reed Paper Group in the U.K., and with Ferranti-Packard in Toronto.

Grant intends to uphold the existing tenets of the I.P. Sharp Publications Department, namely to produce manuals which are informal and easy to read, while being thorough, precise, and technically accurate.

Dallas

Bill Paxson has been appointed branch manager in Dallas. He has a B.A. from St. Xavier College, Chicago, and an M.A. from Western New Mexico University, Silver City, New Mexico.

Before joining I.P. Sharp as a marketing representative in Chicago, Bill gained experience in teaching and in sales.

Under Bill's capable hands, the Dallas office will continue to be sensitive to our clients' needs.

Houston

David Hopkins is the new branch manager in Houston. He has a B.S. in computer science from Trinity University, San Antonio, and an M.B.A. from Southern Methodist University in Dallas.

David's previous experience includes working for Trinco, the Trinity Computer Center; STSC, Inc., as an applications consultant and marketing representative; and running simulations in APL with Texas E.C.I., an engineering consulting firm.

David has considerable expertise in planning models, budget systems, refinery simulations, and report writers for the energy industry.

Montreal

The new manager of the Montreal office is **Guy Bitton**. He is a graduate in mathematics from Montreal's Sir George Williams University, and has done extensive studies in marketing and operations research.

Previously, Guy was with Canada's largest brewery as senior marketing analyst. He now hopes to put his marketing skills to good use with the Montreal office's diverse client base.

**Call for Papers
APL 83**

APL83, the annual international conference of the APL community, will be held in Washington, D.C., from 10 April to 13 April 1983. This conference is sponsored by the Association for Computing Machinery (ACM) and the ACM Special Interest Group for APL (SIGAPL).

The program committee invites the submission of papers on all aspects of APL, including the following:

- Commercial and Scientific Applications
- Language Features
- Implementations
- Interfaces with Other Systems
- Tutorials
- Using APL in Education
- Teaching APL
- Interactions with Other Languages
- APL Programming
- APL as a Tool for Management

Drafts of full papers (four copies) should be submitted to the program chairman in accordance with the following schedule:

Drafts submitted for review
22 June 1982
Authors notified of acceptance
27 September 1982
Camera-ready copy due for
proceedings 26 December 1982

**Australian APL
Conference**

Further information is available in the workspace 1 *APL83*.

Submitted papers and questions regarding the conference technical program should be directed to the program chairman:

Tom Puckett
2000 Wyoming
Las Cruces, New Mexico
U.S.A. 88001

Telephone: (505) 522-6017

STSC, Inc./I.P. Sharp
MAILBOX Code: *THP*

Questions in other areas should be addressed to the conference general chairman:

Don Link
STSC, Inc.
2115 East Jefferson Street
Rockville, Maryland
U.S.A. 20852

Telephone: (301) 984-5333

STSC, Inc./I.P. Sharp
MAILBOX Code: *LINK*

The first Australian National Conference of APL users is to be held in Sydney, on 3 and 4 February 1983. The APL user groups in Sydney and Melbourne are jointly organizing the conference. The theme is *The Uses of APL in Commerce, Industry and Education*.

The sponsors invite submission of papers. Please send abstracts by 1 August 1982 to:

GPO Box 1425
Sydney N.S.W. 2001

**Boston Office
on the Move**

The I.P. Sharp office in Boston
has a new address:

I.P. Sharp Associates, Inc.
1 Liberty Square
Boston, Massachusetts 02109

New Dial Access Numbers

BOSTON

30 cps (617) 542-5433
120 cps (617) 542-5116

and a new telephone number:

(617) 542-2313

HARTFORD

30 cps (203) 249-9455
120 cps (203) 249-9797

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.
- Please send me information about your courses in

Name: _____

Title: _____

Company: _____

Address: _____

(City) _____ Telephone: _____

The Newsletter is a regular publication of I.P. Sharp Associates. Contributions and comments are welcome and should be addressed to: I.P. Sharp Newsletter, Box 418, Exchange Tower, 2 First Canadian Place, Toronto, Canada M5X 1E3.

Irene Shimoda, *Editor*
Mary Kopfensteiner, *Circulation*

Printed in Canada
May 1982
ISSN 0226 854X

International Offices

Aberdeen

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

Amsterdam

Intersystems B.V.
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

Bristol

I.P. Sharp Associates Limited
5 Whiteladies Rd., Clifton
Bristol BS8 1NN, England
(0272) 732 918

Brussels

I.P. Sharp Europe S.A.
Avenue du General
de Gaulle, 39
1050 Bruxelles
Belgium
(02) 649 99 77

Calgary

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

Canberra

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

Chicago

I.P. Sharp Associates, Inc.
Suite 1736
2 North Riverside Plaza
Chicago, Illinois 60606
(312) 648-1730

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, England CV1 1EA
(0203) 56562

Dallas

I.P. Sharp Associates, Inc.
Suite 1148, Campbell Center
8350 Northcentral 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
Leostrasse 62A
4000 Dusseldorf 11
West Germany
(0211) 57 50 16

Edmonton

I.P. Sharp Associates Limited
Suite 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

Hamilton

I.P. Sharp Associates Limited
14 Hess St. South, Hess Village
Hamilton, Ontario L8P 3M9
(416) 527-3801

Helsinki

TMT - Team Oy (Agent)
Kalevankatu 33 A
Mail to: P.O. Box 452
SF-00101 Helsinki 10, Finland
(9) 0-6946344

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

Kitchener/Waterloo

I.P. Sharp Associates Limited
3 Menno St.
Waterloo, Ont. N2L 2A4
(519) 884-5420

London, Canada

I.P. Sharp Associates Limited
Suite 510, 220 Dundas Street
London, Ontario N6A 1H3
(519) 434-2426

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., 13th Floor
Melbourne, Victoria
3000, Australia
(03) 614-1766

Mexico City

Telinformatica de Mexico S.A.
(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) 221612

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
Mail to: 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
(415) 327-1700

Paris

I.P. Sharp Sarl.
Tour Neptune, Cedex 20
20 Place de Seine
92086 Paris-la-defense
France
(1) 773 57 77

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
Telex: 0097 8473
0097 8474

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

San Jose

I.P. Sharp Associates, Inc.
3028A Scott Boulevard
Santa Clara, California 95050
(408) 727-9446

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

Singapore (Far East H.Q.)

I.P. Sharp Associates(S) Pte. Ltd.
Suite 1003, PIL Building
140 Cecil Street
Singapore 0106
Republic of Singapore
2210494
Telex: RS 20597 IPSAS

Singapore

Singapore International
Software Services Pte. Ltd.
Suite 1501, CPF Building
79 Robinson Rd.
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

Toronto

(International Headquarters)
I.P. Sharp Associates Limited
Box 418, Exchange Tower
2 First Canadian Place
Toronto, Ontario
M5X 1E3
(416) 364-5361
Telex: 0622259

Vancouver

I.P. Sharp Associates Limited
(Agent)
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
Rechte Wienzeile 5/3
A-1040 Wien, Austria
(0222) 57 65 71

Warrington

I.P. Sharp Associates Limited
Paul House
89-91 Buttermarket Street
Warrington, Cheshire
England WA1 2NL
(0925) 50413/4

Washington

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

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 A.G.
Fortunagasse 15
8001 Zurich, Switzerland
(01) 211 84 24

SHARP APL Communications Network

APL OPERATOR VOICE (416) 363-2051

COMMUNICATIONS (416) 363-1832

Local dial access is available in most of the locations listed above and in:

- Alliance • Ann Arbor • Austin • Baltimore • Birmingham • Buffalo • Clewiston (Fl) • Dayton • Des Moines • Des Plaines
- Detroit • Ft. Lauderdale • Geneva • Greenwich (Ct) • Hartford • Hull • Knoxville • Laurel • Leeds • Liverpool • Lyndhurst
- Minneapolis • Oxford • Quebec City • Raleigh • Red Deer • Regina • Santa Ana • Stuttgart • Sunnyvale • Syracuse
- Towanda • Ukiah

Our private, packet-switched network connects with the Value Added Networks in:

- Alaska • Argentina • Bahrain • Bermuda • Chile • Dominican Republic • Hawaii • Israel • Japan • Luxembourg • New Zealand
- The Philippines • Portugal • Puerto Rico • Taiwan

In the continental United States, the SHARP APL Network is interconnected with the Value Added Networks to provide access in 170 more cities, 40 more in Canada, and 5 more in the U.K. In all, with the 90 cities served by the I.P. Sharp Network listed above, SHARP APL is accessible from over 500 places via a local phone call. Please ask at your nearest I.P. Sharp office for a complete list of access points and access procedures. Our private network also connects with the worldwide Telex network via the Rochester, New York and Amsterdam nodes.