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 : atp @ agora . rain . com
Vol. 97- 1; January, 1997

Table of Contents

Salford Compiler and DOS Extender

Anonymous DBOS was mentioned in the preceding issue. Since then, there has been widespread testing thanks to free E-mail distribution by Masahiro Kan of Toshiba Corporation in Japan, and Prof. Laszlo Prikler of T.U. Budapest in Hungary. These two individuals volunteered their services in response to your Editor's call for help in public E-mail of the Fargo list server dated December 5th. Since then, there have been several stories of success, and a few of failure. No one yet has related the new anonymous DBOS to Trident, however. Recall the preceding issue had a story about the successful experimentation of Cornel Brozio, who now is a doctoral student at the University of Stellenbosch in South Africa. Your Editor asked : *"Would new DBOS make such special Trident drivers unnecessary? If so, what are the advantages and disadvantages, if any?"*

The most comprehensive single report of anonymous DBOS trials was from Mr. Kan in public E-mail dated December 27th. This included the following results for five Pentium computers running the Japanese MS Windows 95 operating system Win95J :

MHz	Graphic chip	Computer class	Result of testing Salford Dbos 2.71	Dbos3.5
200	S3-968	Tower	Terminated	Good
166	S3-VIRGE	Tower	Terminated	Good
150	CT-65550	Notebook	Invalid video mode	Good
100	CT-65548	Notebook	Not tested	Good
133	S3-Trio64V+	Desktop	Invalid video mode	Good

"In conclusion, by using DBOS 3.5, TPLOT worked well for all PCs we tested," according to Mr. Kan.

Prof. Prikler described some finer details of the new DBOS in E-mail dated December 16: *"I have found some changes in the intro screen also: 1) Detection of the processor type is more precise; 2) Detection of the Windows version has improved (Win95 is detected now). Some enhancement has also been observed when DBOS was running under Win95. The Device=emm386.exe setting does not seem necessary in the config.sys file."*

Anonymous DBOS Version 3.5 does require more disk space, although the difference is more an issue for the

supplier than it is for the end user. For example, transfers by E-mail will take longer; and the DBOS floppy disk will have less otherwise unused space for newsletters. Whereas DBOS2P71.ZIP that is used to carry DBOS Version 2.71 is 655 Kbytes in size, the new Version 3.5 requires 919 Kbytes (size of the floppy disk archive that is kept in *E:\EMAIL* on Dr. Liu's Pentium).

Win95J, the Japanese-language version of Windows 95, has failed a test for possible support of Salford EMTP and TPPLOT on otherwise-incompatible PCs that are sold within Japan (the domestic market, which is not all IBM-compatible). Masahiro Kan of the Hamakawasaki Works of Toshiba Corporation reported the following in E-mail dated November 5th : *"My colleague finally tested Salford ATP using an NEC PC with Win95J. .... When he ran DBOS.EXE, the PC was hung up. The NEC PC was revealed to be incompatible with Salford ATP even under Win95J."* Unless and/or until some reader can report a contrary finding, it must be assumed that there is a fundamental problem. As with lack of Salford timesharing under Win95, it is difficult to absolutely sure (how does one prove a negative?). But hope has faded.

JDELAY of STARTUP was mentioned in the July newsletter. Already, it has been needed (i.e., that code in OVER12 was not as unnecessary as had been hoped). A simple flashover switch with delayed opening T-delay was affected adversely by the change to OVER12. This was first observed by Dr. Hiroshi Okamoto of Tokyo Electric Power Company (TEPCO) in Japan in E-mail dated November 1st. Your Editor's advice was: *"I believe you can avoid the problem using JDELAY = -55666 in STARTUP. But that would be for existing TPBIG. For those having source code, the better way seems to be the addition of 4 lines: the DO 7352 loop show below. This is located"* immediately below S.N. 7301.

Compiled TACS is described in a separate story. As a result, the speed of compiling and linking has become more important for users, so shall be reviewed. Relinking is the bottleneck for Salford, it should be understood. Salford compilation is so fast that it takes negligible time for any modest-sized chunk the user might supply. But ATP is a very large program, so conventional linking with the user's special control system modeling is a substantial burden. There are two general strategies for minimizing this: 1) use of external text; and 2) dynamic link libraries. The next two paragraphs expand upon these.

LUNTEX of STARTUP determines whether text is to be internal rather than external. This dates to the Schultz Revolution in the fall of 1993. Prior to that, text was always stored outside the program; it was to be found in disk file BLOCKD51.BIN (see reference in STARTUP). Well, making the text internal substantially delays the linking. By omitting this, the following elapsed times are obtained under Win95 on your Editor's 133-MHz Pentium. These are for the best 5 of 6 trials after an MS-DOS window was opened, and DBOS was started. The times to link without text are (in sec): 3.90 3.84 3.84 3.90 and 3.95 seconds. Subsequent simulation speed seems not to be noticeably changed for DC-22d. That new second row of the table in the story about compiled TACS was for use of BLOCKD51.BIN (external text). Relinking with text, and repeating the experiment, resulted in the following simulation times:

```
Compiled :  3.516  3.407  3.407  3.407  3.407
```

The average here is 3.429 --- slightly longer than the 3.418 reported in that separate story. There seems to be no meaningful difference. Three years later, it is difficult to recall what sort of data cases might have been accelerated by making the text internal. Note the importance of disk caching: the 2048 Kbytes that are given to MS-DOS SMARTDrive seem to be effectively holding the required parts of BLOCKD51.BIN in RAM. If this were not the case, the difference would be expected to be substantial.

Creating a dynamic link library COMTAC.LIB saves most of the preceding time of linking. For DC-22d, the times for both compiling and linking are 1.15 0.93 1.37 0.93 and 0.99 seconds. On the other hand, use of such a Salford DBOS library prevents use of the symbolic debugger (a diagnostic tool that is useful in cases of trouble). Whether the gain of 3 or 4 seconds is worth this price is debatable. For developers, it certainly is not.

NOGNU is a new parameter of STARTUP that was added to either approve or prohibit (values zero or unity, respectively) output to disk for the freeware plotting program GNUPLOT as explained in a separate story.

Improvements to Salford TPPLOT

Fourier series accuracy was improved October 25th in response to a question by Ashok Parsotam of Southpower in Christchurch, New Zealand. October 21st, he observed that, for two superimposed sinusoidal signals, resolution was not good. Furthermore, using twice the period (half the frequency) resulted in substantially worse resolution. Several readers provided interesting and useful discussion. But in the end, the trouble was very basic. Your Editor wrote : Both truncation and roundoff errors contributed to the problem. By adding 13 REAL*8 declarations, and changing SIN and COS to DSIN and DCOS, the following results were obtained:

```
 Harmonic   ....      Complex        Fraction of
   number   ....      amplitude      fundamental
        0   ....  0.000000000E+00      0.0000000
        1   ....  1.000000000E+02      1.0000001
        2   ....  6.503260153E-10      0.0000000
        3   ....  1.000000119E+00      0.0100000
        4   ....  2.936040855E-10      0.0000000
        5   ....  4.684128996E-08      0.0000000
        6   ....  1.692655599E-10      0.0000000
```

Considering that the 2000 input points are REAL*4, these results seem remarkably good. The noise is 9 or 10 orders of magnitude below the peak, and that third harmonic is in error only in the eighth decimal place.

NOGNU of integer index 273 is a new binary switch that controls possible output for freeware GNUPLOT (see separate story). This parallels the older NOPOST and NOHPGL switches, but with an important difference: the new NOGNU applies only to **non**-WINDOW plotting whereas the two older switches apply only to WINDOW plotting. Any program dated December 22nd or later is able to produce the new output files TPPGNU.001, etc. A GNU button has not been added to the PAPER command, it will be noted, because this now is viewed as the old and more cumbersome way. PCs are so fast today that it makes more sense to create auxiliary text outputs automatically (if not prohibited by a switch such as NOGNU = 1), rather than force the user to follow his screen plot by a special request. This was the learning experience associated with WINDOW plotting. If PostScript and HP-GL were being added today, they would be produced automatically for non-WINDOW plots, too. According to current thinking, only the EPSON and LJ2 buttons really belong under the PAPER command. They are fundamentally different in that such output can not be done in parallel (in fact, the EPSON and LJ2 outputs are created by the redirection of Salford screen graphics to disk).

News from Outside USA and Canada

A printed copy of the July and October newsletters was air mailed by BPA to each of its 9 primary EMTP contacts on January 17th,. In the preceding issue, the question of need for this paper was raised. Since no one responded, that former service has been discontinued. The October newsletter OCT96.DOC was E-mailed on December 12th to Prof. Bruce Mork in Houghton and Mathias Noe in Hannover for placement on the aFTP servers that they provide. This was as an MS Word for Windows file. The WP 5.1 alternative followed 11 days later. Manual conversion again was used, but was quite different than 3 months earlier. Using Dr. Liu's Pentium at BPA, the real WP 5.1 did not crash once within the DOS window of WinNT that was being used. On the other hand, many more manual changes were required. For some reason, most quotation marks and apostrophes appeared as slugs (solid rectangles), so they required manual conversion. Curiously, fonts did **not** require modification as had been the case for preceding issues. Time Roman was seen in the WP 5.1 output of MS Word for Windows. This was an encouraging sign, anyway. On the down side, page numbers at the bottom took up too much room, so they were deleted. The result is just a line or two short of the exactly 20 pages making up OCT96.DOC

An Indian ATP user group has resulted from the persistence of Prof. Ned Mohan of the University of Minnesota in Minneapolis. As mentioned in the July, 1996, issue, Prof. Hariharan, retired from IIT-Bombay, now is a possible local source of ATP and ATPDRAW for MS-DOS PCs. Details are summarized in the WP 5.1 form letter INDIALIC.ZIP that can be found in the directory **/pub/atp/license/india** of the Houghton aFTP server. In private E-mail dated November 15th, Prof. Bruce Mork summarized his intention to *"add a direct ftp link from the page where all the user's groups are listed so that people can just click on the license document to download."* Three days later, he confirmed: *"I've added Prof. Hariharan's name and address to the web page on user's groups."*

Advertising for ATP-related short courses can be found from the Web page that is maintained by Prof. Bruce Mork of Michigan Tech in Houghton. In public E-mail dated January 8th, he commented about Prof. Dennis Carroll's upcoming 4-day course in Florida as follows: *"I've put the updated course / registration info on the web site. It is accessible at* ***http://www.ee.mtu.edu/atp*** *-- click on "Training". You can then view and/or download the information. Or if you want to download it directly via ftp, it's at* ***ftp://ftp.ee.mtu.edu/pub/atp/course/9703carr.txt*** *Check the web page periodically for updates on seminars given by Dennis Carroll, Ned Mohan, and EEUG."*

South African user group relicensing has been slow as mentioned in the preceding issue. Yet, no fundamental problem is believed to exist. The last public word on the subject would seem to be your Editor's, who reported as follows on November 9th : *"Before new ATP can be distributed to others in that country by the user group at ESKOM (the national electric power authority), modern (post-LEC) licensing is required. Well, a Can/Am-approved form letter is to be sent back to Dr. Warren Levy of ESKOM no later than tomorrow. Yesterday, Dr. Tsu-huei Liu made several modifications to the ESKOM proposal that was dated October 16th. Only margins and fonts now remain for this writer to adjust."* Then 3 days later: *"Can/Am-approved SALICENS.ZIP was sent to Dr. Warren Levy of ESKOM on November 10th. .... Most changes to the ESKOM proposal had to do with a recent change of format. As this writer explained to Dr. Levy : The European idea was to have a 1-sheet form letter. It is possible if margins are minimal and a proportional font is used. We did this for your document, which looks good here on our screens using 10-point Times Roman. I hope you have something comparable. Of course, readjust margins and font in minor ways, should this be required to remain within the 2-page limit (a single sheet of paper)."* The able assistance of Cornel Brozio (see preceding issue) was acknowledged in that October 16th mail, it is noted.

Mainland China (the People's Republic) was the source of E-mail expressing interest in ATP. In a message dated December 4th, Dr Warren Levy of ESKOM in South Africa wondered why **he** had been approached. The attached inquiry was from Prof. Yuan Bin at Tsinghua University in Beijing. From **yuan@pwrs.eea.tsinghua.edu.cn** came the question: *"Would you please tell me where can I download atp-emtp source code, so that I can install it in my Lab."* For the record, the answer is *nowhere*. Source code is not generally available; and, in any case, ATP

materials of any type except newsletters are not available to the general public prior to ATP licensing.

The Latin American EMTP Users Group, meanwhile, has been moving ahead largely on its own. It was gratifying to learn the following on January 6th from E-mail that was sent by Marco Polo Pereira of Furnas in Rio de Janeiro, Brazil: *"We have already a home-page .... and there we have a list of licensed users exactly in the way you mentioned. Just the name of the representative, the company, the city and the country. I can send this part to the ATP server as soon it is required. The home-page was created almost one year ago, and it was completely reformulated last October. Unfortunately it is written in Portuguese, although it contains some documents in its original form, like the 'licensing agreement' and the Can/Am EMTP News, for example. The basic idea was to complement and/or to facilitate the available information on the other ATP sites. .... Main topics are: 1) How to join the Latin American Users Group; 2) Purpose of the Latin American Users Group; 3) Summary of ATP / EMTP History; 4) ATP / EMTP Items (brief description, ATPDRAW, licensing conditions, licensing agreement, Can/Am EMTP News and Primer for Beginners); 5) Latin American Users Group Meetings; and 6) Other sites of interest. .... Some materials can be transferred directly, such as Can/Am EMTP News, ATPDRAW, Meeting Notes, and the Primer for Beginners (maybe these two last ones are examples of information that are not available from other sites). I am finishing the notes of the last meeting that occurred in the beginning of December and just after that I have the intention to make an announcement through the EMTP List-server. Maybe people from Spain and Portugal can have some interest in visiting the Latin American Users Group home-page and find some interesting thing there, who knows? .... the home-page in question has address* ***http://www.furnas.gov.br/atp*** *.... Of course, comments will be welcomed."*

More about Electronic Mail (E-mail)

UUEX and UUDX are Japanese shareware that are being used at BPA to circumvent the recent 1-Mbyte limit on outgoing E-mail. See a separate story for details.

Incompatible UUENCODE that did not involve file segmentation must be reported, unfortunately. In the preceding issue, there was mention of such a possible problem in conjunction with UUCODE by Sabasoft. But that had to do with the marking and linkage of two or more segments (for big files). More frightening was the trouble your Editor had handling the small attachment KFD.ZIP from Prof. Mustafa Kizilcay of FH Osnabrueck in Germany on December 29th. "UU-DECODE 4.13 FOR PC. by Richard Marks" complained as follows: "Line 11: Line too short or too long." AOL seems to have some involvement since the first line of KFD.UUE is the following: "--PART.BOUNDARY.0.19029.emout20.mail.aol.com.85188 8787". Later, there is mention of "< uuencode 95 (v36) by R.E.M. >" Obviously different is line termination on the right. No longer are lines of uniform length when viewed using Vernon Buerg's shareware LIST. There is a jagged right edge. For example, the first 5 lines appear as follows, after removal of the first 25 bytes of each:

```
... 0PH''-4R''')''''=3DW)O;F<N9&%T[5M;;]LX
... 8#*#:3&+&=3DU**3[;PLW$1)C#AR1I+;;9_G5\P/
... QT.>JW380W[<Z:,^OT1M7_BHY0=3D\?Z\%_Q'2
... LP09U^ZVPPX"V!XC"93=3D($?1FG=3DVK(:>?W3K?E
... \&8ZF4R_C*-;E$P?0C1-[\)8#=3D,4,$U\K<'&
```

What is happening, readers? Is it time the user group updates its copy of UUDECODE (on the GIVE2 disk)? Alternatively, is this v36 software perhaps not freeware?

Japan seems to have its own list server for ATP-related information. Not surprisingly, this is yet another initiative of Masahiro Kan of the Hamakawasaki Works of Toshiba Corporation. The story was scheduled to appear in the preceding issue, but it was delayed in a last-minute search for more space. So, the details will be summarized now. Seeking clarification, your Editor wrote to Mr. Kan as follows on October 8th: *"So, your Japanese subscribers can share messages in Japanese language just as the 400 subscribers of the Fargo list server do in English? If so, this is wonderful, and deserves to be explained in a newsletter."* Initially, Mr. Kan was reluctant, explaining that *"there are only 13 subscribers and there are few messages now."* But this was not at all the perspective of your Editor. Recall how slowly Prof. Mork's operation grew during the early years. A year and a half after massive exposure (3 pages of explanation) in the October, 1991, newsletter, Prof. Mork's report of 105 subscribers can be found in the April, 1993, issue. At the time, this seemed like an overwhelming number, even though it applied to the whole world rather than just one country having half the population of the United States. To have 13 subscribers after 3 months or so has to be viewed as a substantial success, and a possible model for others in non-English-speaking countries.

Martin Jones finally seems to have reliable E-mail once again. Of course, he had good service while a doctoral student at Nottingham University in England, where he performed that pioneering conversion of LEC Rule Book files from Lotus Manuscript to WP 5.1 But while at GEC Alsthom, he has had an unbelievable struggle. A message dated November 29th explained the trouble as follows: *"we originally had an X400 link for our comms with India and we got this same provider to provide Internet access (a task which they were never really set up to do). Anyway, we now have a fully functional system and all is ok .... If you wish to mention my email address in the newsletter, then you may as well do so now, as I gave this system a good hammering yesterday without a single failure."* The new address is **jonesms@pcs.gecalsthom.co.uk**

"America Online, posting huge loss, writes off costs of acquiring members" is the headline of a story about AOL's latest financial problems on page B2 of the November 8th

Wall Street Journal. Yes, AOL has grown fast, but at a cost that was not being supported by the business: *"The nation's largest computer on-line service historically deferred those costs over a two-year period instead of expensing them up front -- allowing the company to report quarterly profits and postponing the day of reckoning for its marketing expenditures."* How big was the loss? *"Nearly $354 million after taking a whopping $385 million charge to abandon"* that sneaky accounting practice. Historically, AOL has made money charging by the minute, but that strategy, too, no longer seems viable: *"Last week, the company abandoned its hourly charges by creating a new flat-rate price of $19.95 a month."*

Meanwhile, CompuServe loses both subscribers **and** money. *"CompuServe posts big quarterly loss"* is the headline on page A3 of the November 22nd *Wall Street Journal*. The story begins with an explanation that CS *"set plans to retreat from the on-line consumer market and shut down a family service unveiled less than a year ago."* March 25 is the date shown for the introduction of *"WOW with flat rate."* This was no small initiative, either: *"The company spent a year and $70 million developing and marketing the slick service, launching it at a party hosted by magicians Penn & Teller in a Broadway theater. But family users appear to have played a disappearing act. The service had only about 100,000 subscribers, and had added only 10,000 new customers in the October quarter. WOW used a flat-fee approach, charging $17.95 for unlimited access, but the low pricing appealed more to sophisticated heavy users who went on-line for hours at a time, driving up the company's costs."*

Robert Schultz of NYPA has been given a new home E-mail address, and a better price, according to a message received December 10th: *"North American Internet has just completed a multi-million dollar system upgrade. As a result of this, my email server is changing. The new address is* ***rschultz@ct2.nai.net*** *The good news is that rates have dropped. Unlimited full Internet access is now $17.95 per month. Light usage(7 hours per month) is just $4.95 per month! ($2/hr over 7 hours). Some things do change for the better!"*

Growth of WWW usage is summarized in a graph on page B6 of The Wall Street Journal dated September 26th. Labeling states: *"The Web Index grew about 6% in September as the number of users on the World Wide Web edged up by about the same amount."* To what number? Some 30 million Web users world-wide. Elsewhere on the same page is a story about the location of Web users. The pie chart shows 73% in the US, 11% in Europe, 8% in Canada/Mexico, 4% in Australia/New Zealand. Why the American dominance? Competitive telecommunications companies seem to be the single, most-logical reason: *"In many European and Asian countries, telecommunications monopolies have impeded the growth of the Net .... In many countries the price for access can be as high as $125 a month plus the cost of a long-distance phone call."* The largest per capita use probably is in Antarctica, though, where 0.1% of all Web users are said to be found (do you believe that? Could there be 30K users down there?). No doubt all use high-speed satellite links, too!

European EMTP User Group (EEUG)

The 1996 annual EEUG meeting did occur November 10-12 in Budapest, Hungary. The following is a quick summary of those many papers that were mentioned last time. In the order of presentation, based on a disk file that was received from Prof. Laszlo Prikler on October 25th, paper titles were: 1) Multiconductor transmission line model for multiwinding transformer for very fast transient studies. 2) Simulation of low inductive current switching processes using ATP-MODELS. 3) Corona modelling for attenuation and distortion of lightning surges in overhead transmission lines 4) Computation of transients on high voltage transmission lines 5) Evaluation of existing secondary arc models 6) EMTP/ATP study of transient phenomena for series-compensated line protection 7) Simulation of transients in a thyristor-controlled series compensation scheme using ATPDRAW and MODELS 8) Performance test of the three-phase induction machine model 9) Transient behaviour of a variable speed slip energy recovery system for wind power generation 10) Fault location algorithm testing using MODELS 11) Relay protection algorithm testing with off-line simulation using ATP 12) ATP modelling of electromechanical distance relays. Part I: Modelling Options 13) ATP modelling of electromechanical distance relays. Part II: Case study 14) How to use MODELS-based user-defined network components in ATP 15) Working with MODELS foreign models in windows environment 16) Automatic massive EMTP simulations for relay algorithm testing 17) Future plans and latest developments of ATPDRAW 18) Round table discussion 19) Transferred voltages on the low voltage bus of a 220/15.75 kV step-up transformer loaded by auxiliary transformer 20) Overvoltage protection of the new Croatian 400 kV substation 21) Comparison of measurement and calculation of surge propagation in a 400 kV transmission system 22) Simulation of switching overvoltages of a 110 kV condenser bank by using ATP program 23) ATPDRAW application examples: Overvoltage protection of LV and HV networks 24) Teaching surge arrester protection using the ATP 25) Monte Carlo lightning backflash model for EHV lines - a MODELS application example 26) Using ATP as a source of data for training and testing of ANNs detecting current transformer saturation 27) Point of wave circuit breaker closing established by an expert system based EMTP rules

EEUG membership dues for 1997 are to remain fixed at 1000 DM for companies, and are to be reduced by 20% for universities --- from 500 DM to 400 DM. There were two proposals, and this one passed by the required 2/3rds majority. The proposal that did not pass would have raised industrial dues by 10%, also.

Prof. Laszlo Prikler is to be passed a small amount of money ("*DM 10000 for the year 1997*") to help support his ATP-related services. According to Chairman Kizilcay, "*his tasks are: a) publication of 4 issues of EEUG News in cooperation with me; b) provide some on-line services (by Email) to new members (help to install and to run ATP and related programs); c) Extend Web pages for the EEUG at Techn. Univ. of Budapest.*" This sounds very good. As your Editor has observed many times, Prof. Prikler is a proven contributor to the ATP user community. The institutional support of his activities is important.

The integrity of EEUG accounting is to be insured by two new auditors. This should be boring but important work, recall, in light of the creative accounting that was performed in Leuven, Belgium, prior to the collapse of LEC in 1993 (see July and October newsletters of that year). According to Prof. Kizilcay, "*Harald Wehrend from SEG and Prof. Juergen Pannicke from Fachhochschule Berlin were elected by members as auditors for the fiscal year 1996.*"

Barcelona, Spain, is to be the site of the 1997 annual meeting, and Prof. Juan Martinez of Universitat Politecnica de Catalunya is to be responsible for all arrangements of the meeting. According to Prof. Kizilcay, the proposal by Prof. Martinez was "*accepted by the members. After the meeting, there will be an EMTP short course (2-3 days) in Barcelona.*"

IBM OS / 2 Warp Used by NYPA

IBM's OS / 2 is being used for the support of ATP by Robert Meredith and Robert Schultz of NYPA (the New York Power Authority in White Plains).

WAT4GW is the free Watcom DOS extender that is used at BPA for compilation, linking, and ATP execution before a new translation is shared with NYPA for processing using its more attractive alternatives (OS/2, WinNT, and Win95). Well, the good news about Dr. Tsu-huei Liu's Dell Pentium (see preceding issue) is that WAT4GW seems to be compatible with the DOS shell of WinNT. Salford DBOS is not, but WAT4GW **is,** as first discovered on December 18th.

The Pentium with NT did crash once while debugging a problem with Watcom ATP (next paragraph), and/or the subsequent validation of a new (current) translation. Just based on this number, there is no concern (Salford DBOS crashes more often under Windows 95 on your Editor's Pentium). But the crash in question was disturbing. It occurred while using Mike Albert's shareware FC, which never ended (the reset button had to be pushed). At the time, the floppy disk light was on (permanently), and the DOS emulator of NT somehow corrupted the floppy disk --- not for real DOS, or even the DOS of Win95, but rather for NT DOS use. After restarting, FC and DIR both failed to read the floppy. But, a couple of days later, after Win95 had written to it, the floppy became readable once again using the NT DOS emulator.

Use of LUNIT4 = 4 rather than -4 as distributed by the user group made BENCHMARK DC-32 and 49 fail in unpredictable ways. The trouble was first reported on December 19th, in E-mail from Harald Wehrend of SEG in Kempen, Germany. C-like .PL4 files were not easily made using Watcom FORTRAN, and there may be more such problems if details are changed. Yet, this one trouble was removed. As reported December 24th, 32 and 49 can be simulated as a matched pair using DISK PLOT DATA to toggle the sign of LUNIT4. This one change was believed to be important enough to warrant extra work. Whether others might be is not known. No promises are being made. Watcom ATP users are advised not to tamper unnecessarily with what works as demonstrated by all DC*.DAT data.

Quality time-sharing of OS/2 always was an advantage that Messrs. Meredith and Schultz had cited. So how is the Microsoft competition doing? Prof. Mustafa Kizilcay of FH Osnabrueck wrote as follows about Win95 in E-mail dated December 2nd: "*Watcom ATP is running without problems in the background under Windows 95, but the computer is slowing down significantly. I use a 133-MHz Pentium with 32 MByte RAM. But if I execute a big data case with large T-max, editing files in the foreground becomes very slow ....*" That is your Editor's impression, too, even though ATP is not involved. While editing within a DOS window, the response becomes annoyingly jerky whenever Win95 occasionally and unexpectedly decides to perform some unknown, system-related activity in the background. The time-sharing is lousy, considering how powerful the processor is, and how much RAM it has.

PCPLOT was enhanced by Prof. Kizilcay in order to recognize automatically the universal C-like .PL4 files that can be created by many ATP versions including Watcom ATP. This is small but important progress for Watcom ATP users who can not use TPPLOT (e.g., those who use Salford DBOS-incompatible OS/2 or WinNT). In E-mail dated December 29th, Prof. Kizilcay wrote: "*PCPLOT version 6.50 can correctly read .PL4 files of Watcom ATP .... I will send you the new version to be included on the GIVE2 disk ....*"

DOS under WinNT is inconsistent with DOS under Win95 in important and troublesome ways. Thus began a paragraph in the preceding issue. Now, following some debugging of Watcom ATP using Dr. Liu's Pentium, yet another good example can be cited. DOSKEY no longer always begins with the most recent command as it does under Win95. Within a DOS window of NT, DOSKEY has memory. I.e., it will begin where it was left the last time. The complaint is not that this is bad, but rather that it is inconsistent and aggravating for those who must use more than one flavor of MS Windows at a time.

New User Interface for Compensation

True nonlinear elements, and time-varying elements, are solved by compensation, which implies superposition for the total solution in the linear portions of the network. This is old. What is new is a user-accessible interface which allows the user with a FORTRAN compiler to provide his own models. If all one wants to do is model rotating machinery, why be forced to use inefficient MODELS as in years past? Rather than program in Dube's arcane simulation language, the user can program in much more powerful, ANSI-standard FORTRAN and be rewarded by much faster execution as a result.

Note how compiled TACS nicely complements this just-summarized use of compensation. If the user has a compiler anyway (e.g., for compiled TACS, to simulate control system modeling at warp speed as explained in a separate story), why not exploit the compiler for better nonlinear element modeling, too? Historically, this is the way compensation was generalized: compiled TACS came first. Then came the desire to connect TACS to compensation the way MODELS can be connected (Laurent Dube's Type-94 element). *But why the same way?* your Editor suddenly asked himself. As observed in the preceding issue, some of the development that is associated with MODELS is not at all inherently or naturally associated. In the case of rotating machine dynamics, who really wants either MODELS or TACS, anyway? What the user wants is the machine equations. The fact that MODELS was required for access was an inconvenience rather than an advantage. Well, if the user has a compiler, that inconvenience now can be avoided.

"New, accessible compensation interface ... " was the title of public E-mail that introduced this subject to the general public on December 4th. It introduced the new user interface as follows: For a multi-phase case, the user simply defines a Type-91 branch for each phase, with "MODEL " in the BUS3 field of columns 15-20. The model name follows this in the BUS4 field of columns 21-26. To illustrate for a Wye-connected but ungrounded armature of a 3-phase machine at bus MACH (node EARTH separately would be connected to ground by an R-L branch, typically :

```
C 34567890123456789012345678
91MACHA EARTH MODEL RTIME
91MACHB EARTH MODEL RTIME
91MACHC EARTH MODEL RTIME
```

The model name is assumed to be the subroutine name, too. That is, SUBROUTINE RTIME will be passed the 3-phase Thevenin equivalent of these 3 branches at each time step. The job of RTIME is to solve the machine equations along with the linear Thevenin constraint V = [Z-thev] * I + E-thev

So, ATP data is trivial. But what about code for such a user-supplied motor? What are the choices? First, between the use of actual coil variables and the use of Park's equations with d-q-0 variables. Recall the story about TEPCO S.M. modeling in the October newsletter. We now must ask ourselves: should we still be concerned with the artificial complexity of variable transformation when machine equations can more easily be programmed in coil variables? Who might be willing and able to contribute FORTRAN for such a solution? This writer could easily enough solve all simultaneous equations by brute force using Newton's method; but this probably is a case of overkill. What reader can recommend a more efficient algorithm for typical usage? This writer does not want to reinvent the wheel unless necessary.

Piracy is the primary concern about the new interface. How can commercial exploitation of it best be prevented? Or, do ATP owners even want to prevent such activity? Perhaps commercial exploitation of the new interface should be encouraged. A separate story summarizes a related problem of how one merchant tried to sell his own modeling that would be used by ATP for better (in some sense) simulation. That did not involve compensation; the coupling to ATP was tighter than that. But the idea is the same. Take Prof. Hian Lauw's Universal Machine (mentioned in the separate story) as an illustration. The original mode of U.M. use involved compensation. As a hypothetical illustration, suppose today, a decade and a half later, Dr. Lauw decided he wanted to generalize or otherwise improve his U.M. code within ATP. Should he be allowed or encouraged to do so in the form of a separate commercial product that each end user would need to purchase from him? Think long and hard about it, readers. These are dangerous, uncharted waters that Good Ship ATP could be sailing into.

Generator modeling using phasor Zo, Z1, and Z2 were of interest to Yong Guo of Texas A&M University in College Station. His inquiry by public E-mail dated December 11th asked how, and your editor explained in response that the new compensation model would be one way since Type-51, 52, 53 modeling does not work. The problem is that ATP requires symmetric [R] + j [X]. If Z1 does not equal Z2, this requirement generally will not be satisfied ..... However, one could pull the generator impedance outside the network as though it were a 3-phase, nonlinear or time-varying element. Although the entire 3x3 matrix could be represented this way, in fact only a correction should be required. That is, a Type-51, 52, 53 branch could be used for diag [Z0, Z1, Z1] and compensation would then be used only for the correction diag [0, 0, Z2-Z1] (both matrices in the symmetrical component domain). This is believed to have two advantages. First, for the phasor solution, which would have no compensation, at least generator impedance would not be totally missing (a symmetric approximation should be better than none at all). Second, once in the time-step loop, the use of parallel admittances (Norton equivalents) should guarantee connectivity regardless of what might be connected. That is, there should be no problem with singularity of the 3-phase [Z-thev]. Who is willing to

supply some typical numbers of a simple illustration? Is any internal voltage required? If conventional positive- or negative-sequence execution were desired, note that it could be handled externally using ordinary Type-14 sources. Unless some interested reader can show otherwise, it would seem that only 1 complex number, Z2 - Z1, would be needed as input to the new component. This would be very simple, indeed.

News About TACS and MODELS

This is a continuation of the story having the same headline in the preceding issue. Recall the final paragraph, which stated that *"MODELS is much slower than the above ratios would indicate .... Also, DEC Alpha running Open VMS at BPA and TransGrid displays even more of a handicap for MODELS simulation than your Editor's Pentium."* Finally, *"a plausible explanation by NYPA of the cause, and a new standard of comparison : compiled TACS ."* All of these items shall now be addressed with the exception of compiled TACS, which is the subject of a separate story.

Details of how slow MODELS was for DEC Alpha first were revealed publicly in E-mail of the Fargo list server dated October 22nd. It stated: Alpha running Open VMS at about 133 Mhz (Randy Suhrbier said a Model 3000/400 is involved) is particularly MODELS - handicapped. .... BPA's James Randall took the 3 data files used in Report III (remember, no plot data points at all), and submitted them for batch-mode execution. His seconds spent within the time-step loop tabulate as follows:

```
DC-22d :  4.940   4.920   4.960   4.820   4.940
DC-22e : 61.220  56.090  56.780  57.170  57.610
DUBE22 : 44.720  42.870  43.720  43.400  44.510
```

Using the average of each group:

```
Furst  TACS   of  DC-22d :  4.916     ratio
Furst MODELS  of  DC-22e : 57.774    11.752
Dube MODELS   of  DUBE22 : 43.844     8.919
```

Stephen Boroczky of TransGrid in Sydney, Australia, then showed that even faster Alpha looked even worse for MODELS relative to TACS. His public report on October 24th was for an AlphaServer 2100 4/275 with 275-Mhz EV4 Alpha CPU running Open VMS V6.1. His raw data parallel to the above was:

```
DC-22d :  3.480   3.530   4.030   4.050   3.780
DC-22e : 49.170  51.160  46.100  51.200  50.260
DUBE22 : 36.380  36.120  36.360  36.320  33.830
```

Again using the averages, ratios of MODELS time to TACS time become:

```
Furst  TACS   of  DC-22d :  3.774    ratio
Furst MODELS  of  DC-22e : 49.578    13.137
Dube MODELS   of  DUBE22 : 35.802     9.486
```

As to why the double-speed processor at TransGrid is only slightly faster, Mr. Boroczky wrote: *"This is very strange as we do get roughly double the performance on our other applications."*

External (secondary) cache memory might be slowing some fast computers more than others for MODELS simulation. This was explained by Robert Meredith and Robert Schultz of NYPA (New York Power Authority) in White Plains. Quoting from private mail dated October 23rd : *"it is likely that this extremely poor programming technique .... is 'breaking the cache' on both Intel and RISC systems with the latter suffering the effects the worst. RISC machines are notorious for needing intensive optimization (a la HPUX Fortran's multiple levels of restructuring and optimization) and mandatory large secondary caches (1-2 MB minimum)."* Note this nicely explains the TransGrid disappointment: if the bottleneck is outside the processor, a faster chip will not help much.

These MODELS - to - TACS ratios of 10 : 1 seem bad enough, but the ratios actually are larger. Why? Because the constant time for everything except control system modeling should be subtracted from all of the preceding times. After all, work on the electric network has nothing to do with either TACS or MODELS. As stated in the separate story about compiled TACS, a Watcom profile from NYPA was used to estimate that 60% of the time was spent on the electric network. The percentage must be slightly less for your Editor's Pentium, obviously, since compiled TACS is very fast, but can not require negative time. Taking the average time for compiled TACS (3.418 sec) as the time outside of control system modeling, this is subtracted from the table of the preceding issue to yield :

```
Furst  TACS   of  DC-22d :  2.703    Ratio
Furst MODELS  of  DC-22e : 53.538    19.81
Dube  MODELS  of  DUBE22 : 40.131    14.85
```

Permanent, user-accessible storage should be made available to interpretive TACS users. The request was made by Gabor Furst of suburban Vancouver, B. C., Canada. It was answered on December 4th by E-mail that proposed the following: I can envision a storage device (probably yet another special relay, although this is an unimportant detail) to which you would specify four things: 1) Direction of the transfer (to or from this secondary storage); 2) The name of a TACS variable. This would, in effect, give an address for the start of a block that is to be saved. While TACS variables could be saved one at a time, they also could easily be grouped together and transferred as vectors. For example, (va, vb, vc) being a 3-phase voltage vector. As long as you as a user were careful that the 3 were defined in order, they could be transferred as a block of 3. 3) An address in the new storage. In the simplest case, this would be a single integer index of the new storage, the start of a block to be used. But more generally, it could be any number of integers. Commonly, you might want two subscripts for the storage, with one being the time-step number. Just as a compiler allows ARRAY (I, J), so could my new logic. Of course, you would be responsible for specifying the I and the J, with one typically being ISTEP (the simulation step number). 4) Number of words to be transferred.

The need for MODELS working space increases as a function of simulation time. This is a complaint from a number of users in recent years, and your Editor discussed

the problem as follows in public E-mail of the Fargo list server having date December 11th : *"This writer must ask again, what is the problem? ..... Physically, or from an engineering or mathematical point of view, does overflow make sense? With TACS, this writer does not believe any such overflow is possible. ..... Why should MODELS be any different? Is there a MODELS component with overhead that continually increases with simulation time until eventually resources are exhausted? If so, which one, and what is the need? If not, how does one explain the MODELS inclination to halt when T-max is extended? Has anyone ever had such trouble with TACS? .... Brute force increases of ATP table sizes beyond any reasonable estimate of user requirements is not a satisfying solution. What happens if those tables need to be dumped to disk (e.g., for either Monte Carlo studies for START AGAIN usage)?"* No answers yet from MODELS supporters.

Prof. Mustafa Kizilcay of FH Osnabrueck in Germany did report an interesting, possibly-related MODELS error, however. From his E-mail dated December 29th : *"I used frequently DELAY function within a WHILE loop. If the condition for the WHILE loop is not specified correctly, MODELS tries to access past values beyond the limit as the simulation proceeds."* About the suggestion that the user is partly to blame (his data asked for the excessive delay), your Editor was unsympathetic: *"That is about like suggesting that ATP should not be protected against a branch of zero impedance (it is, of course). TACS has been protected against out-of-bounds requests of the delay function for 20 years. MODELS was supposed to be 'improved TACS,' but in many ways it is 'inferior TACS,' we have learned. Unfortunately, I did not know this a year or more ago. Only since August have I taken a particular interest in MODELS."* For those curious, August was when concerns about unauthorized modifications to ATP (see the separate story) first were raised.

No Unauthorized ATP Modifications !

"Warning : Modification of ATP for commercial purposes is prohibited" was the headline of public E-mail of the Fargo list server dated January 2nd. In this 14885-byte message, your Editor introduced the problem using six paragraphs very similar to the following (remainder of this story):

Problem: One individual who has had access to ATP source code is reported to have proposed the sale of his own commercial software, which then could be used via a special version of ATP (this, too, would be provided by the seller). The proposal seems to have been made clandestinely to an unknown number of companies and/or individuals. One potential customer reported that he was asked not to inform others of the initiative.

Notice: Canadian/American (hereafter, Can/Am) User Group management in Portland, Oregon, USA, regards any such activity as a clear violation of modern ATP licensing agreements such as its own. Furthermore, it is prepared to take vigorous counter measures to prevent such *EMTP commerce*.

Explanation & discussion: It is true the just-mentioned commercial software might not be a part of the ATP disk file itself (for Salford EMTP, TPBIG.EXE), since it might be connected dynamically at execution time (just as Salford TPPLOT.LIB is). But, this would make it no less a part of ATP for purposes of simulation. This is not at all like the pre-processing of ATP data by Hoidalen's ATPDRAW, or the post-processing of ATP results by plotting programs such as Kizilcay's PCPLOT. Such preprocessing and postprocessing are not integral parts of simulation. But the commercial software just mentioned was to have been made an integral part of the simulation itself --- just as, for example, Prof. Hian Lauw's Universal Machine is an integral part, or Prof. Xusheng Chen's three-phase transformer model is. Of course, neither Lauw's nor Chen's work is a commercial product today. New transient modeling would be provided, and the commercial software would be used within the time-step loop of ATP to provide it, so the new modeling would be part of *an electromagnetic transients program*. That is the user group's position. In that the new modeling would be sold, it would constitute *EMTP commerce* whether or not the modified ATP that supported it might supposedly be supplied free of charge. That's right, readers, this seems to be the gimmick that the devious merchant considered in his attempt to avoid legal challenge : distribution of the secretly - modified (to connect to his own commercial software, which was to be sold), unauthorized copy of ATP free of charge. Well, a legal challenge will not be avoided **this** way.

The use of dynamic link libraries (Salford terminology) to package major portions of a large program is relatively new (e.g., the past ten years or so). But this is a common technique today. Programs of PTI's popular PSS / E package for load flow, transient stability, short circuit, and other forms of system planning - type analysis are understood to have been packaged as dynamic link libraries for purposes of running under Salford DBOS on Intel PCs. This is the modern way of linking: a deferral of part of the effort until execution time. The point is this: the fact that a separate disk file stores part of a program makes that disk file no less a part of the program that uses it. In the case of Salford DBOS, dynamic link libraries are connected by ordinary subroutine calls --- just as could be the case if all code were in the same disk file. It is the user group's contention that mention of *a program* in LICENSE.ZIP covers any dynamically-linked libraries or subroutines as well as all classically-linked libraries or subroutines.

Note that nowhere in a modern licensing agreement is there any provision for modification. Any unauthorized change would constitute an unauthorized use of the ATP

materials --- commercial exploitation of the property of others without permission. That is the user group's position, and the user group is prepared to sue in American courts to enforce this interpretation. Also, it is prepared to ATP-excommunicate anyone involved with such commerce, just as it excommunicated LEC during the fall of 1993 (which also turned out to be the fall of LEC, if a pun is permitted).

ATP source code never again will be shared with anyone who fails to promise in writing not to modify ATP for commercial purposes. Details of such an explicit, written requirement now are open to discussion. What is the best language to prevent the situation described in preceding paragraphs? Should the proposed condition be added to ATP licensing agreements, even though not one in 100 users might ever receive ATP source code? Alternatively, should there be a separate licensing agreement for developers who have access to source code? Is there a need to protect unauthorized use of object files, too, as well as source code? After all, entire subroutines can be exchanged by linking or relinking. Possibly this could be done with an executable file such as TPBIG.EXE, too. It was possible using Apollo Aegis because the linker output had the same form as linker input. A clever merchant might know enough to modify the program usefully in this manner, without ever modifying ATP source code. Some really hard thinking and public discussion by software development experts is being encouraged. To be continued (more next time).

Dependent Switching for Ned Mohan

Simple switch dependence was added December 1st upon the recommendation of Prof. Ned Mohan at the University of Minnesota in Minneapolis. Recall that the concept was proposed and explained in the January, 1996, newsletter. However, no programming was actually done until late November when Prof. Mohan raised the subject again. A note in an envelope that bore November 21st postage metering stated: *"I would also like to include the description of switch-coupled diodes and switch-coupled switches. .... Such a switch pair would be very useful in modeling PWM inverters in motor drives."* So, the extension was provided.

The master / slave relationship is simple in general concept, but can be tricky in detail. It applies to any two switches. If the master is to undergo a status change (opening or closing) on any particular time step, then the slave will obey orders from its master if and only if the slave is not already scheduled for its own change. That is, slavery is powerful, but it does not override any natural status change that already has been scheduled at the time bondage is considered. All natural status changes come first, before any switch dependence is considered; and if a slave is involved in a natural status change, its bondage is ignored for that time step.

Each dependence among two switches is defined on a new card that must follow the definitions of the two switches involved, and precede the blank card ending switch cards. For small data cases not involving data modularization, it probably is simplest and clearest to place all such declarations immediately before the terminating blank card. This is what has been done in the illustrations of BENCHMARK DCNEW-17. Although each slave can serve more than one master, present logic permits each master to have at most one slave.

Dependence can be mutual and bilateral, it should be explained. That is, switch K can be both a master and also the slave of switch L thereby ensuring that the two operate together even though the user does not know, or want to specify, which of the two first will undergo switching naturally.

Different modes of operation are allowed for each switch dependence. In the order these were added to ATP, the user is allowed the following alternatives:

1) ALWAYS TOGGLE will change the status of the slave whenever the master changes;
2) ONLY OPEN/CLOSE will close the slave whenever the master opens;
3) ONLY CLOSE/OPEN will open the slave whenever the master closes.

GIFU switches have not been made obsolete by the new switch dependence, it should be emphasized. Most likely it is only for small or simple configurations that the user will know enough to establish the required switch dependencies. For complicated or multiphase networks, the user may not know ahead of time what the switching sequence will be. Further, the sequence might not be fixed as loading or other parameters are changed. So, the availability of GIFU switches remains important. On the other hand, for simple illustrations such as Prof. Mohan's Exercise 10, the lone dependence is clear enough, so is easily specified. By avoiding the GIFU switch (see first subcase of DCNEW-17) this way, the simulation is speeded because table dumping and restoring are avoided. Also, there is no conflict with STATISTICS (Monte Carlo) or SYSTEMATIC usage.

DEC ATP for VAX / Open VMS

VAX/VMS is the original operating system of DEC VAX computers, and Open VMS is the corresponding operating system for DEC Alpha, the high-speed RISC workstations. ATP has been updated for both at least twice since the last report.

July 30th, a VMS translation was performed for the first time in many moons, and BPA's Dr. Tsu-huei Liu began the laborious job of reconciling VMS .LIS files using all

standard test cases --- first with themselves (VMS history), and second with Salford .LIS files (using TRUNCATE followed by Mike Albert's FC). One important discovery was of two sorts of installation-dependence in Salford DC*.DAT data files: 1) Tab characters (believed to have entered disk files from DOS EDIT as explained before); and 2) E-field numbers that were not right-adjusted. Unfortunately, Salford software skips over such lapses of rigor. It is only when the data case is passed to BPA's main VAX/VMS computers that developers in Portland notice the trouble. For example, the E-6 of DCN-16 under-flowed to zero, and E+3 became E+30 (not good).

Yet another observation (actually, a **re**discovery) that led to reform came from Wendell Neugebauer's famous rail gun, which is the third subcase of DC-22. Even using 64-bit floating-point computation and TRUNCATE, this is difficult to compare mechanically (using FC) rather than using human intelligence. New comment cards explain this in order that it need not be rediscovered again. Anyway, the reconciliation using VAX/VMS via Ethernet connection to Dr. Liu's Pentium was completed around August 8th, when Randy Suhrbier was telephoned and informed that files were ready for his repackaging under Open VMS. It is appropriate for all who have interest to express their gratitude since work with EMTP of any flavor, on any computer, no longer is part of his job. Well, Dr. Liu, your Editor, and all VMS ATP users in the vicinity certainly are appreciative.

The first week of October, VMS installation-dependent plotting modules were reworked to includes HP-GL and PostScript output. The subsequent reconciliation of test cases then revealed two important bugs or potential bugs. These were discovered using the VAX debugger, but were traced to code that is used for Salford EMTP, too: 1) there was overflow of a local CHARACTER*24 string; and 2) there was a name conflict with a variable in COMMON. Executable VAX and Open VMS ATP from this October translation were copied onto a TransGrid tape and air mailed to TransGrid in Sydney, Australia, on November 19th. There, the treatment of numbers that are not right-adjusted became a much bigger issue because production data --- not just test cases --- were found to be troubled for the first time, for unknown reasons. This led to an important discovery (next paragraph).

The addition of BN to FORMAT statements was first proposed by Stephen Boroczky of TransGrid in Sydney, Australia, as a standard (i.e., universal) way of avoiding unintended consequences of those blanks to the right of both E-field and integer numbers. In private E-mail dated December 24th, he wrote: *"I believe that BN is a Fortran-77 standard. From memory, the Digital Fortran manual says that this is so."* So, your Editor proposed the idea for consideration in public E-mail of the Fargo list server later that same day (December 23rd in the USA). Confirmation of intended operation was rapidly obtained for DEC Alpha VMS by Bruno Ceresoli of ENEL in Milano, Italy, and for Apple Macintosh by Stu Cook of Just Services in Montreal, Canada. Privately, Robert Meredith of NYPA approved for the Watcom compiler, and Prof. Mustafa Kizilcay of FH Osnabrueck in Germany seconded the theoretical basis: *"BN is a FORTRAN 77 standard according to ANSI, ISO and DIN standards as I checked in the FORTRAN language manual of Univ. Hannover."*

On - Line Records of ATP Licensing

Soon, E-mail should allow any one ATP user to determine whether any other ATP user is licensed. This is a continuation of the story than began in the July, 1996, issue. There was not space in the October issue, unfortunately.

German law ensuring privacy was what had prevented the publication of names of licensed EEUG members, recall. Well, following the fall meeting in Budapest (see separate story), Chairman Mustafa Kizilcay was asked whether the issue had been resolved. In E-mail dated December 2nd, he wrote: *"Unfortunately, I could not recognize this matter during the meeting. But, I had included this issue in the 'order form for members' in the beginning of 1996. The following text was added: 'Check here if you wish to have your name omitted from:*

|_| *The EEUG Membership Directory*
|_| *The list of licensed ATP users on the FTP file server*
|_| *Outside mailing lists'*

Until now, we have not received any negative reply concerning this issue. No reply means the members accept the disclosure of their names." So, rather than mandatory disclosure for all, each member retains the right to prevent this in Europe. The distinction is important.

Commercial mailing or telephoning (including FAX) should not be practicable using any published list of licensed users, it has been agreed. Prof. Kizilcay wrote : *"Frankly speaking, I am not in favour of the disclosure of all information related to licensed users. Name, company, city (without postal code) and country would be enough in my opinion. My main concern is the likely misuse of these data by ATP competitors .... We maintain a data base of total 357 ATP users/members at the moment."* Your Editor responded: *"We agree. The common use will be for E-mail deliveries. For this, a street address would not be required. A full name, company name, city, and country should be adequate. .... By omitting a street address and postal code, the list would have minimal value for unauthorized mass mailing. Also omit the E-mail address since this will be known by the persons involved. For example, take NYPA's problem of verifying that someone who requests Watcom ATP is licensed. That person will already have sent E-mail to NYPA making the request, so the E-mail address already will be known in White Plains. .... Omit all telephone numbers, too. As long as we have full names so there is no ambiguity within a company, we do not need either street addresses or electronic addresses for our intended use."*

Where should lists of licensees be maintained? Your Editor responded as follows December 30th in response to Prof. Kizilcay's offer to send a copy to Portland: *"It is not a question of sending the list to us, but rather of making it available for aFTP transfers (I would recommend Prof. Bruce Mork's Houghton aFTP server). But where? .... Maybe down within the subdirectory for each user group, using a standard name. For example, your licensing summary could be within* ***pub/atp/license/eeug*** *as a file named LICENSES.LIS or .DOC or .TXT (whatever best suggests a text file). Then we would contribute a file of the same name for placement within* ***pub/atp/license/canam*** *; etc. for others. Alternatively, all lists could have different names one level above. Prof. Mork will need to think about this. No question, having all lists in the same directory would ease acquisition by others."*

Fargo E - mail Requires License

Readers are reminded that one must be ATP-licensed in order to subscribe to Prof. Bruce Mork's ATP-dedicated electronic bulletin board **atp-emtp** that is hosted by North Dakota State University (NDSU) in Fargo, North Dakota. The story about three suspicious subscribers last spring is worth documenting for the record.

First, the problem. Looking at the output of REVIEW, addresses ending in **@der.edfgdf.fr** and **@ee.ubc.ca** and **@wapa.gov** were noted. For those who do not know, these apply to EDF in France (affiliated with EPRI for purposes of EMTP development and sale), UBC in Vancouver, BC, Canada (home of DCG / EPRI advisors and contractors), and WAPA of suburban Denver, Colorado --- a DCG member. Since the individuals who allegedly were associated with these addresses could not be found in any ATP licensing records to which your Editor had access, on May 29th he sent each a form letter requesting clarification. In part, the form letter stated the following (next 3 paragraphs):

.... neither you nor the organization providing your E-mail could be found in computer files that are supposed to summarize ATP licenses of North America and Europe. These are ATP licenses that have been issued by either the Canadian/American EMTP User Group (Can/Am) or the European EMTP-ATP User Group e.V. (EEUG, which has headquarters in Osnabrueck, Germany).

If your current usage of ATP information is covered by a modern .... license, please provide details in order that our records can be corrected. I.e., to what organization or individual was the ATP license in question granted, what is the mailing address and telephone number of the licensee, who granted the license, and when was it granted?

On the other hand, if your current usage of ATP information might not be covered by a modern ATP license, it is strongly suggested that you quickly either: 1) become ATP-licensed .... or 2) desubscribe from the Fargo list server. It should be noted that subscription to Prof. Bruce Mork's E-mail forum does require ATP licensing because this forum involves the exchange of ATP information. ATP information is not in the public domain, and may only be received by, and exchanged among, ATP-licensed users.

End of quotation from form letter. None of the recipients responded, but it **was** noted that each promptly disappeared from the list of subscribers.

GNU ATP Runs Under Linux

Lack of INTEGER*2 support by GNU FORTRAN g77 that runs under Linux was mentioned in the April, 1996, issue. Fortunately, this shortcoming did not last long. Masahiro Kan of the Hamakawasaki Works of Toshiba Corporation reported in E-mail dated November 5th that g77-0.5.18 allows 2-byte integers. BPA's Walter Powell has since confirmed this. The development is believed to be important not so much for the memory saving itself, but rather because it allows more commonality of source code with other compilers including Salford's.

GNUPLOT could be executed from within ATP as first pointed out to developers in Portland by Prof. G. Corwin Alexander of Oregon State University in Corvallis. This is documented in E-mail dated December 5th, and the concept was proven shortly thereafter. This is the spark that started the current interest in GNUPLOT, and all of the work of the following story. After proving execution from within ATP, using CALL SYSTEM to execute the Linux command to run GNUPLOT, your Editor quickly remembered how computer expert David Szymanski had recommended the use of Unix pipes for plotting during the late '80s. Not only was a pipe used with those old Sun Unix workstations at BPA, but SUBROUTINE PIPLOT ("pipe plot") remains in the UTPF today. As a result, Prof. Alexander's idea was quickly modified to keep GNUPLOT external to ATP, and to be fed by a Unix pipe. BPA's Walter Powell then wrote some C code to demonstrate how a pipe could be used for output by a FORTRAN program. At this point, the GNUPLOT revolution for ATP was well under way! For details, read the following story.

Correct execution using all standard **dc*.dat** test cases was mentioned in the preceding issue. It is worth summarizing some of the limitations and assumptions. First, UNFORMATTED **.pl4** files were used. Whether C-like files are possible using g77 still is not known. The Watcom approach of using ACCESS='DIRECT' with RECL=1 for the header has not yet been tried. Of course, C-like files could be created using real C just as Szymanski did 9 years ago. Because the C compiler is

free, there is no serious objection to this efficient solution other than your Editor's natural inclination to avoid a second language whenever possible. Second, about graphics, there is no CALCOMP PLOT output to the screen. Instead, HP-GL, PostScript, and GNUPLOT output are to be relied upon, with the latter displayable immediately via a second process that is connected by a pipe. Third, SPY **does** work in the batch mode as illustrated by DC-56 and 57, although without SPY PLOT vector graphics. Fourth, environment variable ATPDIR functions correctly. The user must remember that this is upper case, however --- unlike all file names, which are lower case (as held by an exclamation point on cards such as $OPEN that might reference them). Fifth, the **dc*.dat** test cases that will be distributed to others are just the ones that differ from Salford EMTP data files. Sixth, there is no glitzy HELP window at the start of execution, just as there are no SPY windows; all output goes to the screen (the window of the X Window System in which ATP execution begins, if in fact such windows are being used). Windows are not necessary, of course. Like old MS Windows that ran on top of DOS, so The X Window System runs on top of Unix, making it voluntary rather than mandatory. Seventh, some file types of DOS are repeated for Linux use even though there is no special meaning. For example, **run.bat** executes all standard test cases --- and executes them efficiently. Whereas your Editor's Salford execution at 133-MHz under Win95 takes about 5 minutes, the Linux execution requires 6.5 minutes. But, as your Editor explained to Mr. Kan in E-mail dated December 3rd, this is good: *"The Linux computer is 90 MHz! Walter discovered this by rebooting and looking in the right place. .... That ratio of 6.5 / 5.0 = 1.3 whereas the ratio of clock rates is 133 / 90 = 1.48 Further, this ignores the multitasking capability of Linux. Your advocacy seems to have been vindicated."*

GNUPLOT for New Universal Plotting

GNUPLOT is a freeware plotting program that can be downloaded from the Internet by aFTP transfers. It comes in many flavors for the many computers and operating systems of interest, so it offers nearly universal plotting on both monitors and printers. ATP for Intel-based MS-DOS computers (Salford EMTP) began producing GNUPLOT output for CALCOMP PLOT (batch mode plotting) beginning December 20th, and any other translation can easily be modified to do likewise if such output is not yet provided. A Beginning reference for Web surfers seems to be **http://www.cs.dartmouth.edu/gnuplot_info.html** on the campus of Dartmouth University in Hanover, New Hampshire (appropriately, the *live free or die* state).

Version 3.6 of GNUPLOT is required for compatibility with ATP output, it is important to note. This is because Version 3.6 supports inline data for curves whereas the preceding Version 3.5, dating to 1993, did not. Masahiro Kan of the Hamakawasaki Works of Toshiba Corporation in Japan made developers in Portland aware of newer Version 3.6 in E-mail dated December 20th. Attaching the HTML instructions, he observed that 3.6 *"has many improvements and extensions."* Yes, indeed, it was rapidly noted that the critical one for universality and unification was in-line data. Previously, using Version 3.5, it was necessary to place data points of each curve in a separate disk file, and the name of this file then would appear between quotation marks of the plot command. An example having two curves follows:

```
plot "curve1.dat", "curve2.dat"
```

To indicate in-line data, each file name is replaced by a single minus sign. The data follows in order (curve 1 comes first, then curve 2, etc.), with (x, y) pairs of each curve terminated by the letter "e" in column 1. Without use of this feature, it would not be possible to package output in a single file parallel to the input data file (next paragraph).

NOGNU is the new parameter of STARTUP that was added to control whether or not ATP produces output for GNUPLOT. In its basic form, this parallels NOPOST for PostScript and NOHPGL for HP-GL : value unity will prohibit such output (i.e., no GNUPLOT). As for approving output for GNUPLOT, most systems should accept zero as a request for a disk file parallel to that of the PostScript that is supported by Robert Meredith of New York Power Authority (NYPA) in White Plains. So, for those who follow the recommended usage of KTRPL4 < 0, a single output file of type .GNU will parallel the input data file. Values of NOGNU other than 1 or 0 may be used to indicate other output alternatives such as pipes (important for all Unix computers). But since alternatives to disk depend on computer, they will not be covered in this general story.

GNU PLOT is the ATP plot declaration parallel to PRINTER PLOT (beginning in column 3) that is used to carry arbitrary GNUPLOT commands of the user. ATP will generate the plot command, and its associated data, automatically. This includes labeling of the curves, the axes, and title. But what about commands associated with a bounding box, or logarithmic scaling, or points without connecting line segments, or output for a printer? Any text to the right of column 11 (the blank that follows declaration GNU PLOT in columns 3-10) will be passed to the .GNU file without modification. This allows the user to change things before each vector plot, and specifically, to override ATP assumptions of a grid and plotting with line segments.

Reference signals are easily added to ATP graphs by manual editing of the .GNU file, it should be emphasized. This is a strength of GNUPLOT : it provides analytical representations of many functions including the obvious sinusoids of power systems. For example, one could add a sinusoid to the right of the previous illustration:

```
... "curve2.dat", 50*COS(377*X-1.57)
```

Logarithms, trigonometric functions, Bessel functions, and hyperbolic functions are all there; and arithmetic can be

performed using operators of C plus the ** of FORTRAN (for exponentiation), and parentheses for grouping.

GNUPLOT is not covered by the famous GNU General Public License (see the October, 1994, issue), despite the name. From a 13 Sept 1993 disk file 0FAQ by Thomas Koenig at **ig25@rz.uni-karlsruhe.de** : *"The authors of gnuplot are: Thomas Williams, Colin Kelley, Russell Lang, Dave Kotz, John Campbell, Gershon Elber, Alexander Woo and many others. The following quote comes from Thomas Williams : I was taking a differential equation class and Colin was taking Electromagnetics, we both thought it'd be helpful to visualize the mathematics behind them. ..... Gnuplot is not related to the GNU project or the FSF in any but the most peripheral sense. Our software was designed completely independently and the name 'gnuplot' was actually a compromise. I wanted to call it 'llamaplot' and Colin wanted to call it 'nplot.' We agreed that 'newplot' was acceptable but, we then discovered that there was an absolutely ghastly pascal program of that name that the Computer Science Dept. occasionally used. I decided that 'gnuplot' would make a nice pun and after a fashion Colin agreed. .... Gnuplot is neither written nor maintained by the FSF. It is not covered by the General Public License, either. However, the FSF has decided to distribute gnuplot as part of the GNU system, because it is useful, redistributable software."*

PostScript and HP-GL can be produced by GNUPLOT, it is important to note. I.e., GNUPLOT can provide yet another way to produce such output corresponding to batch-mode plotting. Quoting from that 0FAQ file: *"Many word processors can use Encapsulated PostScript for graphs. This can be generated by the 'set terminal postscript eps' command. Most MS-DOS word processors understand HPGL (terminal type hpgl)."*

Users of MS Windows now have an alternative to Ghostview (see story in July, 1994, newsletter) for the free display of ATP batch-mode plots. If GNUPLOT is started from the DOS window, an MS Window will open, and this has pull-down menus named File, Plot, Expressions, Functions, General, Axes, Chart, Styles, 3D, and Help. In the left-hand corner will be the usual *close* button, and in the right the down and up arrows for making the window either full-screen or an icon. On the right is a scroll bar. When a graph is produced, this will be in a second window. This is Version 3.5 patch level 3.50.1.17 dated 27 Aug 93 as left at BPA by Taku Noda (see July, 1996, mention). At that time, the program was running under Windows for WorkGroups, but today it runs under Windows NT on Dr. Liu's Pentium without any changes. Of course, for use with the new .GNU files, Version 3.6 is required.

Compiled TACS Speeds Simulation

Compiled TACS is the much faster, and also more flexible, way to simulate control system modeling. It was mentioned in the preceding issue, but there was not space to explain it then. The following material is close to what was explained in public E-mail dated November 2nd and 9th.

COMPILED TACS MAKE is a new special-request word that will result in the creation of FORTRAN code to represent the user's TACS data. This involves a translation from TACS language to FORTRAN language. Then, after compilation of modified code, ATP is relinked to create an executable TPBIG.EXE that can be used in either one of two possible modes: 1) The old familiar interpretive mode for TACS; or 2) the new compiled mode, with hard-wired logic, as selected by the special request word COMPILED TACS USE .

This procedure, which is being named *compiled TACS*, is not yet complete. But enough has been done to demonstrate feasibility and efficiency using DC-22d --- that SVC (static var) modeling from Gabor Furst of suburban Vancouver, B.C., Canada. Whether interpretive or compiled, answers are identical. Only the time spent in the time-step loop is different. Obviously, compiled TACS is faster --- **much** faster.

Supplemental variables are the dominant strength of compiled TACS. At the top of page 3-17 of the ATP Rule Book, these are called *the algebraic and logical FORTRAN expression*. Well, after ignoring columns 1 and 2 (which contain either 88, 98, or 99), and indenting the name to column 7, the result is legal FORTRAN, generally. For Mr. Furst's SVC modeling of DC-22d, no other changes were required. The user-supplied TACS data becomes in-line FORTRAN that will be executed if the COMPILED TACS USE request has been made.

Supplemental devices show less improvement from compiled TACS because devices always have been compiled. Yet, there are meaningful gains for several reasons: 1) Offset subscripting (by Bob Eifrig, to variably dimension TACS during the late '70s) has been largely eliminated. 2) Some DO-loops (e.g., to sum device inputs) are eliminated; 3) Some IF-statements (e.g., to check if there is a named threshold of a relay, or to distinguish a level-trigger) have been eliminated. 4) All computed GO TO statements have been eliminated. 5) All device code now is inline; there is no reuse by looping or module CALLs. Only two CALL statements are required by TACS on each time step, and one of these is a block transfer at the end (CALL MOVER). The other simply provides entry to TACS, which is in a single subroutine.

S-blocks (Laplace transfer functions) remain interpretive in the sense that parameter values are stored in data rather than code. So, for example, the numerical values of gains or time constants can be changed without recompilation. On the other hand, DO-loops and IF-statements that are part of the sparse elimination and backsubstitution no longer exist; and offset subscripts (again due to Bob Eifrig) have been eliminated.

Sources were the final TACS components to be modified. Although improved ease of variable access was the primary motivation for the work on sources, noticeably-increased speed of simulation was the pleasing result for Gabor Furst's SVC modeling of DC-22d. Note the resulting completeness: every single TACS variable now is assigned across an equal sign. The user has access to each and every one of them, should he want to change the way any one of them is computed.

The speed of compiled TACS has been measured using Gabor Furst's same standard SVC simulation as in the separate story about TACS and MODELS (refer to this for background). The new second row below displays the best 5 of 6 successive trials within a new Win95 window of this writer's 133-MHz Pentium:

```
Interpreted : 5.989  6.099  6.044  5.989  5.879
Compiled :    3.571  3.407  3.407  3.352  3.352
```

The averages of these 2 rows are: 6.000 and 3.418 sec, respectively. Clearly, the latter must be barely above the time for everything else. A Watcom ATP profile from NYPA was shown in the original announcement, and from this it was estimated that the electric network is roughly 60% of the burden. Well, 3.418 is 57% of 6 seconds.

The flexibility of user-supplied FORTRAN is believed to be more important than the speed gain of compiled TACS, it should be explained. Note carefully that TACS supplemental variables and devices are nothing other than function evaluations. For each, an output variable must be assigned. How the user does this is his business. COMPILED TACS MAKE produces one alternative, but the user is not forced to follow it. He can modify this FORTRAN however he wants. If he wanted to use function CALLs rather than in-line code for devices, such modification could be made. If the user wanted to load a matrix, invert it, and calculate eigenvalues, that, too, would not be difficult. The user has the full flexibility of the FORTRAN compiler at his disposal. Double subscripting (I, J), extra local storage, special I/O, and complex variables all are available to the compiled TACS user at no cost to ATP programmers because such features come with any compiler. The user has obvious access to all supplemental variables, devices, and TACS sources, so user-supplied source code can be quite arbitrary. There would seem to be no need for any special interface since points of access are obvious (TACS variables can be recognized by their 6-character names).

T-start > 0 for TACS sources can not be allowed for use with COMPILED TACS MAKE because each source must be in use on step 1 when FORTRAN representing the user's data is created. This is an important restriction of the present coding. It is the consequence of a convenience: the scattering of additions among existing code, which made for easier programming and debugging. All additions might later be gathered together in one enormous block, to remove the restriction on T-start if it represents a serious annoyance. But until then, look for warning messages like:

```
===  Warning!  TACS source "XXXXXX" has
     T-start, T-stop =, (1P, 2E11.3) ATP
     now modifies so active on step 1.
```

This is a warning about the interpreted solution only. It does not apply to subsequent simulation using COMPILED TACS USE (such a subsequent solution should correctly honor user-requested T-start and T-stop of sources).

Extra storage for compiled TACS is required for code only --- **not** for data, it is to be emphasized. Data remains in the variably-dimensioned working space of List 19. TACS data storage is unchanged.

Garth Irwin of the Manitoba HVDC Research Center must be credited with subliminal inspiration of the concept of the preceding paragraph. As explained in public E-mail of the Fargo list server dated June 28th, EMTDC was demonstrated at BPA by Mr. Irwin. The idea of function evaluation for EMTDC control system modeling seemed clear enough. EMTDC comes with a library of functions (or are they the more general subroutines, which would permit multiple outputs?), and if the user wants others, he is able to create and compile his own. That is what your Editor learned, and probably is an important reason his thinking recently turned to compiled TACS once the numerical inefficiency of MODELS was discovered and documented (see separate story).

Availability of a FORTRAN compiler that will support ATP is a critical requirement for the use of compiled TACS, of course. The free GNU compiler is believed to be important because it should extend compiled TACS to those who either can not, or will not, purchase a better compiler (e.g., Watcom as used by NYPA or Salford as used by your Editor). Even within companies, GNU g77 might be important. Whereas few companies would think twice about purchasing one compiler, how many companies will purchase FORTRAN compilers for every computer of ATP interest? So, see the separate story about GNU ATP for Linux. This has taken on added importance; and after Linux, that free GNU DOS extender clearly is worth trying, for those who would rather keep things simple.

Logical variables pose the greatest incompatibility for compiled TACS. The average user may not know that these involve floating-point numeric storage within Mr. Dube's TACS interpreter. Obviously, in FORTRAN, the variable type LOGICAL should be used. How this can be done should be shown in the next issue. This is via the AA and BB declarations (in columns 1 and 2).

Advocates of C might wonder about the continued mention of FORTRAN --- the original computer programming language. The answer is simple: ATP is a FORTRAN program, and program variables are stored in COMMON. These variables are immediately available to user-supplied source code if that code is written in the same language. If C were to be used instead, there would be added complexity of sharing variables. So, if other things are equal, FORTRAN is preferable.

Separate SUBROUTINE COMTAC is the new location of compiled TACS , which has been isolated from interpretive TACS for simplicity of use by others. The principal ATP module of the time-step loop is OVER16, and this now will CALL COMTAC rather than interpretive TACS if and only if the data case involves a COMPILED TACS USE request. Considering memory requirements, the use of interpretive TACS adds only if size of the code to represent the user's modeling exceeds size of the unused code that performs interpretation. For small data cases not having great diversity, compiled TACS actually represents a saving of memory compared with interpretive TACS. On the other hand, as the size of TACS data grows to infinity, so does the required memory for its representation using compiled TACS.

TACS is not the only modeling that might be *compiled* rather than interpreted, of course. Persons having interest in real-time simulation (e.g., for relay testing) might consider compilation of the electric network in order to speed the simulation. Certainly the average electric network source is just as amenable to compilation as the average TACS source was, so simple sources look promising. But thereafter, exploitation seems to become substantially more difficult. For example, because of switching, the repeat solution [Y] V = I of the electric network is not fixed. Whereas TACS has a fixed matrix for the S-blocks, [Y] might change on each time step. Another example: nothing like TACS supplemental variables exists on the electric side. Recall supplemental variables were the strength of compiled TACS. Well, no such obvious inefficiency as the interpretation of supplemental variables is to be found on the electrical side. Rather, one finds thousands of lines of complicated FORTRAN that already have been variably-dimensioned (no gain by removing offset subscripting of Bob Eifrig, either).

Florida and Other Short Courses

Prof. Dennis Carroll of the University of Florida in Gainesville is scheduled to give his 4-day ATP short course March 11-14, 1997, according to E-mail dated December 30th. Prof. Carroll wrote: *"We will use pretty much the same format as last year, with the addition of ATPDRAW."* To refresh the memories of readers, this means that the small format is to be advertised, and Drs. Kurt Fehrle (a consultant from suburban Philadelphia, Pennsylvania) and Yin Yuexin (of Global Engineering in suburban Atlanta) will assist Prof. Carroll with both the lectures and laboratory work. Each student will be paired with a fast PC (486 or Pentium) throughout the course, which will be given in Prof. Carroll's own small laboratory unless enrollment forces movement to the larger departmental lab down the hall. Readers who want to learn ATP usage *by doing* are advised to give this Florida offering serious consideration.

Addendum January 1st: A check of Prof. Bruce Mork's Houghton Web page provides a nice, short summary that includes the following (remainder of this paragraph): Enrollment will be limited to allow every registrant the continuous use of a Pentium-based PC for the entire course. The course begins with introductory material for persons with little or no EMTP background. Sessions progress to more advanced material for the analysis of power system transients. This will include the utilization of EMTP models for transmission lines, circuit breakers, transformers, surge arresters, generators, motors, power electronics and associated control systems. It will also include actual case studies in the field of power system simulation. Use of the newly developed ATPDRAW graphical interface program will also be demonstrated in several of the applications. For more details, please contact :

Dr. D. P. Carroll
TEL: (352) 392-0918
FAX: (352) 392-8671.
E-mail: dcarr@admin.ee.ufl.edu

"The trans-Atlantic IPST'97 may have run into such trouble in Seattle ..." This was seen in a paragraph last time about problems managing the costs of distant courses. Of course, IPST is not a course, but rather a biannual conference as explained in the October, 1995, issue. As for financial challenges of IPST in Seattle, how do these prices sound (from E-mail of Thomas Grebe on November 27th): *"Events that could be sponsored (in whole or part) include: 1. an opening reception on Sunday night (approximate cost - $3,000); 2. seven refreshment breaks (approximate cost - $1,200 each); 3. four lunches (approximate cost - $3,000 each); 4. shuttle bus from campus housing to hotel (approximate cost - $1,000 for week); 5. museum tour / dinner cruise (approximate cost - $8,000)"* The missive Subject was *"IPST '97 Sponsorships,"* and the E-mail was addressed to several persons associated with user groups or EMTP commerce. Aside from Electrotek Concepts and your Editor, these were Prof. Hermann Dommel of UBC, Dr. Atef Morched of OH, and Dennis Woodford of **daw@hvdc.ca** (presumed to be one of the EMTDC people in Winnipeg, Manitoba).

Mr. Grebe explained the need for IPST sponsorship as follows : *"Due to unforeseen problems with on-campus meeting facilities, we have had to move the conference to the Red Lion Hotel in Bellevue. As expected, this change of venue has affected our budget. The current registration fee is anticipated to be near $500. We would like to attempt to reduce this fee by soliciting outside sponsorships. We are in the process of contacting Seattle City Light and Puget Power & Light. However, it seems unlikely that in today's utility environment, we can expect very much assistance. That is why we are contacting each of you. In many ways you represent a large portion of the group that will attend, and as such, we feel that it would be appropriate that you been given an opportunity to contribute to the effort."*

For those unfamiliar with Seattle-area geography, Bellevue is the expensive suburb to the east, across freshwater Lake Washington; and the Red Lion is not among the cheaper area hotels! IPST **had** to be moved there, and the registration fee might be around $500? Wow (logistics have really degenerated since the pioneering beginning by Prof. Correia de Barros in Lisbon, Portugal, two years ago)! Maybe Prof. Mohan should have approved the $5K/day rate of the New York Hilton, and then asked Grebe and friends to cover the shortfall (joke)! Needless to say, the user group did **not** respond --- either to this formal request or the preceding, apparently-inadvertent mailing, which had *"Subject: RE: begging for dollars..."* The first request included the following reference: *"ps - what would you think about the respective user's groups ????? (scott's pocket) - too commercial? possibles: EMTP (Atef); ATP (Scott); PATH (us); EMTPDC (dennis). maybe each one does a lunch??"* The PATH as used here is believed to come from Electrotek Concepts (see preceding issue); it most likely is not the train out of New York City (Port Authority Trans Hudson)!

Speaking of Electrotek, an announcement from the power-globe list server caught the eye of BPA's Robert Hasibar. The Subject was: *"BMI, Electrotek Parent Company Acquires Dranetz and Daytronic."* This was dated December 18th, and attributed to Bob Zavadil at address **bobz@bmi.electrotek.com** The story from Santa Clara, California, reads: *"WPT began expanding its data acquisition and analysis business in 1994 with its acquisition of Electrotek Concepts, premier supplier of engineering services and information systems for the electric industry."* Yeah, right, such as that original, quarter-speed simulation of DC-1 under old OS/2 (joke)! Anyway, a change of focus at Electrotek is not surprising.

Hoidalen Improves ATPDRAW

ATPDRAW is the graphical data assembler from Hans Kristian Hoidalen of Trondheim, Norway. In that new, improved versions continue to be released, the story of last time continues. The following report was provided by author Hoidalen in public E-mail of the Fargo list server dated November 20th: *"Based on feedback at the EEUG meeting (especially from Laszlo Prikler) in Budapest an upgrade of ATPDRAW has been made. Besides, we are right know working on a Windows version of ATPDRAW so I would like to request all ATPDRAW users to send us a list of suggestions for the new program. All suggestions will be considered, but we can not, of course, guarantee anything. Most of all, we are interested to know about : 1) discovered bugs; 2) suggestions for new components (1 or 3 phase); 3) clipboard / printing requirements; 4) mouse operations and shortcuts; 5) design; 6" help file system; ...*

An updated version (3.24) of ATPDRAW is now available at ***ftp.ee.mtu.edu*** *(anonymous login). The file AD324EXE.ZIP is the zipped ATPDRAW.EXE file. Old users of ATPDRAW can just download AD324EXE.ZIP and unzip it. New users must down-load ATPDRAW.ZIP and INSTALL.BAT and perform a complete installation followed by downloading AD324EXE.ZIP for an upgrade."*

Mouse operations are an area of major reform, it would seem. Hoidalen wrote: *"The left double click to get the input window resulted in a poorer operation of the program since objects often became selected instead. New mouse operations (again): 1) Select object: left click; press Shift key to add. 2) Unselect object: left click in open space. 3) Move object: left click and hold. 4) Draw connection (new): left (short) click on node; left to place; right to cancel. 5) Resize connection: left click and hold on node (often necessary to select connection first). 6) Get object input window (new): right click on object. 7) Get node input window (new): right click on node. 8) Rotate object: right click on selected object. 9) Get group number window: left double click on selected object. To summarize: A) Edit operation: left button (except rotate (right)); b) Input windows: left button (except group number (left double))."*

Hide is a new feature about which Hoidalen wrote: *"A hide property has been added to objects. Hidden objects are gray on the screen and are not written to the ATP file. The hide property is set by a button in the object input window or in the group number window."* Finally, there was mention of an error: *"Move text error removed. The old error occurring when moving and object label or node name overlapping other node names / labels is removed."*

Tamir Orbach of Kim Development USA Inc. was supposedly creating a commercial product to compete with ATPDRAW, recall (see stories in the April and July, 1996, newsletters). Who can provide a status report on what was supposed to be an ambitious development effort? How many copies of this software have been sold? How many of the promises of future development have been fulfilled on schedule? Does Prof. Laszlo Prikler of T. U. Budapest in Hungary still plan a review for EEUG, or has he perhaps decided that such an expenditude of his time no longer is worth the effort?

Missing Induction Motor Data

Missing data for induction motor modeling of ATP was mentioned by your Editor in public E-mail dated December 4th. This followed the call for FORTRAN to represent such devices using the new, user-accessible compensation interface (separate story): But shortages of machine data remain a problem in either domain. This writer reiterates his public advice dated 20 June 1996. If one could easily and reliably convert nameplate data into transient data for the U.M., this writer would be happy to

write the conversion routine himself. Who has the formulas, and what are the assumptions? What rotating machine expert believes this to be a promising idea? How much error is involved for cases of practical interest? The concept has been talked about since the early 1980s, but has some serious limitations, obviously. What machinery expert will summarize these for the general public?

Note carefully that it does little good to have a new ATP model if data for its use is generally unavailable. For large synchronous machines, those concerned with transient stability have been gathering data for decades. Data for the Type-59 model is not generally a problem, then. But for various induction motors scattered around the world, data is less well known. What is the consensus of experts as to what the ATP user community should be doing in cases where parameters are unavailable? Recall Gabor Furst's recent advice on the subject:

```
> Date: Fri, 22 Nov 1996 09:30:24 -0800
> There is no unique way of calculating these
> parameters.  Remember that a machine of the
> same rating, speed, inertia etc. could have
> vastly different values of synchronous
> reactance., transient reactance etc.
```

Does the user really want to be guessing about what he does not know? The obvious possible consequence is GIGO (garbage in, garbage out).

A possibly-useful published reference was contributed by Bruno Ceresoli of ENEL in Milano, Italy: G. J. Rogers, D. Shirmohammadi, "Induction machine modelling for electromagnetic transient program," IEEE Transactions of Energy Conversion, Vol. EC-2, No. 4, December 1987. This public E-mail dated December 9th was quickly answered by Case Muskens in Edmonton, Canada, who wrote: *"I have used the method outlined in the attached paper and found it useful for developing workable machine models. The procedure can be easily implemented in a spreadsheet program. There are some typographical errors in the equations (eqn 20) and the reader should refer to a previous paper (in the bibliography): 'An Aggregate Induction Machine Model for Industrial Plant' by G. J. Rogers, John Di Manno, Robert Alden IEEE Transactions on Power Apparatus and Systems, Vol PAS-103, No.4, April 1984. I have found the difficulty in getting good estimates for the rotating mass model more difficult i.e. type of load characteristic, load inertias, damping coefficients etc."* Finally, Gabor Furst concluded: *"Both references are indeed valuable. I used the reference given by Case Muskens for about six years now in my motor starting program very successfully."*

More accurate modeling as a function of frequency is yet another challenge of induction motor modeling. Gabor Furst wrote the following at the end of a message dated December 4th : *"The problem with induction motors is the nonlinear behaviour of the rotor circuits, with double cage and deep bars, and the correct modeling of skin effect as the slip changes. In my studies I either neglect this or use some fudging. Could be an interesting project for a Ph.D. student."*

1 - Mbyte E - mail Limit at BPA

A 1-Mbyte size limitation on outgoing E-mail seems to be in effect at BPA , as reported in the October issue. So, the segmentation of Salford EMTP archives such as GIVE1.ZIP is required. It has not been easy. That first beta tester was Prof. Laurie Snider in Hong Kong, and he clearly proved that not all UUDECODE programs are compatible. Specifically, the one used successfully by Prof. Laszlo Prikler in Budapest came from Theodore Kaldis whereas the one in EMAIL.ZIP came from Richard Marks. For multi-segment files, they are **not** equivalent, unfortunately.

Robert Meredith of NYPA (the New York Power Authority) in White Plains summarized his own experience with multi-part messages in E-mail dated November 12th: *"Some people have had problems; others not, even though I do it the same way each time. I believe the people who have had problems have incompetent mailers that unsuccessfully attempt to uudecode the partial mailing and simply abort, rather than delivering the partial message. I have decided not to bother with multipart mailings any more. There were too many bookkeeping problems; too many times I had to remail one piece and keep track of what pieces had been delivered successfully. I've posted atpwnt.zip on my personal ftp site to provide access for those who cannot provide an ftp site for me to deliver to."*

Masahiro Kan of Toshiba Corporation in Japan has proposed an attractive alternative to use of UUENCODE and UUDECODE by Theodore Kaldis. This was summarized by your Editor in public E-mail dated November 12th. Mr. Kan wrote: *"uuex - uuencoder ver. 2.5 95/01/29 .... This software is freeware. .... Using uuex, we can generate uuencoded files whose maximum size are arbitrarily determined. .... After receiving these E-mails, you can extract GIVE1.ZIP by simply invoking the following command, if you named the received E-mails as give1.uu1/give1.uu2. C>uudx give1.uu* You do not need to delete the E-mail headers etc."* Prof. Prikler's procedure, too, allowed retention of the E-mail headers. However, it required an intermediate COPY to concatenate the two files .... If the average recipient does not now possess Kaldis's UUDECODE, there really is no reason to prefer it over this UUDX alternative. Avoidance of a big COPY is nice, and the use of a wild card to point to the two pieces is even better. Finally, the sender could avoid using commercial UUCODE to perform the segmentation. So, when new Salford DBOS was sent from BPA to Budapest and Tokyo on December 12th, it was this Japanese UUEX that was used to segment DBOS.ZIP

Nonetheless, segmentation is a pain that makes no sense in these modern times of large disks and high-speed data paths. BPA's Gerald Lee agreed, and he talked to persons in the computer establishment about the problem of ATP developers and users. He supplied a name and telephone

extension of someone who might grant an exemption. This was Linda Grogan on Ext. 4216. Unfortunately, when details finally were discussed with her by Dr. Tsu-huei Liu on January 2nd, the solution seemed to be too cumbersome and uncertain to be worth the effort. Rather than a general exemption, or an exemption for usage in the middle of the night, all that was offered was consideration on a case-by-case basis. I.e., prior approval would be needed, and even this seemed not to be guaranteed. *"Thanks but no thanks"* was the reaction of your Editor. Mr. Kan's procedure using Japanese UUEX seems easier and more sure for the BPA end. It is an acceptable interim solution until the computer establishment at BPA becomes more reasonable.

Miscellaneous Intel PC Information

Dual-processor PCs seem to be coming just as computer expert David Szymanski predicted (see preceding issue). The January issue of *Computer Shopper* carries a large story with title *"Low-cost dual-processor desktops debut"* beginning on page 72. The key acronym is SMP, which stands for symmetric multiprocessing. Just as Szymanski had said, *"The cost of doing [SMP] is low compared with prior chip sets, because [SMP support] is built into the Pentium Pro."* But what about the software? Is the extra hardware much good without software that exploits it? Not much, apparently. So, no one seems to know when the individual consumer might be interested. But the server market should be affected immediately: *"Compaq, Dell, and others have recently announced four - processor Pentium Pro servers with prices starting around $20,000. This is well below the $100,000-plus cost of a typical RISC/Unix server."* Of course, Windows NT is being used (*"Windows 95 is strictly a single-CPU operating system."*). As for Pentium Pro sales, the bar chart on page 73 shows rapid growth. For the 4 quarters of this year, estimated shipments in thousands are 125, 175, 400, and 1300. But prices remain high. Intel is not yet ready to abandon the lucrative Pentium market, which it continues to dominate. But it will happen eventually. *"Intel will take the Pentium Pro from a $500 price point .... all the way down to a $100 or $200 chip. At that price, ... adding a second Pentium Pro becomes an obvious upgrade whether or not the user really needs it."*

An overheated Pentium was the problem reported by Robert Meredith of New York Power Authority (NYPA) in White Plains. In E-mail dated October 9th, he explained how he used ATP compilation to isolate the problem: *".... I have been trying to change every memory access timing, removing secondary cache, interchanging memory pairs, etc, in an attempt to get trouble-free operation on my 120 MHz Pentium. I even installed WNT and the Watcom WNT compiler version, and installed another disk drive, changed the location of the temp directory and swap file, only to get the same 16 or so failures in ten runs. So it wasn't just OS/2 or the disk drive. I was already researching purchase of a new motherboard when Schultz suggested I try another cpu. So I took the 90 MHz from work home, rejumpered the board for 90 MHz and --- all worked OK. Then I rejumpered my own cpu for 100 MHz, instead of 120 and it too worked OK. No failures in ten compilation runs. In comparing Intel's cpu heat sink .... to my own, I've decided that my heat sink is inferior and I intend to replace it, in an attempt to get back to 120 MHz or even 133 MHz operation. So I plan to order an oversized one, with lots of fins and a larger fan .... I just thought you'd like to know that repeated compilation of ATP in an operating system with protected memory (to generate the error stop), combined with enough RAM and secondary cache (512 K) to keep the cpu busy, was the only way to detect this cpu problem."*

NeXT Software has been purchased by Apple Computer in what looks like a move to replace the aging Macintosh operating system. According to a short account on page A28 of *The Oregonian* dated December 26th, *"Apple ... said late Friday that it would buy NeXT for $400 million in cash, stock and assumed debt, and would rehire Apple cofounder Steven Jobs as a part-time consultant."* Whether Motorola CISC (e.g., 68040-based Quadra models) will be an orphan is an obvious subject of speculation. PowerPC is the clear target. *"The first version of the new OS won't run on Macintosh machines using older processors. ... Apple plans to support its customers with two new versions of System 7, its existing Macintosh operating system, during 1997."* This is worth watching. The competition for Microsoft is about to change in an interesting way.

Miscellaneous Small Items

Free printed copies of the 20-page July and October newsletters are scheduled to be mailed to subscribers of the US and Canada on January 21st. The end (this is the final issue that will be mailed free of charge by the user group). From now on, individual users are expected to pick up their own copies electronically (by aFTP transfers from the Houghton server or its Hannover mirror). If users want paper copies, from now on they can print their own.

$INCLUDE usage was confirmed using a maximum of 30 bytes of file name prior to August 13th, when this output at the start of a .LIS file was increased to a maximum of 80 bytes. At issue are messages such as the following:

```
--- Pass 1.  Card = 7.  Ready to open $INCLUDE =
```

The request to extend that trailing name on the right (not shown here) came from BPA's James Randall, who had so many bytes devoted to DEC VMS node, directory, and subdirectory names that nothing meaningful was left for the file name itself! The change is universal (i.e., users of all computers benefit from the extension).

Lower-case labeling of CALCOMP PLOT output was first seen September 13th as explained in the preceding issue. That was for the vector plot itself. But what about the associated data-card interpretation? After exclamation

points ("!") were removed from DCNEW-15, the .LIS file no longer showed the lower case of some input text. This was corrected December 28th when SUBR28 was changed so that the actual input data card was seen to the right of column 51. A collateral consequence was the appearance of some exclamation points to the right of text of other test case outputs such as DC35.LIS Previously, although lower case was held, the "!" was erased, and was not seen in the .LIS file. For plot cards, it seems better to show the user exactly what he supplied as data.

Seymour Cray died at the age of 71 on October 5th of injuries suffered in an automobile accident. John Markoff of the New York Times News Service wrote a fascinating summary that can be found on page A18 of *The Oregonian* the following day. The way Cray, father of the super-computer, beat IBM during the early '60s is memorable: *"The move to Chippewa Falls (Wisconsin; Cray's home town about 100 miles from Minneapolis) turned out to be important for Control Data. Cray's small team designed the CDC 6600. Introduced in August 1963, it had a speed of 3 million instructions per second. ... Frustrated, IBM Chairman Thomas J. Watson Jr. wrote a memo to his staff noting that the 6600 team totaled only 34 people, 'including the janitor,' and asked how IBM had let such a small team offer the world's most powerful computer."* There also is the story of Cray's 5-year plan for CDC, written only to satisfy bureaucratic demands for which Cray had little sympathy: *"Five-year goal: Build the biggest computer in the world. One-year goal: Achieve one-fifth of the above."* As for the IBM Chairman, this *Junior* is understood to be the son of the pioneer for whom IBM's famous research center (see Robert Schultz's mention in the April, 1996, issue) was named.

Some (not all) printed output of JMARTI SETUP has the frequency scaled by 2 * Pi. Prof. Laszlo Prikler of T.U. Budapest in Hungary described this as the *oldest ATP bug* in E-mail dated October 7th: *"... the frequency scale on the left and the right sides of the tabulated output list is definitely wrong! .... If you ask for a frequency scan from 1 Hz to 1 Mhz, the list starts with 6.2832 Hz and ends with 6.2832 Mhz. .... Interesting, eh? Luckily the JMARTI fitting has not inherited this bug ... The frequency scale seems correct on the left side of the printer plot and the Zc values in that three-column list is exactly the same as that in column 6 of the RLGC list. .... So I think the core of the code is correct."* Yes, Dr. Tsu-huei Liu's correction consisted of leaving the heading label unchanged, and dividing the numbers that had been printed by TWOPI .

POSTPROCESS PLOT FILE (PPF) was extended to allow TACS HYBRID data cases as explained in the July, 1995, newsletter. The second subcase of DC-46 illustrated correct operation, albeit for a very simple case. It was the 9th of December when BPA's James Randall supplied data that was more general, and which illustrated overwriting of the output vector by a signal from the PPF input file. The following day, use of offset KCOUNT was added to modules TACS3, TSTACS, and POSTPS in order to avoid overlaying the partial loading of BVALUE that might occur in OVER16 (switch current output).

Kizilcay frequency dependence, used to model network equivalents, comes from Prof. Mustafa Kizilcay of FH Osnabrueck in Germany. For some time, he has known that operation could be incorrect if his F-dependent branches were not positioned first in order of data input. A large illustration was provided by E-mail on December 29th, and your Editor proposed a simple correction to ENTRY HANN14 the following day.

Acrobat Reader is the name of a freeware document viewer that can be obtained from Adobe Systems as detailed in a story last time. Recall Dr. Glenn Wrate of Michigan Tech in Houghton recommended it along with associated Portable Document Format (PDF) files. Well, January 2nd, BPA finally obtained its replacement copy of Corel WP Suite 7. Why *replacement*? Months ago, a copy had been supplied, but it was for Win95, and it did not quite work properly on the Pentium of BPA's Dr. Tsu-huei Liu, which runs NT. More good work, Bill (joke)! As for how the new Corel WP worked, look for a report next time.

$INCLUDE files should be capable of mathematical processing of the arguments. This idea was suggested by Jeff Peggs of Virginia Power in Richmond. Your Editor responded as follows in E-mail dated December 18th: *"It is an idea that goes back years. I lost track of it then because I could not decide how much math was needed. Since then, I have worked on TACS .... Would TACS supplemental variables satisfy you for the evaluation of dependent parameters?"* For those who do not understand the concept, consider a trivial illustration: a balanced, 3-phase source. If the angle were to be an argument, in fact three of them would be required, for the three Type-14 sources of the module. But this is wasteful, since only one of them is independent. It would be much more convenient for the user to input just one angle (e.g., for phase 'a'), and have the program calculate the other two (e.g., the angle of phase 'b' by subtracting 120 degrees, etc. for 'c').

"csv format" was asked about by Dr. Mike Ennis of S&C Electric in Chicago, Illinois. His public E-mail of the Fargo list server on December 18th asked if there was *"a way to convert ATP output files to csv format in order to pass them to other applications?"* As Harald Wehrend of SEG in Kempen, Germany, observed, *"I don't know the special CSV format (but would like to know)."* Who does know, readers? Is there interest by others?

Lightning backflash was the Subject of public E-mail from Prof. Bruce Mork of Michigan Tech in Houghton. On November 23rd, he announced: *"Gabor Furst has just contributed an ATP data case for Monte Carlo simulation of lightning backflash on a high voltage transmission line."* Look for **ftp://ftp.ee.mtu.edu/pub/atp/dcase/backfl.dat** Yes, and also look for an alternative without MODELS .