Can / Am EMTP News

Voice of the Canadian / American EMTP User Group

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

Authorized by Co-chairmen :
Dr. W. Scott Meyer, Editor
Dr. Tsu - huei Liu
E - mail : thliu @ bpa . gov
Vol. 99 – 2 ; April, 1999

Table of Contents

Salford Compiler and DOS Extender

Newer Salford DBOS is available from the Internet free of charge according to Masahiro Kan of Toshiba Corporation in Japan. His E-mail from the Fargo list server on January 4th explained: *"I found a newer version of Wdbos.386 on the web page of Salford Software ([download -software]). The time stamp of wdbos.386 distributed with dbos3.5 is 96-07-07 whereas that on the Salford Web page is 98-06-10. It is described as free. The URL is* ***http://www.salford.co.uk/*** *It worked on my desktop PC, and I plan to test it using a troubled notebook PC at work."* The first success story using the new DBOS originated with Raul Bianchi Lastra of Universidad Nacional de La Plata in Argentina, who reported use with a *"Toshiba Satellite 2500CDS, and apparently it works fine."* This report was relayed by Mr. Kan on January 4th along with the explanation that *"the old Wdbos.386 had trouble"* with this new notebook PC.

Age of the program being used finally can be documented precisely and automatically in the header of each .LIS file having width 132 columns. *"Source code date is 27 December 1998"* is the replacement for the former imprecise *"Program is no older than ..."* when KOMPAR has value zero or one. This is as seen during first testing two days after Christmas. The date now being displayed is the date of translation, not the date of linking (the latter associated with possible redimensioning). Note use of four digits for the year. Although this might appear to be independent of century-ending confusion, in fact it is not. The four digits are constructed in FORTRAN from the two digits returned by DBOS TIME8@. Yes, Y2K should be handled correctly, but Y2100 will not be. What reader worries about DBOS use 100 years from now?

NUMC0B stores the number of variably-dimensioned COMMON blocks that hold program tables for simulation. Beginning January 4th, this variable is being carried in deck LTURBO along with the original content, which was created to communicate the total size of all COMMON block sizes as part of the Schultz Revolution (see the July and October, 1993, issues). As originally codes, only the Salford translator provides the linkage automatically. Of course, effectiveness of the Schultz compression depends on the data being processed, so it is unpredictable. Overhead is only known upon conclusion, with one word needed to summarize the burden for each COMMON block. The same fixed amount of working space (currently 87 words) can be used to store the actual starting and ending pointers associated with any one common block. This is the reason for continued use of a small amount (200 cells) of List-29 table space: minimization of CALLs to Mr. Kan's C code. Robert A. Schultz's initials are associated with several variables, and overflow protection has been provided. For example, the overflow message reads as follows: *"Halt in C-like RUNOUT. Overflow JARRAY(200) at KRAS = xxx."* If any user ever notes this complaint, he is advised to contact program developers immediately.

Value -6666 for KTRPL4 became the preferred alternative on February 26th, if this works. It **does** for the first three computers tested, and most likely will for most others. Meaning of the new -6666 involves two points: 1) there is a minus sign, so any .PL4 file is to be named in parallel with the input data file; and 2) the appropriate absolute value is to be determined internally by Salford EMTP. No longer is it necessary for the user to supply this number, which indicates the disk of interest (e.g., 3 for C:). The request for the change came the day before in E-mail from Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. Somehow, your Editor thought of a solution this year whereas he obviously did not around 1990 when the original code was developed. Then, your Editor must have been too far into the DBOS forest to see DOS trees. This year, the possibility of using DOS DIR instead of DBOS FILES@ seemed obvious and natural enough. The former indicates the disk letter following *"Directory of"* in the header of DIR output whereas the latter was missing such information if there was no explicit specification of the disk as part of the name of the input data file. As initially coded, there is reliance upon the character string stated, followed by one blank, followed by one capital letter (A through Z) that indicates the disk. If Bill G ever changes this, or if there is use of some old DOS version not having precisely this, recognition should fail and execution should die. On the other hand, nothing is lost. The old values for KTRPL4 should continue to be honored as before. The new value is an extension; it should not modify the support for any other possible values of KTRPL4.

Fortran 95 first was considered for ATP by Masahiro Kan of Toshiba Corporation in Japan. His E-mail dated January 4th summarized ease of acquisition for testing: *"The web page of Salford is ... **http://www.salford.co.uk/** From the 'Download' pane, you can download the 30-days evaluation version of FTN95."* Three days later, there was a message indicating apparent satisfaction: *"I am running Salford FTN95 under Win95. It's very nice to use. I am running Plato, which is an Integrated Development Environment, from within which it is possible to compile, link and execute Salford programs. For Win3.1(x), another Salford IDE editor (SIDE) can be used. It is a Win32 application that can be run under Windows 3.1(x) (Win32S), Windows 95 and Windows NT. A more primitive Win16 version is also available. SIDE can be used with any or all of Salford compilers (Win32 and/or DBOS based)."*

Improvements to Salford TPPLOT

MODELS variables of FREQUENCY SCAN use were not properly treated by the CHOICE command prior to correction on December 21st and 22nd. This was a continuation of the work associated with .PL4 files for such data, as mentioned elsewhere. The special multi-part CHOICE display required substantial work to accommodate the single-part MODELS variables, which disrupted multi-part uniformity. For a row that mixes multi-part current and a MODELS variable, only the *Magnitude* or *Real part* labeling will be used for each MODELS variable. The one used will be the top row, whatever that may be (it depends on the .PL4 file). On the other hand, for a row with nothing but MODELS variables, the multi-part structure that is used for electrical variables is replaced by the familar, normal, single-part output. As before, variables are selected by clicking the mouse in intuitive fashion. Testing involved wide08 .PL4 files DC52AW.PL4 and DC52BW.PL4, FORMATTED files DC52AF.PL4 and DC52BF.PL4, and C-like files DC52AC.PL4 and DC52BC.PL4 files. If anyone still cares about UNFORMATTED (no longer recommended), he can try these himself. Here the "A" was for data DC52_3.DAT supplied by Prof. Kizilcay whereas the "B" involved an added branch current.

RELAY32 is the name of the illustrative .PL4 file of widexx type that can be found in the TPPLOT archive. Prior to the February 13th addition of one line to POSOUT of HARDCOPY, attempts to plot this file would hang TPPLOT immediately after connection. This was for one output part. For 2 or 4, the extension of the preceding paragraph worked well. Fortunately, no copies of the modified version were given out to the general public, it would appear. The .LIB file in D:\TPPLOT on Dr. Liu's computer was dated November 7th, and it worked correctly. Later, following substantial testing of changes, the TPPLOT archive was updated on February 25th.

BITS=15 or more for COMTRADE output failed prior to February 22nd. Change was made in response to the observation of Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. On February 19th, he observed: *"When testing COMTRADE with TPPLOT, I*

sent BITS=16, but TPPLOT aborts. I don't know if the measurements have such resolution, but TPPLOT contains a bug." Investigation was inconclusive. All who use COMTRADE are believed to use 2-byte integers for the MODE=BINARY alternative. So, internal storage was INTEGER*2, and this seems to prevent use of BITS=16. Is any user bothered by this limit? No doubt about it, a 4-byte alternative could be added, but this will not happen unless substantial practical use is anticipated. As for 15 bits, this is doable without changing any vectors. Only two new scalars were used. The trouble was traced to computation of 2**15, for which 2-byte storage is one bit short (32767 can be stored whereas 32768 can not). Replacement by 4-byte integers solved the problem as demonstrated using BITS = 15 with the 10K-step file MORK.PL4 that was mentioned in the October, 1997, issue. Both BINARY and ASCII alternatives are compatible with BITS=15, it was found. Who needs more, and why?

The new C-like .PL4 file format (see separate story about concurrent plotting) was honored in minimal fashion beginning March 7th. No, concurrent plotting should not be possible using Salford DBOS, as far as your Editor knows. That was not the reason for the work. Rather, some way to verify the new associated ATP source code was needed, so TPPLOT simply was selected as the first viewing tool. Recognition of the new file type is automatic.

The limit of 20 curves per plot was not observed prior to March 12th, when protection was added in response to a warning by Orlando Hevia. He was right, of course: tables are dimensioned 20. If 21 curves were selected, execution would die. That was prior to added protection that accepts only the first 20 requests. This can be demonstrated by @TWENTY where disk file TWENTY is a new addition to the archive that demonstrates the use of both the NAME and the # command as a means of variable selection. While it seems highly unlikely that a single, realistic, ordinary plot ever would involve 20 curves, window plotting could reach this limit without difficulty. It is important to note, however, that real plots are involved. The limit of 20 should not apply to either the RELAY or the COMTRADE command to produce alternative output files. The limit without a plot is much larger, with 200 being a promising guess. A third plot from TWENTY illustrates 32 outputs for COMTRADE (every signal).

News from Outside USA and Canada

Mexico no longer enjoys the same lower prices as the USA and Canada for ATP materials ordered from Dr. Kai-Hwa Ger. It was during the first week of March that your Editor was surprised to learn of the change from what had been explained in the October, 1994, issue. In a note dated March 2nd, Dr. Liu explained: *"Mexico is* ***not*** *included in North America as far as the ATP order form is concerned. Air mail cost to Mexico is much much higher than US / Canada."* Subsequent research by her showed that the first order at the higher, overseas prices was dated 6 October 1997. The order form had been revised on August 13th so that Mexico no longer was listed with Canada at the domestic rate. Anyone wanting to verify postal rates within the USA is referred to **www.usps.com** (USPS stands for U.S. Postal Service). Good luck. The snail mail people do not seem to handle Web sites any better than normal mail. Your Editor reached the home page easily enough, and entered the rate calculator. After selecting Mexico as the country, and 3 pounds as the weight of a package, the waiting began. Some five minutes later, another window opened to indicate "Proxy error."

IPST'99, the 1999 International Conference on Power Systems Transients, is to be held in Budapest, Hungary June 20-24. Laszlo Prikler of T.U. Budapest, the EEUG Deputy Chairman, is Conference Secretary. His progress report, issued February 1st using the Fargo list server, concluded as follows: *"I would like to call your attention to the latest update of the conference web site at address* ***http://www.vmt.bme.hu/ipst99*** *The Advance Program and the Registration Form are now available in Acrobat PDF format. The early registration deadline is March 31."*

Osaka University in Japan seems to be the present source of Japanese-language E-mail about ATP. This was learned March 13th by reading a message from Masahiro Kan of Toshiba Corporation: *"The Japanese mailing list which was served by our company's mail server was forced to stop on the occasion of the move of the mail server. Dr. Tsuyoshi Funaki kindly took over the server, and now the Japanese ATP mailing list is served by Osaka University's server."*

More about the Internet and E-mail

Security of ATP-related FTP file storage at Michigan Tech was summarized by Prof. Bruce Mork in E-mail dated January 5th. Responding to someone who had inquired, Prof. Mork explained: *"If you're a licensed ATP user, I can supply you with the password. Note that the directories license, atpdraw, and canam are not password protected."* It is understood that /license is for form letters of the various user groups, /atpdraw is for the graphical data assembly program by Hans Hoidalen, and /canam is for newsletters of the North American user group. All of the mentioned items are available free of charge to anyone, whether ATP licensed or not, it is to be noted.

MS Outlook 98 is the replacement for MS Mail that began operation on Dr. Tsu-huei Liu's Pentium at BPA after she left for home on December 21st. According to MS advertising, *"Outlook can be configured to look like the familiar MS Mail by selecting View from the menu and deselecting the Outlook Toolbar and the Folder List. This will be the default setting for new Outlook users."* According to a BPA form letter left by those who effected the change, the result might be chaos: *"If you hear of co-workers getting error messages attempting to send mail to*

you, may be because they are replying to an old message you sent or they have your old mailbox name stored in their Personal Address Book." As this paragraph is continued December 23rd, E-mail remains unusable. In fact, NT remains unusable because the password seems to have been changed. Dr. Liu was stranded at home yesterday by snow. Conclusion: the great E-mail blackout ended during the morning of the 28th when Dr. Liu next returned to work and solved the log-in problem by guessing her modified password (what are the odds of being able to do that?).

Consistent ordering of "Save as" for a highlighted, consecutive group of mail messages is the only obvious way MS Outlook seems better than the old MS Mail that it replaced. During recent months, in response to pressure from the BPA post office, this technique has been used to save all mail messages of a month in one file. For example, DECR.LIS stores all received messages for December, and DECS.LIS stores all sent messages. When first used during March on mail dating to December and January, content of the resulting files seemed to be ordered chronologically. The storage for previous months, produced using older MS Mail, was not.

The Columbian is Vancouver's principal newspaper. It does not amount to much by big-city standards, but the editorial staff certainly **does** know how to communicate with the public. A sidebar on the editorial page is entitled "TO REACH US," and the first and fourth of four ways listed are: *"E-mail: Electronic mail is best for all submissions. Send it to editors@columbian.com .... Web: Recent editorials and links to other opinion sources can be found on the World Wide Web at www.colunbian.com"* This Web address also appears at the top of the front page, immediately below the paper's name.

Yahoo would seem to have joined Juno and Hotmail as prominent suppliers of free E-mail. Recall that the latter two names were explained in the October, 1997, issue. Perhaps it has never before been mentioned publicly, but BPA's Dr. Tsu-huei Liu is a satisfied Juno customer for E-mail to and from her home. This provides a free replacement for her relatively expensive CompuServe (most recently, AOL) subscription of years past. Now, Internet giant Yahoo seems to have joined the fight for Internet customers by offering free E-mail. This first was seen by your Editor in public E-mail of the Fargo list server dated February 16th. The message from Basset M. Bendalla ended as follows, below a horizontal separator line: *"DO YOU YAHOO!? Get your free @yahoo.com address at http://mail.yahoo.com"*

Repeated attempts to deliver E-mail that bounces is worthy of note. During years past, a 4-hour message was common. This would inform the sender that attempts to deliver his message had failed during 4 hours, but that attempts would continue for 5 days. Well, the numbers have been reduced. As first noted March 11th, the new message is: *"451 <pekaric@uns.ns.ac.yu>... I/O error: Error 0. Warning: message still undelivered after 1 hour. Will keep trying until message is 1 day old."* The one hour certainly represents progress, but what about abandoning attempts after one day? One can easily imagine unavailability during most of a weekend, while there might be no maintenance.

Auctions are the specialty of ebay.com which is run by eBay Inc. of San Jose, California. According to a story on page E5 of *The Columbian* dated March 2nd: *"sellers log on to a Web site and post items they want to unload, and buyers bid for them … Currently, eBay has nearly 1,300,000 items for sale in more than 1,000 categories. … eBay was started in 1995 by California programmer Pierre Omidyar, who set up a website to help his girlfriend sell Pez dispensers. … Today, eBay is a retail giant worth more than $7 billion. More than 39 million items have been put up for sale on the site since its inception, and more than 144 million bids have been placed. Each month, 600 million people log on to peruse the site."* Question: with private parties rather than companies involved, how does a buyer know that a seller is trustworthy? Answer: it is not easy! This is a problem for Internet sales just as it is for telephone sales. *Caveat emptor*.

European EMTP User Group (EEUG)

The annual EEUG Meeting for 1998 was held November 9th and 10th at the Hotel Krystal in Prague, The Czech Republic. Organized locally by the Faculty of Electrical Engineering of Czech Technical University, the meeting was attended by 48 persons, and was followed on November 11th by a day-long ATP course on frequency domain harmonic analysis.

Local culture provides added attraction for EEUG meetings in different cities. About Prague, Prof. Kizilcay reported: *"An interesting and well-prepared sightseeing tour was organized between 09:30 and 16:00 for the participants who arrived a day or more ahead of time."*

The January issue did not contain a report on the Prague meeting because, at the end of the year, your Editor had no quotable account. This time, the story can be written thanks to E-mail from EEUG Chairman Kizilcay dated March 30th.

Technical sessions were held throughout November 9th and during the morning of November 10th. A total of 15 papers were presented in four sessions. In alphabetical order, titles were: 1) Access to a database of power systems recordings using the WWW; 2) Application of ATP to teaching of electric machines and apparatus; 3) ATPDraw for Windows, Version 2; 4) Calculation of very fast transients caused by disconnector switching in GIS; 5) Comparison of measured and modelled voltage-time characteristic of insulator strings; 6) Comparison of surge simulation by Marti and Noda line model with field test results; 7) Fundamental frequency equivalencing of series capacitors with MOVs in a

transmission line using ATP's MODELS; 8) Iterative harmonic and interharmonic analysis on multi-driver systems based on single drive time domain EMTP-ATP simulations; 9) Modelling and simulation of FACTS devices using EMTP; 10) Modeling of traveling waves reflection on cable termination; 11) Mutual impedance between overhead and underground cables; 12) Remarks on modeling of three-winding transformers using saturable transformer component; 13) Signal analysis in converter-fed induction motor drive; 14) Using ATPDraw as the operating shell of the ATP-EMTP simulation system; and 15) Voltage dips simulation in MV distribution networks by using an improved ATP circuit-breaker model.

About possible future work, Prof. Kizilcay wrote: *"In the discussion session ..., development of the user shell ATP Control Center and further development of PCPLOT for Windows were supported by the members."*

The meeting of EEUG members began at 14:00 on the second day with 23 in attendance. Elections seemed to be the most important business --- to select management for the next four years (1999-2002). *"The six members of the new Executive Board were elected by secret ballot in the following order: 1) Chairman: Prof. Dr. Mustafa Kizilcay of Fachhochschule Osnabr ck in Germany; 2) Deputy chairman: Assistant Prof. Laszlo Prikler of the Technical University of Budapest in Hungary; 3) Secretary: Mr. David Bailey of Merz and McLellan in New Castle upon Tyne, UK; 4) Treasurer: Mr. Harald Wehrend of SEG in Kempen, Germany; 5) Member: Dr. Luciano Tonelli of CESI in Milan, Italy; and 6) Member: Dr. Thor Henriksen of SINTEF Energy Research in Trondheim, Norway. All elected persons accepted the results, and agreed to cooperate …. Prof. Kizilcay, reelected as Chairman, stated his intention not to head the association after this second term of office ends in 2002."*

Never forget LEC accounting (newsletters dated July, 1993, and later)! Keeping track of EEUG money is important business if only as a matter of self-protection. *"Dr. Lutz Hofmann of the University of Hanover in Germany, and Dr. Sven Demming of BEWAG in Berlin, Germany, were approved unanimously by the members as auditors to examine the EEUG accounts for 1998."*

Some 22 students attended that short course on the 3rd day, which was taught by both Prof. Kizilcay and Gabor Furst (the latter from suburban Vancouver, B.C., Canada). Prof. Kizilcay noted: *"Interest in harmonic analysis was great. Mr. Furst introduced newly-implemented features of HFS whereas I handled older FS and its recent modifications. In the theoretical part, I summarized frequency-domain models of power system components, and illustrated the use of Frequency Scan for this purpose. Mr. Furst, as a senior consultant with enormous experience in the field of power system operation and analysis, reviewed time-domain and frequency domain methods of harmonic analysis .... Several examples were executed using a notebook PC to illustrate harmonic analysis using ATP."*

Watcom ATP for MS Windows

Optimized compilation was described in the preceding issue, which mentioned a local paging file of 100 Mbytes. This did not last long, however, it should be documented. First, there was a doubling from 100 to 200 Mbytes after 100 proved inadequate for EEUG-sized usage having 400 coupled coils. That was some weeks earlier. But even this staggering overhead did not represent the end of the story. December 4th, working on yet another expansion (this one to allow 3-phase representation) for Bernd Stein of FGH in Mannheim, Germany, the 200 Mbytes were exceeded. First, the DOS window was closed and reopened, but this did not help. Next, WinNT was shut down and restarted, but neither did this cancellation of past memory leaks help. So, BPA's Dr. Tsu-huei Liu again expanded her virtual memory of disk E: (which includes the virtual D: that is used for Watcom linking in D:\WATNT), this time from 200 Mbytes to 250 Mbytes. Linking was successful. But think of the waste.

The .PL4 file for FREQUENCY SCAN having 2 or 4 output parts (as opposed to one part for magnitude only) and MODELS variables should be mentioned later in this issue. That correction applied to Salford EMTP only, however. Then, on January 13th, Prof. Mustafa Kizilcay of FH Osnabrueck in Germany reported comparable trouble using a new Watcom version. So, Watcom modules were massaged by your Editor and Dr. Tsu-huei Liu January 14th and 15th. Use of RECL=1 for the Watcom C-like .PL4 file necessitated substantial changes to Salford ATP, unfortunately. A full day's work was involved, and new WATAIL was created in the process.

Why do we continue to struggle this way? This was the question that your Editor posed of Dr. Liu upon conclusion of the work of the preceding paragraph. Instead of such repeated work to support Watcom ATP, why not instead buy a new Salford compiler for MS Windows and drop support for Watcom ATP? Recent experimentation with F95 by Masahiro Kan of Toshiba Corporation in Japan has provided added incentive for such thinking. Note how times have changed. When Robert Schultz selected Watcom around 1994, it offered the advantage of linking for a variety of Wintel platforms that included the then-preferred IBM OS/2. Since then, interest in OS/2 seems to have disappeared (see the July, 1997, issue), and a single executable version for 32-bit MS Windows should satisfy all Wintel users. So why the continuing struggle to maintain Watcom ATP?

Variables NPOWER, MODHFS, and IHSPL4 were added to C-like .PL4 files for the first time on January 15th. This was in order that Salford TPPLOT or any other ATP-aware plotting program could make intelligent

decisions about variables such as time (or frequency), power and energy, and multiple output parts (possible for frequency scans of any type).

The .PL4 file was being destroyed unnecessarily after an error was detected within the time-step loop. This problem first was called to the attention of developers in Portland by Massimo Ceraolo of the University of Pisa in Italy. The subject of his E-mail dated December 18th was "Auto-killing of pl4." His data encountered switching trouble: "Logic trouble within SUBROUTINE SWITCH of overlay 16." Without concurrent plotting using his PLOTXY (see separate story), Mr. Ceraolo might not have thought anything of the missing .PL4 file. But prior to the error, he had been plotting signals, so he knew that the .PL4 file had existed, and that Watcom ATP somehow must have destroyed it during the termination. Mr. Ceraolo's complaint was verified in Portland, and correction was made easily enough 2 days later. Before, the .PL4 file was trashed if it had not been connected by the user in data (using $OPEN). Now, both this condition and zero value for miscellaneous data parameter ICAT are required for deletion during an error termination.

The new 3rd of 3 subcases of DC-41 is mentioned elsewhere in this issue. After successful testing first with Salford, and second with GNU djgpp, it was a surprise to discover that Watcom failed to connect the disk file DC41C.PCH as required when first tested on February 8th. Yet, a data error was involved: the marker preceding the in-line comment was missing prior to correction that same day. What did Watcom do? The DOS command DIR DC41* revealed a disk file having a name that ended with the comment text (*"Output the just-created branch cards to a disk file"*)! That's right, embedded blanks and all were part of the disk file name. This is a problem about which the Watcom ATP user had better be aware.

New, exceptional disk files RUNCT46A.BAT and RUNCT46B.BAT will be seen within RUNCTACS.BAT beginning February 9th. The CT or CTACS refers to compiled TACS, of course, and the 46 refers to the same DC-46 that has been troubled in the past (see the July, 1997, issue). Restoration of C-like DC45.PL4 is the reason for the change. For some still-unknown reason, this .PL4 becomes corrupted during use, so the second or later use requires refreshing the .PL4 file using a copy that was set aside as usage began. Not nice, but such is life using Watcom ACCESS='DIRECT' to provide C-like .PL4 files until some C-language programmer might switch to real C for Watcom ATP as Masahiro Kan did for GNU ATP (see the October and January issues).

C-like .PL4 files for Watcom collided with in-line use of supporting programs (see mention in separate story) the first time this was tested. Such use is illustrated by the new 3rd subcase of DC-41. Not wanting to modify data specially for Watcom (one alternative, to add $DEPOSIT use), it was decided on February 9th that accommodation instead would be made in Watcom REQUE3. The user will note that his output file is UNFORMATTED rather than C-like, and will have the fixed name PLOT.PL4 regardless of what the user might have intended. Later, parallel naming might be provided, if there is sufficient demand. But for now, fixed naming seems adequate. The user easily can avoid the internal change by using an UNFORMATTED .PL4 file rather than a C-like .PL4 file, if the distinction is important to him.

START46 is a special copy of STARTUP that is used to allow all subcases of DC-46 to complete normally. The name will be found within Watcom RUN.BAT as modified on February 12th following about one full day of debugging by BPA's Dr. Tsu-huei Liu and your Editor. Corruption of input DC45.PL4 was occurring within Watcom CLIKE4 as plot data was being written to unit 4 on the final time step. Yet, the input file was connected to unit 63 (no mistake about this)! The Watcom compiler seems somehow to confuse the two random-access uses --perhaps due to optimization as summarized in the January, 1998, issue? Confusion can be removed by changing data, however --- by using wide08 for the .PL4 file of DC-46. This is the reason for the special STARTUP file: a change from C-like to wide08 for DC-46 only.

Recall TSTALL, which is supposed to slow output as first explained in the July, 1998, newsletter. For some unknown reason, use was found to hang execution for Watcom ATP using Dr. Tsu-huei Liu's Pentium at BPA. Internally, there was reliance upon SLEEP.OBJ which seems to have been produced from SLEEP.ASM (Intel assembler) of Robert Schultz. After confirming the problem a second time on March 5th, CALL SLEEP was replaced by the universal, time-wasting, nested DO loops that preceded SLEEP. The internal loop limit of 100 was replaced by 500 to reflect progress by Intel during recent years (Pentium speed as opposed to 486 speed).

The new C-like .PL4 file of DCNEW-28 first was demonstrated to be compatible with Salford TPPLOT on March 14th. The following morning, both Watcom TPBIG and an archive of Salford DCN28.* was sent to Pisa. Even though data is small, a variety of signals is involved: node voltage, branch voltage, branch (actually, switch) current, branch power, branch energy, and TACS.

DEC ATP for VAX / Open VMS

The death of DEC VMS ATP seems to have been prematurely assumed (*"The DEC ATP ship is sinking"* can be found in the October issue). The first response to the call for help in the July issue was received on January 6th: a long E-mail message from Stephen Boroczky of Network Development (*"note new section name --- yet another effect of restructuring"*) at TransGrid in Sydney, Australia. Mr. Boroczky began with an explanation for his

delay: *"Having taken the opportunity over the Christmas break to catch up on my reading, I have discovered that your call for help on VMS has gone unanswered."*

History seems to be what dictates continued use of VMS by TransGrid. Mr. Boroczky explained: *"The VMS version of ATP is our main production version of ATP. The PC version only has a minor role in our organisation. This is because data preparation and post processing is normally done with our own in-house software on VMS."*

Technical expertise of TransGrid should be guiding further enhancement of VMS ATP. As your Editor summarized in his response on January 10th, the former DEC code was close to being abandoned: *"There has not been any significant VMS ATP development for several years. The departure of Randy Suhrbier ended all real development. During recent years, all Dr. Liu and I did was test a new VMS translation once or twice a year, typically. That was for Bob Hasibar. But with Bob's passing last July, all interest seems to have ended. … We might never have tested another VMS translation, had you not responded. No one else, either inside or outside of BPA, had expressed any interest in the continuation of VMS ATP."*

Two separate and different possible VMS solutions to the detection of ATP environment variables were explained in detail by Mr. Boroczky: *"Symbols are most like environment variables and you would use calls to LIB$GET_SYMBOL and LIB$SET_SYMBOL in your FORTRAN code to get and set symbols. ... The syntax is: .... Logicals, on the other hand, are more intimately tied to the operating system than environment variables are in either Unix or Windows. A logical would be the ideal way for you to set up the ATPDIR functionality. ... In fact, in the LEC days, we would reference the STARTUP file in an external directory by just defining a logical STARTUP to point to the startup file. ...."*

Transformer Tap Changing

This story is drawn from three E-mail messages of Prof. Bruce Mork's Fargo list server between November 4th and the 5th. It serves to outline the procedure used to model tap-changing transformers for purposes of ATP simulation.

The original inquiry was from Sedlak Alex, who explained that *"the model can be quite rough, (i.e., no saturation) as the objective is to investigate performance of the tap-changing controllers. I've been through the ATP Rule Book, but haven't been able to find anything quite right to do the job."* True, true. Furthermore, prior to this story, newsletters were not of much help, either. But good advice quickly was provided by Christian Collombet of Schneider Electric SA in Grenoble, France. He provided the following good outline: *"I think I would do the following things: 1) connect in parallel as many transformers as taps; 2) drive the connections of the transformers with controlled switches; and 3) control the switches with MODELS for example. Don't forget that for each tap the transformer has a ratio, a resistance, and a reactance."* Finally, your Editor responded, confirming in some detail the advice of Mr. Collombet. The remainder of this story is drawn from that third and final communication on the subject.

The modeling of a tap-changing transformer in EMTP is a very old problem for which there is a very old solution: switching. Unless and/or until someone can explain why switching might be inadequate, this writer is inclined to do nothing using code. A special model could be developed, but it seems to be unnecessary. Until proven otherwise, data should be adequate.

About newsletters, a special case, representing a degenerate application of the general technique, was found in the July, 1997, issue. Interested readers are referred to the story entitled *"Transformer Fault Simulation,"* which was inspired by the work of Gabor Furst of suburban Vancouver, B.C., Canada: *"The winding to tank (ground) fault can be modeled by a simple autotransformer, with its lower voltage tap being the faulted point in the winding. The transformer can then be modeled using BCTRAN (or the regular saturable transformer model)."* This is a special case in that a single switch was used, and no TACS or MODELS logic was required to control it. For a tap-changing transformer, one would have as many windings as there are taps, plus two.

Perspective of the user is interesting. One has dealt with equations of tap changing for so long that one thinks of the component as if the transformer changes. In fact, parameters remain fixed. It is the tap that changes, and this is modeled by switching that is outside the transformer itself. Anyway, this is the proposal: have the ATP model mimic the physical device rather than its mathematical equivalent. As for any switching of an inductive circuit, it is recommended that an old contact not be broken before a new one is established. Dependent switches come to mind as a possibility, but should not be needed because of the simple nature of the switching that approximates movement of the tap. Unlike the case of thyristors, one should always know which switches are to be open, and which are to be closed, and any time. This would be the logic of TACS or MODELS, and it should be simple enough. Associated switches would be TACS- or MODELS-controlled.

The stated desire for a crude model means that the saturable TRANSFORMER component might be used. This is easier than using a supporting program such as BCTRAN, with its output that is a full matrix. It should be both easier and clearer to use the saturable TRANSFORMER since data is out in the open and easily changed. The turns ratio is obvious by inspection, and an essentially-ideal transformer of any number of windings

can be constructed easily. Of course, if there are 98 taps, one will be using 100 coupled coils, but this should not be a problem. Dr. Hiroshi Arita of Hitachi Ltd. in Japan already has reached that huge 400-coil limit that was mentioned in the April, 1998, newsletter, and this would have roughly 16 times the burden. One hundred coupled coils is nothing, these days! By using the saturable TRANSFORMER, one keeps the burden of data linear (100 input lines), and he avoids full matrices (10000 nonzero entries). Data input should be much easier. It almost seems as if one is obtaining something for nothing. Is this **too** easy (*"If it sounds too good to be true, it probably is"*)? If so, what is the complication or contradiction?

Concurrent Plotting of PLOTXY

PLOTXY is the interactive plotting program from Massimo Ceraolo of the University of Pisa in Italy. For background on this software that runs under MS Windows, see newsletters beginning with the July, 1998, issue. PLOTXY has taken on added importance recently because of its ability to display ATP waveforms as they are being generated. This is equivalent to SPY PLOT of Apollo days. Yet, there are important differences. As noted briefly in the January issue, using generic terminology, this allows concurrent plotting of ATP simulation without the need for SPY. Just as for Szymanski's independent, parallel execution of SPY a decade ago, the operating system provides the time-sharing (this was not the case with Apollo). Of course, the operating system has changed over the past 11 or 12 years: (Ceraolo uses MS Windows whereas Szymanski was using Unix System V, Release 3).

$OPEN with STATUS=APPEND (for an illustration, see DC-49) or OVERLAY has been modified for C-like .PL4 file usage. This was the first change to .PL4 file practices, and it was made January 8th --- prior to any work on the new alternative format. The reason was to inform any concurrent plotting program such as PLOTXY that simulation is ongoing as opposed to complete (the conclusion based on DC32.PL4 prior to connection for more simulation using DC-49). Changes are well-commented in the Salford version of CIMAG2 (used for experimental C-like handling of LUNIT2 prior to later use by Watcom or GNU versions).

A new .PL4 file type was created to allow PLOTXY and all other plotting programs better access to information about an ongoing ATP simulation. The detail of the preceding paragraph was just the first of many that now are different. As suggested by Mr. Ceraolo, the user sets NEWPL4 = 2 to request the new alternative. Only C-like is allowed for this case, and this is automatic, regardless of other parameters in STARTUP.

PISA.TXT is the disk file that documents details of the new .PL4 file type. It is secret in that informational content must not be disclosed without permission of the author (your Editor), and a copy is available only to one who has a need to know. A copy was to sent to Pisa on March 10th along with a compatible version of Salford TPPLOT (see second story). A compatible version of PLOTXY then was received in Portland five days later (talk about speed)! More next time, when there is more space.

Higher - Order Pi Circuits

ATP data for 400 coupled coils is available from the secure FTP storage at Osaka University in Japan. In a message from the Fargo list server dated March 12th, Dr. Tsuyoshi Funaki announced availability of disk files from Dr. Hiroshi Arita and Mr. Takahide Matsuo of Hitachi, Ltd. Two Web addresses were given, with each one ending in … /from_japan/bnchmark As Dr. Funaki explained, such innovative data represents yet another, and a new sort of, contribution from Japan to ATP users everywhere.

Prior to control of all dimensions, use of NMAUTO = 1 resulted in overflow of List 7 that was not properly protected. NAME03 was location of the addition.

Brain - Damaged MS Windows

Trouble with GNU POCKE1 is mentioned elsewhere in this issue. The first time trouble with DC-22 was investigated by your Editor with Dr. Liu watching, the DOS window quickly disappeared and an error window mentioning Dr. Watson appeared. Ten or so seconds later, your Editor made the mistake of clicking on the icon of the DOS window at the bottom of the screen. The screen then quickly went black and the keyboard became locked (no response to **Alt-Tab**). Only the hardware reset button seemed responsive. Not a nice experience, this WinNT termination of erroneous ATP execution. The second and later time (as the error was investigated using special diagnostic printout), your Editor modified his reaction. By quickly clicking on the "cancel" button within Dr. Watson's error window, the DOS window was rapidly recovered, and rebooting was avoided. The first time, maybe NT choked while trying to save to disk that 250-Mbyte swap file that was necessitated by Watcom ATP linking? What a pain.

The Windows NT paging file itself --- a disk file --- might be corrupted. It was December 27th that BPA's Brian Furumasu offered this suggestion as your Editor described recent troubles with Watcom ATP linking. Mr. Furumasu casually asked whether the paging file had been deleted after NT was restarted. It had not. But the idea seemed promising, so was not forgotten. When the next trouble with Watcom linking occurred on January 12th, BPA's Dr. Tsu-huei Liu tested the idea, and it seemed to work. I.e., for the first time, linking was successful with a **smaller** paging file than the one with which linking had failed. Previously, sizes always were increasing.

The term *memory leak* first can be found in the April, 1997, issue. Finally, news of recognition and correction by MS has been received. The following is from MS-certified computer expert David Szymanski, who sent E-mail from Internet address **david@szy.com** on January 13th: *"I have also been re-reading our correspondence from July of 1997 .... The Microsoft Knowledge base returned over 300 articles when I searched for 'memory leak'. Most of the articles describe programming tips to follow so your software won't create 'memory leaks', but, the following article shows they can also be caused by running older 16 bit programs."* Yes, the article, *"last modified on 06-01-1998"* (in this country, June 1st) describes the symptoms well: *"In cases where 16-bit applications are launched repeatedly, system resources may be depleted because of a WOW memory leak. Symptoms of the leak condition will be popup error messages indicating 'low virtual memory' and potential system crashes if memory resources are severely depleted over a long period of time. The problem can be recognized by monitoring MemUsage under Windows NT Task Manager or by running Performance Monitor ...."* But what about the solution? The article states: *"To resolve this problem, contact Microsoft Technical Support to obtain the following fix, or wait for the next Windows NT service pack because there are several dependencies necessary to completely resolve this issue."* Szymanski also provided advice about disk fragmentation, although this seemed to be a minor problem compared with fatal memory leaks.

Inability of WinNT to control DOS was mentioned in the OS/2 story of the January, 1997, newsletter. More than two years later, the operation has been observed again --- this time using a different floppy disk, and with no obvious relationship to Watcom. After copying 4 files onto a floppy using the real DOS of Dr. Liu's 486, Mike Albert's FC would hang in a DOS window of her Pentium running NT when A: was accessed. This was February 11th, and your Editor repeated the trial several times. Even "DIR A:" would hang a DOS window, so Mr. Albert can not be blamed, obviously. Normally the NT Task Manager (which appears in response to **Ctrl-Alt-Del**) could kill the process, but eventually even this failed. The floppy disk light remained lighted continuously, so the hardware reset button was pressed. Subsequently, it was discovered that removing the disk from the drive accomplished the desired function, however, after a substantial delay (maybe 10 seconds). This was found to be the most efficient way of killing the task: simply push the mechanical button of the disk drive. When this was done with the Task Manager window open, your Editor watched the task disappear automatically. Additional information the following day: this time, editing and copying of files at home using Win95 did not correct the disk for NT use. The first attempt to use the floppy involved saving of a mail message to disk using new MS Outlook 98. This resulted in an error window with the message: *"The floppy disk in drive A: is not formatted or has been formatted for a Macintosh."* Yeah, right (what a system, this top-of-the-line OS from Bill G!).

"Here NT can be interpreted as No Trabaja (Not Work)." This insight about NT use in Spanish-speaking countries came from Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. In E-mail dated February 21st, he also mentioned the well-known trouble that General Motors had selling its Chevrolet Nova down there: *"(No va = not go) --- bad name for a car!"* Indeed! Nova, in turn, brings to mind Radio Shack, which was purchased by Tandy. During the early '80s, Tandy Radio Shack offered a PC having the name TRS-80. Although a decent machine, it became known as *Trash 80* among PC users of the USA (not even a foreign country). As Mr. Hevia observed a couple of days later, the search for a name that is politically correct and not negative somewhere in the world must be a difficult job.

The MS Word file of this newsletter can be carried from Win95 at home to WinNT on Dr. Liu's Pentium at work, but not vice versa. This is the latest of Bill G's planned obsolescence, as discovered using the January issue. Normally, everything but the final conversion to PDF is done at home, so it is not necessary to take a .DOC file home. But around the end of the year, the need for minor change was discovered at work immediately prior to the PDF production. So, changes were made, and the file was saved as usual as a .DOC file. But the result was huge, for some unknown reason:

```
… 290,304  12-30-98  6:08a  JAN99.DOC
```

Worse, a copy was found to be incompatible with Word 7 at home. While the file could be read, this was only as text; it was refused as a .DOC file. Later, when your Editor eventually found the time, he compared MS-DOS text output of the two programs, and adjusted the .DOC file at home to match the one at work. Prior to renaming, the result at home ended up with believable size as follows:

```
… 149,504  02-28-99  5:31p  january.doc
```

Are there any other good reasons not to use Windows NT? What a pain, this continuing incompatibility of MS with itself. At least this no longer is a problem of the reader, who was switched from MS to Adobe PDF for other reasons (see a story of the July, 1998, issue).

A "Reply" to E-mail from William Bloethe of Sargent and Lundy bounced March 1st. Apparently Bill G's Outlook 98 seems allergic, for some unknown reason. Origin of the incoming message was clear enough. As saved to disk, this was: *"From: William G. Bloethe [bloethe@wwa.com]"*). But somehow only 'bloethe' was extracted by MS Outlook 98 to address the reply. Curiously, the missing remainder was mixed with a BPA address to produce a defective address at the end of the complaint from the post office: *"<wwa.com@bpa.gov>... User unknown"* The experiment was repeated 3 times over several days, with consistent results. A "cc" always was added, using the same address as stored in the address book. The "cc" always arrived whereas the principal copy bounced quickly (a matter of seconds).

A 200-MHz Pentium Pro-based PC with 128 Mbytes of RAM and a 4-Gbyte hard disk replaced Dr. Liu's 133-MHz

machine having 32 Mbytes. This was March 10th. Whether the added RAM will put an end to troubles with MS Windows remains to be seen. Your Editor is hopeful.

Estimate Actual Table Sizes

ESTIMATE ACTUAL TABLE SIZES (EATS) is a new request word that instructs ATP to dimension itself according to needs of current data rather than needs as declared by the user in LISTSIZE.DAT (VARDIM data). The modern implementation of this old idea began following research on dynamic memory allocation by Masahiro Kan of Toshiba Corporation in Japan (see separate story). Actual coding of new SUBROUTINE ESTIMA began January 15th as a 3-day weekend began.

As with Partial Table Dumping which first was reported in the preceding issue, details remain secret. However, the philosophy is simple, is easily understood, and probably could not be protected from exploitation by others, anyway. So, it now will be described. Since the early '80s, data first is assembled by the evaluation of $INCLUDE, and then it is ordered by processing of "/"-card requests for sorting. Upon completion, data is ready to be loaded into program tables. But are the tables ready for the data? Sometimes not. Big data sets frequently overflow one list size or another the first time they are considered, resulting in a KILL = 1 error termination. Even if such a fate is avoided, some tables may have been sized substantially larger than necessary, resulting in a waste of virtual address space. This translates into real memory, and possible paging to disk, whenever tables are dumped and/or restored (e.g., for Monte Carlo simulation or START AGAIN use). Such problems or waste could be avoided by sizing program tables just a little larger than necessary to hold the user's data. That is the concept or hope of Minimal Table Sizing.

Much more easily said than done, of course. The idea is decades old, but had not before been implemented for several reasons. First, substantial work is involved (coding is far from trivial). Second, the expansion of data case size without obvious bound is a relatively recent phenomenon (e.g., read in the January issue about the 20 thousand Argentine nodes and branches required for Orlando Hevia's LISTSIZE.20K). Third, dynamic memory allocation for Wintel versions of ATP (not David Szymanski's Unix version of a decade ago!) was not seriously considered prior to Masahiro Kan's recent successful experimentation. If it takes time to create memory dynamically (it certainly did for Szymanski using Unix System V, Release 3), there is increased incentive to minimize table sizes.

Initial results were communicated to others on January 19th. About COUNT.LIS your Editor concluded: *"note sizes in the heading before and after:*

```
Old: ...  230338 INTEGER words.  ...  752  900
New: ...  168470 INTEGER words.  ...   40   26
```

The saving of virtual address space is between a quarter and a third (168K instead of 230K). Thus far, only lists 1 and 2 are adjusted automatically. Also, there is automatic recognition that TACS and MODELS are not being used." Switches of List 6 were the 3rd table to be adjusted automatically. As first observed on January 23rd, this decreased the total to 143153 words as 16 replaced the original limit of LSWTCH = 120. Nonlinear element Lists 9 and 10 were the next to be adjusted. January 24th, the original 225 and 480 were reduced to 14 and 16, respectively; and the total was reduced to 135516 words. Then came the work of BPA's Dr. Tsu-huei Liu, who provided code to count all distributed line models. January 25th, Lists 8, 20, 21, and 22 were reduced from original values 5250, 2580, 300 and 1050 to 128, 10, 28, and 10, respectively; and the total was reduced to 104431 words.

To be continued. The story is just beginning. Storage using just two working vectors was redesigned with the collaboration of BPA's Dr. Tsu-huei Liu on January 20th.

Evaluating Analytical Functions , IV

A sawtooth waveform on the electrical side was mentioned in the January issue. Since then, Orlando Hevia has supplied additional creative uses of the pocket calculator to define Type-10 sources. On January 26th, that 1st subcase of DC-22 was expanded once again. Yes, the predicted pulse train is there, along with a rectangular waveform. As Mr. Hevia shows, these are easily generated analytically from a reference sine wave. The time has been extended to two cycles, and an added CALCOMP PLOT displays all 4 signals for easy comprehension.

Superposition of two or more such Type-10 sources at a single node was not possible prior to the correction of POCKET on January 29th. Once again, it was Orlando Hevia who accurately reported the trouble following your Editor's assurance that superposition should work. But it did not due to a structural peculiarity of your Editor's pocket calculator. Each Type-10 source involves a name, of course, and two or more at the same node led to confusion about which one was which during evaluation. The correction allows up to 9 sources at the same node, and it should be fool proof as long as the analytical function of each later Type-10 source is independent of the value of each source that participates in the superposition. A negative offset of 1/4 was added to DC-22a to illustrate such use. This was by means of a second Type-10 source at node PULSE. Order does matter, however: sources involving superposition are assumed to be contiguous.

Variable Dimensioning of ATP

An odd List 7 was allowed within ATP until January 23rd, when it was forced to be even by the addition of VARDI2. At issue were the two different locations for

table-sizing code --- the original in separate program VARDIM, and later the same formulas in ATP itself. The latter began some five and a half years ago to support dynamic dimensioning (see the July, 1993, issue), but were not identical to the original. Specifically, they did not force List 7 to be even, and the 4th and 5th subcases of DC-22 were found to be using odd value 285. Following the change, the odd 285 is being rounded up to even 286, and "Total size of LABCOM tables" is two words larger than before (74738 vs. 74736 INTEGER words).

LWORK is the name of the variable that begins the 4th data card of VARDIM data. It is unusual in that it dimensions working space for all supporting programs. Prior to changes on March 1st, there was no protection against overflow of the associated storage during another usage, however: normal simulation. For very large data sets, overflow during simulation might occur as first reported by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. Not only did Mr. Hevia report the troubled simulation, he also learned by himself that larger LWORK solved the problem. He was right: there is minimal use during simulation. The problem occurs when the usual 240K words is small compared with tables used for simulation. Ten or 15 years ago, this was unthinkable, so no protection was added at the time of first usage. But today, tables used for simulation may reach several million words, and data finally has been found for which the 240K words is inadequate. About the overflow error stop, it will use KILL = 1 as expected, and will mention List 31. Although perhaps not labeled this way in the Rule Book, the numbering is logical enough. If there were more than 30 lists, the first number on the 4th card would be number 31.

Linux and Programs It Supports

The cost of Linux from Red Hat Software was mentioned in the preceding issue. It should be emphasized that many users pay nothing, or next to nothing, however. This was explained by Orlando Hevia in E-mail dated January 14th: *"Here, a group of persons acquire the Red Hat CD-ROM and reproduce it. The cost is around $5 per copy (each CD)."* This is completely legal, too. Three days later, Mr. Hevia clarified: *"I asked a friend about ... legal conditions: he says that the CD-ROM may be copied without legal violation. ... Linux can be downloaded from Red Hat, via FTP, or directly installed from FTP, without legal restrictions. The advantage of buying the CD-ROM from Red Hat (I suppose that Slackware or Cheapbyte are comparable) is the on-line technical support. But the best support is from a nearby friend with experience (my own reliance upon Gustavo Courault). Generally, a Linux expert is generous with his knowledge. If one Linux user wants to update his kernel, a good manner to reduce the time is to contact another user in the neighborhood. In Santa Fe, I was the first to obtain Red Hat 5.1, but a month after this, another user obtained Red Hat 5.2. He offered to me a copy, but I did not want to reinstall the entire system in the absence of real need."*

Recall letter S replaces the asterisk (star, *) of other ATP versions for possible use as the third argument of RUNTP as explained in the April, 1997, issue. Curiously (and dangerously), unintended use of an asterisk seemed to result in the deletion of all .LIS files within working directory D:\GNUNT on Dr. Liu's Pentium at BPA. This was December 18th, using djgpp within a DOS window of WinNT. It seems that DEL %3LIS was responsible, with %3 being the asterisk. This is the only possible explanation Dr. Liu and your Editor could imagine. To make sure the experience is not repeated, any asterisk in the input line will be noted by GNU ATP, and execution will be halted with an error message: *"Halt in GNU CIMAG4. Asterisk (star, *) has been detected. Use letter S instead."*

List 5 of the 2nd subcase of DC-10 has had value 133 for years using Salford or Watcom ATP. But as first noted by BPA's Dr. Tsu-huei Liu on December 17th, GNU ATP showed value 137 even though the columns of dT-loop output were identical. Why? Matrix [Y] had different nonzero elements, apparently because of roundoff error. Might this be another case of confusion between internal 80-bit and external 64-bit results? For background of this detail over which LEC stumbled years ago, read Robert Schultz's insight in the January, 1994, issue. There seems to be no error in ATP source code. Cells 31 and 34 of TX should be equal, and they almost are. But not quite. When subtracted, they result in a small difference. Printing the numbers using E30.22 prior to [Y] formation in LAST14:

```
TX(31) = 2.64119186179894035147 73E-04
TX(34) = 2.64119186179894089357 84E-04
diff = 5.42101086242752217003 73E-20
```

Note that the two numbers differ in the 17th decimal digit, which is about the limit of 64-bit representation. The same two cells of TX printed by Salford EMTP show exact equality and precision limited to 64 bits:

```
TX(31) = 2.64119186179894089000 00E-04
TX(34) = 2.64119186179894089000 00E-04
```

Failure to update GNU POCKE1 resulted in erroneous execution of DC-22 and DCNEW-26 that was diagnosed on December 19th. But understanding occurred only after considerable difficulty. To ensure that the learning experience is not repeated, a special error stop has been added to GNU POCKE2 to trap code that is stale. Quick payoff: this was noted during testing on March 15th!

Variable NUMC0B should be mentioned in the opening story of this newsletter. Initially, only the Salford translator had been updated. But after stumbling over lack of the variable more than once following a new GNU translation, the decision was made to ensure that the trouble never would be repeated. For both Watcom and GNU, the change to the translator was made February 7th.

Faster starting of GNU ATP for djgpp can be traced to a December 26th report from Masahiro Kan of Toshiba Corporation in Japan. Huge table sizes are being used by

some, it must be remembered, and this has led to objectionable delays in starting as djgpp seems to page the entire virtual image. Mr. Kan documented relief as follows: *"I noticed the size of JARRAY ... is big. 4929580 integer words corresponds to about 20MB. I modified the code for main00.f and newmods.f by using the malloc() function of C. Malloc() allocates memory dynamically, so there is no initialization at startup. The startup time was reduced greatly."*

This was the critical, inspirational experiment and report. It identified the problem: huge JARRAY. But f2c was being used to convert otherwise universal ATP FORTRAN to C. As a less-disruptive alternative, developers in Portland first attempted to avoid use of JARRAY by table dumping to disk rather than RAM. This is the same change that recently helped Watcom (see explanation of trouble using SUBROUTINE MOVE in the October, 1998, issue). This reversed the development for VAX as reported in the April, 1994, issue (see discussion of JSDISK). While use of disk generally was successful, execution failed for the simple but important MEMSAV = 1 and START AGAIN usage of DC-32 and DC-49. Too bad. We nearly had a big improvement at essentially no cost.

Yet, a more complicated alternative existed : C-like logic that had been developed for Salford EMTP some 7 or 8 years ago. That was before use of MS disk caching (SMARTDrive), which provided universal acceleration of sequential I/O involving disk. In place of Salford DBOS, it was Mr. Kan's C-language routines, originally written to provide C-like .PL4 files, that were used to effect the table transfers for GNU ATP. Unfortunately, operation failed initially for DC-32 because neither your Editor nor BPA's Dr. Tsu-huei Liu knew enough about C to appreciate a temporary limitation of the coding. But Mr. Kan immediately recognized the problem, once symptoms were described to him. Following reception of a generalized copy of CLIKE.C on January 30th, the struggle ended easily and quickly enough. February 3rd, the last of standard test cases had been verified. Curiously, following the modification of I/O, one old data case (DC-13) required splitting whereas another one (DC-24) that previously had been split could be recombined into a single file. This is a reference to the well-known trouble djgpp sometimes has keeping track of the use of I/O units (e.g., see the April, 1997, issue). In any case, the project ended happily. Dr. Liu E-mailed a complete new GNU ATP package to Mr. Kan on February 5th after all compiled TACS data cases, too, had been verified.

How much quicker will the starting be? Generally, the answer will depend on both available RAM and limiting program dimensions. It should be noted that operational dimensions of LISTSIZE.DAT do not affect the speed. Improvement will be dramatic if previous paging to disk now can be confined to RAM. This was what Mr. Kan seemed to note at the end of his report dated February 11th: *"4. I observed the startup time of ATP djgpp was greatly reduced. The previous 10-15 seconds was reduced to 1-2 seconds."* But the ratio is highly variable. Consider the two limiting, asymptotic cases: 1) As dimensions approach infinity, paging of djgpp should be cut in half; and 2) As dimensions approach zero, there generally is no data to be paged, so no difference.

Superposition of Phasor Solutions

The automatic creation of sources that were used during previous phasor solutions became available December 15th, with operation illustrated by the 14th subcase of DCNEW-26. Comment cards explain when this labor-saving feature is used, and when it is not. The 14th subcase, of 16 total, ended the illustrations of superposition with an example involving three harmonics: 50 Hz, 75Hz, and 100 Hz for a series connection of constant parameter, Jmarti, and Semlyen modeling of distributed lines. The smooth result of period 40 msec (50 Hz), illustrates coexistence of all line types except Noda.

ATP Licensing Problems

EPSRI is the acronym used by your Editor some 16 years ago to indicated the Chinese EPRI outside of Beijing, China. This was explained at the bottom of page SHOW-4 of the Vol. XII *EMTP Memoranda* dated 26 January 1983: *"I might briefly mention problems with 'Science' or 'S' in the name of the Chinese institute. Tsu-huei assures me that the Chinese name has such a 'Science', and our first correspondence with Peking used this (actually, the abbreviation 'Sc.') in the return address. But recently, it has been dropped from the official name, and the official acronym is just EPRI. But since America already has an EPRI, use of this could only lead to confusion. ... I prefer to revert to the original EPSRI, which is both unique and logical (a literal translation of the Chinese)."* The distinction shall be continued in this newsletter.

EPSRI is a second institution of the world that has required the separation of noncommercial ATP use from commercial EMTP development. Recall that the University of Wisconsin in Madison was the first, as explained in the July, 1993, newsletter. The need for another separation was suspected following Prof. Hermann Dommel's recent disclosure of commercial collaboration by Lin Jiming (see mention in the October newsletter).

Student Guo Jian is the person who ably arranged the limited use of ATP by others following Can/Am user group licensing of him as an individual. First, in E-mail dated February 1st, he summarized EPSRI organization : *"... EPRI, China is composed of ten departments, several corporations and a graduate school which are all relatively independent entities for each other."* About the problem of EMTP commerce, he clarified: *"My advisor, Prof. Wu Zhongxi, works in the Computer Application Department. Prof. Lin Jiming works in the Power System Research*

Department .... he is the only person to have the license of that version of EMTP which can not be used even by other groups in the Power System Research Department ...." Clearly, the graduate school of EPSRI had an ATP licensing problem, if Lin Jiming was associated with it.

Three days later, your Editor and Dr. Liu responded : *"It sounds as though we could license the Computer Application Department for ATP use free of charge. Your Prof. Wu would be the primary contact (i.e., she would sign the form). This assumes that Prof. Lin Jiming, and anyone else who might have participated with him on work for DCG and / or EPRI on DCG / EPRI EMTP, has no connection with the Computer Application Department, and therefore would be excluded from access to ATP materials by such a license. ... Later, if there is interest in ATP by EPSRI personnel outside the Computer Application Department, other boundaries might be discussed. But for now, let's begin with just the Computer Application Department, which is understood to involve no person who has been involved with EMTP commerce. That is the necessity : organizational separation from anyone who could not be licensed to use ATP free of charge."*

How might corporate segmentation affect existing prohibitions against free ATP licensing? This issue first was raised explicitly by Peter Dick of Ontario Hydro (OH), who telephoned BPA from Toronto, Ontario, Canada, on March 4th. Just as BPA is being split into two halves, so it would seem that Ontario Hydro is being split --- perhaps into three or more pieces. Of course, as a single, cooperating organization, OH was denied free access to ATP because of its membership in DCG. But what will happen if there is a split into two or more separate, non-cooperating parts? This is an interesting question that is worthy of thought by everyone. The OH problem is the opposite of the ABB problem. Recall a merger masked the originally-separate companies ASEA and Brown Boveri --- important because only the former was known to be involved in EMTP commerce.

Cepel of Rio de Janeiro, Brazil, is **not** licensed to use ATP. This was the conclusion of the Can/Am user group, as contained in E-mail to Marco Polo Pereira of the Latin American user group on January 12th. But there remains the possibility of an ATP pardon as explained in the January, 1992, and July, 1995, issues. More next time.

Comings and Goings

Akihiro Ametani of Doshisha University has explained the status of ATP-interested faculty at his school. Specifically, your Editor had inquired about Naoto Nagaoka, whose name can be found in numerous issues of this newsletter over the past decade. Professor Ametani's E-mail dated January 11th clarified the status as follows: *"Naoto will become a full Professor. Presently he is an Associate Professor. The following steps apply to faculty of my university: 1. Research Assistant; 2. Lecturer or Assistant Professor; 3. Associate Professor; and 4. full Professor. I was promoted to the rank of full Professor 14 years ago. Now, Naoto is being promoted. So, my power systems analysis laboratory has two professors, and Dr. Y. Baba from Tokyo University will join us as a new Research Assistant. This should give me a little more time to concentrate on my job as Dean."*

Prof. Yoshihiro Murai of Gifu University in Japan died of some illness on Friday, February 5th, 1999. This sad news was received four days later in E-mail from Dr. Tsuyoshi Funaki of Osaka University. *"What a pity that he should have died so young."* Readers who do not know about Prof. Murai's most famous contribution to ATP are referred to the January and April issues of this newsletter during 1996. Understanding of the need for *dynamic current redirection*, provided today by GIFU switches, will endure as Prof. Murai's memorial in ATP.

Tom Varilek and Prime Computer were names from the past as seen in E-mail from Stu Cook dated February 17th. Mr. Cook and his company, JUST Services, are to be moved from suburban Montreal, Quebec, Canada, to Rideau Ferry, Ontario, during the last week of March. About Prime (or PR1ME as used in advertising of the day), Mr. Cook explained about an old manual: *"EMTPLT Manual version M34. The inside cover calls it the ElectroMagnetic Transients Interactive Plotting Program. The first line of text says that it is for Prime computers. It seems that it could use character displays, Tektronix 4014 display and could output to a Calcomp compatible plotter. ... This last page is signed 'Tom Varilek, Jan 29, 1981' ..."* At the time, Mr. Varilek was with Minnesota Power in Duluth, Minnesota. But when last heard from some years ago, he was south of Minnesota, in Iowa. This followed loss of his job as a Vice President of sizable Iowa Power in Des Moines, following a corporate merger (he was on the losing side, apparently; turmoil in the industry continues).

Who has publishable information concerning the whereabouts of former BPA contractor Laurent Dube? From two different sources, it has been learned that the Oregon Coast no longer is his home. In fact, the Pacific Northwest no longer is.

Florida Resort Course during March

A much better color projector is noteworthy progress from Prof. Dennis Carroll's ATP short course during the 2nd week of March. In E-mail dated March 10th, he described *"much brighter and more colorful presentations than in previous years with our old LCD projection panel. So room lighting was really not an issue, since we can clearly view the presentation even in broad daylight."*

A peculiarity of students this year is that each had access to ATP prior to the course. There were no new users. In fact, most seem to have brought ATP with them on notebook computers since Prof. Carroll reported: *"that*

gave us a lot of fun trying to get all the various hardware and software coordinated at the beginning of the course."

Attendance was on the low side this year (12 students), and half came from outside the United States. Two were from Italy, two were from Australia, one was from Norway, and one was from Caribbean Puerto Rico.

Power and / or Energy Outputs

The 4-punch in column 80 of branch or switch cards has been used to produce power and energy outputs in place of the voltage and current outputs of a 3-punch. Now, as part of the first general change in a quarter of a century, more flexibility is being offered. The user no longer is forced to choose between the pair of voltage and current outputs (v, i) and the pair of power and energy outputs (p, e). Finally, power and energy have been freed from the voltage and current that are used to compute them. Of the 4 variables, the user is free to request any combination for output. Although unavailable to the general public prior to January 28th, complete success first was demonstrated Christmas day when four less-affected modules were updated.

Massimo Ceraolo of the University of Pisa in Italy provided inspiration for the generalization of power and energy outputs. His E-mail message about alternative .PL4 files ended with the following idea on December 15th: *"The branch quantities that can be outputted are of four types (current, voltage, power, energy). The number of bits that a hexadecimal digit can describe is 4. If the 80th column of a branch card is viewed as a character between 0 and F, it could ask for any combination of the four quantities. ... Is there (or will there be in the future) a need for such an extension? I don't know. What's your opinion?"*

The need for continuity has prevented a change from the old output alternatives of 0 through 4 to such a rational, new, hexadecimal system of requests, however. As Mr. Ceraolo observed, 0 through 3 would require no change, but the old 4 then would become C (decimal 12, equal to 8 plus 4). This would be using the right-most bit (bit number 1) for current, bit number 2 for voltage, bit number 3 for power, and bit number 4 for energy. I.e., old data with a 4-punch in column 80 would require change. Alternatively, new data would require some such declaration as HEXADECIMAL COLUMN 80 to distinguish it from the old. But this change would represent too much of a burden, with the cost being paid forever by persons assembling data. Instead, it seemed better to map the 16 choices into column 80 punches less rationally, preserving the meaning of existing 0-4. This is what has been done, as detailed on comment cards of the 6th subcase of DC-37 which has been modified to illustrate the new capability. Power consumed by the nonlinear element has been added to the preceding output of current. To prevent copying by commercial competitors, details remain a secret of ATP design that will not be detailed in newsletters (which may be read by persons who are not ATP-licensed).

Lack of one voltage or one current (but not both as ordered by a 4-punch) represents an economy in the sense that storage for plotting (e.g., in a .PL4 file) is lessened compared with usage that preserves all four variables. However, it should be explained that work is required to remove the unwanted voltage or current from the output vector. If any increase in elapsed time for the time-step loop ever might be noticed (typically it will be too small to measure), this is one possible reason. Deleting either both of them, or neither of them, allows avoidance of the extra effort that is required to delete just one of them. This is a peculiarity of the implementation.

More storage is required to provide the new output choices, although only a small amount. MAXPE is the limit of VARDIM List 18, and the multiplicity of this has doubled. That's all. Following the change, three times default dimensioning results in *"Total size of LABCOM tables = 230338 INTEGER words."*

TEPCO Improves S.M. Model

Dr. Hiroshi Okamoto of Tokyo Electric Power Company (TEPCO) supplied developers in Portland with important changes to Type-58 and 59 S.M. modeling as explained in an E-mail message dated January 6th.

Comment cards in data cases DC-25, DC-26, DC-47, DCNEW-11, and DCNEW-20 have been changed as a result of the S.M. modifications from TEPCO. Changes to the last of these are both minimal and illustrative: *"Answers change 10 February 1999 following the massive changes from TEPCO (Tokyo Electric Power Company) in Japan. See April newsletter."* Noteworthy by their absence are the two subcases of DC-53, for which the solutions are unchanged in spite of Type-59 S.M. modeling, both saturated and unsaturated. For posterity, old solutions are being saved both at BPA and on your Editor's home computer as DC*SAL.992 disk files. Also, comment cards that document solution numbers of years past (maybe dating to Apollo?) are **not** being changed.

Changes to results are very small, generally. The only exceptions are some variables of DCNEW-20 --- not surprising when one remembers that this simulation demonstrates instability of Type-59 modeling. Typical of changes is the torque TQGEN of the first subcase of DC-25. Limits of the PRINTER PLOT are unchanged (only one interior number has changed), and the value at the 500th time step, and extrema (of which the times are unchanged), compare as follows:

```
          Final step      Minimum       Maximum
   Old:  4.269726268  -1.02756122  6.372899605
   New:  4.2697321    -1.02756092  6.372882218
```

For engineering purposes, signals are effectively the

same. Why signals are changed at all is not known. Sometimes there is no obvious reason (sometimes it just seems to happen, using the Salford compiler).

Convergence of the single-phase load flow of the 3rd subcase of DC-26 is substantially affected. Recall this is the famous 5-bus problem from Stevenson's text book. Whereas NEKITE = 164 iterations sufficed in the past, 394 iterations now are required to achieve the same tight convergence.

NEW LOAD FLOW is the way a user requests the new load flow modeling of TEPCO. Illustration first was provided on February 19th in the form of a new 2nd subcase of DCNEW-20. Of course, the data was supplied by Dr. Okamoto; your Editor knows next to nothing about the new alternative as it first is being made available to others. Continuation on March 18th : Dr. Okamoto now is preparing an English-language summary of changes. The plan is to publish this explanation in its entirety next time. Look for it.

Parameter Variation Studies by PCVP

Assignment of a negative constant was in error prior to correction on January 10th. The trouble was illustrated by data that was attached to an E-mail message from Mustafa Kizilcay of FH Osnabrueck in Germany. On January 5th, the EEUG Chairman provided an illustration such as PHASEB____ = -120 about which he explained: *"A negative number assigned to a variable is recognized as zero."* Your Editor responded: *"This trouble was quickly confirmed and corrected. The line having S.N. 8390 had two errors. The subscript was bad, and a minus sign was forgotten."*

Numbers were not being right-adjustment as intended when less than maximum precision was required. This was the 2nd of 3 separate and distinct errors that were corrected on January 10th following complaint by Prof. Kizilcay. Blank spaces in the output of all standard test cases that involve $PARAMETER --- DC59, N19, N25, and N26 --- changed as a result. Numerically, results were unchanged, but only because either F-field numbers were involved or system-level software provided protection. As the professor observed, with scientific notation (E-field numbers), there might be a potential problem were it not for that nifty BN reform of Stephen Boroczky (see the July, 1998, issue).

Character strings were not being handled properly when the constant on the right was shorter than the variable name on the left. This was the 3rd of 3 separate and distinct errors that were corrected on January 10th following the complaint by Prof. Kizilcay. As the professor observed, consequences of a shorter right-hand side were severe: DBOS interrupted execution with an error message like "illegal data in field."

Two-dimensional parameter variation is illustrated by the 11th subcase of DC-59 as mentioned in the January newsletter. Included was explanation that the .PCH file *"will include all results."* As detailed on comment cards that were added February 24th, this is true if the data is extracted and executed from a separate disk file (the normal case) with no name following the $PUNCH request. But as stacked after 10 preceding subcases, with its own special name for punched output, it is only output of the 2nd of the 2 passes that end up in the punch file. The .LIS file is complete, however.

Integer data symbols are another creation of Prof. Martinez. The original request came in E-mail dated February 9th: *"We have tried to use $PARAMETER to change the number of runs using PCVP. It can be done, but PCVP only accepts integer numbers, and $PARAMETER only produces real numbers."* Well, that was prior to February 19th, when the need finally was understood in Portland. Observe that numbers of ATP data are mostly floating point, and the original use supported these. However, some numbers are integers. If a user wants the pocket calculator to change one of the integers, he needs an integer data symbol rather than a floating-point data symbol. As Prof. Martinez explained, encoding must be integer rather than floating-point. An illustration is provided in the 11th subcase of DC-59. On February 24th, the old data was modified by the addition of integer data symbol DCD to replace the previous fixed value of one for the number of decades in columns 60-62 of the frequency card of LINE CONSTANTS data.

Both $PARAMETER use and the pocket calculator are described in the same chunk of the computer stored Rule Book. This is disk file H01-L as first made available to others on February 19th when pages were switched in the printed Rule Book. This was immediately prior to a printing in Portland and subsequent reproduction in Florida for Prof. Dennis Carroll's resort course. A WP 7 copy was sent to EEUG the following day.

Publishing Programs and Viewers

Both spelling checkers and grammar checkers have their limitations. Consider the lack of one letter (an "a") in public E-mail from South Portland, Maine, on January 22nd. *"I am an electrical engineer working for a small insulting group in Maine."* This sentence receives a clean bill of health from MS Word 6 on your Editor's home computer!

A Spelling checker is bad enough when used on the language for which it was intended. It can become downright dangerous when used on a foreign language. The best illustration was received February 8th from Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. First, your Editor had asked about the following message from his Spanish-speaking country: *"Liu, y want to express my gratitude for the program you sent me. At this moment it is working correctly, if y need*

something Y will ask ...." Mr. Hevia explained the use of "y" instead of "I" as follows: *"This is because of the spelling check ..., surely: the 'I' or 'i' is not used in Spanish as a separate word. The spelling checker replaces the 'I' or 'i' ... by the 'Y' or 'y' (the equivalent in Spanish of the English 'and'). … Mr. XXXX must be using the spelling checker automatically, and with a Spanish dictionary."*

Frequency Scans and Harmonics

The printed heading for MODELS variables of FREQUENCY SCAN use was treated in a paragraph in the January newsletter. Recall Prof. Mustafa Kizilcay of FH Osnabrueck in Germany first reported trouble. Well, that trouble extended to .PL4 files, too, as subsequently reported by him in E-mail dated December 20th. Correcting the .PL4 problem was the preoccupation of December 21st and 22nd --- nearly two days of work. Yet, it must be stated that ATP itself was less than half of the problem (see mention in the separate TPPLOT story).

Year 2000 Compliance of ATP ?

History provides guidance for dealing with this cursed Y2K problem. Thanks to BPA's Jon French for the following anonymous letter from Rome about 2000 years ago: *"Dear Cassius: Are you still working on the Y zero K problem? This change from BC to AD is giving us a lot of headaches and we haven't much time left. I don't know how people will cope with working the wrong way around. Having been working happily downwards forever, now we have to start thinking upwards. You would think that someone would have thought of it earlier and not left it to us to sort it all out at this last minute. I spoke to Caesar the other evening. He was livid that Julius hadn't done something about it when he was sorting out the calendar. He said he could see why Brutus turned nasty. We called in Consultus, but he simply said that continuing downwards using minus BC won't work and as usual charged a fortune for doing nothing useful. Surely we will not have to throw out all our hardware and start again? Macrohard will make yet another fortune out of this I suppose. The money lenders are paranoid of course! They have been told that all usury rates will invert and they will have to pay their clients to take out loans. It's an ill wind ... As for myself, I just can't see the sand in an hourglass flowing upwards. We have heard that there are three wise men in the East who have been working on the problem, but unfortunately they won't arrive until it's all over. I have heard that there are plans to stable all horses at midnight at the turn of the year as there are fears that they will stop and try to run backwards, causing immense damage to chariots and possible loss of life. .... Anyway, we are still continuing to work on this blasted Y zero K problem. I will send a parchment to you if anything further develops ...."*

"Statement about Year 2000 Compliance" of ATP was the good idea of EEUG Chairman Mustafa Kizilcay. His cover letter for the first draft explained as follows on January 24th: *"With such a ready text, we will save some working time. If someone asks about Y2000, he/she will be sent this statement without any further explanation. We can put it also on FTP file servers and the JAUG secure WWW pages."*

"Le bug de l'an 2000" seems to be the way French refer to the Y2K problem. This title was seen at the Jean Levant WebSite in Amiens, France. Is this authentic French, or is this Franglais (considering *bug*)? English words have been creeping into colloquial French during recent decades.

Computers that prepare paychecks might exhibit Y2K problems, of course. It is not obvious that employees would be the losers, however. The following was received from BPA's Walter Powell on January 29th. It takes the form of a memorandum from management dated January 1, 2000: *"Dear valued employee : Our records indicate that you have not used any vacation time over the past 100 year(s). As I'm sure you are aware, employees are granted 3 weeks of paid leave per year, or pay in lieu of time off. One additional week is granted for every 5 years of service. Please either take 9,400 days off work or notify our office and your next pay cheque will reflect payment of $8,277,432.22 which will include all pay and interest for the past 1,200 months. Sincerely, Automated Payroll Processing"*

Branch Data Input Restructured

TO SUPPORTING PROGRAM is a new request word that allows in-line (as opposed to separate, preceding, off-line) execution of a supporting program to produce needed data. The idea was received in Portland on January 26th in E-mail from Prof. Juan Martinez Velasco of the University of Catalunya in Barcelona, Spain. Much more radical thought was involved, but this was a foundation tool that seemed capable of implementation without too much trouble, so it was selected for initial experimentation. New 2nd, 3rd, and 4th subcases of DC-41 illustrate operation for both LINE CONSTANTS and JMARTI SETUP. All 3 subcases involve the same 138 miles of 500-kV line from John Day to Lower Monumental (BPA substation names). For those who want to derive data, and then use it in the same data case, these are illustrations. Of course, the K.C. Lee derivation of the 3rd subcase is fast, so is well suited. Much slower is the JMARTI derivation, so the 4th subcase is less practical --- even for execution using a Pentium. Operation should be universal, and it began January 29th with the addition of module SUPORT to provide service for ATP branch cards. Can any reader envision the need for other types of data?

The saturable TRANSFORMER always has allowed current output of the magnetization branch and/or winding number 1. There never has been provision for current output of the 2nd or later winding. Nonetheless, any column-80 punch apparently was ignored correctly during

years past whereas recent versions (after branch data input was restructured) resulted in an unusable .PL4 file. The first report of trouble came from Francisco Gil Garcia of Schneider Electric in Grenoble, France. On February 4th, he wrote: *"I have the latest Watcom version of ATP (11/98), but I have found several problems using it. .... a datacase which works properly in ATP10 but which doesn't work properly in ATP (11/98)."* Later that same day, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, observed that the problem had nothing to do with Watcom, and was in fact due to an extraneous 1-punch in column 80 of the data card for winding 2 of a saturable TRANSFORMER. Your Editor agreed with Mr. Hevia that protection was warranted, and later that same day he added two lines to INTRAN for this purpose. Henceforth, any use of columns 79 or 80 on the 2nd or later winding card will be ignored by erasure, in effect. Whether or not the user will see the offending punch will depend upon whether or not an in-line comment is involved (details are less than ideal, although protection should be fool-proof).

A DO loop in ATP data was another of those radical ideas from Prof. Juan Martinez (see a preceding paragraph). More precisely, an ordinary DO loop of FORTRAN is the simple way that your Editor chose to implement the professor's more general request for looping. Illustration is provided by a new first subcase of DC-58 that was added February 6th as the original data was forced downward to become the 2nd of two subcases. Anyone having interest is advised to read data comments carefully. Included is new numeric output of the pocket calculator, which can be used to build A6 node names that are a function of the DO-loop index (i.e., the names are numerically serialized). In its simplest form as illustrated in DC-58, an alternative to CASCADE LINE (see the numerous illustrations in DC-9, and mention in the April, 1998, newsletter) has been provided. This assumes uniform line sections. But much more generality is possible. Parameters *could* be made a function of section number (the DO-loop index), for example. Another advantage compared with CASCADE LINE is availability of all branch cards that are created by the loop. I.e., the feature can be used as an ordinary supporting program to create branch cards that then are accessed by $INCLUDE for subsequent simulation that avoids re-creation.

Selective branch outputs of a JMARTI line might be erroneous. The trouble first was pointed out to developers in Portland by Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. This was in E-mail dated February 1st. There would seem to be similarity to the just-mentioned saturable TRANSFORMER problem since an ATP Rule Book from a decade or more ago clearly warns about non-uniform punches of column 80. From the top of page 4D3-7: *"Variable K of column 80 can be used as usual for branch output requests, only with the added restriction that all phases should be keyed equally whenever branch current (or the derived power and energy) is involved. This applies to values one (for current), three (for both voltage and current), or four (for both power and energy)."* February 9th, BPA's Dr. Tsu-huei Liu and your Editor modified the logic of branch data input in INDIST to enforce the uniformity of output for all distributed models. While not necessary, this was simplest solution. It should be fool-proof for Semlyen modeling of any number of phases and for all other models having 9 or fewer phases (phases numbered 10 or higher escape the modification logic). Whatever punch the user keys in column 80 of the first phase automatically will be extended to later coupled phases. This is illustrated by several data subcases, for which output signals are unchanged despite the following changes to data: 1) 1st subcase of DC-41 (add 1-punch to column 80 of 3rd phase of Semlyen line; 2) 2nd subcase of DC-41 (add 3-punch to column 80 of 2nd phase of the K.C. Lee line; and 3) 1st subcase of DCNEW-4 (erase the "1" in column 80 of the 2nd phase of the JMarti line). These changes were made February 15th.

TACS Control of Unloaded S. M.

Trouble with TACS control of the torque of an unloaded generator was mentioned by Dr. Ali Moshref of PowerTech in Vancouver, B.C., Canada. Public exposition of the problem on February 3rd mentioned zero power, voltage, and frequency. But a previous inquiry, made privately, had mentioned only zero power, which seems to be a more common and practical case. Lack of load makes normalization difficult. Recall a scaling factor is required for TACS control: *".... it fails if I want to have no power in steady-state and start loading the machine."*

Your Editor's earlier private response was made public on February 7th. It included the following advice and an offer: *"That is an interesting point. Yet, there always are losses, so some small power must be expended to supply these. You have a lot of precision (about 16 digits), so I suspect in practice there would not be a problem scaling the power of an open-circuited machine. .... I have never heard of this complaint before. Even if there were no mechanical damping (perhaps this is legal, and a useful idealization), there is always some electrical loss. The windings always have resistance (assuming you are not modeling a superconducting generator), and a power system generator typically will have an attached transformer that will draw magnetizing current. There will be resistive loss due to this even though the breaker on the high side is open. This is why starting of generators is so dangerous: the transformer is connected --- always. If you model the transformer realistically, there should always be armature current, and hence loss. Having written that, I would be willing to modify the program to bypass scaling provided expert users agree that the change would be useful. .... If a consensus is reached publicly that un-scaled power would be a useful extension, we should be able to make the change without much trouble."*

That concerned just the zero power of *spinning reserve*. The public inquiry mentioned more, and your Editor's

response ended as follows: *"About zero voltage and frequency, this writer would not be qualified to make such a radical change. Support by the experts at TEPCO in Japan would be required. But is such a change practical? For large thermal units of the electric power industry, starting from zero speed might require a day or more, so this would represent one long EMTP simulation! While it may be true that hydro generators or smaller generators can be started more rapidly, how practical is such simulation? Readers are reminded that the Type-19 U.M. remains available for any such unusual configurations that might occasionally be of interest. Furthermore, thanks to Prof. Hian Lauw's later work, the U.M. will accept Type-59 S.M. data with minimal changes. If one switches from Type 58 or 59 to Type 19, does the problem remain? If so, this writer will work on it."* No response to this final question has yet been received from anyone.

Very low frequency use of the Type-58 and Type-59 S.M. models was confirmed by expert user Gabor Furst of suburban Vancouver, B.C., Canada. This was in response to a follow-up question by Dr. Moshref on February 15th : *"I would like to start SM58 or SM59 from almost near zero conditions (say for example, voltage = 1% and speed = 1%)."* Mr. Furst reminded readers of his Fargo-publicized explanation nearly two years ago. This was summarized in a story entitled *"Inadvertent Energization of Generators"* in the July, 1997, newsletter. Mr. Furst summarized relevance to the present question as follows: *"There is no difficulty in starting a machine using SM58 / 59. The machine is started at near rest, say 1Hz speed as noted by Ali Moshref, and a very small voltage, could be a few volts. Reactances should be referred to 1 Hz if per unit, if engineering data is used. This was discussed in detail in a posting some two years ago by this writer and Bruno Ceresoli of ENEL. The machine can then be started using the TACS control option as described in the Rule Book on page 8-21, Rule 2. The user will run the machine without TACS control first and note the TACS TQGEN value which will be very small, say on the order of 1e-10. Then, using the reciprocal in TACS, or some very large multiple, a reasonable starting torque can be applied for starting the machine. It should be noted that in Rule 2 above, the reference to the external mechanical power may be misleading, it is external torque."*

Partial Table Dumping

Minimal initialization is a related reform that was discovered during the early days of work on partial table dumping. This idea, too, is simple. Dating to pre-ATP days, entire vectors were zeroed at the start of execution. Of course, not much memory was available during the early days, so such use was not a problem. For example, use of the BPA CDC-6500 for EMTP studies involved a memory partition of between 32K and 64K words (close to a quarter to a half a megabyte) of core storage until usage was dropped in 1979 (a DEC VAX-11/780 was installed during March of that year). Nothing approaching Meredith's 14 Mbytes, and its associated waste, could be envisioned in those days. But the problem is upon us today. As recognized for years, zeroing portions of memory that never will be used to store data **was** wasteful. At least one such vector initialization was removed during the overhaul of branch input (see the April, 1998, newsletter), and 18 more were removed October 14th at the end of a concerted drive to eliminate unnecessary initialization. With the most common class being List-2 branch vectors, the saving is huge. List 2 has size 3000 cells in LISTSIZE.DAT and 6000 cells for EEUG dimensioning. Furthermore, unlike table dumping, unnecessary initialization slowed every simulation. There was no way to avoid it.

Drifting past history of List 8 is an old phenomenon that was eliminated around October 20th. Each time step, data representing any one line moved downward one cell within List 8. Why? No good reason that your Editor could imagine. After all, history for a Type-53 transport delay in TACS never moved within its List-19 storage as coded by BPA contractor Laurent Dube. Static history seems conceptually clearer than drifting history, so, after decades of drift, electrical history finally has been confined. Answers are unchanged, of course.

The *partial* of partial table dumping varies from table to table and from vector to vector. In words and theory, the concept is simple: just dump what is being used. But in practice, it is not always easy to know what is being used. Care is required. A good illustration is provided by the Type-58/59 S.M. (synchronous machine) as written / maintained by Tokyo Electric Power Company (TEPCO) in Japan. It was easy for your Editor to detect when no S.M. modeling was involved, so the omission of all such dumping and restoring in case of non-usage was easy. But what portion of the six COMMON blocks is being used when there is one or more Type-58/59 model? Because your Editor is unqualified to answer this question, he asked Dr. Hiroshi Okamoto for assistance on October 20th.

List 19 storage of TACS is worthy of special consideration because of widespread use. As with the Type-58/59 S.M., the saving is perfect and complete if TACS is not being used. But if TACS is being used, no improvement is offered. The user should understand why. It is curious how perspective has changed over the past two decades. The present-day problem with TACS is rooted in the minor miracle that was performed by Robert Eifrig two decades or more ago: variable dimensioning by offset subscripting. TACS author Laurent Dube had used fixed vectors, so overflow or waste was common, and was a serious problem, prior to the Eifrig reform. Yet, because of the offsets, data of these TACS vectors now are scattered throughout the storage of List 19.

List-19 TACS and the List-25 U.M. are special cases that deserve special explanation. There is similarity in that only the total working space is limited, and within that total, the user is able to allocate component tables using

one of the ABSOLUTE declarations. For illustrations, see ABSOLUTE TACS DIMENSIONS in DC-18 and ABSOLUTE U.M. DIMENSIONS in DCNEW-1. The ABSOLUTE declarations serve to space (separate) vectors within the total working space. If unnecessarily large, holes are created. Yes, if these are significant, Schultz will find them, and remove them. But the service is not free. One pays for it in terms of virtual address space (RAM, if paging to disk is to be avoided) as well as unnecessary indexing. So, the user is advised to minimize his ABSOLUTE declarations, if he wants to improve the efficiency of his table dumping and restoring. Prior to Schultz's afterburner, no such saving was possible. Now, it is, so the user might profitably pay more attention.

MODELS is little improved, it should be explained. Of course, if MODELS is unused, the saving is total. But the *partial* for MODELS by Laurent Dube is limited to the character storage of List 15, which typically is much smaller in bytes than the numeric storage of List 28. This was the simple part, so your Editor quickly exploited it. But numeric storage is more complicated. Lacking documentation of source code (in fact, lacking real source code, too), your Editor was unable to understand enough, during a day or so of study, to simply and effectively restrict Schultz's application to List 28. So, he abandoned the idea of such automatic saving. Nonetheless, assuming one can trust author Dube's memory management, note that the user has the power to eliminate manually most unneeded virtual memory. Using either NEW LIST SIZES or $DEPOSIT, the user can, in data, reduce List 28. This would be done after execution first had been shown to be correct using much larger dimensions. This is comparable to using one of the ABSOLUTE declarations (see preceding paragraph) for the U.M. or for TACS, note.

SPY data may or may not be included with program tables, the user is warned. This is not a new problem, or a new uncertainty. It existed in the past, and it continues to exist. It is mentioned now only in order that the issue be understood, and not be forgotten. Whereas batch-mode data always (well, at least since START AGAIN was introduced) has been stored in COMMON blocks that are part of the tables being dumped or restored, this has not been the case for SPY data. Should any SPY data be dumped and restored? Your Editor is unsure. The interactive nature introduces a potential conflict. If viewed as a tool of observation, SPY should be independent of ATP tables. If the program does not stop, why should the observation of it (e.g., the use of EXAMINE to monitor key program values) be lost, or at least modified, by the restoration of tables? Note that the independence of SPY might be more than philosophical: Running under standard Unix more than a decade ago, computer expert David Szymanski of Wattsburg (suburban Erie), Pennsylvania, offered multiple copies of SPY running as separate programs. This did **not** stop when ATP stopped! If more work ever is done on SPY (if and when decent windows, not to be confused with Bill G's Windows, are used), there should be additional consideration of SPY data relative to table dumping and restoring.

Interactive Plotting Programs

"GTPPLOT and Salford TPBIG" was the subject of E-mail from the Fargo list server dated January 6th. In this, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, explained that there was a potential conflict between use of his GTPPLOT and the Salford DOS extender DBOS: *"GTPPLOT may abort with DBOS installed, due surely to conflict between DBOS and CWSDPMI. This conflict can be avoided using the following script ...."* In this way, *"DBOS and CWSDPMI are not run at same time, and the conflict is avoided. The time to load DBOS is small."*

IEEE PES COMTRADE output of GTPPLOT is being worked on by author Hevia. His public E-mail dated February 17th explained: *"I am adding the COMTRADE command to gtpplot. But I have not the standard IEEE C37.111-1991. By trial and error I obtained an ASCII file ..."* Later that same day, after another idea that might assist users of Salford DBOS-incompatible MS Windows NT, your Editor offered cooperation: *"Alternatively, the TPPLOT code to support COMTRADE could easily be added to GTPPLOT, most likely. As long as use is non-commercial, this writer has no objection."* About lack of the IEEE standard, your Editor claimed: *"This may be an advantage. The book is not easily read .... It also is ambiguous in more than one respect. That is why this writer copied the work of BPA's Randy Suhrbier rather than spend much time studying the IEEE standard. Finally, recall concern about how some manufacturers might be reading the standard (see the October, 1995, newsletter)."*

Miscellaneous Intel PC Information

The U.S. government's antitrust case against Microsoft continues, and is being followed by the news media as a minor story (nothing like the O.J. trial, or Pres. Clinton's latest scandals). The front page of *USA Today* dated November 10th contained the following summary of the interesting connection to Intel beside a picture of Bill G: *"Microsoft effectively threatened to cripple Intel and force the chipmaker to stop developing software and backing Microsoft Rivals, an Intel executive testified Monday. And consumers have been hurt by Microsoft, led by CEO Bill Gates, Intel's Stephen McGeady said in the trial's most damaging testimony."* The story itself can be found on page 1B, with a sidebar explaining that this is day 13 of the trial.

Pentium III is to be released during the last week of February. This according to a story on page C1 of *The Oregonian* dated February 18th. The headline is: *"550 MHz chip near release, Intel says."* The story mentions a plan to spend $300 million on advertising. About the new chips, *"Intel has included a serial number feature ..."*

According to critics, this *"could allow tracking of a user's digital footprints on the Web."* Anyway, availability of Pentium III should be good news for buyers of Pentium II (bargains always can to be found below the cutting edge).

Miscellaneous Small Items

Gabor Furst's data generator for induction motors was mentioned in the April, 1997, issue, and is illustrated by the 2nd subcase of DC-15. Since that installation, your Editor has done nothing. But author Furst apparently has worked with his separate code, and there has been work on the ATP implementation by Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina. His E-mail dated November 16th accompanied substantial changes. It began as follows: *"I was testing INDUCTION MOTOR DATA, and encountered two problems: 1) If you change the data, the results are the same; and 2) If you set the default data in a file, the results are different than the results using blank lines. I send to you the INDUCT.F file with the corrections I added. With this, both results are the same, and the user can set his own data."* Mr. Hevia's corrections entered the UTPF on December 20th.

$OPEN and $CLOSE shared the same error message (*"Error with disk file OPEN or CLOSE ...."*) prior to separation on January 13th. The problem was as follows. $OPEN typically would follow $CLOSE with no data separating the two. It sometimes was difficult for the user --- even for this program developer, who had access to source code, and understood it well --- to be sure which was which. So, the single message was split in order to end the possible confusion. The associated two lines of error will mention either OPEN or CLOSE, but not both. Thereafter, all output should be ignored, as ATP attempts to halt execution in a civilized fashion (not always trivial, considering location of recognition the error down within SUBROUTINE CIMAGE).

$BLANK COMMENT was explained in the October, 1993, issue. But use was not illustrated in any test case prior to an addition to DC-6 on February 20th.

SET VARIABLE of conditional data assembly was mentioned in the July, 1998, issue. Unfortunately, Salford DBOS failed to set the associated environmental variable as hoped, and Watcom ATP was no better (see the October, 1998, issue). A believable explanation finally has been offered by someone who knows more about such matters than program developers at BPA. The following was received January 12th from Stephen Boroczky of TransGrid in Sydney, Australia: *"I **suspect** that the problem in DOS is that by performing a system call from within the program, the program first creates a new 'process/command' space in memory for it to execute the command. This is OK if the program affects things that are outside its process space, but any changes made to the process are lost when it returns to the calling program."*

TSTALL replaced PIXPUN in STARTUP beginning April 13th. So began a paragraph in the July, 1998, issue. The October issue then explained protection against data values that exceed 0.5 seconds: *"... the monitor will be notified instantly of the reason for subsequent delays."* Unfortunately, as discovered, corrected, and explained in public E-mail dated January 29th, the warning was not instantaneous. Rather, it, too, was delayed prior to a correction to universal RSTART.

The DATA command of SPY did not operate properly in interactive mode (i.e., through the keyboard) prior to correction on February 20th. This is the date the problem was described in E-mail of the Fargo list server from Dennis Long of the Omaha Public Power District in Nebraska. As your Editor explained later that same day, trouble seemed to date to that increase in the limiting length of file names. Currently, the storage is 132 bytes long (see the July and October issues of the newsletter last year). A new variable was required for this, and the SPY DATA command had not been connected to it. ... Note mention of *interactive*. Batch-mode use of SPY is demonstrated by standard test cases DC-56 and DC-57, which always have been simulated correctly.

"KOMPAR = 1" was found in the interpretation of comment lines for any positive value of the control parameter. This was prior to correction on February 5th. At issue was fixed text rather than numeric serialization. Apparently the text being used was years old, dating to the original implementation when KOMPAR was allowed only two values (0 or 1). Following correction, a positive value will be mentioned rather than value one.

A delta connection of compensation-based elements results in a singular Thevenin impedance matrix [Z-th]. Subscribers to the Fargo list server were reminded of this fact on January 6th following a question about Type-94 elements and MODELS by Marjan Popov, writing from Delft in The Netherlands. Your Editor observed: *"The 3x3 Thevenin impedance matrix for any 3 delta-connected, compensation-based elements will always be singular ... Voltage vectors V-ab, V-bc, and V-ca are never linearly independent; Kirchhoff requires that they add to zero. If a matrix of order 3 is required, there is a fundamental and irreconcilable problem. If it has not yet been tried, this writer would recommend breaking the delta by the insertion of a small measuring impedance."* Later that same day, Orlando Hevia of Universidad Tecnologica Nacional in Santa Fe, Argentina, quickly confirmed success of the technique: *"Some hours ago I sent this solution to Mr. Popov. I inserted a resistance of 0.001 Ohm in the triangle. With this, the case runs correctly."* Of course, the resistance can not be arbitrarily small. Mr. Hevia described troubles with a much smaller value, and concluded with the observation that *"the user ... must investigate the adequate value of this resistance."* True, just as for any current measuring branch (as opposed to a switch).