Can / Am EMTP News

Voice of the Canadian / American EMTP User Group

Table of Contents

Salford Compiler and DOS Extender

READ PL4 HEADER was removed as a special request word for C-like .PL4 files on August 7th. This was the three-day (Labor Day) weekend of work preparing for GNU C-like files, when it was discovered that no standard test case DC*.DAT used the request. So, entire REQUEC was removed for Salford, leaving an unsatisfied external. The structure is being left to allow use by other compilers. After all of these years,   it has been forgotten which compilers required the request word for which type of .PL4 files,   and why. If any DBOS user stumbles over the change, he is asked to contact program developers with an explanation or illustration of his use. It would be easy to restore the C-like service, and this will be done, if need can be demonstrated soon. For file types other than C-like, it should be mentioned that nothing has changed.

About illegal library function usage (see story in the preceding issue), no changes have been made to Salford EMTP. Note that DBOS ver. 3.5 provides a good error message. It also documents the location (subroutine name and hexadecimal address) of the infraction,   and will provide a traceback of locations through calling subroutines if the user presses the down-arrow key. In a word, Salford DBOS ver. 3.5 performance is exemplary.   Even the performance of ver. 2.66, dating to April of 1992, was exemplary.   Time and exposure to newer competing products only makes the old Salford compiler look better.

<Loaded from command line: "DISK"> illustrates the first of three new lines of output to the screen that confirm parameters of batch-mode (RUNTP-like) usage. The idea comes from Robert Schultz and Robert Meredith of the New York City area. Quoting from Meredith's E-mail dated

September 20[th], this is done in order *"that users do not have to wonder whether to respond or not to the last of the prompts that sit on the screen while the case runs: Send desired disk file name for LUNIT6 [<CR>, -R]:"* The preceding two prompts typically pass so quickly the user has no time to think, of course. It is after the third prompt (for the output file name) that nearly all of the time of a typical simulation is spent. The new output has been a Watcom ATP feature for some time (Schultz's initial's and the date 4-12-96 are seen on the line of code). It became a Salford feature on September 23[rd] in order that Watcom could share module CIMAG4 with Salford. This eliminated the necessity of maintaining a second copy of what has grown to hundreds of lines.

The DIR command was a victim of the reform that allowed longer file names (see story in preceding issue). It remained broken until October 6[th] when it, too, was modified to handle 132-column names. This was the day following the first complaint of trouble. *"Illegal file name. Try again ..."* was the wrongful complaint reported by Prof. Juan Martinez Velasco of the Polytechnic University of Catalunya in Barcelona, Spain in E-mail dated October 5[th]. Whether DBOS actually will return names in excess of 80 characters has not yet been tested; but at least now the source code should be ready.

Those long file names (see preceding issue) never were tested using Salford DBOS, although the code seemed to execute correctly for the usual short names. That was using ver. 2.66 and 2.67 of DBOS. Upon switching to ver. 3.50 compilation, linking, and execution on October 24[th], it was found that execution sometimes would end prematurely --- always near the start of any second or later subcase of simulation. The problem was zero rather than blank fill of the uninitialized end of ANSI32L in SYSDEP. This trouble already was explained in the July issue (the only difference this time was 80 rather than 32 bytes as the start of the garbage).

Really big ATP networks have been simulated, even without large numbers of coupled coils (see *high order* in recent issues). Such cases make the superiority of Salford DBOS obvious. One request was from Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. He wrote as follows on November 30[th]: *"I received a requirement from a user (Eng. Eduardo Lalla, surely the biggest user in Argentina), about a bigger TPBIG. I sent to him the gnu TPBIG, ... but there is a problem with the time to load: about a minute for each run."* So, you Editor quickly redimensioned Salford EMTP using Mr. Hevia's dimensions (preserved at BPA in disk file LISTSIZE.20K, which is used by VARDI20K.BAT). The 20K reflects the maximum number of nodes or branches. Total table size is staggering: *".... LABCOM tables = 10694972 INTEGER words,"* which is just under 43 Mbytes. Yet, Salford EMTP simulated DC-6 without any apparent hesitation. About possible table dumping, your Editor warned Mr. Hevia: *"God help the person who tries to dump tables ....*

Beware of STATISTICS, etc. .... Well, if you had 64 Mbytes of RAM, and this all were available to DBOS, I guess you could dump tables. But beware of Bill G and the games he will play with Windows memory management."

Improvements to Salford TPPLOT

AUTO SCALE is the keyboard equivalent of a mouse click on the AUTO button of the *Multipliers* heading within the CURVE window. This service began November 7[th] when your Editor realized that a user without a mouse was deprived of an easy way to display both power and energy on the same graph. The new command should be honored, even though it will not be found in any of the colored, pull-down menus. In order that it not be forgotten, it will be added to the HELP file, however. Even users with mice should find occasion to use the new command. Sometimes it is easy to forget limitations of a mouse, which are severe: mouse clicks are interactive, so are a pain for repeated operations. For example, if one wanted to process a family of 100 .PL4 files, typically one would bury his plot commands within a batch file and then execute this using the @ command. Whereas mouse clicks could not be placed in such a batch file, the new AUTO SCALE command could be.

The EXTREMA display was tied to the SMOOTH tolerance in a curious way prior to correction on November 7[th]. This was the result of work with power and energy following an inquiry from Dr. Peter Naumovich of Cooper Power Systems in Olean (located about 100 miles south of Buffalo), New York. Before Cooper, the company name was McGraw-Edison. Energy is unusual in that it is monotone non-decreasing. Furthermore, for typical real signals, values in jules are much smaller than power in watts, or voltage in volts, or current in amperes. As a result, the energy signal might become lost among the larger signals. Of course, it is common for the maximum to occur at the final point, and for the user to begin with unity scaling (assuming 2 or more curves). If this happened, and if all previous points had been discarded as being visually redundant (using the initial, default scaling of unity), the final point was not being tested for extrema. So, both the minimum and the maximum were zero, and the curve plotted as identically zero --- even after automatic scaling (see preceding paragraph). This weakness was overcome by forced application of the final point to the extrema vectors. Using an old program, the user has two alternatives for avoidance of the problem: 1) zero the smoothing tolerance; or 2) alternatively use FACTOR to manually scale the curves believably.

News from Outside USA and Canada

A 252-page, soft-bound book about EMTP --- and specifically and conspicuously the ATP version --- was

received at BPA in a mailing envelope from Doshisha University postmarked Oct 9[th]. The book is a work of the Japanese EMTP Committee, published by Ohmsha or Ohm (its English symbol or logo), with the introduction written by JEC Chairman Akihiro Ametani. Your Editor can understand the many figures and diagrams and graphics, which contain some English. As for the text, Dr. Liu can read the Chinese (if not the Japanese) characters. On the cover is the Internet address **http://www.ohmsha.co.jp/**

Spanish and Portuguese are alternative languages for the publication of EMTP news. For background, see the April, 1998, newsletter. Prof. Juan A. Martinez Velasco of Universitat Politechnica de Catalunya in Barcelona, Spain, made the following announcement in public E-mail that was received early in the morning of November 16th (not to be confused with the futuristic message date): *"A new issue of the EMTP News in Spanish and Portuguese is already available. It can be downloaded from the following web site: http://iitree.ing.unlp.edu.ar/estudios/caue/revista.htm This issue closes the 1998 volume -- also called 'Volume 0' -- which has been edited without a fixed periodicity. As usual, the new issue includes three papers which deal with different ATP applications. Editors started this business with the aim of learning from experience .... From the next volume on, (we hope) four issues will be released every year on March, June, September and December. As mentioned in the previous message, those interested in publishing their work can submit their documents in either PDF format, or in a popular text processor. A file with 'Instructions for Authors' can be downloaded from the web site."*

Indian ATP use remains confused, with the latest indication being an E-mail message to Prof. Bruce Mork of Michigan Tech. Dated December 3[rd], Ashok Maitra of B. E. College (D.U), B. Garden, Howrah 711 103, wrote from a Hotmail address. He began his inquiry about hysteresis modeling as follows: *"Our institute and I myself are members of Indian EMTP User Group. We are also licensed user of EMTP software. We sent drafts to Mr. I. S. JHA, Secretary, Indian EMTP User Group on 03.8.98 for renewal of our membership for 1998, and to obtain password for accessing ATP software through Internet, but no acknowledgement or reply so far been received in this respect. ..."* Prof. Mork responded with a brief reference to Prof. M.V. Hariharan in Bombay (see the Januay, 1997, issue). Before some words about hysteresis, your Editor made a general observation: *"If you or anyone else knows of any other group of Indian users of ATP, that information should be mentioned in our newsletter."*

More about the Internet and E-mail

There is a 72-byte limit on lines of E-mail as received by MS Mail running under Windows NT on Dr. Tsu-huei Liu's Pentium at BPA. As a result, 80-column EMTP data lines that might be included in messages generally will be broken,

unfortunately. This problem is well known to many, but is ignored by some who request help using the Fargo list server. On October 20[th], your Editor reminded readers of the problem once again under a headline that read *"Periodic advice about broken data lines."* The following day, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, explained how he avoids such difficulties: *"I am using Pegasus mail program, and the option Wordwrap long lines can be disabled, and the file saved, without such modifications."* On the other hand, neither Mr. Hevia nor any other respondent explained how to avoid the wraparound using MS Mail.

"Amazon.com: Books, Music & More!" is the title of the Web page used by the pioneering Internet retailer. Anyone needing to check the name of a book, or its authors, is advised of this resource, which offers a search engine. This is better than a card catalog for the electronic book store. Millions of titles are included, and the user is able to search for key words (more powerful than simple alphabetical sorting of titles). Prices are attractive, with discounts from prices of the publishers clearly stated.

Notices of undeliverable mail occur from time to time. For an address that is valid, typically a 4-hour warning is to be expected. But what about the following explanation from BPA's post office: *"Subject: Undeliverable Mail. Date Sent: Monday, November 30, 1998 12:14PM. .... The following recipient(s) could not be reached: Gabor Furst (SMTP:gfurst@istar.ca) That user's mailbox is busy."* Who has ever heard of a mailbox being busy?

All IEEE papers are available from IEEE according to Marjan Popov, who on November 17[th] wrote public E-mail from TU Delft in The Netherlands. *"You can visit http://www.ieee.org/power There you can find information about ordering of all IEEE publications from conferences, transactions and proceedings."*

"The World Wide Web and high energy physics" is the title of a story by Bebo White in the November, 1998, issue of *Physics Today,* a publication of the American Institute of Physics. This article traces present WWW use from the Internet, which was begun *"in the late 1960s by the US Department of Defense, which wanted to develop a communications network that could survive a nuclear war. In particular, messages had to be able to reach their destinations, even if significant communication routes were destroyed. The solution, the 'packet switching' network, was designed so that the particular route a message takes to reach its destination is unimportant. ... The original Internet was named Arpanet, after its sponsor, the Advanced Research Projects Agency."* But that was E-mail and TCP/IP for file transfer. It did not provide hyperlinks for documentation. This was a European extension that came from the physics establishment rather than the defense establishment. *"CERN, located on the Franco-Swiss border near Geneva, is one of the most successful high-energy laboratories in the world."* The subtitle of the

article explains: *"The Web --- what you may now be using to buy an airline ticket or look up Miles Davis's discography --- began life as a tool for improving the flow of information at CERN."* Fig. 1 shows *"Tim Berners-Lee, the Oxford-educated physicist ... who invented the World Wide Web."* Fig. 3 shows *"the actual NeXT computer that ran the original World Wide Web server and browser at CERN."* This dates to November of 1990. It is quite a story, and another good illustration of how necessity drives invention: *"In the beginning was --- chaos"* (of information exchange).

Unlicensed subscribers to the Fargo list server remain a concern as this issue goes to press during mid-December. Look for details in the following issue.

Watcom ATP for MS Windows

The DIR command in response to the opening prompt was mentioned in the first (Salford) story of this issue. Upon consulting the corresponding Watcom code, it was noted that DIR was not supported. Yet, GNU ATP for Linux offered directory output (although not interactive selection), and this seemed preferable to nothing. So, DIR output was added to Watcom ATP on October 8[th].

Optimized compilation was described in the January, 1998, issue, which reported use of **/ot /on /5 /fpc /fp5**, with the final three of these a replacement for Robert Meredith's **/ox** (full optimization). This resulted in just over 9 seconds for time spent inside the time-step loop of DC-1. At the time, this seemed good. But then Meredith published his own benchmarks on October 8[th], and these were clearly faster. As reported in the October newsletter, there was an increase in the local paging file size to 100 Mbytes. This was the only known change since a year ago. Anyway, October 9[th], optimized compilation as recommended by Meredith (using /ox) succeeded at BPA, and all test cases except DC-56 were handled correctly by the resulting Watcom TPBIG. But what about DC-56? Dr. Tsu-huei Liu had the idea to omit **/ox** from just the compilation of that chunk of fortran that handles SPY HELP, and this resulted in correction of the DC-56 solution. This is the significance of new COMP6.BAT rather than the previous COMP5 as explained on rem (DOS comment) cards.

Protection against illegal arguments of library functions might be one reason to avoid Watcom compiler optimization. BPA's Dr. Tsu-huei Liu located details among disk files that accompany the ver. 11 compiler. One bit of writing has headline *"What you should know about optimization."* The first paragraph contains the following warning: *"when you are using the OM option to generate inline math functions no argument validation will be done for the intrinsic math functions such as 'sin' or 'cos'. Consider the needs of your application and select the optimization options that best meet your requirements."*

The 11[th] and final subcase of DC-59 illustrated two-dimensional parameter variation, with PCVP providing one of the dimensions and LINE CONSTANTS itself providing the other (frequency). This was handled correctly by Salford EMTP when data was in a separate disk file, but it was wrong as positioned in DC-59. The trouble first was recognized during Watcom ATP testing of the same code. For Watcom, there was no mistaking the symptoms of trouble: execution died when first tested October 26[th]. Later that same day, MAIN00 was found to be missing the cancellation of MAXKNT (which indicates a PCVP study, in this case involving a supporting program).

Two Type-59 S.M. models are connected to TACS in large standard test case DC-47, and this simulation suddenly failed for Watcom during the final week of October. Whether the cause was a change of compiler options (that use of /ox for optimization), or the removal of unused COMMON blocks (see story about Schultz's afterburner), is not known. In any case, Watcom ATP died in SSTACT, and this was traced to subscripting of data that was not properly initialized. So, simple avoidance was applied on October 30[th]. The patch is used by all computers, apparently to good effect. Watcom ATP execution has been restored to normal (within roundoff differences), and the Salford and GNU solutions are unchanged.

UNFORMATTED .PL4 files might never be created by the average Watcom ATP user because C-like are available, and C-like offer the advantage of portability. As distributed from Portland, Watcom STARTUP requests C-like. But the UNFORMATTED alternative is available, and was the choice of Dan Durbak of Power Technologies (PTI) in Schenectady, New York. He uses his own plotting program, and the logic of this assumed UNFORMATTED .PL4 files. He also uses a different compiler. Surprisingly, Mr. Durbak found compatibility for small files. But for large files (e.g., DC1.PL4), he found that Watcom ATP was adding unexpected end-of-record marks. Presumably Watcom ATP would read these correctly (e.g., if used with REPLOT), but his non-Watcom compiler did not prior to mandatory expansion of the record size to 32 Kbytes on November 12[th]. This is progress in that the previously unknown and unpredictable has become predictable.

News About TACS and MODELS

TACS and Type-59 S.M. modeling form a dominant part of large DC-47. It was during the final week of October that the solution was found to have been incorrect --- for a long time, and for all computers. Watcom ATP died during the execution, and this was traced to subscripting in SSTACT that was erroneous. The connection to the Type-59 S.M. was involved, and all computers previously seem to

have been using a zero subscript, which, not surprisingly, accessed a zero value. But suddenly Watcom accessed garbage. Whether this was due to a change of compiler options (that use of /ox for optimization), or the removal of unused COMMON blocks (see story about Schultz's afterburner), is not known.

MODELS use with FREQUENCY SCAN is illustrated by the 3rd subcase of DC-52, and the loop heading should reflect this usage. For each MODELS variable, the six letters "MODELS" should be seen in the first name of the pair of identifying names. Unfortunately, this was not the case prior to correction on October 27th. In E-mail earlier that same day, Prof. Mustafa Kizilcay of FH Osnabrueck in Germany alerted developers in Portland to the problem. He was using different data (COPENH4.DAT), but similarity to DC-52 was unmistakable. The problem was not noticed in DC-52 output because branch flows are requested, and as a result, the column heading was inconspicuous. Not so for Prof. Kizilcay's data. For the record, only 2 lines differ:
Incorrect lines of DC52.LIS :

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

Corresponding correct lines:

```
   Step   F [Hz]      GEN       LOAD      MODELS ..
                                          VRE1   ..
```

The FREQUENCY SCAN data of DC-52 produces output of just variable magnitudes. For such data, the correction of the preceding paragraph seemed to be correct. But for multiple (2 or 4) output parts associated with either polar, or rectangular variables, or both, the printed heading was in error. This first was reported by Prof. Kizilcay, in an E-mail message dated October 30th. Printed variable headings (but not the following variables themselves) incorrectly were repeated for MODELS variables prior to correction later that same day. Now each MODELS variable never has more than one part.

ABSOLUTE TACS DIMENSIONS allows the user to size TACS tables arbitrarily within the total available working space of VARDIM List 19. The fourth of eight table sizes is for the number of TACS sources, and this is confused by the fact that 11 such sources are supplied internally and automatically. Recall names such as TIMEX for simulation time and PI for the geometric constant. The user's size is supposed to include space for these, but what if he forgets and requests a number less than 11? It was found that ATP **was** protected, provided data defined at least one source of its own. Otherwise, there was no protection, and results were erroneous prior to the addition of special error checking to SUBR1 on October 30th.

Line and Cable Constants

The mixture of overhead and underground cables is one of the latest innovations of Prof. Akihiro Ametani of Doshisha University in Kyoto, Japan. Attached to a letter to

BPA's Dr. Tsu-huei Liu dated October 2nd was a 4-page report entitled *"Report of the recent modifications of Aki Cable Parameters (June to Oct. 1998)."* Near the top of page 3 is the section heading *"5. Mutual impedance between overhead and underground cables."* At the end of the section, Prof. Ametani wrote: *"I am going to present the above modification and calculated results in comparison with the accurate result at the EEUG'98 meeting in Prague, The Czech Republic."* The following modification was entitled *"6. Accurate formula of space impedance (self term)."* The professor may now be a dean (see the July, 1998, issue), but he seems able to remain productive in the area of research that made him internationally famous.

Higher - Order Pi Circuits

Real data for 400 coupled coils were received from Dr. Hiroshi Arita of Hitachi, Ltd. on October 23rd. For background, see the April, 1998, issue. Full-size data was received in the form of HIT6.DAT, which consisted of 31884 lines and required 2613 Kbytes to store. Yet, formatting continues to expand such data unnecessarily. Dr. Arita has not yet switched to the special C-like alternative (HIGH ORDER PI CIRCUIT, or HOPC in abbreviated form); he continues to use conventional Type-51,52, ... branches, which have become overpowering. There are 3736 branches, of which only 400 are the famous high-order Pi. Next issue, more statistics might be provided.

RENUMBER WITHOUT COUPLING (RWC) is a new request word that was introduced October 24th in order to avoid the burden of transient renumbering. Illustration began the following day with the addition of a new 5th subcase of DC-5 (see comment cards for further explanation). In case the reader does not understand the need, consider preliminary results of that 400-coil example from Hitachi. As reported by your Editor in E-mail that same day, the burden of renumbering is reflected in the second component of time ("Seconds for overlays 6-11"), and this was overpowering: 112.527 sec for your Editor's 133-MHz Pentium. The preceding data input required only 15 seconds (using $LISTOFF to suppress nearly all output), and the following preparation for the time-step loop only 41 seconds. Transient renumbering (there was no phasor solution) clearly was the bottleneck. Using RWC, the 112 sec was reduced to 3.132 seconds. Furthermore, there was no noticeable change in resulting simulation speed. Storage of List 5, which includes triangularized [Y], was unchanged. Typically it **will** change, of course. For normal data, John Walker's Scheme-2 renumbering is of value. If it is largely avoided, simulation speed generally should suffer. But not for the 400 coupled coils of HIT6. Why? Perhaps each node already is connected to every other node, so order of elimination might be immaterial. About the burden on List 23, the saving is striking, although small, when viewed from the perspective of List

5. As reported in case-summary statistics, the figures are 3664 and 82216, respectively, for use with and without RWC. More details are ready for the next issue.

Bernd Stein of FGH in Mannheim, Germany, was mentioned in the October newsletter. His data, too, continues to grow. He reported progress in E-mail dated November 9th: *"The data case now encloses 3121 nodes, up to DU2664 automatically serialized, 6570 branches and about 19000 data lines after inclusion of all DBM data. The square of the maximum number of coupled phases, LSIZ26, is still 5476, as we do not yet model coupling between the three phases."*

Brain - Damaged MS Windows

John Kappenman is mentioned in a later story. He, too, seemed to be a victim of Bill G's incompatibility with himself. Kappenman wrote: *"I am having a problem getting my hands on a printable licensing form."* This is the first complaint your Editor can recall about MS Word not importing the user group's standard WP 5.1 disk file. At BPA in Portland, using Word under WinNT on Dr. Liu's Pentium, there is no problem. But a copy of the same file sent to Duluth was rejected. Why? About software being used, Mr. Kappenman wrote: *"I did try it on both the old license agreement as well as the new. Both are WP5 ... I can at least get an old WPD file to open though it is garbage in the translation ... still no go. I am operating in Win95 with a Word version from Office97. Like I mentioned before, I can open and translate most WPD files, but apparently not this WP5 file."* To solve the problem, we in Portland exported as a Word file, and sent the result. End of difficulty.

Corrections to This Newsletter

Note that Prof. Hermann Dommel's contribution to the preceding issue was **not** published under the present title. It was **not** submitted along with any claim that your Editor's writing was in any way in error.

Prof. Dommel **did** apologize, but not for his misleading association of BPA with the error, it is to be noted.

Prof. Dommel's characterization of DCG/EPRI EMTP version 1.1 as used by TEPCO as *"a very old version"* is revealing. The professor indicates that *"corrections were made ... in EMTP96 (Version 3 ..."* But, according to EPRI's own projection (see the October, 1996, issue), this was not available to members two years prior to Prof. Dommel's writing. In round numbers, of the 14 years of DCG/EPRI EMTP development, the error had been corrected for two years at most. Rather than refer to a version that is uncorrected as *very old*, it might be more accurate to refer to the corrected version as *very new*. Of course, this would be more embarrassing to both the professor and his sponsor. It also makes more preposterous the association of BPA with the error. Not only did BPA have no special responsibility during the mid-80s, it has supplied to DCG or EPRI no EMTP work since then, as far as your Editor is aware. The error clearly is DCG's.

Variable Dimensioning of ATP

The Type-50 SCE generator model preceded the Type-59 model of Dr. Vladimir Brandwajn, recall. Although the SCE code itself was removed by Dr. Brandwajn during the early '80s, a remaining trace has lingered on without anyone noticing or caring all of these years. Your Editor is referring to dependent List Size 53, which sized compensation current vector CURSUB using both the number of nonlinear elements (List 9) and the number of S.M. components (List 17). But the Type-59 code does not use compensation as the SCE model once did. On the other hand, the U.M. does. Rather than replace List 53 by List 9, your Editor decided to replace List 17 in the formula by an estimate of the number of U.M. machines based on default proportions. This was done October 18th when it was assumed that one machine exists for each 112 cells of List 25 space (see DC34.LIS, which assumes the default value of NUMFIX = 3 machines).

Several COMMON blocks were found to be unused, so were removed. This was the most surprising discovery during research about table dumping and restoring: COMMON blocks that had not been used for years. The heading that begins the .LIS file changed as follows on October 17th: from *"size of LABCOM tables = 233316"* words to 230308 for 3 times default dimensioning. This may be the first decrease in history.

Evaluating Analytical Functions , III

Exponentiation is requested by the ** of FORTRAN, and it was supposed to be a supported function, of course. But as Gabor Furst observed in E-mail dated September 30th, exponent value 2 was handled correctly whereas any other value seemed to result in a temporary error termination. Apparently the general case never had been tested. Once existence of the trouble was known, correction was found to be easy enough (the following morning).

Exponentiation A ** B was not properly handled for negative A prior to October 12th. Working with Gabor Furst who had interest in the subject, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, wrote that your Editor's assurances were theoretical rather than actual: *"Yes, but this is in FORTRAN, not in PCVP. I tried (-1) ** KNT and ATP aborted with some error message like 'real exponent for negative base' from DBOS."* So, your Editor modified logic to check for non-negative A.

Should A be negative, ATP will continue with special logic provided B is an integer. If B is not an integer, there will be an error message.

$PARAMETER definitions were limited to a single line for each variable that was being replaced. This was prior to October 17th when intermediate variables first were allowed in response to a request from Mr. Furst. Now, an arbitrary number of intermediate variables is allowed, permitting much more complicated formulas. For an illustration involving one intermediate variable, see the 5th subcase of DCNEW-25. The user is warned, however, that all intermediate variables are erased upon use of the following variable of substitution. So, if one intermediate variable is used by two or more variables of substitution, it must be repeated before each. Later, if this is found to result in substantial repetition, the rule might be modified. But, for now, it is convenient. As usual, the capability evolves (bottom-up design, as with most creative engineering).

An induction motor model prompted an increase in the number of intermediate variables from 15 (the original limit) to 25. This was a couple of days later. It did not take long for limits of the preceding paragraph to be recognized as inadequate. Writing to Orlando Hevia and Gabor Furst on October 20th, your Editor finally could report success handling Mr. Hevia's induction motor data: *"I ... have been debugging my way through the pocket calculator. ... I modified the code to preserve intermediate variables after each substitution variable. ... Preceding intermediate variables remain defined ... Next, I found lack of protection against overflow of temporary variables. This morning, I allowed only 15 (it seemed like a lot at the time). ... clearly, I had not anticipated heavy duty usage such as yours. So, I added protection against overflow at the same time I expanded the storage to 25. Next, I found that xxxxxx was an assembly language verb produced by my compiler. But it was missing in the code to execute. ... So, I now add the missing execution code (no problem once the failing is recognized). ... Good deal, now execution runs to completion. ... Let me know what more might be wrong, if anything, and I will work on it."*

One or more blanks separating a parameter name from its following equal sign were not generally permitted prior to October 22nd. The preceding day, Gabor Furst had noticed that column boundaries for R and L of a series R-L-C card were not being properly observed. This was for his induction motor model within an HFS study. The error was traced to an extra blank left of the equal sign of a substitution variable. But blanks provide useful visual separation, and should be encouraged. So, the logic of S.N. 1094 and below was added to MATDAT in order to ignore any number of blanks preceding an equal sign.

The Type-10 electrical source card carries an analytical expression to the left of familiar T-start in columns 61-70. Unfortunately, an easy error for the user to make was a formula that was too long. If the formula spilled over into T-start, the result generally would be unreadable, and execution would die (for Salford EMTP, with DBOS opening a debugger error window). Beginning October 28th, ATP will check for legal T-start and T-stop, and will end the subcase in civilized fashion (KILL = 236) if an error is found. This is illustrated by a new 2nd subcase of DCNEW-19 (the former 2nd has become the 3rd). While few users care about Type-10 sources, the improvement is believed to be noteworthy as a test of procedures that might be applied for checking of other, more popular data.

A sawtooth waveform on the electrical side (i.e., not in TACS or MODELS) was the contribution of Orlando Hevia. His E-mail dated November 10th began as follows: *"I send a file with a sawtooth source."* For your Editor, this was a complete surprise. Mr. Hevia used the pocket calculator to define a Type-10 source, as first mentioned in the October issue. That is, rather than use sawtooth logic that is built into the program, Mr. Hevia supplies his own analytically, in data. To be sure that this creative idea is not forgotten, Mr. Hevia's centered waveform has been added to the 1st subcase of DC-22. Might there be more such contributions? Maybe. Mr. Hevia wrote: *"I think that with some work, a rectangular pulse train (as TACS 23 source) could be simulated."*

The numerical interface of the pocket calculator was changed November 11th in order to remove a conflict with other sources. This was found to be necessary for the just-mentioned modification of DC-22. Surprisingly, two signals of one other standard test case changed in the process: node voltage COS and associated current (COS, TERRA) of the 1st subcase of DCNEW-19. Clearly, the signals have been improved (previously, they were wrong after the first time step).

Continuation lines of FORTRAN are distinguished by blank columns 1 through 5 and nonblank column 6. This capability was added to your Editor's compiler on November 14th in order to solve a problem posed by Orlando Hevia. His E-mail the preceding day explained: *"I tried the pulse source, but I can't reduce the number of lines to two."* First, the Type-10 analytical source allowed just columns 11-60. Then, when more space was required, CONT. in columns 61-65 was provided to indicate the desire for a single continuation line. But even this was not enough to satisfy Mr. Hevia's creative mind, it would seem. So, expansion to an indefinite number of lines now is being provided. This is independent of CONT. which will continue to be honored for historical reasons. Yet, it is not required. FORTRAN-like continuation is being allowed, and is illustrated in the 3rd subcase of DCNEW-19.

$$ (a double dollar sign) was selected to mark an intermediate variable within $PARAMETER data. By definition, an intermediate variable is one that is used to define later data symbols, but is not itself a data symbol. Such a symbol will not be searched for and replaced in data cards that follow the $PARAMETER data structure, then.

Illustration is provided in the 5[th] subcase of DCNEW-25, although this use is trivial and artificial. The need came from Orlando Hevia during mid-October, as he was constructing a frequency-dependent model of an induction motor for Gabor Furst's HFS use. In a disk file named GABMOT.DAT, this data uses 20 intermediate variables followed by 5 data symbols. Sample lines are scheduled for illustration in the associated Rule Book section.

Simulation speed for TACS first was measured on December 4[th]. This involved the same data as for the original TACS ASSEMBLY LANGUAGE (TAL) as summarized in the January, 1998, issue. Once again, we use the best 5 of 6 consecutive trials of the 60-variable, 100K-step MATHTACS.DAT. But this time Salford DBOS is run under DOS of Win95 rather than within a DOS window of Win95. We now can add a row to the preceding tabulation of time-step loop seconds using your Editor's 133-MHz Pentium:

```
Dube's original  TACS,  60 variables:  106.21
TACS assembly language, 60 variables:   18.33
New built-in compiler,  60 variables:   18.41
Original TACS, one dummy variable :      8.46
```

About time to compile, this additional burden would be part of the initial timing figure, which measures data input. This has not changed significantly. Again taking the best 5 of 6 figures gives 2.08 seconds. No big deal (your Editor's compiler is plenty fast).

Linux and Programs It Supports

C-like .PL4 files now are offered by GNU ATP thanks to the efforts of Masahiro Kan of Toshiba Corporation in Japan. For background, see the story entitled "Szymanski C-like .PL4 structure" in the preceding issue.

Your Editor did proceed as planned, modularizing and testing using Salford library functions. But then came Mr. Kan's surprising response. Quoting from E-mail dated September 9[th]: *"I looked into the source code, and I believe the work I should do is to write the following functions in GNU C ..."* The following day, your Editor concluded: *"Why did I do the work? If you can satisfy the Salford library routines, we do not need my work. Too bad I did not see this earlier. Yes, if you can supply Salford-like routines, we can use Salford ATP FORTRAN. We do not need the neutral interface that I wrote. Yes, we can do away with the extra layer of interface. This is perfect. I like your idea very much. I think Schultz would be proud of you. This is more than a convenient (i.e., easier) way to proceed, this is a **better** way to handle the problem. Somehow, I was tied to Szymanski's procedure of a decade ago. Of course, that was before Salford. Today, Szymanski probably would do exactly what you have proposed."*

Mr. Kan worked quickly, attaching his work to E-mail dated September 13[th]: "I wrote Salford compatible C-like functions in GNU C. .... I also attach testclk.f and makefile which are used to test these functions." After verification using Mr. Kan's small test program, GNU installation-dependent modules were modified by the addition of Salford-like code to handle C-like files. Then came verification of test cases, changing all usage from UNFORMATTED to C-like where appropriate. The job was completed without the discovery of a single error in Mr. Kan's code, as best your Editor can recall. To conclude, GNU ATP today allows C-like .PL4 files as demonstrated by any of the numerous standard test cases.

GTPPLOT is Orlando Hevia's interactive plotting program as explained in the July, 1998, issue. Well, it began as an alternative to Salford TPPLOT, which will not run under MS Windows NT. But it had the advantage of reliance upon a graphics library that was available for many platforms, so extension to completely different Linux was just a matter of time. This important milestone is documented by E-mail from the author to Masahiro Kan of Toshiba Corporation in Japan on October 24[th]: *"I send the source code of GTPPLOT for Linux, the makefile, the binary, the .ini and .hlp files. Some work remains .... (Mr. Hevia lists 5 things). .... Please send your opinion, if you can run the program."* The following day, author Hevia added: *"For Linux, the best option for hard copy are the Postscript files. I don't know if hp2xx is available for Linux."* Three days later, Mr. Kan responded with a favorable evaluation: *"It also worked fine at home without the libdislin.so.7 file. Gtpplot/Linux is ready to publish? Anyway, you did a great job!"*

Three days later, Mr. Hevia acknowledged the importance of Mr. Kan's support for Linux. He also related interesting details about his recent acquisition: *"Thanks, with the help of a good friend! By the way, the RedHat Linux 5.1 took 3 days to travel from the USA to Santa Fe! The package cost $49.95 and included Installation guide (2 CD-ROM + 1 CD-ROM as catalogue of Linux products and demos)."* Delivery by Federal Express added $32, and the Argentine government added $35 in taxes. So, Linux in Argentina is more expensive than in the USA, but still a very good deal. Consider the alternative from Bill G: *"Compare with Windows + Fortran + Graphics package + ... !!"* About use by others: *"I send the document lgtpplot.doc in word 6.0 for Windows 3.1 format. Please, review the text ... I converted Word to PostScript, and PostScript to PDF, with very good results using Linux. The PDF format is nearly universal. Can/Am EMTP News could be printed using Linux, too."*

"WordPerfect for Linux" was the subject of E-mail from Masahiro Kan. Dated October 27[th], this explained: *"The following is an abstract from http://www.corel. com/news/1998/october/linux.htm : Corel Announces No Charge Linux Download Special version of Corel WordPerfect8 for Linux Personal Edition to be available at no charge ... Information about the download version will be available from http://linux.corel.com/ . Interested users can preregister at this site for the application, which is expected to be available in November."*

8

Compiled TACS & Assembly Language

Secrecy surrounding TAL has not yet been relaxed. There are two schools of thought, and no final decision has yet been made about the question. Recall that secrecy was declared in order to prevent commercial exploitation. That was explained in the January, 1998, issue, before work on a compiler had begun. Once the free compiler becomes available to the general public, any window for commercial exploitation should have closed. So, according to this line of thinking, secrecy could be relaxed. But there is another viewpoint. Once a free compiler is available, what user needs to see TAL? It will be used internally, but no longer will be required externally. Should user-supplied TAL survive availability of the free compiler? Is the burden of compilation large enough so that users will want to save the assembler? Answers are not obvious.

The Type-68 Meredith hysteresis device was described in the April, 1998, issue. It was not supported by compiled TACS until September 28th, however, when operation first was tested. New disk file DCN25A.DAT was added to the set of compiled TACS illustrations to ensure that Meredith's model will not be overlooked again. Prior to modification of the code, the Salford compiler warned that variable A was undefined. For completeness, DCN21O.DAT is a second disk file that was added at the same time. In that execution halts prior to entry into the time-step loop, compiled TACS really is not involved, and little is demonstrated. But, TACS **is** involved, so DCN21O (the 15th subcase of DCNEW-21 corresponds to letter O) was added just to prove that it, too, can be handled without difficulty.

About AUT5 (Meredith's benchmark mentioned in the April newsletter), compiled TACS did not make a great deal of difference. Meredith's Type-68 device already had saved most of the inefficiency. Following that reform, devices dominant the TACS usage, and the difference between compiled and not-compiled devices is not great. Disk file AUT5.DAT allows the addition of a third row to the April table. Times for the time-step loop, in sec:

```
Dube supplemental variables :  72.802
Meredith Type-68 device      :   9.011
Compiled TACS of this        :   6.209
```

COMPILED TACS MAKE (CTM) finally is illustrated by standard test cases. A new 6th subcase was added to DCNEW-25 on October 8th for two reasons, actually. Yes, it illustrates CTM (important). But it also documents the basic pocket calculator formulas that were used many times to illustrate how slow Dube's MODELS and TACS can be for such elementary mathematics. Six equations are involved. When these are copied 9 times, one has the 60 variables mentioned in the January, 1998, issue. About possible replacement of CTM by COMPTACS=MAKE, the reader is advised not to try. An extraneous error

termination resulted, due to the stacking. Since it seems unlikely that CTM ever would be appended to the end of another subcase for practical use, the trouble might not be worth the effort to resolve.

BEGIN SUPPLEMENTAL ONLY and END SUPPLEMENTAL ONLY are two new declarations to accommodate standard FORTRAN, which begins in column 7. To understand the motivation, consider some TACS history. During 1984, Chinese visitor Ma Renming made several important improvements to the TACS code that originally was written for BPA by contractor Laurent Dube. Probably few readers recall the separation and ordering of data by TACS class, with each terminated by its own blank card. Thanks to Mr. Ma's creative work, this no longer was necessary in the "M39." version dating to July of 1984. But the price that was paid for such progress was reliance upon a unique type code in columns 1-2 for all TACS data. The card requesting TACS output variables became type 33, etc. Well, standard FORTRAN, which uses columns 1 through 5 for a statement number, collides with this assumption. Standard FORTRAN can be allowed only if some means other than the type code of columns 1-2 is used to determine TACS data type. The request words that began this paragraph are one such means, and they always works. There is no limit on the number of times such switches can be used, either, or where they can be placed within TACS data. They work for TACS data much as $VINTAGE or $UNITS does for branch data. They are needed for standard FORTRAN if the equal sign is not preceded by one blank and followed by one blank. This latter rule is simple, and probably will be used more even though it is not perfectly general.

Superposition of Phasor Solutions

The proper way to handle multiple-frequency, sinusoidal excitation of the steady-state, phasor solution is to use superposition. Or is it? For the first time in two decades, minor progress has been made. Superposition of phasor solutions has been demonstrated for linear networks of series R-L-C branches. But where do we go from here, and why? What are practical uses? How much more effort is warranted to handle other branches? What about nonlinear elements? Your Editor had no shortage of questions.

DCNEW-26 has been given a new 5th subcase to document the first successful superposition, which occurred November 8th. ATP-licensed readers who have interest are encouraged to study comments of this data, and to verify smoothness of the solution. The screen plot using Salford EMTP is perfectly repetitive and smooth, and the HP-GL output looks similarly perfect using WP 7. The .LIS, .DAT, and .HPG files that document operation were sent to Prof. Mustafa Kizilcay and Gabor Furst in Prague early the following morning. This message had *"Subject: Another new type of F-scan."* Your Editor concluded as follows: *"There are lots of questions that need to be*

thought through. You, Gabor, and your short course seem like a good forum for the discussion."

Within each phasor solution, two or more frequencies are allowed as long as they do not overlap. This is no different than it was in years past. It is the user's responsibility to ensure that subnetworks involving different frequencies are not connected. The new SSOMIT switch (see mention elsewhere in this issue) might be of assistance in this regard. Alternatively, if there **is** a connection, the two different frequencies now can be separated by placement in different phasor solutions. The user has complete control over which phasor solution involves which excitation. Although most users probably will order the solutions so frequency increases, in fact there is no such need. Order is arbitrary except for one detail: All phasor solutions preceding the last one are added to the last one for continuation into the time-step loop. As a result, the final excitation must include sources for the time-step loop. All previous groups are used only for the associated phasor solutions. The final group serves an additional function, so is fundamentally different. Perhaps it is more useful to have the final solution include the power frequency.

Additional storage is required for the new superposition, and an existing working vector of List 13 has been used for this. There should be protection against overflow, of course, although such trouble seems unlikely for networks that are not huge. Storage does depend on the number of nodes and branches, but thus far is independent of the number of harmonics or phasor solutions involved.

Integer miscellaneous data parameters IDOUBL and KSSOUT are important to the user of superposition. About the former, value zero will result in no output of the connectivity table, of course. On the other hand, value unity will result in just one copy of the table, which will be seen as part of the first phasor solution. For the second and later solutions, such constant output is suppressed automatically. As for KSSOUT, this is no different than for ordinary FREQUENCY SCAN (see the illustrations of DC-51 and 52). If nonzero, beware of voluminous output.

Success superimposing two frequencies on a multiphase Pi-circuit occurred November 12[th]. A new 6[th] subcase of DCNEW-26 documents this higher-frequency (400 and 800 Hz), balanced, multiphase use that involves that famous 3-kHz line section from DC-3. One end is excited by balanced, 3-phase voltage sources while the other is connected to ground by 3 equal R-L branches that draw plenty of current. The period is 2.5 msec, and agreement of unknown voltage nodes for the first two cycles is obvious. The amplitude of each harmonic is unity, and the .LIS file shows the following for unknown voltages:

```
Step   Time    TRANA     TRANB      TRANC
  0    0.0    .6411782  -.655262   .0054071
100   .0025   .6410699  -.655039   .0052891
200   .0050   .6410606  -.655017   .0052750
```

Nonlinear reactors having type code 93 (true nonlinear)

or 98 (pseudo-nonlinear) were made compatible with superposition on November 16[th] and 18[th]. Prior to this, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, included such elements in his simulations, and seemed to be producing results that were useful for engineering. However, if one looked closely at time steps 1 and 2, discontinuities were obvious. Mr. Hevia's use of AVERAGE OUTPUT was an indication that initialization was imperfect. Well, no longer. A new 9[th] subcase of DCNEW-26 demonstrates a smooth start despite high harmonic content in a nonlinear reactor. Actually, there are two of them. There are two independent, parallel circuits, with the Type-93 model used in one and the Type-98 model in the other. Both produce identical answers. Of course, the reactors are not saturated (saturation is a different problem that has nothing to do with multiple-frequency excitation, and is not helped by it).

A comment card was allowed immediately before a declaration of phasor superposition beginning November 20[th]. Before that, there was no problem if all comments were destroyed (NOCOMM = 1). But the user who retained his comments, and placed one immediately before a declaration for superposition, was treated badly: ATP failed to recognize the declaration.

The burden on List 13 did not correctly reflect use of SUPERIMPOSE prior to correction on November 29[th]. In fact, usage was not being considered at all. So, if there was a plot, it was only the burden of plotting (the original meaning) that was reported. If there was no plot, as for the 9[th] subcase of DCNEW-26, the undefined indicator -9999 was reported. Following correction, this changed to 24.

An error in the initialization of JMarti line modeling for a second or later subcase that is unexcited was removed December 3[rd]. Although unrelated to use of superposition, the error was discovered during such work. For that single-phase example in the 25[th] subcase of DC-9, answers changed only very slightly. The first output variable is a line voltage, and it now has maximum 713185.9 (compare with 713186.0 yesterday). A change to SUBR1 provided correction.

Stu Cook Uses Apple Macintosh

STATUS='OLD' was added to the OPEN used by $INSERT as a result of Absoft testing. Previously, there was no protection against a missing file. In effect, without an explicit STATUS declaration, UNKNOWN was assumed, and this resulted in successful opening of an empty new file rather than an error as intended. Other compilers may well behave comparably (thus far, the statement remains useful). So, the change was made to CIMAGE on December 11[th] as changes from Mr. Cook were processed in Portland. This prefaced retranslation.

Absoft is the first compiler not to return exactly unity

as the cosine of zero. Apparently all other compilers of current usage treat zero as a known special case, thereby avoiding possible numerical error. Absoft does not, as explained in E-mail from Absoft Tech Support dated December 8th. Mr. Cook had inquired about the problem as follows: *"... the DCOS function doesn't return a true value of unity. Rather the value returned is one bit less than unity. The following program demonstrates this ..."* The Absoft response was both helpful and reasonable: *"The PowerPC doesn't have built in hardware transcendental functions. DCOS is calculated by a software routine supplied in the runtime library. .... The runtime library SIN/COS routines are implemented as rational approximations. The test code we use (from Argonne National Labs) shows an average error of 1 bit or less. There aren't many values that you can give the cosine function that have an exact result in any case, but 0.0 certainly is one of those few."* As to why this created an ATP problem, optimal encoding was involved. Curiously, this number that is one bit short of unity failed to satisfy the D9 .LT 1.0 test that dates to the beginning of ATP. A correction from Mr. Cook involves one trivial instruction, no arithmetic, and should be independent of computer. This was installed in FLTOPT on December 11th.

Comings and Goings

The University of Pittsburgh in Pennsylvania is one place where EMTP has been used for decades. The Rule Book dated September, 1980, had a PDP-10 page near the beginning. Although only the computer specialist, Prof. J. Tom Cain, was mentioned on that page x-k, Prof. R. Gerald ("Jerry") Colclaser seemed to be the main engineering user. After 15 or more years of non-contact, E-mail from Jerry Colclaser was received September 15th. This referred to *"a graduate student in my power system analysis class. I introduce the students to ATP and would like to get the latest version."* As has happened so many times during recent weeks, BPA's Dr. Tsu-huei Liu responded by supplying login information for those password-protected Japanese storage sites on the Internet.

John Kappenman remains in Duluth, but no longer is an employee of Minnesota Power. This change was learned in E-mail from **jkappenma@aol.com** dated November 2nd. Metatech Corporation is the new company name, and circumstances are interesting: *"I have recently left Minnesota Power for the world of consulting. ... Metatech is a California firm with offices in Santa Barbara and Albequerque, New Mexico. They have about 20 full-time employees in scientific and engineering areas with specialties in difficult electromagnetic coupling problems. They were nice enough to allow me to open a Duluth office which will be developing products for power systems applications. We are currently specializing in geomagnetic disturbances: impacts on power systems and storm forecasting services. We will also be working in the lightning area in which I will be*

using ATP from time to time. My office is only about 2 blocks west and one block north of the MP offices, so no big change in that routine."

Florida Resort Course March 8 - 11

Prof. Dennis Carroll's ATP short course in Florida is scheduled for March 8-11, with faculty unchanged from last year. This according to E-mail of the Fargo list server dated December 15th. Listed faculty consist of Prof. Dennis Carroll, outside experts Drs. Kurt Fehrle and Yin Yuexin, and graduate student Carlos Mata.

Course location has changed drastically. From the campus in Gainesville, the course has been moved to *"the small resort town of Cedar Key, which is located on the Gulf of Mexico, about 50 miles west of the University. A block of guest rooms will be available for course participants at the Cedar Cove Beach and Yacht Club."* This is an environment conducive of diligent and uninterrupted study (joke)?

Price is interesting. Although the basic fee is higher, substantial discounts are offered: *"The base course fee is $1350. .... There will be a $200 discount for registrants bringing their own portable PC, and an additional $100 discount for early registration prior to February 1, 1999. University students will also receive a $200 discount."* For information, send E-mail to **dcarr@ece.ufl.edu**

Branch Data Input Restructured

An error associated with column-80 punches of branch cards was removed by restructuring of GETBUS use on October 15th. The trivial little example that was being improperly handled was MRVIC2.DAT as received earlier that same day from Masahiro Kan of Toshiba Corporation in Japan. The name of this file is a reference to the person first reporting trouble: Jovan Mrvic at the Institute of Nicola Tesla in Belgrade, Yugoslavia. The combination not being handled properly was a constant-parameter, 3-phase distributed line followed by a Type-92 nonlinear element with 4-punch in column 80. In turn, the nonlinear element this was followed by a series R-L-C branch with a 1-punch in column 80. Prior to correction, there was one output too many in the .LIS file and the .PL4 file was obviously broken (TPPLOT refused to plot because some time values were negative). To ensure continued correct operation, MRVIC2.DAT was added as a new 9th subcase of DCNEW-24.

Parameter Variation Studies by PCVP

POCKET CALCULATOR VARIES PARAMETERS (PCVP) is a request word that will result in a family of

simulations. For background, see the beginning of this story in the preceding issue. The feature evolves.

Statistical tabulation of extrema is a concept familiar to those who perform Monte Carlo simulation. As illustrated by standard test cases DC-24 and DC-40, extrema not only can be tabulated, extrema also can be saved on disk for later combination and tabulation. Well, this concept of STATISTICS studies was extended to PCVP studies on August 29th, and the 4th subcase of DCNEW-25 was enhanced to illustrate such output. There are no new rules of data input since all the old ones of STATISTICS use continue to apply. Included among alternatives are the use of FIND and EXCLUDE to locate extrema that are not the worst (the only ones revealed by the extrema of extrema output).

Batch-mode plotting was added to the PCVP data structure on August 30th. How practical or popular this will prove to be is not known, however. One problem is that ATP screen graphics can not be scrolled today as they could using Apollo workstations of a decade ago. In any case, graphics now exist, giving data the same structure as a STATISTICS data case that includes a base case solution. Plot cards are to appear after the blank card ending node voltage (and possibly other) output requests, and before any request for statistical tabulation. Provided the user has not disabled plotting with negative IPLOT, the interpretation of plot cards will be seen in the .LIS file following the output of extrema. If the user has requested a PRINTER PLOT, this output, too, will be seen in the .LIS file. But, for the CALCOMP PLOT alternative, nothing more will be seen on the .LIS file, although plots will be seen on the screen if BOTH is used. For the user who disables plotting, nothing of plot cards or the blank card terminating them will be seen in the .LIS file. For an illustration, see the ever-changing DCNEW-25.

Use of PCVP for frequency scans began September 12th. This was in response to Gabor Furst's reaction to the preceding use for time-domain simulation. Quoting from his E-mail dated September 11th: *"This raises an exciting possibility. Why not use it in HFS? We can then vary the value of a resistor as a function of harmonic number h using a user defined function."* The message from your Editor the following day began as follows: *"I have put in about 6 hours, adapting PCVP to loop over frequency. I have my first results, and will attach them. Note that I vary both frequency and inductance using the pocket calculator. I do believe there is going to be yet another type of F-scan for you to illustrate at your course in Prague!"*

DCNEW-26 entered the set of standard test cases on September 19th in order to document three different types of PCVP usage related to phasor solutions. By design, all three have comparable output. The first subcase creates its own loop over frequency whereas the second relies upon a familiar FREQUENCY SCAN declaration, and the third upon a HARMONIC FREQUENCY SCAN declaration, to do the looping. This is in addition to a PCVP declaration. Plotting as a function of frequency is permitted, of course, and the associated .PL4 file can be saved for later use by another program. For the F-scan user, the past 3 months have brought an important expansion of capability.

Supporting programs are the third class of PCVP studies. Operation began October 2nd, with only details differing from usage of the other two classes. A new 10th subcase has been added to DC-59 to illustrate use with LINE CONSTANTS to vary the height of a transmission line. About differences, there is printout control IOPCVP, which has only limited effect as initially coded. Of course, full output requires no explanation. But reduced output is not reduced as much as for the other classes of PCVP use. Only printout outside the supporting program in question will be suppressed. About the .PCH file (punched output), DC59J.PCH will contain the card images for all passes concatenated in a single file. I.e., such output is unified. As usual, comment cards document which geometry produced which set of branch cards.

Two-dimensional parameter variation is possible for LINE CONSTANTS, and probably also for CABLE CONSTANTS or CABLE PARAMETERS, if one of the two parameters to be varied is frequency. This innovative idea came from BPA's Dr. Tsu-huei Liu. Upon seeing the just-described variation of conductor height, she suggested the switch from a single, fixed frequency to a scan of frequency as used by the JMARTI SETUP fitter. It worked perfectly on the first try, and is illustrated by a new 11th subcase of DC-59. Note that variation of the raw data (geometry) is the outer loop whereas variation of frequency is the inner loop. Both the .LIS file and the .PCH file will include all results in this same order, with comment cards of the latter clearly delineating the start of each new pass of the outer loop.

Any STATISTICS or SYSTEMATIC data case is incompatible with the PCVP declaration. This observation about your Editor's programming came from BPA's Dr. Tsu-huei Liu during debugging of PCVP use with LINE CONSTANTS. It is true: at least one variable (KNT) is used by both features, but with different meanings. Is this limitation worth removing? It seems to your Editor that simultaneous use would be impractical. It would be placing too many eggs in one basket. By itself, STATISTICS or SYSTEMATIC is enough of a challenge, generally. What reader disagrees, and why?

Timing of PCVP use with a supporting program passed through simulation formulas prior to correction on October 25th. Sometimes it died on DC-59, with division by zero time. This was using BPA's shared, 90-MHz Pentium across the hall. Why the trouble was not seen at home on your Editor's 133-MHz Pentium, or on Dr. Liu's 486 at BPA, is not known. Salford DBOS is newer, of course.

If one symbol contains another, the shorter one must

follow the longer one rather than precede it. This was a new restriction beginning October 23[rd], when the illustration of DCNEW-25 was changed. This became necessary following permission of one or more blanks preceding the equal sign. When such a blank was part of the symbol, it distinguished MILLIH from MILLIHENRIES, for example. But once this distinction ended (once trailing blanks were ignored, as they should be), the shorter symbol would be located where the longer was intended. In fact, both would be found if the shorter had been declared first. ATP now checks for such conflicts, and will issue a message such as the following: "++++ Notice. Symbol 1 is contained within symbol 2. Swap these two and try again."

Use of FMTPL4 = widexx with a PCVP loop over simulation resulted in missing heading lines of the 2[nd] and later .PL4 files. I.e., the .001 file was correct and complete whereas the .002, etc. files were missing the first 3 lines. This was prior to corrections in TSHEAD and PLTFIL late on December 3[rd]. Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, first reported the problem in E-mail early that same day.

ATPDraw from Hans Hoidalen

Trailing blanks of punched cards might sometimes cause trouble with ATPDraw. This was detailed in three public E-mail messages dated October 7[th], with the first coming from Jovan Mrvic at the Institute of Nicola Tesla in Belgrade, Yugoslavia. This gentleman first provided the surprising explanation that punched cards produced by Salford EMTP have no trailing blanks (through column 80) whereas other program versions (e.g., Watcom ATP) do. Following research, your Editor explained the difference later that morning: *"Yes, Salford punched cards are missing trailing blanks (on the right). But full 80-column card images seem to be written. It would appear that DBOS or DOS is removing trailing blanks. The code clearly orders the output of all 80 possible bytes. It is proposed that source code be modified to strip trailing blanks from each punch line. Salford does not need this help, but it would appear that other program versions would be improved by it."* As no reader objected within the allotted 24 hours, the change was made. All punch files DC*.PCH created by standard test cases then were compared (before vs. after) using Mike Albert's freeware FC, and all were found to be identical for Salford.

"Statistical Outputs with ATPdraw" was the subject of public E-mail from Philippe Bergin of Schneider Electric in Grenoble, France. He inquired on October 22[nd] about possible support by ATPDraw. In his response the following day, author Hoidalen confirmed trouble handling such details: *"ATPDraw handles statistical output rather poorly. Only output of node voltages is possible with no BASE control. Special request STATISTICAL DATA is not supported. It looks like the user has to change the ATP*

file manually. I have just looked into the problem of statistical and systematic switching, since I'm introducing the concept of a master and slave switches. However, at this moment I see no immediate solution to the statistical output problem. I am planning to release a new version of ATPDraw around the EEUG meeting in Prague. This new version 1.5 will remove several bugs and introduce some more general components, smoother handling of MODELS, direct execution of ATP++ etc."

Too long a constant-parameter, distributed transmission line was a problem for Watcom ATP prior to the addition of special protection on November 19[th]. Without any indication of location, *"*ERR* KO-06 integer overflow"* was the uncivilized message that terminated program execution. This was seen by BPA's Jules Esztergalyos after phasor branch flows, so he submitted the case for analysis. Dr. Tsu-huei Liu provided explanation later that same day: *"The problem is with the …. branch cards contained in your data. … The first line mode has a travel time of 6173 sec!"* Divided by a step size of 1.E-6, the 32-bit integer limit of 2.048E6 was exceeded. Execution still ends, but now, an explanatory error message will precede a civilized halt. As background, it should be explained that *"ATPDRAW overflow"* was the Subject of the original complaint. That's right, it illustrates how data assembly isolates the user from ATP data. Mr. Esztergalyos later explained: *"ILINE in columns 51-52 defaulted to a number 1 instead of a 0 that requires an input of the surge impedance Z in column A and propagation v in column B."*

"The User's Manual for ATPDraw version 1.0 for Windows is now available on the ftp.ee.mtu.edu server." This according to author Hans Kr. Hoidalen, who issued a public announcement using the Fargo list server on December 7th. Continuing: *"The manual is on a zipped pdf format for Acrobate Reader 3.0 or later … The manual contains elementary information on how to use ATPDraw, but also information for more advanced users along with several useful examples. The sample circuits in the manual will be distributed with ATPDraw."*

A parallel announcement the same day was about ATPDraw itself. Author Hoidalen provided an impressive explanation from which the following has been extracted as a brief outline: *"Version 1.3 includes several extensions and useful updates, as will be documented below, but is not totally debugged yet. … Version 1.3 of ATPDraw for Windows is capable of reading the .MOD file (a text file starting with MODEL <<ModelName> and ending with ENDMODEL) directly, examine its input, output and data variables, and suggest a support file on the correct format. … This new module of ATPDraw is activated if the user selects a file with extension .MOD under MODELS in the Component selection menu. … Version 1.2 of ATPDraw also supports RECORD of model variables. …. Direct execution of ATP was removed from ATPDraw when introducing the Windows version 1.0. This was due to possible conflicts with DBOS. Now this feature has been*

reintroduced in a new and powerful way. The user is allowed to specify several programs to execute and what to send as parameter when calling these programs. ... All the older circuit objects of version 1.0 are supported in the new version, but some of them have been removed from the Component selection menu and replaced by other more general components. ... The general 3-phase transformer object supports both 2 and 3 windings, three types of coupling for each winding (Y, D-lead and D-lag) and the option to switch between high or low homopolar reluctance (TRANSFORMER or TRANSFORMER THREE PHASE), by clicking on a 3-leg button. The icon of the transformer object adapts the user settings. The TACS transfer function object supports both fixed and named limits. .... For the new statistic/systematic switches the concept of independent/master/slave has been introduced. ... The Type94 MODELS objects are handled in a very special way. ... Harmonic Frequency Scan is now supported by ATPDraw. Some new load models have been introduced in ATP and are supported by ATPDraw as well. A new harmonic source is also introduced. ... Handling of transposition for objects having several 3-phase nodes is improved. ... The icon of each component can be edited individually. ... The high precision option ($VINTAGE,1) has been extended and improved. ... All models are written directly in the ATP file instead of using $INCLUDE. ... TACS Device 50 has received an extra variable (max. acceptable change). ... The ATP|Settings dialog box has changed a lot. ..."

Frequency Scans and Harmonics

Nonlinear frequency dependence within HFS first was used realistically by Gabor Furst, who reported success in E-mail dated October 2nd: *"I attach the file DUMF.DAT which is a semi-realistic HFS file with some of the sources made near zero, to be able to use PCVP. I have a line from GEN to TRAN with 2 +j6.1 ohms and a load from TRAN to TERRA called RESIST. This resistance is varied by the Funk-Hantel formula. It works very nicely. I checked the results manually, and they are correct. .... I think we have something really neat now."* The following day, Mr. Furst's data was added to DCNEW-26 as a new 4th subcase. Neatness continued to improve, too. Important progress of that morning included elimination of the need for all harmonics (note that only numbers 1, 3, 5, 7, and 11 are involved). That afternoon, Mr. Furst's request for polar output, POLAR OUTPUT VARIABLES, was activated (previously, any of the three alternatives to magnitude-only led to a later error).

Supercaps (i.e., super capacitors) possess frequency dependence that might sometimes require representation. This detail from the cutting edge of harmonic analysis was learned from Gene Porter of Power Profiles, Inc. in Knoxville, Tennessee. The following summary was received in E-mail that was dated October 5th: *"If you have some time, visit* **http : / / www1.psi.ch / www_f5_hn** */Electrochemistry/supercaps.html to see the frequency response of a few supercaps. A practical example we investigated was the use of supercapacitors on the dc side of a pwm drive to increase the ride through capability of the drive during sags and such. A boost regulator was used in conjunction to further increase the ride through capability. The simulations showed that it worked but we were skeptical about the model. In the end we suggested that a frequency dependent model of the capacitor was necessary for accurate analysis. My portion of the work with Ferraro Oliver and Associates & Technology Insights ended and I don't know what has happened since."* Well, ATP would seem to be ready. The founding father of HFS, Gabor Furst, assured Mr. Porter that *"we have now the ideal tool for such frequency dependent simulation. The ATP PCVP is now working fine. I will write to you in more detail ...."*

Limits of LINE MODEL FREQUENCY SCAN (LMFS) were removed October 26th following a report of trouble from Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. For background on LMFS in ATP, see the January and April, 1992, issues of this newsletter. Two sorts of unintended and unprotected limits were exceeded by Prof. Kizilcay's data, which used 100 points per decade. This was an industrial-strength scan never before attempted, it would seem. The two limits, exceeded for the first time, were for: 1) local storage of floating-point numbers; and 2) rows of the LMFS comparison table. Changes were confined to module LMFS11, with local storage replaced by the List-28 (floating-point) storage of MODELS, which henceforth is prohibited to LMFS data. It is unlikely that this large cache (126K words for LISTSIZE.BPA) ever will fill, but if it does, an error message and program halt should protect the user. As for rows of the output table, they now are unlimited.

The warning message about nonlinear elements being either linearized (for inductors) or ignored (for resistors) during a phasor solution should be familiar to all. Designed to precede the output of a single phasor solution, it was found to be highly disruptive and annoying as part of repeated output of FREQUENCY SCAN or HARMONIC FREQUENCY SCAN. The problem first was reported by Prof. Mustafa Kizilcay. So, for these cases (FS and HFS), the message was suppressed on October 30th.

The output of current in an unexcited branch was reported to be missing by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. In E-mail dated October 2nd, Mr. Hevia reported: *"The case is an HFS with a filter for the 7th harmonic isolated by a set of switches with T-close = 1.0 ... Yes, gtpplot fails. ... The Salford and gnu djgpp versions produce the same result. If you delete the 7th-harmonic filter branches, the problem disappears. I think that the same error is in the .LIS file."* The code was corrected later that same day, and an unexcited branch was added to the 3rd subcase of DCNEW-21 to make sure such branches continue to be handled

correctly in the future. The trouble was traced to conflicting demands of the new output by columns (like the time-step loop) as opposed to the original (early '70s) output that was limited to branch flows. Today, one can have both, of course. The original program output preceded not only all frequency scans, but also the use of computer monitors at BPA. Real paper was being printed, and branches that carried zero current represented considerable waste. So, such output was suppressed. While no problem for the table of branch flows (positive KSSOUT), this omission upset recent tabular output of the .LIS file, and .PL4 files of all types (the error was in fact independent of plot file type). There is not yet any provision for the omission of zero signals of the .LIS or .PL4 files. If any reader can imagine a workable procedure, the idea will be considered. Meanwhile, FS and HFS solutions will disable the long-established omission of output of unexcited branches. If new DCNEW-21 is used with an old program, execution will die in plotting due to the inconsistent .PL4 file. Also, the column heading of (NONE, TERRA) will have no numbers below it (except for extrema, which come from different code).

Frequency dependence of saturable TRANSFORMER models first was demonstrated by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. His E-mail dated November 13[th] had *"Subject: Funk-Hantel functions for parameters."* It explained: *"I tried a case with the Funk-Hantel functions for transformer parameters (resistance and inductance). I .... applied the parameters to the TRANSFORMER fields directly, not to external branches. I changed only the primary branches. The case ran OK. The transformer parameters can be varied as a function of frequency in a very easy manner. This PCVP is a wonderful tool, if one knows how to use it!"*

Quotation marks may be used to distinguish EMTP symbols from data symbols defined by the user. This was an early decision associated with $PARAMETER data of the Powell pocket calculator as explained in the April, 1997, issue. Quotation marks should continue to be honored, but are they necessary? Your Editor thinks not. Although some risk of false identification is involved, ATP should be capable of telling the difference between the user's data symbol names and EMTP symbol names. All user symbol names will be searched first. Only if a name has not been defined by the user will ATP tables be searched for it. Quotation marks ceased to be a requirement, and were removed from the 1[st] subcase of DCNEW-19, on November 28[th].

Changes to Universal ATP Graphics

PostScript output changed in small ways September 28[th] in response to experimentation by Robert Meredith of the New York City area. In E-mail dated September 20[th], he supplied changes to three modules. Most adjustments are related to the spacing of text: *"It works better with lower case descenders ... I found that there is an inconsistency in how plotted title lines are spaced and how the legend lines are spaced."*

The header of a widexx .PL4 file might be mildly corrupted if load flow (FIX SOURCE) is involved. Attached to E-mail dated October 26[th], Prof. Mustafa Kizilcay of FH Osnabrueck in Germany supplied disk file DC25_1.DAT to illustrate the problem. An extra, unintended, and corrupted 4[th] heading line of FS or HFS, not being used by the data, was being produced due to a variable conflict between KTAPER in FXSOUR and MODHFS used elsewhere. Correction was made later that same day.

MODHFS was mentioned in the October issue as a possible appendage to .PL4 files. On December 12[th], there was a modification to the circumstances under which such information is appended. The creation of a .PL4 file now is required, although one or more comment is not for the case of simulation (as opposed to frequency scan). The first to request such improvement of rules was Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, although nothing was done until Massimo Ceraolo of the University of Pisa in Italy wrote with a question related to more work on his own plotting program PLOTXY (see the July, 1998, issue).

Year 2000 Compliance of ATP ?

The alleged year-2000 problem, popularly referred to as *the Y2K problem*, first was mentioned in the October issue. Since then, several collaborators have reported nuisance inquiries that reveal the mindlessly stupid fashion with which some (most?) bureaucracies seem to be responding to the alleged threat. The humor continues. At the risk of overlooking some good report, several that come to mind will be summarized.

First, Robert Meredith of the New York City area shared his own response to a Y2K inquiry in E-mail dated October 5[th]. An inquiry about Y2K compliance of Watcom ATP had been received from the headquarters building of Northern States Power Company in downtown Minneapolis, Minnesota. A questionaire was attached, and prior to the Y2K questions, there were others such as: 1. *"Briefly, what does the app do? 2. Is it NT 4.0 compliant? If not, when is an NT 4.0 compliant version coming out? What is the cost? 3. How frequent are upgrades? When is the next one? 4. Is it network-compliant? 5. Is it meterable? ...."* More tolerant than your Editor, Mr. Meredith **did** answer the NSP questions quickly. He also concluded with a summary (great writing): *"The program is developed entirely by volunteer effort and the user assumes all risk of using it. If you find it unsuitable for your use, do not use it. It is probably Y2K compliant, since it does not use dates. It does calculate run time statistics by comparing elapsed times since 1970, measured in hundredths or thousandths*

of seconds. I am not personally aware of when that counter may become a problem or whether it might overflow into anything significant. The program will undoubtedly be revised thousands of times before then anyhow, so its current status is somewhat irrelevant. If it breaks, it gets fixed. Excuse me if this sounds blunt. But if your first question is 'What does the program do?' you are obviously not a licensed user of the program and may not be entitled to any more detailed response. I am forwarding your note to the user group, however. Perhaps they wish to respond further." The user group's answer: no.

Laszlo Prikler of the Technical University of Budapest shared his own response to Claudio Hernandez in E-mail dated November 12[th]. About your Editor's mention of DBOS, Prof. Prikler advised: *"If you are using Watcom ATP under NT, the hypothetical problem with alphabetical sorting of DBOS mentioned in October issue of the Can/Am EMTP News does not effect you."* But why? Your Editor responded as follows on November 15[th]: *"The DOS in a DOS window of NT is different? Maybe it shows 1998 for the date? Still, when this becomes 2000, the two least-significant digits will be 00, and these would be used to name the .PL4 file. Or would they? No, we do not have the space. We have been saved by lack of space! Yes, the year-2000 problem should be the same as the year-1990 problem. We only use one year digit, so we have the problem every 10 years. Who can forget the havoc created by the year-1990 problem (joke)? Really, all this concern about Y2K seems stupid and wasteful to me. Those who inquire typically seem not to recognize the difference between programs that require a current date and programs that do not. I have yet to hear from one person I respect who believes Y2K will bring major disruption."*

Consider some of the questions that Prof. Prikler was asked. Question 1: *"Will you certify that your product (ATP) will perform without disruption for all dates of the next century, ..."* Your Editor's answer: *"No. The infamous 1990 problem already has been explained. This might plague users well into the next century."* Question 2: *"Is there, today, an upgrade, or newer version that can be offered and that is year 2000-secure? (specify)."* Your Editor's answer: *"I think this comes closer to the real reason so many persons have Y2K interest: upgrades can be sold."* Question 3: *"Will you describe the consequences of the year 2000-problem?"* Your Editor's reaction: *"The person who asks the questions does not seem to know what the problem is. He does not even know if there will be a problem. So, he is easily frightened into purchasing a Y2K-compliant upgrade."* Question 4: *"Advice on how to test (if needed) for the year 2000 problem."* Your Editor's reaction: *"These people do not know how to set the calendar ahead on their computers? Meredith should be in the Y2K consulting business."*

Sellers are not the only ones who exploit Y2K hysteria, however. Some (probably many) buyers are profiting from the opportunity as well. Y2K readiness seems to be a

nearly-perfect justification --- just one step below national security. Consider the following explanation, which was received November 19[th] from a source who probably would prefer anonymity: *"Sometimes there is a few minutes' delay before the processing is done. Sun6 is only a sparc1+. I'm getting a replacement server (for Y2K 'readiness') that will be much faster, and have more memory and disk space. .... I'm also getting version 5.5 of Framemaker, so it is Y2K 'ready'. .... I also bought X-designer version 5 to put on Verdi for Y2K 'readiness'. That way we can say that the Y2K ready package is in place. If XXXX doesn't use it, that's his Y2K problem."*

"Stupid hardly begins to cover it! Company policy is to ask all vendors whether their stuff is Y2K compliant, file the response for future lawsuits, then ignore the answer and try to test it. This will be the ambulance-chasing lawyers' full employment act for the next ten years. Anyone who answers is sticking a neck out to be sued." More great writing. Unfortunately, there can be no attribution (this contributor, too, prefers to remain anonymous).

Data Sorting by Class

ATP data sorting by class first was reported to be incompatible with ATPDRAW-created data on September 28[th]. In his E-mail complaint, Prof. Juan A. Martinez Velasco of the Polytechnic University of Catalunya in Barcelona, Spain, provided several examples illustrating trouble using /INITIAL for the specification of initial conditions. *"It is also obvious that we know how to solve this. But the main problem is when we use ATPDRAW. Then the file is created by the graphical preprocessor. Of course a hand modification (or some other tricks) can always solve this problem."*

Beginning September 30[th], SZBED is used as a flag to choose between old sorting logic and new, simplified sorting logic. If the user does nothing, his data will be sorted by the old logic, of course. But if a minus sign is appended to normally-positive SZBED, new sorting logic will be used. It really is a question of old and new rather than right or wrong. Each has its own philosophy, and different limitations. Neither will handle all possibilities correctly. The new is offered as an easy way to solve the problem posed by /INITIAL and any other class of data that has no blank card associated with it. This is what complicated the old logic: data between the blank card ending sources and plot cards. Included is possible load flow, initial-condition, and output-variable data.

The new sorting logic requires that all data following miscellaneous data cards be sorted. The user is not allowed to sort only part of his data. For example, note that just /BRANCH and /SWITCH declarations are used in DC-8. The new logic will fail on such partially-sorted data. The new logic requires that all data be sorted if any is sorted. This is the key to simplicity, which allows (but does not

require) the omission of all blank cards. Blank separator cards need not be supplied by the user because they will be generated internally by the program. Such ATP-generated lines will be easily recognized by their inimitable English-language explanation.

NEWSORT and OLDSORT are requests in data for use of the new, or the old, sorting logic, respectively. The first of these is illustrated by a new 4th subcase of DCNEW-25 (the former 4th of 4 subcases became the 5th of 5 on September 30th). It is assumed SZBED is positive for correct handling of all test cases by RUN.BAT. This preserves continuity with the past. It is the way STARTUP is being maintained by the user group: no change. For additional explanation, see comments in the new data, which is a modification of Prof. Martinez's FILE6.DAT illustration.

NORUN is the switch that allows the user to inspect data after sorting, recall. This was described in a story in the January, 1995, newsletter. Use is highly recommended if problems are suspected. Remember that an ATP error message generally will be wrong if data sorting is wrong. I.e., GIGO. To understand what has happened, the user should look at the output of sorting as a complete file.

Use of Initial Condition Cards

A new KILL = 237 error termination was added to prevent the manual specification of initial conditions for a constant-parameter, distributed line when two circumstances are satisfied: 1) there is no phasor solution; and 2) the simulation time TMAX is less than the travel time of the line. Previously, execution would continue, but the solution was wrong because line history had been omitted in order to save space. Since this logic has been in place for more than 25 years (and probably forever), it is being left unmodified. But an error stop is being added November 20th to warn the user of the problem.

For more information about the manual specification of initial conditions, see the separate story about /INITIAL and data sorting by class.

JMarti line modeling did not handle manually-supplied initial conditions properly in some cases. Weakness of the logic, which dates to the early '80s ("M35." idents) was discovered October 22nd while working on partial table dumping (see separate story). This provided initial inspiration for strenuous work on initial condition cards.

List-8 storage of past history has been saved in certain situations. This first was demonstrated November 18th, before solutions were correct. Using manual specification of initial conditions for a JMARTI line, Dr. Tsu-huei Liu showed a decrease of List 8 from 82 cells to 41 cells.

Frequency in Hz may be specified on the right of the first

node voltage card of initial conditions. In the absence of a phasor solution, the program had no other way of knowing the frequency of the solution. That dates back to year one, before ATP. But ATP **does** have another way: variable STATFR of STARTUP. So, if frequency as read from the first node voltage is not positive, the value of STATFR will be assumed. Interpretation of the first Type-2 (node voltage) card will show this value if and only if the frequency field of columns 39-53 is non-blank (not non-zero). That is, the .LIS file of test cases will be unchanged. Several involve initial condition cards, but only for lumped elements, which do not require frequency. Output of these remains unchanged.

DCNEW-27 is a new data case that is being devoted to the verification of initial condition cards for different distributed-parameter transmission line models in their most basic form (without the complication of /-card sorting, for example). This began November 22nd, following success using a single-phase, distributed model. Comments of the data describe significance, as complexity is increased from single- to multi-phase, and from constant-parameter to Jmarti and Semlyen models. In all cases, just a few (typically 20) time steps are taken, and these approximate a continuation from the sinusoidal steady state. All modes are strongly excited (operation is not balanced) in order to be sure that nothing is overlooked. Smoothness of the solution is documented on comment cards for those who might not otherwise know the answer.

All linear branches were interpreted the same way prior to improvement on November 24th. Now, transmission lines are distinguished by their own interpretation that includes the display of the fourth, previously-missing floating-point number (the imaginary part of I-mk). See DCNEW-27 output for many illustrations of this improvement.

JMarti frequency dependence was not originally provided with manual ac (phasor) initialization, it is curious to note. Instead, there was manual dc (zero-frequency) initialization. Rationale behind this peculiarity has long been forgotten, and dc initialization never was used, as far as developers in Portland know or can remember. If any reader knows otherwise, or would like to make a case for the importance of true dc (as opposed to low-frequency ac) initialization, it is requested that he do so using the Fargo list server. For background of first use, consult page IONM-1 onward of the Vol. XI *EMTP Memorandum* dated 3 October 1981. Meanwhile, the dc capability has been changed to ac, and first illustrations of DCNEW-27 (see subcases 3 and 4 for single phase and 3-phase, respectively) involve use at the power frequency. But later one will find use of frequency .01 Hz in the 8th and 9th subcases, and this should be good enough for engineering. Who might disagree, and why?

Multi-frequency initialization began Thanksgiving Day. For years (see page MSPR-6 of the Vol. XI EMTP Memorandum dated 5 August 1981), the phasor solution

has allowed multiple frequencies provided each was isolated from all others in a disconnected subnetwork. But there never was a complementary extension to manual initialization. Today, there is, as first illustrated by the 6th subcase of DCNEW-27. Read data comments for details about how the program knows which branch is excited by which frequency (an important detail, obviously).

Noda frequency dependence is not involved in any of the extensions mentioned in this story, unfortunately. This is a consequence of the difference of language (C rather than FORTRAN). If manual specification of initial conditions later might be required for Noda modeling, programming in C will be required. This provides yet another illustration of the disadvantages of mixed-language programming. Today, this complication is being avoided.

Partial Table Dumping

Partial Table Dumping is an extension to Turbo Table Dumping by Robert Schultz of the New York City area. For background on Schultz's quantum leap forward, see newsletters beginning with the October, 1993, issue. Partial or selective table dumping began early in the morning of October 12th when MODELS tables were the first to be ignored for any data case that involved no MODELS usage. This was quickly followed by the addition of logic for TACS, the U.M., etc. In colloquial terms, partial table dumping is being referred to as *Schultz's afterburner*.

Schultz handled all variably-dimensioned COMMON blocks, and therein lay the problem as program dimensions increased without obvious bound. Recall the writing of Robert Meredith in the preceding issue: *"LABCOM for me is 3639004"* which covers more than 14 Mbytes. Meredith also advocated use of maximum dimensions (the removal of LISTSIZE.DAT), which means that the entire 14 Mbytes would be handled by Schultz. This would be very slow for any computer having 16 Mbytes or less of RAM. Storage that for smaller dimensions could be confined to RAM would spill over onto disk (there would be massive paging). This also might happen on a computer with much more memory (e.g., 32 or 64 Mbytes) if Bill G's top-of-the-line MS Windows NT is being used. Recall the summary of Daniel Durbak's temporary problems in the preceding issue. In an attempt to speed execution, the memory of his PC was doubled to 64 Mbytes, but even this was less than satisfying. As your Editor wrote in public E-mail dated October 8th: *"The issue is not how much RAM the computer has, but rather how much RAM Bill G is allowing the process of interest (here, Watcom or GNU ATP) to use. Maybe Meredith knows enough to control his resources under WinNT. It is not obvious that others have the knowledge."*

Partial table dumping prohibits the application of Schultz's logic to those portions of ATP tables that are known not to be in use. Although details are secret, the concept is easily understood. Why dump and restore what is not being used? Not only is this demanding of memory, it also wastes time (unnecessary indexing in RAM may be very fast, but is not free). It is more efficient to apply the programmer's knowledge to omit such unused regions of memory from consideration. Upon reading of the idea, pioneering table archivist Schultz seemed highly pleased. On October 12th, he wrote: *"I just received and read your email. Bravo. Clean and simply based on using one's brain! The only pitfall might be the special exceptions you mentioned -- in which case … Incidentally, I got to touch an actual jet fighter engine at the San Diego Aerospace Museum, fully equipped with afterburners. I have pictures of it at home. Nice experience."*

PARTIAL TABLE DUMPING (PTD) is a new special request word that is required for use of the new logic. Initially, the user is being forced to declare his desire explicitly. This is to prevent unintended use while debugging is underway. Later, if and when usage might be understood to be extremely reliable, the need for explicit declaration could be removed. One might later add a switch to STARTUP. Only time, and lots of practice by small burrowing animals (guinea pigs), will tell.

To be continued next time. A lot more already has been done and written, but publication space is in short supply.

Interactive Plotting Programs

Availability of GTPPLOT was announced by Dr. Tsuyoshi Funaki of Osaka University in Japan. *"New Gtpplot is uploaded"* was the Subject of his public E-mail dated November 3rd. The message was noteworthy because of the inclusion of free Unix: *"The new GTPPLOT is released from Mr. Hevia. This time, the Linux version is added. Also, a 32-bit version (for WinNT) is released from Mr. Kan. All of them are placed at the password-protected server, http://atp.pwr.eng.osaka-u.ac.jp/~support/common/ gtpplot"*

"Shaft fatigue estimation with GTPPLOT" was the Subject of public E-mail dated October 3rd. In this, author Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, explained that *"The procedure with GTPPLOT is very similar to TPPLOT. But the method is different. The method is described in [1]"* (an IPST'97 paper by Argentine authors J. A. Nizovoy, J. L. Alonso, A. C. Alvarez and L. M. Bouyssede). *"The authors of this paper developed an auxiliary program (named PERVIDSM) to estimate the fatigue, reading the torques from the .lis file (the ATP output) of each case. The program was compiled with constant parameters, corresponding to the machine studied, and was verified with a commercial program of a European consulting firm. .... GTPPLOT reads the torque values from a .PL4 file, and shaft data from an auxiliary file .... The shaft section data may be in metric or English units."*

PLOTXY comes from Massimo Ceraolo of the University of Pisa in Italy as explained in the July, 1998, issue. December 14th, BPA's Dr. Tsu-huei Liu and your Editor verified use of this separate interactive plotting program under Windows NT for plotting of an ongoing simulation. This development, which promises to be important as a replacement for SPY PLOT, should be the subject of a separate story in the next issue. It also should justify a new, improved .PL4 file format.

Miscellaneous Intel PC Information

The Asian financial crisis, now a year and a half old, has resulted in dramatic reductions in prices of PCs in many other countries such as the USA. High-end workstations of unprecedented power have become amazingly affordable. Consider the following October quotation to BPA by a local company named Affordable Technology. For $1818, one can purchase a 400-MHz Pentium II PC having 128 Mbytes of RAM, a 6.4 Gbyte hard disk, fancy output card (8 MB ATI EXPERT@WORK PRO/AGP), floppy disk, keyboard, mouse, 100-Mbyte Iomega ZIP drive, CD-ROM (32X LITEON), sound (card and speakers), 17-inch monitor (.26 PROVIEW DIG.), and Windows NT.

That was at the high end. For the low end, consider an offering on page 8 of the November mail order catalog from MEI/Micro Center of Columbus, Ohio. Look what $518 will buy: 233-MHz Cyrix MediaGX MMX processor, 29.5 Mbytes RAM, 14-inch monitor, 24X CD-ROM, 2.1-Gbyte hard drive, sound card and speakers, 33.6 modem, and MS-style keyboard. Amazing. The $500 Internet computer has arrived, and it is not a specialized, limited appliance at all. Rather, it is a full-power, Pentium-based PC! About other prices, a 17-inch, .28-mm Philips Magnavox color monitor costs only $289 if purchased with a computer that has none ($309 otherwise). A 19-inch version is $529 ($549 otherwise). How about a refurbished Epson Stylus 400 color printer for $100? Or a 4.3 Gbyte Seagate hard disk for $147? Computers and their parts are not worth much, these days. For more information, the catalog prominently recommends that a reader: *"Visit our new Web site at* **www.mei-microcenter.com"**

Miscellaneous Small Items

Massimo Ceraolo of the University of Pisa in Italy continues to provide exemplary support for PL42MAT (the utility that will convert .PL4 files to MATLAB format). The subject arose September 29th when Adeoti Adediran of Texas A&M University in College Station reported trouble during heavy-duty use: *"The size of the files is about 20 Meg and the number of variables is above 90."* Mr. Ceraolo's advice the following day included the following: *"Do you use the latest version of the program?*

(3.1a) You can find it at the ATP ftp sites. For large files, it is highly advisable to use wpl42mat.exe instead of pl42mat.exe ... I would be very much surprised knowing that the problem is related only to the file size. However, if this is the result of your tests, let me know."

Frequency spreading is the procedure that is used to determine possible overlap of phasor solutions having different frequencies. The operation begins at sources and expands through branches and closed switches, keeping track of which node is excited by which frequency. If a collision (overlap of two different frequencies) ever might be noted, ATP should halt with a KILL = 222 error message. That is the theory. In practice, the logic was imperfect prior to a correction in OVER8 on November 16th. Any older program will not produce a correct solution to the 10th subcase of DCNEW-26, which was created at the same time. Note clarification of the line that precedes the error message. This is new.

SSOMIT is a special request that is the opposite of SSONLY (see the October, 1994, newsletter). ATP recognition began November 6th, following the composition of E-mail to Prof. Akihiro Ametani of Doshisha University in Kyoto, Japan. The combined use of 50 and 60 Hz in Japan had been mentioned as a problem in E-mail dated November 1st, and your Editor and Dr. Liu proposed the following solution: *"What you write is true: any one subnetwork is limited to a single frequency. Thus, one can not mix 50 and 60 Hz. However, one subnetwork can have 50 Hz and the other 60 Hz. ... Thyristors must not provide a connection of the two subsystems, however. With all thyristors open during the phasor solution, this should not be a problem except for leakage paths as are provided by snubber circuits. Yes, snubbers do create a complication, but this could be avoided in one of two ways. First, one might be able to ignore snubbers only during the phasor solution by using small enough TOLMAT. Yet, this is a little risky: one would be disabling his natural protection. A better approach might be for us to ignore certain branches during the phasor solution."* For an illustration, see the new 4th and 5th subcases of DC-10. One more cell of List-7 storage is required, following the modification to recognize SSOMIT.

BEGIN NEW DATA CASE (or BNDC in abbreviated form) is supposed to appear at the bottom of a data case as a bound. This is recommended practice. Among other virtues, BNDC at the end allows recovery from KILL-code errors, providing a location for the execution of a following data subset to begin. But many program users do not bother with such niceties, and one finally was burned badly due to a recent program addition (the HFS code). Daniel Durbak of PTI (Power Technologies, Inc.) in Schenectady, New York, made public his wound in mail of the Fargo list server dated October 9th: *"None of the three versions are working at all for me when I attempt to run a good old FREQUENCY SCAN (like DC51.DAT). My simulation 'ends early' with the following message: Halt in*

OVER4. Minimum Type-14 source frequency is neither 1 (harmonic #) nor equal to power freq. Has anyone else had a problem with this?" It does seem that some sort of frequency scan was a necessary ingredient for such trouble. As discovered later that same day, standard DC-51 itself was not troubled. Lack of BNDC at the end was a prerequisite for trouble. Since all standard test cases are properly terminated, testing in Portland did not reveal the weakness. Well, October 10[th], protection was added: one line initializing MAXCRD immediately before DO 1758 in OVER1. This modification permitted ATP to execute correctly using Mr. Durbak's data.

Complexity of data gathering and validation were illustrated by a public E-mail exchange on October 25[th]. First, Kenji Takahashi of Toshiba Corporation in Japan posed a question about the induction motor constant H as used in Gabor Furst's supporting program INDMOT (see April, 1997, issue). There seemed to be ambiguity (two equations were given), and the confusion was confirmed by author Furst later that same day. He explained: *"In dynamic calculations the value of H and the value of WR**2 are unique; they are calculated the same way by all manufacturers in the world. Some European manufacturers calculate GD**2 differently, with one value being twice the other. I too have come across this inconsistency, which is the reason why I did not put GD**2 as an option into INDMOT. Your best bet is to go back to the manufacturer, and ask them for the H value. If this is not possible, the only thing you can do is to calculate H both ways and see which value is reasonable for the given machine. In most cases you will find that it is the 1.37 multiplier which gives the correct conversion."*

$DEPOSIT was mentioned in the October, 1995, newsletter, and is illustrated by use near the top of DC-6. Operation remains generally correct, it would seem, since no user yet has complained. However, one concurrence of unlikely circumstances was discovered October 9[th] as Masahiro Kan's famous benchmark case TEST_ATP was modified by the addition of $DEPOSIT use. Dimensions already had been adjusted within data (NEW LIST SIZES), but that eliminated only one of two changes that once were required prior to timing. The second was a change of KOMPAR from the usual value 4 (for program developers) to unity. An attempt to add this assignment to data led to an infinite loop within the time-step loop, when storage of plot points within RAM (240K words) had filled, and SPY was being used to flush RAM to disk. The trouble occurred because: 1) SPY symbols did not include KOMPAR and 2) the plot cache had filled. The first problem was solved October 10[th] when SPY symbols were updated for the first time in years (the defective storage dates to October of 1996). Although work was limited to EQUIVALENCE statements of BLKCOM, many changes were made. Included should be most names that have been mentioned in newsletters the past 2 or 3 years. Some old symbols were deleted while many new ones were added. The total number of SPY symbols increased from 820 to 877. But

separate from this update, operation was corrected by the addition of spelling protection to SPYING earlier that same day. After all, even if the symbol list in SPY were complete, there would remain the problem of misspelling by the user.

Zero LU6VRT will send some diagnostic output to the .LIS file rather than the .DBG file. This was rediscovered again by mistake on November 22[nd]. The differences (comparison of all DC*.LIS output with stored known solutions *.SAL) were surprisingly large, so were worked on for an hour or so. In most cases, changes were made to force diagnostic output back to the .DBG file when LU6VRT is not positive (i.e., when output is not being cached in RAM prior to being written to disk). But in two cases, non-blank lines of production output that belonged in the .LIS file had been hidden in .DBG files. The first was a warning message from Prof. Hian Lauw about U.M. use in DCNEW-16: "Warning : Unbalanced electric power network elements or sources ...." Second, there was parameter tabulation of TEPCO's Type-58 S.M. --- optional output that follows the phasor solution. The latter has been properly converted to lower case and removed from the code for reasons of multilinguality. At the time of the original installation (see the April, 1997, issue), your Editor was busy, so he postponed the conversion. Then, apparently, he forgot about the need for more work. The output of DCNEW-20 now is substantially larger, readers will find.

SMOOTH SATURATION USING TANH was less than properly illustrated by the final subcase of DC-13 prior to correction on November 14[th]. As reported the preceding day by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, labeling of the FLTINF value is truncated on the right. This is due to the truncation request of the preceding subcase. So, full, normal, 132-column output was requested at the start of Mr. Hevia's data in order to cancel the truncation.

The computation of power and energy, requested by a 4-punch in column 80, always has been noted in the heading of printed output of the time-step loop (the .LIS file). There were two extra lines of output, if applicable. But missing prior to October 25[th] was any indication of the number of variables involved. This was a trivial but informative recent addition. For example, see DC-16.

A *"permanent magnet synchronous motor driven by an inverter"* was of interest to Andrea Moratto, who submitted public E-mail from **itric.mail.abb.com** *"Does anyone have suggestions for modelling of this kind of motor?"* he asked. This was October 23[rd]. Later that same day, Prof. Chee-Mun Ong of Purdue University responded with a reference to his book: *"I have something on that in Chaps 7 and 10 of 'Dynamic Simulation of Electric Machinery', Prentice-Hall, 1998, ISBN 0-13-723785-5, that might be useful to you."* But who might be willing to contribute data illustrating such ATP simulation?