
Can / Am EMTP News

Voice of the Canadian / American EMTP User Group

Publishers and Mailers :

Drs. Kai - Hwa Ger and Tsu - huei Liu
3179 Oak Tree Court
West Linn, Oregon 97068
United States of America

Authorized by Co-chairmen :

Dr. W. Scott Meyer, Editor
Dr. Tsu - huei Liu
E - mail : [thliu @ bpa . gov](mailto:thliu@bpa.gov)
Vol. 98- 3; July, 1998

Table of Contents

Salford Compiler and DOS Extender	1
Improvements to Salford TPPLOT	3
More about Electronic Mail (E-mail)	5
European EMTP User Group (EEUG)	5
Watcom ATP for MS Windows	6
Line and Cable Constants	6
News About TACS and MODELS	7
Japanese ATP User Group (JAUG)	7
DEC ATP for VAX / Open VMS	8
Relay Testers to Read .PL4 Files	8
Higher - Order Pi Circuits	9
New Cascading of Line Sections	10
Brain - damaged MS Windows	10
Frequency Scans and Harmonics	10
Machinery Saturation Questioned, II	12
Corrections to This Newsletter	14
GNU ATP for Linux and DOS	14
Compiled TACS speeds simulation	14
Chen on saturable TRANSFORMER	14
Branch Data Input Restructured	15
Stu Cook Uses Apple Macintosh	15
Publishing Programs and Viewers	16
Changes to Universal ATP Graphics	17
Conditional Data Assembly	18
Power Company Politics and Religion	18
Interactive Plotting in MS Windows	19
Miscellaneous Small Items	19

Salford Compiler and DOS Extender

C-like input of the UTPF is used by the Salford translator. This reform began March 28th in an effort to

speed execution. Before documenting this success, consider a fringe benefit : **Tab** characters no longer are tolerated. The problem has been summarized before, beginning in the October, 1993, newsletter. **Tab** characters would go unnoticed in Portland because Salford would convert each to 8 blank bytes upon FORMATTED input. While this was nice and civilized, it resulted in delayed trouble using other compilers that were not so nice (e.g., Apollo at LEC during 1993). So, in order to maintain purity, it is a good idea to avoid the helpful Salford logic. The switch to C-like input does this. As a result, output FORTRAN for Salford will be just as illegal as it would be for Apollo in Leuven. The compiler will choke, and the program developer will know immediately that he has a problem (most often caused by MS-DOS EDIT for reasons that never have been explained).

Speed of the Salford translator **has** increased with the replacement of normal FORMATTED input by C-like input. Running as usual under Win95 DOS on your Editor's 133-MHz Pentium, elapsed times for complete translation of the 133191-line UTPF were timed as follows:

Old (formatted input) : 155 sec
New (C-like input) : 51 sec

Needless to say, this is considered to be an important gain as it typically will be realized every day or so. Even more important should be the larger gain at BPA using Dr. Liu's 486 / DX2. After parameters of Smartdrv were changed from the original /X (who knows what this meant?) to "2048 2048" (as used by your Editor for years), a full translation took only 66 seconds. This is to be compared with an original 5 minutes or so.

Interactive use of BOTH or DISK end in program termination (i.e., the DOS prompt). But why? The suggestion to provide the interactive user with a return to

the opening prompt of the program was made by Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. E-mail dated October 18th credited Prof. Peter Schegner of the University of Dresden with first observing an error in such behavior using Watcom ATP. Whereas Salford EMTP correctly ended in the DOS prompt, Watcom ATP returned to the opening prompt but in erroneous fashion. After fixing that error with Watcom ATP, the idea to return to the opening prompt rather than the DOS prompt was floated and accepted. The idea was excellent, but implementation was not as simple as it seemed upon first thought. After several hours of hard work, your Editor gave up.

The Diamond Stealth II S220 (Turbo) PCI video card recently acquired by BPA's Fred Elliott is believed to be DBOS-incompatible. In public E-mail dated April 15th, Mr. Elliott asked: *"Does it require any special setup?"* Five days later, Mr. Elliott provided details of his observed trouble: *"When I try to use TPPLLOT I get error messages such as the following General protection exception at User/7F911504 In routine GRAPHICS_BOARD_RECOGNISE@ I have three different sets of drivers for the card (two from Diamond and one from an independent developer) and have gotten the same results in each case. I tried to find information on either Diamond's or Rendition's web site that would enable me to use the DBOS custom configuration option without success. PCPLOT and WPCPLOT work correctly since they do not depend on DBOS."* Your Editor knows of no insight that might have been contributed by anyone else; the problem remains unresolved, as far as is known.

Long file names were mentioned in public E-mail from your Editor dated March 11th. Robert Meredith of the New York City area had noted the truncation of \$INCLUDE file names as they were echoed to the .LIS file. For example: *"--- Pass 1. Card = 40. Ready to open \$INCLUDE = d:\steurer\nst\emtp\atpdraw\u"* For ease of printing, lower case is used in this paragraph whereas directory and file names normally would involve capital letters (assuming KINSEN = 1). At the time, your Editor had claimed this was just a cosmetic problem of output, and should not affect execution. But later experimentation proved otherwise, so four segments of the UTPF were corrected accordingly on April 6th. A 32-byte limit was involved, but has since been increased to 80 bytes: variable CLIK32 became CLIK80. This has been tested using DC8.DAT within a sub-sub-sub-subdirectory, which was used in a \$PREFIX definition within the file. Consider output of the long path, beginning with the \$INCLUDE evaluation: *"--- Pass 1. Card = 16. Ready to open \$INCLUDE = \dumtest\test1\debug\subdir1\subdir2\dc8incl1.dat"* Line 2 of the .LIS file shows the .PL4 file name on the right. For 132-column output, this has been expanded from 14 to 35 bytes, and the least-significant part of any remainder will be omitted. So, this appears as follows: Name of disk plot file is test1\debug\subdir1\subdir2\dc8.pl4 For 80-column output, the limit of 14 bytes remains. Finally, there is the input data file name as

documented on the first line of interpretation : *"|C data:f:\dumtest\test1\debug\subdir1\subdir2\dc8.dat"* This does not approach the new 80-byte limit, but it does show that the old 32-byte limit is gone. If any reader has trouble with longer paths, he should send his summary to developers in Portland. About fundamental limits, Robert Schultz has offered the figure 255 for WinNT in a private message dated February 27th. What the figure for DOS of Win95 might be is not known. Yet, this all is a little academic as long as no user ever runs into trouble (the limit of 80 should suffice for a while).

Local 80-byte character variable ANSI32L is used to hold the name of the .PL4 file for purposes of connection by OPENRW@ in SYSDEP. This was following the reform of the preceding paragraph, which was modeled after earlier work on Watcom ATP by Robert Schultz and Robert Meredith of the New York City area. Curiously, the initial change worked perfectly for compilation, linking, and execution using Salford Ver. 2.66 software. But execution failed for BPA's Dr. Tsu-huei Liu using Version 3.5 software for all 3 phases. Whereas the old software seemed either to blank fill, or to ignore, unused bytes to the right of byte 32, the new software properly treated such bytes as part of the name, and complained that the file name was illegal. This was April 8th, when the logic of the preceding paragraph was improved to remove garbage from otherwise-undefined bytes of the character variable.

Fred Elliott of BPA telephoned to explain why **Ctrl-Alt-Del** does not lead to the new Win95 DOS on your Editor's computer. This was shortly after availability of the January newsletter, which posed the question. In E-mail dated April 8th, hard copy of the explanation was provided: *"On machines such as yours that are set up for dual boot (Win95 and a prior version of dos/windows) F4 is the shortcut key to boot to the prior version. On systems that are not dual boot, the F4 key does nothing during startup. What you should do instead is use the F8 key which brings up the boot menu and allows you to choose 'command prompt only' (menu item 6 on my machine) to boot to dos 7. Then you will get the dos prompt, new EDIT, etc. and no windows."* It is true. The **F8** key leads to *"startup menu"* of which choice number 7 is *"Safe mode command prompt only."* This is the desired new DOS, your Editor has confirmed. On the other hand, AUTOEXEC and CONFIG seem to be ignored, so behavior is not as expected. The first output after your Editor sends "7" is: *"Windows is bypassing your startup file."*

Program NOTABS now purifies the entire UTPF automatically. This is a continuation of the preceding paragraph. If **Tab** characters are noted, correction is required; and NOTABS is by far the easiest procedure if more than just one or two isolated **Tabs** are involved. NOTABS is a Salford program that scans the entire UTPF. If a **Tab** is found, a corrected .TAB file in parallel with the .SPL will be created, and such creation is reported to the screen. For example, as run on March 28th, it found 5

offending segments the first time it was used:

```
Tab found. 199). Byte = 667 SNSAVE =\UTPF\HEADFS.SPL
Tab found. 240). Byte = 1199 SNSAVE =\UTPF\GUTS28.SPL
Tab found. 246). Byte = 42 SNSAVE =\UTPF\NEWPXX.SPL
Tab found. 250). Byte = 29073 SNSAVE =\UTPF\LINEXX.SPL
Tab found. 256). Byte = 390 SNSAVE =\UTPF\GRIDXX.SPL
```

The first number of each line is the segment number, and the byte location in the file is for the first **Tab** only. There may be many more, and usually there are (see the July, 1996, newsletter). Using Mike Albert's freeware FC, any one resulting .TAB file should compare perfectly with the .SAL, but file size will be larger by 7 bytes for each **Tab** that was recognized and converted. Execution is very fast, too: NOTABS takes 12 seconds using your 133-MHz Pentium when no **Tab** actually requires conversion.

TSTALL of STARTUP was explained in a separate story later in this issue. Not surprisingly, implementation is installation-dependent, with Salford DBOS, which offers SLEEP@, used for initial development. Value 0.1 works well. It slows output enough so even the most uncoordinated of users will have time for CTRL-S to pause and study. Smaller values seem little different (DBOS ver. 2.66 used by your Editor seems unable to awaken much quicker, unfortunately).

A .PL4 file that is named using the date and time of day rather than the input data file has the potential problem of a name conflict. This assumes no file version numbers as used by DEC VMS or the original CDC. For operating systems that allow long names, resolution is limited to one second, which probably is adequate for most practical usage. But MS-DOS poses a much bigger problem. With file names limited to 8 bytes, resolution is no greater than one minute. Those who execute with KOMPARE = 0 no doubt have seen many messages such as "**** WARNING. The normal plot file name = F:84181754.pl4' already exists. One second is being added to the time by installation-dependent SUBROUTINE SYSDEP, and the check will now be repeated." About this, two observations might be made. First, only non-blank bytes of the file name now are shown (here within apostrophes rather than the actual quotation marks). This illustrates a new general capability that should be used in other places that involve file names: display through the final nonblank byte only, within a limit of 80. The capability is universal, and has been added to accommodate longer path and/or file names. Until someone complains, this is limited to 80 bytes. Second, the message mentions one second, but in fact one minute is used by Salford EMTP running under MS-DOS.

Old DC-9 data, which involves the no-longer-supported CASCaded PI feature (see preceding issue), was found to overflow the swap file of DBOS using ver. 3.5 compilation and linking. This was traced to variable NPHCAS having garbage (a very large positive number), which resulted in attempted zeroing of gigabytes upon exit (on the way to an error message). It is not known why ver. 3.5 linking does not automatically zero NPHCAS as ver. 2.66 linking does. Anyway, manual initialization was

added at the top of CASCAD on May 13th after the discovery.

Improvements to Salford TPPLOT

Bar charts that result when NOBARS = 0 (see preceding issue) are limited to a single variable. This is because screen plotting is 2-dimensional. There was no protection against a request for two or more, however, prior to the addition of it on April 30th. Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, had reported the problem, prompting the addition of protection. Command file FURST was modified at the same time to illustrate such usage.

Dummy variable names were seen in the CHOICE output that documented content of a widexx .PL4 file, even though the three header lines were present, and had the correct structure. This was for time simulation (not F-scan), as reported by HFS author Gabor Furst in E-mail dated April 28th. The trouble was this: TPPLOT had remained unchanged while details of the header lines created by ATP had been modified. In the past, the first line of the file carried *Second* whereas now it carries *Time* (followed by *in Second* in lines 2 and 3). Beginning May 1st, both the older and the newer file formats will be recognized.

Many more variables of a widexx .PL4 file needed to be handled. Another observation of Gabor Furst was that TPPLOT died while trying to handle files with many variables. Your Editor investigated using a file that had 132 variables, and found (and removed) several hidden limits. Describing this, he wrote: *"In one place, only 255 bytes were searched to find the right edge. I made this 2550. Later, the program died after failing to encode the number of variables (132) using two decimal digits. After that, I found a format that allowed a maximum of 20 columns, which I increased to 300."* So, the theoretical limits of 300 variables and 2550 bytes of total width finally might be fully usable, it is hoped.

Disk file names POSTxx.LIS and HPGLxx.LIS have been removed. Since the beginning, these were the default names that were offered to the user of POST or HPGL (subcommands of the PAPER command of the yellow menu). But WINDOW plotting was added later, and it was different. Rather than ask the user for a file name, it generated and used file names TPPPOST.xxx and TPPHPGL.xxx automatically if not prohibited by the NOPOST or NOHPGL controls. Whereas the old files never were overwritten, the new always were (xxx always began at 001). The two schemes were quite different philosophically, and program developers now decidedly favor the second. So, as part of a complete switch to WINDOW-like logic, file names and refusal to overwrite are being changed first. New output consists of a single line

that notifies the user of the name being used. Immediately following the POST command, the program will output a message such as: "Note: disk file TPPPOST.001 will hold PostScript copy now beginning."

Integer INDPAP having index 65 is a repackaged control for the use of POST (a subcommand within PAPER of the yellow menu) to produce PostScript output of non-WINDOW and non-bar (e.g., Fourier or FS or HFS) plots. For years, the choice of paper size came from the keyboard following a prompt such as: "Choose plot size in inches [1 = 8.5x11; 2 = 11x17; ...]" However, this was more a burden than a solution. Program developers never knew of a user who would change the response between plots. So, on May 5th, the decision about paper size was moved to TPPARAM.DAT, which controls the plotting environment. The default value will be unity, obviously.

While developers do not like to change TPPLLOT data requirements after years of established use, the preceding two modifications are believed to have minimal impact on the masses. If batch-mode data immediately following POST is not removed, two extraneous lines will simply be rejected as unrecognized input in response to the TPPLLOT command. That is for batch-mode use. For interactive use, the issue does not arise.

The CHOICE display of years past continues to be used for cases involving frequency scan (either FS or HFS) **provided** only magnitudes of the phasor solutions are involved. For either 2 or 4 output parts of each variable, a new output display has been implemented. The following illustration is for the two output parts of polar representation (magnitude and angle in degrees), for a file having 2 node voltages, 2 branch voltages, and 2 branch currents:

Type-4 entries (node voltages) :

GEN	LOAD
5 Mag	7 Mag
6 Angle	8 Angle

Type-8 entries (branch voltages).

SWIT	LOAD	LOAD
1 Magnitude	3 Magnitude	
2 Angle in deg.	4 Angle in deg.	

Type-9 entries (branch currents).

GEN	SWIT	SWIT	LOAD
9 Magnitude	11 Magnitude		
10 Angle in deg.	12 Angle in deg.		

Note that variable numbers (for the # command) now are below the associated variable name or names, rather than on the same line. Otherwise, operation is similar, and should require no special explanation. The new CHOICE display became effective May 19th.

The new CHOICE display just described should always be seen if C-like .PL4 files are being used. But for lower-performance alternatives that include FORMATTED (if FMTPL4 is nonblank), it may or may not be seen. The difference has to do with a possible (recommended) preceding TIMESPAN command. In the absence of such a command, the old CHOICE display will be seen. This

is no real problem; it is just as usable. However, the nature of the output variables is not explicitly seen (instead, as in the .LIS file, each entry will be repeated 2 or 4 times).

MWINDO is the parameter that, if positive, requests automatic, uniform allocation of windows for WINDOW plotting. Background is provided in the October, 1993, newsletter, and JULES1 through JULES6 illustrate usage for one through 6 windows. But what about 7 or more? The first person to report unsuccessful use (a crash of TPPLLOT, in fact) was Bruno Ceresoli of ENEL Research outside Milano, Italy. So, your Editor plugged the leak on May 25th. At the same time, more than MWINDO curves will be allowed. Originally, the number of curves was reduced to the number of windows. But since then, a second screen (2 pages for the screen plot) was allowed as explained in the July, 1996, issue. Now, operation is as follows. If the user attempts to plot with MWINDO equal to 7 or more, it will be reduced to 5. The number of curves is unchanged. So, for example, suppose the user has selected 9 curves and had set MWINDO = 9. There will be a warning message about the reduction ("MWINDO is too big. Reduce it to reasonable 5"), then one screen filled with the first 5 curves and another (after the user sends <CR>) that uses the first 4 of its 5 available locations.

No grid was to be seen on the second page of a window plot that required more than one screen to display all curves. This was the observation of Mr. Ceresoli in E-mail dated June 1st, following the previous, more-serious correction. Grids were restored later that same day, with the change applying to cases of manual as well as automatic allocation of windows. At the same time, the limit MWINDO on automatic allocation was raised from 6 to 7. Readability using VGA is marginal, of course, but this is relative. Seven windows look worse than six by about the same amount that six look worse than five; etc. Who would like to see eight (joke)? If publishing programs can (and do) display pages so small they are unreadable, we can display vector graphics with less resolution than printer plots!

Zero JDENSE is a flag indicating that the user wants to vary the number of curves in each window. This was a later suggestion from Mr. Ceresoli, after learning about positive JDENSE usage. Quoting from his E-mail dated June 2nd : *"I will give you another suggestion, based on Randy Suhrbier's VMS program PLOT. What about a syntax such as 2,1,2 ... which would indicate three windows, with 2 curves in the first, one in the second, and two curves in the third (it is assumed that 5 variables had previously been selected)? This should be more flexible, and yet not too complicated"* Following the plot command (e.g., GO), the program will prompt for the number of curves in each window: "SEND # OF CURVES IN EACH WINDOW [<CR> to abort] :" If the input is blank, JDENSE will be set to unity. But if the input line is non-blank, free-format integers will be extracted. In order, these specify the number of curves in each window. If there are fewer integers than windows, the final integer will be assumed to

apply to all remaining (higher-numbered) windows. Work was completed June 3rd, and the following day BRUNO was added to the TPPLOT archive to illustrate usage for 8 curves of HHMMSS spread over 6 windows on 2 pages.

Labeling all curves in each window has proven to be difficult for cases involving two or more curves. Your Editor decided the simplest procedure was to use the old label for the first curve, and then leave any additional curves of the window unlabeled. Dr. Liu's first reaction to the display of BRUNO was to suggest that there is room for a second label, to the right of the one used for the first curve (e.g., "V-bran BUS2 BUS3"). True, but this is only because there is no FACTOR or OFFSET usage, which would place two numbers to the right of the curve names. Rather than change the existing display, the second and later curves in any window were left unlabeled.

More about Electronic Mail (E-mail)

Robert Meredith of the New York City area is to be thanked for managing Agora's E-mail for your Editor following BPA's evacuation of Skyport on January 28th. On February 24th, he arranged for Agora to forward mail to BPA, and de-subscribed from the Fargo list server. Shortly thereafter, he forwarded all mail that already had accumulated. During the final half of March, Agora mail had slowed to a trickle (maybe an average of one message per week) so loss of the service on April 1st seemed not to inconvenience many ATP users.

About 320 million Web pages now exist, and this number is expected to increase ten fold *"in the next few years, according to Steve Lawrence of the NEC Research Institute."* This in an AP story by Paul Recer that was printed on page B1 of *The Oregonian* dated April 3rd. There is mention of a story in the journal Science having this same date. *"The amount of information on the Web overwhelms even the most sophisticated efforts to sort it all out and there may be huge numbers of pages existing in an electronic shadowland never seen by humans."*

Statistics about Internet use, and what they mean to business, are difficult to agree upon. A *Wall Street Journal* story, found on page G1 of the April 19th issue of *The Oregonian*, makes this point. For example, 62 million Americans are estimated to use the Internet, and of these, *"almost 25 percent were Internet newcomers in 1997."* This is impressive, but it remains small compared with television. *"Online use reaches 23 percent of homes up from 17 percent a year earlier."* For advertisers, this is nothing compared with television, which reaches 98% of homes. Also, business has not yet learned how to advertise effectively using the Internet. Nearly every business is there, but the use is not obviously profitable. It is quite unlike television advertising. About searching, *"Web users spend an awful lot of time looking for places to go ... Among the 25 most frequently visited outposts on the Web, a staggering nine sites ... are search engines*

or directories Yahoo!, the No. 1 destination, drew about 32 million people in March." Business is attracted by the disposable income of Internet users, however: *"About 43 percent of adult Web users have a college degree or higher ... The average Web user makes about \$55,000 a year That's more than double the average U.S. individual income ..."*

AOL was used effectively by Prof. Mustafa Kizilcay during his recent visit to Japan. Your Editor thought he was sending mail to Prof. Kizilcay in Germany, but in fact the message was received in Japan! Dated May 25th, E-mail from the EEUG Chairman explained: *"Thank you for your two messages, which I received in Kyoto via AOL. In almost every big city of Japan, AOL is accessible via local call. So, I could download your two Emails without any problem. Because this information is important, I printed it ... (I brought a notebook printer, Canon BJC-80, which communicates through infrared interface with the PC. I do not need to carry a parallel port cable.)"* As for purpose of the trip: *"The EMTP workshop and annual Japanese EMTP user group meeting will be held this coming Thursday and Friday in Nagoya, respectively."*

Priceline.com has been advertising on Portland radio stations as follows: tell us when and where you want to travel, and how much you are willing to pay. So, your Editor connected to www.priceline.com Clicking on the *"our company"* button of the home page, he found an official summary of the revolutionary idea: *"Priceline.com is the world's first buying service through which consumers name the price they're willing to pay. priceline finds sellers willing to meet buyers' needs and price, thus facilitating thousands of buyer-driven transactions each day. In April 1998 priceline debuted with its service enabling consumers to name their own price for leisure airline tickets. In the summer, it will expand its buyer-driven concept to offer new cars, and in the fall, consumers will be able to name their own price for a number of financial services."* At the bottom of the page, one learns the following: *"Headquarters: Stamford, Connecticut. Employees: 100+ Ownership: priceline.com is privately held by individual and institutional investors."* So, what possible disadvantage is there, for the consumer? Plenty, it would seem, if he does not know real value. Just as at any auction, the customer must honor his bid. I.e., if priceline finds the requested ticket, the customer must make the purchase. Nonetheless, the concept is intriguing. Might priceline provide real competition for travel agencies? It seems to be another activity that is ideally suited for the Internet. Obviously, both computers and speed of communication are essential ingredients.

European EMTP User Group (EEUG)

"EEUG Meeting 1998 and ATP Course" was the Subject of public E-mail of the Fargo list server dated May 22nd. The annual meeting for 1998 is to be held on

November 9th and 10th, and a course on frequency scans is to be held the follow day, on November 11th.

Prague in The Czech Republic is to be the site, with Dr. Tomas Hasman of the Faculty of Electrical Engineering of Czech Technical University heading the local organizers. Internet address **hasman@feld.cvut.cz**, was given for this first of two names at the bottom of the announcement. The second name was that of Prof. Mustafa Kizilcay of FH Osnabrueck in Germany, the EEUG Chairman.

Elections make the upcoming meeting exceptional: "*The Members' Meeting makes up the second part of the EEUG Meeting. Four years after foundation of the EEUG Association, the term of office of the present Executive Board will expire at the end of 1998. A new Executive Board consisting of six members will be elected by the EEUG members at the Members' Meeting in Prague.*"

About technical presentations: "*The main part of the meeting is allocated for the technical sessions Two types of contributions may be submitted for presentation at the technical sessions: full meeting papers (20 minutes presentation) and short contributions (10 minutes presentation). A single page abstract of your technical paper should be submitted before September 4, 1998 by normal mail, fax or E-mail to Prof. Dr. Mustafa Kizilcay, Chairman of the EEUG. The language of the meeting is English.*"

"Frequency-Domain Harmonic Analysis using ATP" is the title of the one-day short course this year. Teachers are expected to be Gabor Furst (see separate story about HFS) and Prof. Kizilcay. About content, the announcement states: "*In this course, the new feature HARMONIC FREQUENCY SCAN of ATP will be presented, and different applications of FREQUENCY SCAN will be demonstrated.*"

Further information including registration forms for the meeting and course are available from the Web page of EEUG, which has address www.vmt.bme.hu/eeug

Watcom ATP for MS Windows

Robert Meredith of the New York City area visited BPA on Sunday May 10th in order to provide advice about Windows NT and the Watcom compiler that is supported using this environment. He arrived at the back door of Dittmer around 10:30, following a 3-hour drive down from Seattle in a rental car; and he did not leave until 18:30. Like colleague Robert Schultz nearly 3 years earlier (see the July, 1995, issue), Mr. Meredith provided exceptional support to program development at BPA, and his efforts were greatly appreciated. For the record, the NYPA business trip to Seattle had nothing to do with electromagnetic transients. Seattle is simply the location of Cegelec ESCA, which offers classes associated with its

software (widely used to monitor and control power systems). EMTP-like programs are not yet involved there.

Noda frequency dependence for Watcom ATP was handled by Robert Schultz of the New York City area. As explained in the October, 1995, issue, Noda modeling differs from all other ATP modeling in that it was coded in C language rather than FORTRAN by Dr. Taku Noda. Although execution is correct today, there was trouble initially. At first, execution of Watcom ATP using DCNEW-14 failed in NODA3A during data input. About the language interface, Schultz wrote as follows on May 2nd, following correction of the problem: "*Watcom string handling from C to Fortran is complex -- a Fortran string descriptor structure is needed. Going from Fortran to C is trivial.*" Stated in general terms, code in C that is workable with one compiler (e.g., Salford) generally requires changes in order to work properly with another (e.g., Watcom). The interface between the two (any two?) languages has not yet been standardized, unfortunately.

Meredith's multiphase saturable transformer model is mentioned from time to time, so status of this work is worthy of a summary clarification. The model is years old, and execution finally is illustrated by a standard test case (DCNEW-25). Furthermore, documentation of use already exists in the form of a 390-Kbyte, 114-page, PDF-format document entitled "*EMTP modeling of core-form transformers by magnetic circuit analysis, including finite sectioning.*" But this has not yet been released to more than a few insiders. During his visit to Dittmer on May 10th, Mr. Meredith explained that he wants the feature to be usable by the general public before he releases writing about it. But usability by others will require work on Meredith's related supporting programs, which generate ATP data. According to current thinking, Meredith's supporting programs will be added to ATP using some yet-to-be determined request word or words. Your Editor volunteered his services for the job, of course.

Line and Cable Constants

CABLE CONSTANTS and CABLE PARAMETERS of ATP are the contribution of Prof. Akihiro Ametani of Doshisha University in Kyoto, Japan. Recently, Prof. Ametani has been promoted to Dean of the Library and Computer Information Center.

CABLE PARAMETERS has never been connected to SEMLYEN SETUP. Whether it ever will be seems unlikely. In any case, the user who attempts such impossible use will be treated better than in years past. Beginning June 9th, any such attempt should result in a clear KILL = 60 error message. After referring the user to CABLE CONSTANTS, it adds a warning: "*Semlyen modeling of ATP involves two serious approximations for cables: 1) constant transformation matrix [T]; 2) only 2 exponentials.*" About timing, it might be explained that

S. Navaneethan at the University of Strathclyde in Glasgow, Scotland, reported trouble in public E-mail dated June 3rd. What he was subjected to was KILL = 182, which complained about an incorrect number of frequency cards. First Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, and then BPA's Dr. Tsu-huei Liu, explained that the message was misleading. Six days later, the error termination was corrected.

News About TACS and MODELS

That Type-68 TACS device from Robert Meredith (see preceding issue) had its input data interpretation changed on May 5th. The author suggested "hysteresis element" as a replacement for "transformer model" that had first been used by your Editor. This can be seen in DCNEW-25.

Supplemental variables with non-blank column 80 were mishandled by all program versions except GNU ATP until May 9th. Then, upon the recommendation of Masahiro Kan of Toshiba Corporation in Japan, a patch was added to the UTPF. This had been tested only in the GNU ATP translation that Mr. Kan maintains and distributes to others. The story began in Japan, with Dr. Yuan Bin reporting trouble using data that had been accepted by DCG / EPRI EMTP. Your Editor's trial correction dates to February 4th (the *.MEL files).

An indexing limit of 10K was removed from TACS on July 1st in order to handle data from Humberto Henao, an Assistant Professor at the University of Picardie in Amiens, France. Recall this is the school to which Gerard Capolino was promoted, as summarized in the July, 1995, issue. However, communication was with the University of Oviedo in Gijon, Spain, where Prof. Henao currently is "on leave." While your Editor is not authorized to describe the engineering involved, he can observe that the data is voluminous. The input data file consists of nearly 12K lines, and few of these are comments. This has to be the mother of all TACS data cases --- by far the biggest your Editor ever has seen. It may well be the first TACS data to reach the limit of 10K, which was increased to 100K (should hold us for a while longer, anyway).

Japanese ATP User Group (JAUG)

BPA's Dr. Tsu-huei Liu addressed some 65 delegates at a workshop of the Japanese EMTP Committee (JEC) in Tokyo, Japan, on March 13th. The presentation and meeting were exceptional for two reasons. First, like EEUG last November (see the January issue), JEC reimbursed BPA for travel and subsistence expenses of the trip. Second, a special nondisclosure and non-commercialization agreement was signed by anyone in attendance who was not already

ATP-licensed. As in Barcelona last November, ongoing ATP development was discussed and summarized.

"*DSPATP -- plotting of ATP PL4 file*" is the title of an 8-page paper by Masahiro Kan of Toshiba Corp. The date in the upper-right corner is 13 March 1998, and *DSP* probably is short for *display*. Point 2 explains that "*Dspatp is copyrighted by Masahiro Kan and can be used by ATP licensed users*" free of charge. Point 1a indicates compatibility with both C-like .PL4 files produced by others and also UNFORMATTED .PL4 files produced by GNU ATP. Continuing: "(b) Postscript output is supported. The generated .ps file can be displayed and printed by Ghostscript. (c) Dspatp32.exe (32-bit version) works under Win95/NT." Point 4 explains that "this program uses Gnuplot for displaying waveforms; so Gnuplot should be installed." Later pages include screen displays that illustrate many plots.

Japanese continue to provide and enjoy the best of ATP-related Internet services. As of March 13th, there had been some 6000 visits to the "Home Page of ATP-EMTP for Japanese and East Asians," some 2000 accesses to the "ATP WWW Board (Japanese)," and some 600 accesses to the "ATP WWW Board (English)."

"*New ATP textbook from Japan*" was the Subject of public E-mail dated May 25th. Writing from Osaka University, Dr. Tsuyoshi Funaki, the Secretary of JAUG, summarized yet another Japanese contribution to ATP users of the world. This was "*a new primer handbook of EMTP and power systems. The textbook, entitled 'Highly Sophisticated Electrical Power Systems,' was written by Dr. Eiichi Haganomori while he was at the Tokyo Institute of Technology as an Invited Professor (he has since returned to Toshiba Corporation). It treats essential topics of power engineering, and provides sample EMTP data files for each of them. ... You can obtain a free copy of the textbook from the following URL: <http://atp.pwr.eng.osaka-u.ac.jp/~hagi/index.html>*" Also, <http://atp.pwr.eng.osaka-u.ac.jp/~hagi/hseps.zip> is a second location being used, according to private E-mail from Masahiro Kan dated July 22nd. "It was called from the preceding URL."

Dr. Liu had the pleasure of participating in a meeting with the board members of the newly-formed Japanese ATP User Group (JAUG). The complete list of Japanese user group management follows:

Chairman: Dr. Hiroshi Arita of Hitachi Ltd.
Vice-Chairman: Mr. Masahiro Kan of Toshiba Corp.
Secretary: Dr. Tsuyoshi Funaki of Osaka University
Secretary: Dr. Yasuhide Kashiwagi of Kisarazu National College of Technology
Advisors: Dr. Akihiro Ametani, Dean of Library and Computer Information Center of Doshisha University; and Prof. Naoto Nagaoka of Doshisha University
The official approval process occurred later, on May 29th, at the annual JEC general meeting, which was held in

Nagoya during May. EEUG Chairman Mustafa Kizilcay attended this meeting, and wrote a report that could be used for further description next time.

Envoy files of the EMTP Theory Book were put on the password-protected FTP site at Osaka University in Japan on April 6, 1998. Quickly, these were replaced by a better alternative -- Adobe PDF -- as described in some detail in a separate story.

DEC ATP for VAX / Open VMS

VAX/VMS is the original operating system of DEC VAX computers, and Open VMS is the corresponding operating system for DEC Alpha, the high-speed RISC workstations.

Use of BN in FORMAT statements was first proposed by Steven Boroczyk of TransGrid in Sydney, Australia, as explained in the January, 1997, newsletter. This was to protect the DEC VMS ATP from a change of compiler or operating system. With BN, formats would tolerate ATP data of I-field or E-field nature that might not be right-adjusted as they should be. From the beginning, the Rule Book has warned about the need (see Section I-H). However, as the years have passed, more and more compilers or operating systems seem to be ignoring blanks on the right of all numeric data fields, so data that is less than universal goes undetected (and seems to have accumulated in the case of TransGrid). Then, if and when that data ever might be used with a less-tolerant system (even from the same manufacturer, in the case of TransGrid), a major problem results. Well, finally, the numerous changes were made by hand using EDIT. The archive ALL985.ZIP is used to store the old files in case they later might be needed for some reason. In all, 33 disk files are involved, and these total 2530 Kbytes following modification. Work was performed off and on April 3rd and 4th. Output was ignored, as was Laurent Dube's MODELS --- the latter because all of its input is believed to be free-format. Input from a FORMATTED .PL4 file also remains unchanged. Purely alphanumeric use was ignored, as were FORMATS that decoded nothing more than a type code from columns 1-2 (using I2, of course). Finally, I1 was ignored for this consideration, since an extra blank is not possible if the field width is one.

Higher - Order Pi Circuits

This is a continuation of the story about high-order Pi circuits in the preceding issue. An outgrowth and generalization, \$INSERT, is treated as well.

Special utilities MAKEPI, MAKEFORM, and MAKEC --- all .FTN files --- were written to test both the two new HOPC alternatives and also the old Type-1,2,3... data input

of card images. These programs make the ATP data files to represent a regular, full matrix of the appropriate type, for any order. Although they have no physical meaning, for the first time, really high order Pi circuits have been created, and this data has been successfully read and stored by ATP. Recall that nothing more than order 50 was possible previously because type code 51 began another branch type (coupled R-L). Three sizes were tested. In Kbytes:

Number of coils	100	200	400
<hr/>			
Old card image size:	384	1528	6095
New formatted size :	312	1239	4937
New C-like size :	122	485	1930

Differences of speed for data input of these 3 files are informative. The following ATP execution times in wall-clock seconds were obtained using your Editor's 16-Mbyte Pentium and Win95 DOS (not a DOS window of Win95, which would be shorter of memory). Execution time is total elapsed time to start the program, input the Pi-circuit, and then halt in response to \$STOP. The best 5 of 6 repetitive trials were averaged to produce the times shown:

Number of coils	100	200	400
<hr/>			
Old card image time:	5.82	12.56	43.08
New formatted time:	3.56	4.67	9.26
New C-like time :	3.19	3.30	3.69

If one is dealing with order 100 or smaller, probably the difference between the old 5.82 and the new 3.19 would not be much of an issue. But note how this changes for order 400: 43.08 vs. 3.69. This difference is hard to ignore.

Expansion of the .LIS file represents a portion of the just-shown 43 seconds, it must be acknowledged. The figure shown includes interpretation of those 80K input data lines. The .LIS file was enormous: 10188 Kbytes. Yet, this extra output, which is avoided by HOPC, is not the dominant cause of the slowdown. To demonstrate this, \$LISTOFF was used. Then the 43.08 seconds decreased to 36.44 (and 12.56 decreased to 10.89, and 5.82 decreased to 5.28). This confirms what has been known for many years: Salford DBOS performs I/O very efficiently.

Flushing of the output file from RAM to disk was added to SUBR29 on February 27th. This becomes an issue for large data sets that might fail for some reason (e.g., overflow of the input card buffer, which has limit LIMCRD) prior to there being any program output. The general principle is not new: failure during \$INCLUDE processing or /-card sorting is particularly confusing because it occurs prior to normal output. In the past, such problems would be solved easily enough by running such output to the screen. Lack of a disk file did not matter, since one could read any error messages from the screen. But what if the data file is too big to be shown on the screen? This was the new problem encountered while testing HOPC. Following improvement, a small .LIS file will appear when LIMCRD is overflowed (this was an actual case of practical interest).

\$INSERT is an alternate \$INCLUDE that offers both limitations and possible advantages. Whereas \$INCLUDE

is evaluated at the start of execution, \$INSERT is ignored until the program actually needs input from the associated disk file. For this reason, \$INSERT does not add to the burden on the card cache (having size LIMCRD, recall). Neither is it involved in "/"-card sorting, then. For cases that have no arguments, and no "/"-cards, \$INSERT offers an equivalent, dynamic connection. For large files, there can be a substantial saving as a result. This is the connection to HOPC: inspiration. It was while working on HOPC that the idea for \$INSERT occurred to your Editor. Unlike HOPC, \$INSERT is not limited to Pi-circuits. On the other hand, it is limited for formatted data.

The best example of hoped-for savings using \$INSERT might be James Randall's assembly of data for use with DEC VMS at BPA. There may be a lot of Marti branch data that could be handled by \$INSERT rather than \$INCLUDE, thereby reducing the VMS-expensive burden on LIMCRD. If James Randall last used 100K for LIMCRD, there is a lot to be saved: 8 Mbytes of virtual address space. The installation-dependence of \$INSERT is minimal: only the familiar Salford's STATUS = 'READONLY' requires change for some other compilers. To illustrate operation, the 2nd subcase of DC-4 was modified on March 27th. New disk file DC4BINCL.DAT contains the nonlinear element that previously was contained in the main file.

List 26 is the square of the number of coupled phases, and for years it has been limited to 3000, which allowed 54 conductors. Such storage within disk file LISTSIZE.BPA was increased to 10K on April 20th in order to process the first real (as opposed to artificial) high-order data, which was supplied by Dr. Hiroshi Arita of Hitachi. This data had order 95. Ability to handle 100 coupled phases now has become standard. But 400? Maybe later, but not yet. For now, special dimensioning is required for anyone wanting more than 100. Of course, Dr. Arita was informed that a special version allowing whatever limits he wanted would be created and sent to him upon request. This was in E-mail dated April 23rd.

Relay Testers to Read .PL4 Files

Details of .PL4 file structure are being made available to ATP-licensed companies that write software to evaluate relay performance. This decision, following public discussion using E-mail, was reached during June. There should be greater convenience for ATP users of such software, who no longer will be forced to use either the alternative widexx (which is both numerically inefficient and wasteful of disk space) or IEEE PES COMTRADE (which requires conversion from the .PL4 file using a program such as TPPLOT).

Robert Hoschek of Omicron Electronics in Altach, Austria, is the person who initiated the modern

consideration of the subject. His public E-mail of the Fargo list server dated June 10th asked: *"Is there a document about the C-like PL4 file format available?"* Later that same day, Prof. Laszlo Prikler of T.U. Budapest in Hungary explained that *"it is not available for the general public. The PL4 format is a kind of ATP secret. ... No one has the general right to share it with other users."*

"Relay tester to read C-like .PL4 files?" was the subject of public E-mail from your Editor late on June 19th. It explained that, *"unless some subscriber can think of a good reason not to allow this, the company should be provided with details (disk file PUGET.LIS) on Tuesday, June 23. The interim has been devoted to public discussion by any subscriber who might have concerns."*

What was new about the latest request? *"Comparable software seems to be offered by Doble, which has been mentioned in several newsletters. ... Yet, there is a fundamental difference that readers should understand. Doble will not read C-like .PL4 files, as far as this writer knows. Doble reads COMTRADE, and this is public (via IEEE PES) rather than secret. For the first time, it is proposed that we provide the detailed specifications for .PL4 files to such companies. Although the request came from one European company, obviously any competitor would be treated comparably provided it is ATP-licensed and in good standing with developers. The information would remain secret, but that secret would be shared with makers of Doble-like software who wanted to read .PL4 files directly."*

The supplier's software would read .PL4 files, but would **not** create them. This is an important point. Specifically, the following was stated: *"We do not intend to offer any information about the PL4 format to the customers or anybody else outside our software development group. We shall not ship anything which is part of the ATP information material or ATP licensed programs. We do not provide any tool to the customer to let him create PL4 files, it is up to the user to have them or get them somewhere."*

About BINARY COMTRADE, your Editor asked Mr. Hoschek: *"Do you suppose your binary is compatible with Doble's binary? As explained in the newsletter, our BINARY output of TPPLOT is not necessarily right, but it does now seem to be compatible with Doble."* Has any reader driven other (non-Doble) software with a BINARY COMTRADE file created by TPPLOT? If so, what happened? Your Editor has not heard of any such usage, which remains of interest. Some incompatibility seems inevitable. Remember that request from Dallas-based Multi-amp in the October, 1995, newsletter?

Ajay Koliwad of Powertec Industries provided the only public discussion from the industry in question. His message dated June 29th began with the opinion that *"it would be very beneficial to people involved in protective*

relay transients to be able to read the PL4 files directly into some kind of playback system. Our company had approached WSM about this, a few years ago. We eventually used the LIS output, and later the COMTRADE output from TPPLOT. The disadvantage here is that both options create ASCII files, which are unwieldy in size." In response, your Editor explained that newer BINARY COMTRADE, unlike the original ASCII COMTRADE, certainly has no problem with size : "This is because COMTRADE BINARY uses 2-byte integers whereas .PL4 files use 4-byte reals. If file size is most important, stick with COMTRADE."

But apparently convenience, rather than file size, is the dominant consideration. Users want to avoid the separate step of conversion to BINARY COMTRADE using TPPLOT. Use of C-like files is a compromise. The C-like alternative is more efficient and compact than widexx while being less compact than BINARY COMTRADE. For big files, two factors of two approximately separate the three common alternatives. As a general rule, widexx requires about 8 bytes per number whereas C-like requires exactly 4 bytes and BINARY COMTRADE exactly 2 bytes.

New Cascading of Line Sections

The REPETITION loop was extended to include transposition in order to satisfy academics. Practically, transposition is expensive, so is used sparingly (e.g., not at all in the east where distances are shorter than out here in the western USA). But academics such as Profs. Alexander or Hevia want to be able easily to illustrate how frequency of transposition affects transients. Prior to the extension of April 25th, each transposition required a separate request card. The REPETITION loop extended back over any possible preceding lumped elements, but not any possible preceding TRANSPOSITION card. By adding the qualifier INCLUDE TRANSPOSITION IN LOOP, this shortcoming is rectified. There now are four different interpretations of a REPETITION request, then. The line always is present, but transposition may or may not be involved, and lumped elements may or may not be involved in the loop.

DC-9 was the old CASCADED PI illustration, and this has changed. Before this old data will be found at the end of the new file, 25 illustrations of the replacement CASCADE LINE will be noted. Familiar among these is the 22nd, which provides an alternative, newer, more-compact way to specify the same data as used by DC-3. Those 17 copies of the initial Pi-circuit are requested using REPETITION requests with multiples 5, 6, and 6, in order. About the 26th and final subcase, this was slightly modified so as better to illustrate the power of the new modeling. The final Pi-circuit was noted to be very nearly double the length of preceding ones. Perhaps this was the intention, since most numbers are exactly double. A few are a little different. Anyway, the final segment

was replaced by two single-length sections, and this permitted illustration of the new request INCLUDE TRANSPOSITION IN LOOP, which modifies the normal REPETITION request. Comment cards describe details.

Brain - Damaged MS Windows

"Oops! Windows 98 crashes in high-profile demo" is the headline of an AP story that was carried on page 1C of *The Denver Post* dated April 21st. BPA's Robert Hasibar is thanked for noting this during his work away from home. The story is more irreverent than the report noted by this writer in *The Oregonian*. High profile is right: Bill Gates himself, at spring COMDEX (Computer Dealers Expo) in Chicago was involved. *"The system, scheduled to be released in June and retailing for about \$89, collapsed when a Microsoft employee attempted to plug in a scanner with his boss standing alongside. I guess we still have some bugs to work out,' he noted ruefully. That must be why we're not shipping Windows 98 yet.' He said Windows 98 will be far simpler to use than the popular Windows 95 version, which has 150 million copies in use."*

Frequency Scans and Harmonics

Gabor Furst revealed to the general public the existence of HARMONIC FREQUENCY SCAN (HFS) using public E-mail of the Fargo list server. This was just prior to release of the April newsletter, in a message dated April 19th. The announcement had "Subject: HFS - Harmonic Frequency Scan." Prior to a detailed summary, author Furst provided the following introduction: "Harmonic Frequency Scan (HFS) is the new ATP Request for performing harmonic analysis. HFS is a companion procedure to Frequency Scan (FS). FS performs a series of phasor solutions using user specified 14 type sources, by incrementing the source frequencies in a linear or logarithmic fashion. The amplitude of the sources is kept constant. HFS on the other hand, performs a set of phasor solutions on sources entered by the user, but does not change either the amplitude or angle of these sources, nor does it increment their frequencies. HFS therefore is a procedure which is identical with the methods used in all commercial harmonic analysis software known to the writer. This option is available in TPBIG versions Power quality has become a major preoccupation of transmission planners in the last decade and its importance is increasing with the proliferation of power electronics applications in the various industries and in the homes. Harmonic analysis is perhaps the most important branch of power quality studies. Harmonic analysis with ATP has previously been performed in the time domain. While this is a perfectly legitimate approach and in some ways even superior to the generally used approach for harmonic analysis, it has a number of drawbacks. These are discussed in the HFS Report to be placed in the Fargo

Server archives shortly. HFS features will be described in the Server report referred to above. The following is a brief summary ..."

Prof. Mustafa Kizilcay, the EEUG Chairman, was one of a handful of individuals who had been kept informed about HFS development prior to the public announcement. Prof. Kizilcay had the idea of a one-day course on the subject of F-scans (FS as well as HFS) in conjunction with the upcoming annual meeting in Prague, The Czech Republic, this coming November. See the EEUG story.

Gabor Furst's later detailed writing about the use of HFS was announced by Prof. Bruce Mork in public E-mail dated April 28th: "A documentation file, 2 sample data cases, and a utility that generates 6- and 12-pulse rectifier harmonic sources is included."

Originally, node voltages were the only possible outputs for FREQUENCY SCAN (FS) and HFS data, and these were made into four currents for plotting purposes. The second of the two associated 6-character names was used to indicate which component of the complex variable was involved: 1) mag for magnitude; 2) angle for the angle in degrees; 3) real for the real part; and 4) imag for the imaginary part. Thus 4 real branch currents were used to represent one complex node voltage. In FS, this was used for the past two decades or so --- until reform finally arrived on March 16th. Today, arbitrary outputs are allowed: node voltages, branch voltages, and branch currents, with each having its natural variable type. Fundamental to the reform was a switch from the special, fixed FS printout to printout that is modeled after that of the time-step loop. But before explaining the new printout, let's briefly illustrate the old format. From DC-51:

```
mag:LOAD      angle:LOAD      mag:MID      ...
real:LOAD      imag:LOAD      real:MID      ...
--- Output vector for step number 1.  Freq ...
1.0000000E+02  0.0000000E+00  9.3571522E-01 ...
1.0000000E+02  0.0000000E+00  8.7556297E-01 ...
```

Here the precision was fixed, and use of E-field numbers was mandatory. Finally, there was forced separation of all polar variables (magnitudes and angles, which came first) from all rectangular variables (real and imaginary parts). If MODELS variables were involved, they provided a third class, with each class beginning on a new line.

Magnitude-only is the new default output choice for both FS and HFS data. As long as the average user is willing to ignore phase, it is both simpler conceptually, and more efficient computationally, to suppress output of the other parts of the complex solution (i.e., angle, real part, and imaginary part). Printed and plotted output variables then have the same familiar form as for time simulation (the most common use). Time in seconds is replaced by frequency in Hertz, of course. Otherwise, output has the same form, so requires no special explanation. To illustrate, consider the left side of the beginning of frequency-scan loop printout for the 6th subcase of DCNEW-21:

```
Column headings for the 4 output variables ...
First 3 output variables are electric-n ...
```

```
Next 1 output variables are branch cur ...
Only the magnitude of each variable is output ...
Step F [Hz] LOAD GEN LOAD LOAD
          GEN          TERRA
1      50.  .846733  1.0  .532018  .16935
2     100.  .727727  1.0  .685867  .145545
3     200.  .846733  1.0  .532018  .16935
```

Only the fourth line is all-new, and peculiar to F-scan. It explains that the default choice of magnitude-only is in effect. Yet, it really is not needed. It should be immediately obvious to any user that this is the case. Compare the preceding with following paragraphs.

BOTH POLAR AND RECTANGULAR is the request for all 4 parts of each complex variable. This corresponds to the mandatory output of the past two decades or so. However, the printout looks very different, and the .PL4 file has different structure. Take the illustration of the preceding paragraph. The first two columns (for step number and frequency) would remain unchanged. But thereafter, each column would be replaced by four columns. The names at the top of each of the 4 columns would be the same. For example:

```
Step F [Hz] LOAD LOAD LOAD LOAD GEN ...
          GEN GEN GEN GEN ...
```

Within each adjacent quadruplet, the first half is polar (angle follows magnitude) and the second half is rectangular (imaginary part follows real part). This will be documented in a line immediately preceding the lines shown. For an illustration, see the first subcase of DC-51 or 52.

POLAR OUTPUT VARIABLES is the second of three requests for an alternative to the default choice of magnitude-only. Instead of four columns of printout for each output variable as explained in the preceding paragraph, there will be only two. Magnitude and angle retain their same relative positions (magnitude precedes angle in degrees).

RECTANGULAR OUTPUT VARIABLES is the third of three requests for an alternative to the default choice of magnitude-only. This is like the preceding POLAR except that real and imaginary part replace magnitude and angle.

The plotting of complex output variables has changed substantially because variable type and naming now is different than it was. For batch-mode plotting, the actual variable class (4 for node voltage, 8 for branch voltage, and 9 for branch current) now is used. No longer does an extra name appear with each variable on the plot card. Instead, the desired components must be selected ahead of time by means of a special F-SCAN COMPONENTS declaration. This is illustrated by DC-51 for the simultaneous plotting of both magnitude and angle. Each plot contains a maximum of 4 curves, and each can involve any one of the 4 possible parts of the complex variable. Selection remains in effect until the user changes it by another such declaration, or until a new data subcase begins. Of course, if the default choice of "magnitude-only" is in effect, then F-SCAN COMPONENTS has no use, and can be ignored.

RMS and THD values will be reported following each vector plot of ATP just as they have been for years following a FOURIER ON plot. Each vector plotting request results in a Fourier-like bar chart of harmonic composition. Yet, there are important differences. Recall the Fourier display referred to the preceding harmonic table. No such table exists in the case of FS or HFS, so there is no concern about whether an adequate number of harmonics has been chosen. For FS or HFS use, the bar chart will always display all harmonics available, and all of them will contribute to the numerical computation of RMS. All except the fundamental frequency (the starting frequency of the scan), that is. The associated text should make this clear: *"Each of these two statistics is produced by taking the square root of the sum of the squares of magnitudes of the scan after the beginning (power or fundamental) frequency."* This represents a fundamental difference from what is reported for Fourier use. The text explains: *"Rather than being the rms value of the total signal, this is the rms of the distortion (the signal consisting of second and higher harmonics only)."* According to author Furst, *"this is in accordance with IEEE Standard 519."*

Arbitrary precision and inter-column separation are principal benefits of the change to printout resembling that of the time-step loop. Recall it is the PRINTED NUMBER WIDTH declaration that controls such formatting of the printout. The optional calculation of variable extrema is yet another benefit, with the same integer miscellaneous data parameter MAXOUT used to request either use or avoidance. Rather than "times of maxima" and "times of minima," however, the user should see "F [Hz] of maxima" and "F [Hz] of minima."

LINE MODEL FREQUENCY SCAN (LMFS) deserves special comment because it was accorded special treatment during the overhaul of F-scan output. Yes, LMFS is F-scan as FS or HFS are, but LMFS is fundamentally different in that it involves special output. The decision was made not to modify LMFS output at all. Associated code was moved to a new module LMFS11 as part of the ongoing campaign against spaghetti code, and several simplifying assumptions were made at the same time: 1) no output of phasor branch flows ever will be requested (i.e., it is assumed that KSSOUT = 0); 2) no switches are involved; 3) no MODELS data; and 4) no rotating machinery is involved. As long as usage does not change, these should not be a problem (it is to be noted that the illustrations of DC-51 and DC-52 do adhere to these restrictions). Work was completed March 25th.

The frequency dependence of lumped resistance presently is represented in just the simplest way imaginable: linear interpolation using two user-supplied data points. Yet, smoother or more general alternatives are expected. Quoting from Gabor Furst's announcement : *"A more sophisticated frequency dependent representation will be available in the future using the Funk-Hantel equation. (ref. Dr. Kizilcay's EEUG paper at the 1994 Lyngby meeting).*

See the HFS Report for details." Another alternative, suggested by Gene Porter of Power Profiles, would be to linearly interpolate on a table that stored an arbitrary number of points. As this issue goes to press, no final decision has yet been made.

Machinery Saturation Questioned, II

This story might alternatively have been published under a title announcing more DCG / EPRI misinformation and/or disinformation. It began that way. But it also provides some resolution to a story in the October, 1997, newsletter, and this explains the II in the title.

Was BPA, not DCG or Ontario Hydro, the source of a Type-59 S.M. error in DCG/EPRI EMTP? This creative and misleading suggestion was observed in public E-mail of the Fargo list server dated March 19th. Paolo Mattavelli of the University of Padova in Italy provided context as follows: *"I would like to share with you a reply I got from the EMTP Technical Support Group regarding the SM59 model (saturation problem). Actually, one of my previous messages went unintentionally also to the Power-globe list and this is the reason why I got this INTERESTING email."* There followed a copy of what was alleged to be a message from Prof. Hermann W. Dommel, the former DCG / EPRI Technical Advisor. In its entirety, this stated:

"Dear Mr. Mattavelli, your correspondence with Ontario Hydro about synchronous machine saturation was forwarded to me because I was responsible for keeping track of bugs and of bug corrections in the DCG/EPRI EMTP some years ago. There was indeed a bug in the saturation representation in the code received from BPA years ago, and DCG/EPRI made a major effort to correct it. The corrections were made in the Hydro Quebec Research Institute by Mr. Jiming Lin of EPRI/Beijing, and after the corrections, type 59 and universal machine type 19 gave the same results. This may not help you directly because you do not have the DCG/EPRI EMTP, but at least you know that there was indeed a bug in what came from BPA, and I suppose that same bug may still be in the ATP."

Why was BPA named in this explanation? *Code received from BPA?* Why did Prof. Dommel not write something such as *code developed by DCG members?* Or, *code written by my former student?* This would have been more accurate. At the time, all DCG members supposedly were cooperating on a common, shared, public-domain version of EMTP. The obvious implication of naming only BPA is association of BPA with the error. It must be explained, however, that any real responsibility is highly unlikely. At issue is not the source of the computer file, but rather the source of the Type-59 S.M. error --- obviously. The appropriate question to be asked is: What person or company or organization was responsible for Type-59 S.M. changes during that period? Prof. Dommel certainly should

know, since it was his former student, Dr. Vladimir Brandwajn, who performed all development and validation work on the Type-59 model. This was documented in considerable detail in *EMTP Memoranda* of the period, and Prof. Dommel was on the mailing list. Let's look at one. From the middle of page FMOE-5 of a Vol. XIV memo dated the 2nd of January, 1984: *"As for Vladimir, a clarification is in order. He works on-site at BPA 16 days/month, under a contract with IREQ, which in turn passes the expense along to DCG ... Yes, Vladimir is right in the middle of a complete rewrite of Type-59 S.M. code ..."*. Dr. Brandwajn is the same person who was listed by Prof. Dommel as co-author at the start of the S.M. chapter (8) of the Theory Book. Finally, Dr. Brandwajn contributed progress reports to the *EMTP Newsletter* of which Prof. Dommel was Editor. For example, see the November, 1983 issue, which contained a story entitled *"Investigation and improvement of long term stability for the Type-59 synchronous machine (S.M.) model"* on pages 54-57. Prof. Dommel certainly should know who was making Type-59 S.M. changes at that time. It was **not** BPA.

About money and contractual control, Dr. Brandwajn then was an employee of Systems Control in Palo Alto, California. An overhaul of the Type-59 S.M. code was under way, with DCG --- not BPA --- in charge, and with Hydro-Quebec/IREQ --- not BPA --- handling the paperwork. Without any formal obligation, BPA donated supporting services that included office space and a computer terminal. But that was all. Prof. Dommel failed to mention this key detail: DCG fingerprints are all over the code in question. So are DCG member Ontario Hydro's, since Dr. Brandwajn was an OH employee for years prior to leaving Toronto for California. To name BPA but not DCG or Hydro-Quebec/IREQ or Ontario Hydro makes no sense if one really is interested in accountability.

Remember that famous error with Marti's transformation matrix more than a decade ago? This was the great discovery of LEC's Guido Empereur. Prof. Dommel has a history of misleading others about both DCG/EPRI EMTP and also BPA's EMTP. For a summary of this earlier circumstance where Prof. Dommel was caught confusing EMTP users, read the top 2/3rds of page 36 in the March, 1988, issue of *EMTP News* published by LEC. The latest disinformation about saturation bears obvious similarity. Specifically note how it ends with vague speculation: *"I suppose that same bug may still be in the ATP."* Why should the sneaky professor suppose anything of the sort? We now are in year 15 of ATP development, and he **never** has been granted access. Prof. Dommel has negligible knowledge about ATP. Yet, to an ATP user, he offers an opinion about a detail of ATP operation. Were circumstances not serious, this would be comical.

Saturation in ATP has undergone one important correction years ago, it must be noted, and Prof. Dommel **was** allowed access to a report of this progress. Readers who have forgotten the great detective work of Prof. Juan

Martinez in Barcelona, Spain, are referred to a story about it in the January, 1993, newsletter. Prof. Dommel mentions that saturation of the Type-59 S.M. was corrected in DCG / EPRI EMTP, but he fails to mention a comparable correction to ATP. Do you suppose this is a convenient oversight? Or is it honest ignorance?

Prof. Dommel also fails to consider the possibility that more than one error might be involved. This is a more serious mistake in logic. Is the professor operating in a dream world? Anyone who works with huge, real programs knows that all of them involve errors, and the last error never will be found (Lubarsky's Law). If one error might have been found in one program some time during the '80s, it requires an enormous leap of faith or professorial naiveté to believe or suggest that this might be the same error that was observed in another program a decade or more later. The connection defies logic. Old errors are continually being removed, and new ones are continually being added, as part of ongoing development. Rather than speculate wildly about something he knew nothing about, why did the professor not instead resort to simple experimentation (you know, as in *scientific method*)? Rather than **assume** that DCG / EPRI EMTP was bug-free for the data of interest, why did Prof. Dommel not attempt to demonstrate this contention experimentally? Of course, results of such an experiment might have spoiled both his day and his story.

Dr. Hiroshi Okamoto of Tokyo Electric Power Company (TEPCO) seemed to be the best informed of all who offered their opinions. He seemed to offer the final word in public E-mail dated March 24th, and this provides little support for Prof. Dommel's blind faith in DCG / EPRI EMTP: *"We think that the erroneous behavior of SM59 during saturation, which was pointed out by Dr. Haginomori of Toshiba, is due to the manner of correction of transformer voltage. The current version of SM58 is identical to SM59 in the treatment of saturation effect. So, I recommend that you should use UM instead of SM58 and 59 if the saturation effect is very important in your simulation. Mr. Cao, the author of SM58, has been correcting the bug and testing the new SM58 model. I will send the source code of corrected SM58 to Dr. Scott Meyer promptly after finishing the tests. I don't know that DCG has already corrected this problem. DCG1.1 made the different computation result from ATP SM59 for the data case from Dr. Haginomori. But the result of DCG was erroneous."*

TEPCO contributed Type-58 S.M. modeling to ATP as documented in a story in the October, 1996, issue. These are the people who should know what really is happening. Mr. Cao and Dr. Okamoto independently programmed the Type-59 S.M. equations so that their Type-58 procedure shares the tables used by Dr. Brandwajn's Type-59 model. Today, these Japanese probably know more about details of Type-59 S.M. coding than anyone including Dr. Brandwajn. Testing DCG / EPRI EMTP, they convinced themselves that it did **not** correctly simulate the data of interest. The Dr. Haginomori mentioned is Eiichi Haginomori, who last

September had expressed concern about Type-59 saturation. This was summarized in a story entitled "*Machinery saturation questioned*" in the October, 1997, issue.

Corrections to This Newsletter

Sorry, no one submitted any claim of any mistake, so operation continues in its usual flawless fashion (ha, ha). Come on, readers. Are none of you able to find significant errors? Any issue beginning with September, 1988, is fair game. If a mistake of consequence is alleged soon, it will be discussed in this same column of the following issue.

GNU ATP for Linux and DOS

Larger .LIS files of previous executions were not being completely overwritten by new ones as tested in Dr. Liu's DOS window of WinNT on April 7th. Both she and your Editor were slow to realize that extra output at the end of .LIS files came not from current program execution, but rather from less than perfect operation of the STATUS = 'UNKNOWN' opening of the file. The extra output for the MAKE phase of compiled TACS execution disappeared when the .LIS file was deleted prior to execution. Deletion of the .LIS file is recommended at the top of any batch file such as RUNTP that is used to execute the program.

The ASSIGN statement, which associates a statement number with a variable (e.g., ASSIGN 2542 TO NEXT), is handled differently by the GNU compiler than by Salford or Watcom compilers. The variable being used was a full-length (32-bit) integer, and it had been assumed that the content represented the associated address. This worked for Salford and Watcom. Although the variable needed to be a scalar to satisfy the Salford compiler, its content could be copied to a vector for storage. Later, when ready for use (the jump), it would be copied back out of the appropriate cell of the vector. This worked fine for both Salford and Watcom, but execution of GNU ATP died with a SIGSEGV error. So, in order to avoid such use for normal TACS data cases, a special patch involving -4897 was added to TAC SUP and MATH15 on June 11th. As for TACS ASSEMBLY LANGUAGE, as presently coded, it seems to be fundamentally incompatible with GNU, unfortunately. Yet, since the even-faster compiled TACS is free for GNU users, ATP developers are not particularly worried about the just-realized new limitation.

Taku Noda's frequency dependence was mentioned in the preceding story about Watcom. Similar progress should have been reported long ago for the GNU translation, since the conversion to GNU came first --- before the end of last year. Masahiro Kan of Toshiba Corporation in Japan did the good work, beginning with Salford C code that he obtained from Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. Mr. Kan reported success as follows in E-mail to

Prof. Kizilcay on Dec. 11th: "Today, I succeeded to run Noda Setup data case by ATP/Linux and ATP/DJGPP. It worked! I used VERT500 / SWITCH data case. The modification of the C code was minimal. I believe I can send it tomorrow." Answers to DCNEW-14 are correct here in Portland, and agree with Salford, except for a minor problem of display (the vertical bar in column 51 of interpreted lines is not being defined properly).

Compiled TACS Speeds Simulation

This is a continuation of the story about compiled TACS in preceding issues.

The Type-54 device ("binary pulse delay") involved possible error prior to correction on April 4th. In one place, variable I should have been evaluated numerically, but instead was left as a variable. This was in the NDX1 assignment immediately below IF (D7 .NE. -9999.) THEN. The result was an undefined variable, which caused immediate termination by the djgpp memory manager (some name such as sigsegv was mentioned). This was the good news. For some reason, Salford EMTP did not die. But why did the Salford compiler not warn about the undefined variable? Normally it does. After some investigation, it was found that variable I was used later, in a 2-line DO 6631 loop that initializes the vector of TACS right-hand sides. To conclude, variable I was undefined, but the Salford compiler was not smart enough to recognize the fact. The bad value supplied by djgpp was useful in identifying the universal error that plagued the code of TACS variable DEV54F in DC-19A data.

Industrial-strength testing of compiled TACS began in Portland on May 11th when your Editor first considered a large data case from Masahiro Kan of Toshiba in Japan. With 881 supplemental variables, this was found to be a substantial shock to the logic. I.e., bugs were observed. The following day, corrections were made to segments OVER2, TSTACS, and TAC SUP, and what is believed to be a correct .LIS file was sent to Mr. Kan by E-mail.

Chen on saturable TRANSFORMER

Prof. Xusheng Chen of Seattle University has written a 4-page, IEEE PES-format paper that is scheduled for presentation at the POWERCON '98 conference in Beijing, China, during August. The title is "*On the numerical instability of the saturable transformer component in EMTP*." Prof. Chen's Abstract follows (next paragraph).

"Numerical instability of the EMTP built-in model, the Saturable Transformer Component for the three-winding case, has been the topic of interest of the EMTP community for the past twenty years. This paper identifies the source of the numerical instability of the component through theoretical analysis and computer simulation. A simple

modification of the structure of the component is suggested, and the numerical instability is eliminated. The modification is supported by theoretical analysis."

"*Unstable saturable transformer*" is the title of a story in the April, 1993, newsletter. The initial paragraph reads as follows: "*The saturable TRANSFORMER (request word of columns 3-13) may be unstable for certain combinations of parameters, readers are reminded. An illustration from Jerry Almos of BPA follows. Any theoretician who believes he might be able to shed light on this mysterious and troublesome phenomenon is encouraged to consider this or any other realistic example, and share his insight with others.*" Prof. Chen responded to this call, and is to be commended for his resourceful analysis. Look for a continuation in the next issue.

Branch Data Input Restructured

The overhaul of branch data input reached the general public after April 26th, when six disk files were switched in the UTPF. These were CASCAD, INDIST, OVER3, SUBR3, INNONL, and HANNFD with the six old segments saved as .984 files (archive ALL984.ZIP was sent to Dr. Liu's 486 at BPA).

Noda simulation is validated by standard test case DCNEW-14. Unfortunately, this was neglected during the overhaul of branch input because old (ver. 2.66 or 2.67) Salford FORTRAN compilers continue to be used by developers in Portland, and these are incompatible with C (used by Dr. Taku Noda). Trouble with DCNEW-14 first was recognized by Dr. Liu as she worked on Watcom ATP using C code that recently was contributed by Robert Schultz of the New York City area. This was during early May. Your Editor confirmed fundamental problems on May 13th by experimentation using that shared 90-MHz Pentium at BPA. This offers newer ver. 3.5 Salford software, and it links with Noda object files. Beginning that evening, substantial changes were made to INDIST. But final corrections were delayed until Dr. Liu's return from a week-long trip on May 19th. Finally, DCNEW-14 produces the same answers as before the overhaul of branch input. To better ensure that this continues to be the case, REM text that had disabled the DCNEW-14 simulation in RUN.BAT on the P90 has been removed.

Power and energy output for a compensation-based branch seems never to have been verified by standard test cases. This surprising conclusion is the result of investigation of a complaint from Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. No doubt about it, a 4-punch in column 80 of a nonlinear element was not being handled properly by the new (see preceding issue) code for branch data input. Well, a correction was made to GETBUS on May 30th, and a new 7th subcase was added to DC-37 to ensure that any such error will not again be overlooked.

Noda frequency dependence temporarily was given an added restriction by the new (see preceding issue) code for branch data input, it should be mentioned. This was by accident rather than design, as discovered shortly before May 22nd. This was the date Dr. Liu edited her data (DCNEW-14) so that the special Noda declarations on the first phase began to the right of the BUS4 field (which ends in column 26). But who would have remembered this rule, and who wants to change the user instructions to reflect it? It seemed easier to modify the code, so this is what was done June 2nd --- even though the change (one more IF statement in SUBR3) is out of harmony with general principles of the restructuring. Blame this inconsistency on history and lethargy.

The flux of nonlinear inductors was mentioned in the January newsletter, but not the April issue as predicted. It is time for more information on this subject, which was nearly forgotten by developers in Portland. Fortunately the feature was not forgotten by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, however. In E-mail dated June 7th, he correctly observed that the progress of last December had been lost! Upon investigation, it was found that only a block of code in INDIST was missing --- a block of code that had been added to the UTPF during December, but not to files involved in the update of branch data input. To be sure operation will not be overlooked again, flux was added as a new (additional) output to the 2nd subcase of DC-7. For the branch from XXXX to ground, the column-80 punch was changed from the original value 1 to new value 3. The very first output variable now is the new flux, which of course varies smoothly from the initial value 30.833 (long seen in the initialization).

Stu Cook Uses Apple Macintosh

Two different versions of ATP have been executed successfully on a PowerPC-based Macintosh from Apple Computer. This report of progress comes from Stu Cook of Just Services in suburban Montreal, Quebec, Canada,

Watcom ATP for MS Windows was the first version of ATP to be successfully used on Apple Macintosh by Mr. Cook. This startling news was received in E-mail dated July 2nd, when Mr. Cook wrote: "*We seem to have the Watcom version of ATP running on my G3 PowerBook under Virtual PC emulation.*" The following day, it was learned that DC-1 spent 25.98 seconds in the time-step loop (39.11 sec total solution time). "*My CPU is a 750PPC running at 250MHz. it would seem that the emulator is running at about the speed of a 50MHz Pantium. I'll be letting Connectix, the VPC developer, know about this test.*"

The average heavy user might be disappointed (**only** a 50-MHz Pentium), but your Editor certainly is not. Yes, about 80% of the power is being wasted (25 seconds is 2.5

times the 10 seconds of Salford EMTP at 133 MHz, and 250 MHz is nearly double 133 MHz; $5 = 2.5 * 2$), but the 20% that remains is formidable. For example, it should easily beat any 486 (the DX2 used by Dr. Liu spends about 97 seconds in the dT loop of DC-1). For the average Mac PowerPC user, this should be an important breakthrough --- provided he can tolerate MS Windows 95 under which Watcom ATP is being run.

Pantium of the second paragraph is **not** a typographical error, your Editor learned. Rather, it represents humor from the non-Intel (in this case, Motorola) world. As Mr. Cook explained five days later : “*The reason for spelling it that way is the Pentium power consumption is in the 30-40W range vs. 5W for the G3 PPC, so it pants for breath and needs an auxiliary fan for cooling.*” For background, see PowerPC in the January, 1994, issue.

About Windows 95, Mr. Cook explained that this is real, just as Watcom ATP is real. No imitation is involved in either of these two programs. “*VPC emulates a Pentium and the OS is pure MS. VPC can also run OS/2 and WinNT, it seems, although I haven't yet tried either of these alternatives.*” So, VPC makes Motorola act as if it were Intel. Both MS Windows and Watcom ATP are Intel programs, so both can be supported. Neat.

Next, the Absoft compiler was used by Mr. Cook to provide higher-speed simulation. This is for those PowerPC users who either do not want to give up 80% of their power, or use Bill G's windows, or both. Absoft ATP for Apple Macintosh **does** seem to simulate at a normal, expected rate. Consider DC-1 simulation on that 250-Mhz Apple Macintosh previously mentioned. E-mail dated July 10th indicates 5.5 seconds spent within the time-step loop. Cycle for cycle, this is close to the performance of Salford EMTP on your Editor's Pentium:

Salford: 133 MHz * 10 sec = 1330 MHz-sec

Absoft: 250 MHz * 5.5 sec = 1375 MHz-sec

Here, a smaller number is better. Even more important, Mr. Cook reported on July 13th that his DC-1 answers including extrema agree within limits of round-off error.

Publishing Programs and Viewers

Novell Envoy is yet another alternative format for BPA's EMTP Theory Book. Yet, usefulness is severely limited because of the way WordPerfect (WP) 7 handled Greek letters of the original WP 5.1 files. The story began with public E-mail of the Fargo list server dated April 10th, which had Subject: “*Envoy files of the EMTP Theory Book are on secure FTP site in Japan.*” BPA's Dr. Tsu-huei Liu wrote: “*The Envoy files of the EMTP Theory Book were put on the secure FTP site at Osaka University in Japan by Dr. Funaki on April 6, 1998. ...We have two sets of files for the figures of this Theory Book : higher resolution (300dpi) and lower resolution (75dpi). It is the lower resolution set that was used to generate these Envoy files.*”

An Adobe Acrobat copy of the April newsletter was the next step in the quest for universal, free display of technical writing. The story was told in E-mail dated April 27th, which shared a copy with important contacts who distribute ATP materials via the Internet: Finally, E-mail of the Fargo list server was used on April 30th to inform the general public of this progress. With title “*Abandonment of Novell Envoy in Portland,*” this news release provided text for the following five explanatory paragraphs.

The April newsletter was the catalyst that began the most recent consideration. This writer uses Word 7 within MS Windows 95 for the work, and each issue of recent years has always filled every last line of a full 20 pages. For more than a year, this has been done at home using that Szymanski-supplied Pentium. Others, including Dr. Liu at BPA, typically would see more pages --- most commonly 21 and a fraction pages. No one ever reported exactly 20 pages, and everyone understood that differences of fonts were responsible for minor differences of display using different programs and printers. The problem was serious, and not easily solvable using either Envoy or PDF, because creation of these alternatives was not possible without upsetting the fonts. Envoy could be created if the MS Word file were taken into WP 7, but that would disturb the fonts. Alternatively, Dr. Liu could create PDF via PS output of her MS Word, but that, too, would upset the fonts.

--- until a week or two ago. Then, Dr. Liu observed for the first time that she had exactly 20 full pages on her screen. So, PS was created, and this was carried to the 90-MHz Pentium across the hall. There, Adobe Acrobat Distiller was used to produce APR98.PDF; and this file, too, when displayed by Adobe Acrobat, had exactly 20 full pages. Encouraged, copies were sent to contacts in Europe, Japan, and South America. No recipient reported any change of size.

Laszlo Prikler discovered an unexpected property of considerable importance. Quoting from his E-mail dated Tuesday, April 28, 1998: “*As Masahiro Kan wrote, the new PDF covers exactly 20 pages here, too. It remains true even when it was printed using an A4 laserprinter.*” It would appear that the decade-old discrepancy between American and European paper sizes has been solved by accident! This assumes one paginates using the American 8.5 x 11-inch size, and then wastes about half an inch at the bottom of each piece of European paper. Of course, for viewing on a computer screen, there is no waste.

Masahiro Kan has demonstrated that PDF can be created using free software. He supplied a PDF copy of the newsletter that he had created using GhostScript. Except for minor differences of fonts, it looked very good. For persons who do not have access to commercial Adobe Distiller, GhostScript looks very interesting for relatively simple writing such as newsletters. But output file size was substantially larger: 368 Kbytes vs. 156 Kbytes. Since BPA **does** have Distiller, it intends to continue using it.

While the 156-Kbyte size looks good, and comparable to the MS Word .DOC file, this is a little deceptive. The PDF file that is created by Distiller must already be compressed since application of PKZIP results in little additional saving (size of the .ZIP file is 148 Kbytes).

To be sure there was not something special about the April newsletter, the January issue next was converted and shared with others. Prof. Prikler confirmed undistorted display of this, too: *"Exactly 20 pages on the screen and on paper."*

End of story about the abandonment of Novell Envoy. Actually, Envoy is **not** being abandoned any more than HTML is. At any time, either could be produced as alternative output of WP 7, which continues to be used. Rather than the abandonment of Envoy, it is the addition of PDF that has occurred, for important documents. Future newsletters should be made available by your Editor in PDF format only, according to current thinking. Next time, the informed opinion of Prof. Mustafa Kizilcay on this subject should be shared with readers. He has some concerns.

"PDF files of the EMTP Theory Book are on the secure FTP site in Japan." This was the headline of Dr. Liu's public E-mail dated May 12th. Slightly-edited highlights of this announcement follow: *"The Adobe PDF files of the EMTP Theory Book were checked and put on the secure FTP site at Osaka University in Japan by Dr. Tsuyoshi Funaki ... Appreciation also goes to Prof. Laszlo Prikler in Hungary and Mr. Masahiro Kan in Japan for their efforts in validating these files. To download this file TBPDF.ZIP, one needs the password to log onto this site. ... The free Adobe Acrobat Reader 3.01 (ar32e301.exe) can also be downloaded from this secured FTP site if you don't have a copy yet. With the Acrobat reader, you can read / print these PDF files without need for any special word processor such as WordPerfect or MS Word. We have two sets of files for the figures of the Theory Book: 1) higher resolution (300 dpi) and 2) lower resolution (75 dpi). It is the set of higher resolution graphics that was used to generate the PDF files."*

Robert Meredith of the New York City area made the next major contribution: combination of the different chapters. Dr. Liu's WP 7 and PDF files were separated along lines of Theory Book chapters just as daughter Kwang-yi's original WP 5.1 storage was. We had many separate, smaller files rather than one big one. Your Editor was concerned that a reader might have trouble appreciating and comprehending the set of PDF files as a single work. There was concern that changing from one chapter to another would be artificially cumbersome compared with the same operation in a printed book. Everyone understood the advantage of hypertext links, which allow the user to jump from one document (chapter) to another. While such links are offered in WP 7, they have not yet been used, and would be missing from the PS output in any case. Yet Meredith quickly unified the chapters. Details next time.

Changes to Universal ATP Graphics

Changes to PostScript output (if NOPOST is zero) of ATP were contributed by Robert Meredith of the New York City area. Attached to E-mail dated May 5th, he wrote as follows about his improvements to PSLINE in two places: *"Both address long-standing problems with overflow on X-Y plots. The addition prevents pixel values from overflowing the PS file's I4 fields for excessively negative or positive values. (Excessively negative x values even used to hang the Apollo for X-Y plotting.) The deletion affects where the trace maxima label will be placed. ... Time plots should never need such an adjustment. X-Y plots can easily have traces running off screen at a 45 degree angle. Relocation of the off-grid maxima label only leftward makes no sense, because it is then well above the trace and nowhere near the maxima. Or it will be seen where it belongs when viewed by DisplayNT -- which can display plots as wide as 35 'inches' correctly. Window boundaries are meaningless with DisplayNT."*

LOWERC is a new binary switch in GRAPHICS that was added on April 22nd in order to accommodate the desire of Robert Meredith, a consulting engineer of the New York City area. At issue was the case of text in his PostScript output. Whereas ATP has used lower case since year one (1984) for all internally-generated text such as the units of time, Mr. Meredith prefers upper case. At last the choice is up to the user. As distributed by the user group, the value will be unity, which implies no change (i.e., continued use of lower-case text of batch-mode (CALCOMP PLOT) graphs. If value zero is used instead, internally-generated labeling should change to capital letters for all forms of graphs except any PRINTER PLOT in the .LIS file. As for user-supplied labeling such as the case title or vertical-axis label, the user remains responsible for this. I.e., whatever the user supplies should be displayed.

The 2550-byte limit of widexx was mentioned in the January newsletter. Unfortunately, as reported by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, this dimensioned limit could not be reached in practice due to a subscripting error. In E-mail dated April 30th, Mr. Hevia reported a program objection at about 2020 bytes: *"Overflow fixed 2550 bytes of CUNIT6 in PLTFIL."* Like Prof. Alexander, Mr. Hevia seemed to be thinking of education involving traveling waves: *"I did a test with CASCADE LINE, to make a movie. The 'camera' program runs OK, and TPBIG makes good 'book' but only if the number of photograms is limited. When I tested a file with 99 copies of a line section (the first case of DC9.DAT), TPBIG shows an error message I reduced the file to 49 copies, and then TPBIG ran OK. I am sending you the movie.dat and movie.lis files."* So, the trouble was found and corrected easily enough. May 2nd, your Editor observed: *"I changed N19 to unused N7, and now execution completes normally. But, when I try to confirm the width using LIST, this breaks each line at about 2*

Kbytes (probably 2048). We need a new version of LIST for our industrial-strength usage! Recall the 2550 is arbitrary. If there is a practical reason, it easily could be increased --- at least to 65K for the Salford compiler. So, if you want more width, we can give it to you."

Conditional Data Assembly

IF-THEN-ENDIF statements are allowed as part of ATP input data beginning April 10th. Modeled after ANSI FORTRAN 77, the new declaration allows either acceptance or omission of data lines depending on the content of an environment variable of the user's choice. Any type or class of ATP data can be involved, although for predictable results, any control block that is started in one disk file should also be completed in that same disk file. Users are warned not to have the terminal ENDIF in a \$INCLUDE file that is different from the file that contains the preceding IF-THEN statement.

Salford EMTP for MS-DOS (but not a DOS window of WinNT) was the vehicle used for initial implementation, as usual. Other program versions should require minor modification. Specifically, Watcom and GNU compilers allow the interrogation of environment variables, although details differ from those used by the Salford compiler. Regardless of whether or not the user has set KINSEN for conversion of input data, all bytes of the new control records will be unchanged from what the user supplies in his data file. Ensure consistency in this --- between operating system variable and ATP data!

The environment variable of an IF-THEN statement can have any length that is permitted by the operating system while also satisfying the need to have both variable name and character constant on the same line. I.e., a continuation line is not allowed. As for variable content, it must be nonblank. This is because Salford DBOS indicates blank content for any variable that has not been defined. If such an undefined variable is encountered, the program will halt following an error message. This forces the user to choose explicitly among the two or more data alternatives prior to starting execution of ATP. It is a somewhat arbitrary safety measure that protects the user from unpredictable behavior that might otherwise result.

Indentation of key words IF, ELSEIF, ELSE, and ENDIF is not allowed. Each must begin in column 1. While arbitrary, this requirement came from a desire for faster execution, which requires minimum handling of the IF line. Not only must columns 1 and 2 be used for the leading IF, the following open parenthesis must be in column 4, preceded and followed by one blank. Thereafter, blanks before and after the variable name and the character constant are arbitrary. Finally, the relational operator .EQ. or .NE. is not allowed to have any embedded blank. It is to span no more than the minimum four bytes.

Only the leading IF-THEN and the terminating ENDIF lines are mandatory parts of the new structure. Any interior ELSEIF (any number of these) or ELSE (just one of these, preceded by all ELSEIF that might be involved) is optional. As for the relational operator of any IF-THEN or ELSEIF line, only .EQ. and .NE. are allowed. Control is by means of a character string, and this can only be checked for equality or inequality with the following character constant. This is a fundamental difference from the ANSI F77 definition, of course. No provision for arithmetic, and numeric control of data, is being offered at this time. What reader can see any need? Recall that math within a \$INCLUDE file already has been separately considered (see the story entitled "Evaluating analytical functions" in the April, 1997, issue).

Old \$DISABLE and \$ENABLE are equivalent to weak special cases of the new IF-THEN-ENDIF structure. They are fixed forms of what more generally is variable. They, too, provide conditional data assembly, but are constant, and require manual setting. The new IF-THEN statement is dynamic, and control can be altered merely by changing an environment variable of the operating system. Also, any environment variable can be defined in data by a special SET VARIABLE modification to a normal IF statement. This latter usage (definition in data) has some similarity to the PARAMETER statement of ANSI F77. Unfortunately, Salford DBOS ver. 2.66 seems not to make the changes that it should. Separate 8-line test SET.FTN was used to demonstrate the problem, and it behaves no better using ver. 3.5 (the P90 across the hall). There is no error executing CISSUE@, but neither is there any effect on the environment variable RESISTOR that supposedly is being set. Any experienced user of Salford FTN77 is welcome to the test materials, if he wants to investigate what might be wrong, and why.

What about compilers other than Salford's? Might they work better? For example, what reader knows the appropriate statements for DEC VMS? More than just IF-THEN-ENDIF is involved for VMS, though. Before this latest development should come environment variable ATPDIR, which was introduced in the April, 1996, newsletter. What reader believes that ATPDIR should be honored by VMS ATP? If not, why not? Such an extension will happen only if some informed reader will show developers in Portland FORTRAN code to learn the content of a VMS environment variable. Once that works, developers then would be interested in the way an environment variable can be defined within ATP.

Power Company Politics & Religion

A new Administrator for BPA finally has been selected. In what can only be viewed as a setback for the "we must spend more to protect fish" crowd, Judi Johansen was named in a May 2nd announcement from the nation's capitol. The headline story of the May 2nd issue of *The*

Oregonian reads: "Johansen will take reins at BPA." The story explained that "Johansen has received support from most BPA constituents, with the exception of some environmental groups. ...In its final days, the search for a BPA chief became a political tug-of-war among regional leaders. ... The struggle exposed rivalries among BPA beneficiaries, including environmentalists and public utilities. Oregon's Governor John Kitzhaber seemed to lead the fish crusade, and to have been ignored. "Kitzhaber opposes top candidate to run BPA because of fish issues" was the title of a story in *The Oregonian* the preceding day. As written in the January newsletter, the Clinton/Gore White House seems to have had enough of BPA's wasteful spending on fish. Its limit has been reached.

Interactive Plotting in MS Windows

Several alternatives for interactive plotting now exist for those who either prefer to, or are forced to, work under MS Windows. WinNT probably is the most important case because NT is Salford DBOS-incompatible (so Salford TPPLOT can not be used). Consider alternatives (next):

PLOTXY is the creation of Massimo Ceraolo of the University of Pisa in Italy. Availability to the general public was announced in public E-mail of the Fargo list server dated June 5th. This had "Subject: A brand-new ATP plotting program," and it explained that PLOTXY "has several interesting features, such as:

- *) easy-to-use Graphical User Interface;
- *) full 32 bit (win32) code for very fast operation;
- *) up to 8 plots per sheet;
- *) allows plots from different files on the same sheet;
- *) allows plots versus time, as well as X-Y plots;
- *) factors and offsets and zoom support;
- *) clever automatic axis scaling;
- *) cursor to see values in numerical format.

If there is a sufficient request, I can distribute the program to the whole community (300 kB Zipped, 500kB unZipped)." Then, four days later, more information was received: "I had an instantaneous, abundant, answer. Many people want to try the program. I'm very glad of this, since I spent a lot of time in adapting a plotting program that I built for other purposes for the use with ATP. A few minutes ago I've sent a copy of program and instructions to Bruce Mork in the hope that he will agree to put a copy in the ATP FTP sites. I've personally tested the program with the maximum care, in both Win95 and WinNT environments, but I could not make extensive testing involving other people. Therefore, if you find any bugs, I hope you will be patient and send to me detailed reports. Comments and/or suggestions of any type are welcome."

GTPPLOT comes from Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. It was announced to the average ATP user in public E-mail of the Fargo list server dated June 8th. The Subject was: "GTPPLOT, the GNU version of TPPLOT." No, it is **not**

really Salford TPPLOT compiled by GNU, but GTPPLOT is written in FORTRAN (as TPPLOT is), and it does offer many similar or comparable features. GTPPLOT author Hevia explained: "There is a new option to plot output of ATP. This is GTPPLOT, the gnu version of TPPLOT. The new tool is available from the secure ftp server in Japan, at Osaka University, in the directory /common/gtpplot. The program runs under DOS, with nearly the same capabilities as TPPLOT. But it has some limitations, such as the lack of mouse or cursor keys for user input. The program can read C-like, FORMATTED, widexx, and gnu (not Salford) UNFORMATTED .PL4 files. Each plot can contain up to 20 variables, with each curve having different color, line type, or symbol, with or without grid, and with seven different fonts --- all under user control. The output can be plotted on the screen in different resolutions, and to files in HP-GL, CGM (Computer Graphics Metafile), PostScript, or gnuplot formats. The FOURIER command calculates the Fourier components as well as various power quality indices such as TIF or I.T product (IEEE Std. 519). Furthermore, the program generates plots of transformer derating (ANSI 57.110). Plots of FREQUENCY SCAN and HARMONIC FREQUENCY SCAN can be done as bar charts with more than one variable at a time, using differently-shaded bars. The program is written in FORTRAN 77, compiled with gnu FORTRAN Version 0.5.21, and makes use of the graphic package DISLIN version 6.3. For more details, see the documentation file, which is in Word 6.0 format."

DSPATP is the interactive plotting program from Masahiro Kan of Toshiba Corp. in Japan. With reliance upon GNUPLOT for the screen display, this was formally introduced at the March 13th meeting of JAUG (the Japanese ATP User Group) as explained in a separate story.

WPCPLOT is Mustafa Kizilcay's PCPLOT as ported to MS Windows by EEUG. While not yet available to all ATP users, it probably is widely used by EEUG members. If and when WPCPLOT might be made generally available, operation should be described in more detail.

DisplayNT from Robert Meredith and Robert Schultz of the New York City area does not provide interactive plotting in the sense of the preceding programs. However, it **does** provide an alternative. See the January newsletter.

Miscellaneous Small Items

Unit 6b of The Fontaine no longer should be used as a mailing address for the Can/Am user group. Like the address of BPA's headquarters building, it is a casualty of BPA's ongoing split into halves dominated by engineers (concentrated at BPA's sprawling Ross Complex in Vancouver, Washington, which includes the Dittmer control center building) and peddlers (who will dominate the main building). Unit 6B of The Fontaine was sold on

June 29th (date of the *closing*), so that mailing address, used by your Editor since 1980, no longer is applicable. Your Editor has left Oregon, and today resides within walking distance of work in *Vancouver USA*. This is what one hears on the radio, with *USA* added to distinguish this suburb of Portland from the much bigger Vancouver some 300 miles to the north (in Canada). Any snail mail to your Editor can be sent to the Vancouver post office box in the January newsletter. As for E-mail, there seems to be no need to change again. Continue to use Dr. Liu's mailbox at BPA.

A Transient Stability Program (TSP) might be combined with an electromagnetic transients program such as EMTP in order to exploit the advantages of both. This radical idea is not new, and has previously been advocated privately by another contact. But the idea became sufficiently public in the ATP community when Dr. Yuan Bin, writing from Mitsubishi Electric in Japan, entered an ongoing discussion about the relative merits of TSP and ATP. This was begun by one Jorge F. Gutierrez of Universidad Industrial de Santander in Columbia on May 29th. Whereas others (first Gene Porter, and then your Editor) were recommending TSP rather than ATP for the proposed use, Dr. Yuan Bin expressed need for a combination of the two in some cases. His message dated June 2nd explained : *"It is almost impossible to use EMTP-like program for large scale system. For example, to a system which have more than 60 generators and 100 substations, because the time-step used in EMTP-like program is almost < 70us, maybe you will need more than 1 day to run one time. TSP does not consider the fast transients of the electric network and generator stator winding, and also the harmonics produced by FACTS. I think sometimes the results by the TSP will be wrong Maybe we need a hybrid program which consists of the EMTP model and TSP model to analyze today's power system?"* What other informed reader is interested enough to join the public discussion of this subject?

TSTALL replaced PIXPUN in STARTUP beginning April 13th. New TSTALL is comparable to MOLASS of TPPILOT as reported in the April, 1996, newsletter. Recall MOLASS provided an extra, artificial time delay before each line of output. This was a response to Pentium, which made output to the screen too fast to be comprehended by the human eye --- or trapped by the fingers (CTRL-S). Warp-speed output is fine as long as execution is normal. But if execution crashes, how can a user learn what has disappeared off the top of his screen? He might look in the .LIS file, provided the output buffer had been flushed by the operating system (not generally the case for Watcom ATP as observed at BPA), and provided output was not being cached in RAM as a result of positive LU6VRT. To remove such uncertainty and complications, TSTALL has been added. It is the desired delay time in seconds for each non-blank line that passes through standard output routine TFLUSH. Value zero (no extra delay) should be provided in copies of STARTUP that are distributed by the user group, of course. If the user wants artificial delay, he must edit STARTUP himself. Loss of

former PIXPUN is no problem because screen pixels (in this case, an average value of pixels per unit of distance) were moved to GRAPHICS years ago. PIXPUN has not been used for years. To be continued next time.

Command lines as long as 200 bytes were a reform by Robert Meredith of the New York City area. E-mail dated April 17th documented various changes that he had made to Watcom ATP FORTRAN, and expansion of the limit on length of the command line was one of them. At first, your Editor viewed this as a case of overkill. Who had approached the 100-byte limit then in effect? How was this possible? Mr. Meredith explained: *".... with space for DISK and -R and three separating blanks, the first limits file name length to 45 characters."* This implies a factor of two, which your Editor had not anticipated. No such factor exists for parallel naming of the output file (using a star), but it **does** apply for an independent, separately-specified output name. This is a second file name. So, the expansion is reasonable, and was made April 21st. The Salford code differs somewhat from the Watcom code, however, and Salford EMTP does have another limit. Your Editor wrote: *"We still have a limit of 80 bytes, I see, but that is just for the input data-file name, so should be acceptable."*

Yorick is the name of a language that was of interest to Tom Short of Power Technologies, Inc. (PTI) in Schenectady, New York, USA. Your Editor first learned of Yorick in E-mail dated January 26th, when Mr. Short commented about yet another possible alternative for reading .PL4 files: *"I'd like to write a routine to read it into my favorite plotting program (Yorick - a free Matlab-type language that runs on Unix and Windows.)"* Your Editor encouraged the initiative: *"Let's plan on doing this one way or the other, provided you are interested in sharing the resulting conversion program with other licensed ATP users."* At the same time, the new widexx was mentioned, and Mr. Short agreed that this might provide a temporary, lower-performance solution. About a more efficient alternative, Mr. Short wrote: *"If I attempted the binary import into Yorick, it would not be a conversion program, just a Yorick routine to directly read the binary file. Yorick is good at reading binary files including binary files written on other platforms (it will do the conversions internally as far as byte order, integer sizes, and floating point format), so it might be possible to have this routine read PL4 files generated on other platforms as well. I would be happy to share this with other ATP users."*

Lightning was mentioned in the two preceding issues. Since then, the same Tom Short of Power Technologies has provided information about IEEE involvement using the Internet. Quoting from public E-mail dated May 30th: *"A new mailing list hosted by IEEE on lightning and its effects on the power system has been created. Discussions on line modeling for lightning (including ATP) are encouraged. ... For more information on IEEE mailing lists see <http://majordomo.ieee.org/eleccomm/maillist.htm> or contact Tom Short (ta.short@pti-us.com) who is the list administrator."*