
Can / Am EMTP News

Voice of the Canadian/American EMTP User Group

Publishers and Mailers:

Kwang-yi and Kwang-chien Ger
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. 94 - 1 ; January , 1994

Table of Contents

Salford Compiler and DOS Extender	1
Improvements to Salford TPPILOT	3
News from Outside USA and Canada	4
More about Electronic Mail (E-mail)	4
NYPA Optimizes EMTP Computation	5
Buffer for LUNIT6 Output to Disk	6
LEC Financial and Political Problems	8
NYPA Eliminates BLOCKD51 . BIN	9
BPA : a Government Corporation ?	10
E-mail in Portland : BPA and Agora	10
News about Laurent Dubé ' s MODELS	12
DCG / EPRI EMTP Information	12
Miscellaneous NYPA Changes to UTPF	13
Florida Short Course March 7 - 11	13
Mohan Course : San Francisco, July 23-24	13
Free Waterloo ATP Workshop Feb. 19	14
Real - Time EMTP - like Simulation ?	14
Turbo Table Dumping from NYPA	14
Isolated Corrections from LEC	17
Miscellaneous Intel PC Information	18
Miscellaneous Small Items	19

Salford Compiler and DOS Extender

Dynamic dimensioning of Salford EMTP will not have its limiting list sizes increased from those of TP20 to those of TP99 as had been predicted 3 months ago. This is for distribution to the general public. Almost no one needs the extra factor of five, and it just seemed unreasonable to subject all other users to the additional overhead merely to satisfy the one user in a hundred who

might have occasional need. To understand the added burden, consider the following total job times (from case-summary statistics) for DC-1 simulation that was preceded by rebooting of MS-DOS to ensure each case began with the same history in the SMARTDrive disk cache:

	TP3	TP20	TP99
1 Mbyte	80.165	81.758	87.308
2 Mbytes	70.714	70.275	74.780
3 Mbytes	69.780	68.407	72.802
4 Mbytes	69.780	68.407	72.802

Each row is for a different amount of Extended Memory as reported by DBOS (controlled by varying the memory used by SCROLLIT). Why TP20 is a little faster than TP3 for 2, 3, or 4 Mbytes is unknown. But the final column of TP99 clearly is significantly slower, and extra memory does not help, either. So, nearly everyone should be happy with TP20. TP99 could be created for anyone who needs it. Acknowledgement: Dr. Mustafa Kizilcay of Lahmeyer International in Frankfurt, Germany, was the first to suggest that resources required for dynamic dimensioning might be a little larger. In E-mail dated December 5th, he observed that the swap file used on his home computer required expansion. The preceding experiment using DC-1 then was run to better understand the effect.

No problems have been reported over the past 3 months, and the STARTUP choice using LISTSZ has been eliminated. No longer is there any way to avoid dynamic dimensioning. In fact, JTURBO of the next paragraph replaced LISTSZ in the STARTUP file.

JTURBO is a new parameter of STARTUP that allows the user either to accept (value unity) or reject

(value zero) the new *Turbo table dumping* by Robert A. Schultz of New York Power Authority (NYPA) in White Plains. Recall this is the latest breakthrough to speed the dumping or restoring of tables as required for START AGAIN, etc. Details are provided in a separate, later story. The general public received such code beginning January 10th.

OTG Systems of Clifford, Pennsylvania, no longer is a distributor of Salford products. David Vallance, the Managing Director of Salford Software, notified OTG customers of this change in a form letter dated December 1st --- the date the association ended.

PKZ204G.EXE replaced PKZ101.EXE on the GIVE1 disk of Salford distribution beginning January 10th. This involves the switch from old Version 1.01 to newer Version 2.04 of file compression and archiving software PKZIP by PKWare, Inc. of Brown Deer, Wisconsin. The copy now being made available to others was received many months ago from Dr. Mustafa Kizilcay of Lahmeyer International in Frankfurt, Germany. It was not needed before, so was not previously used. But then along came the restructuring of NYPA (most importantly, Robert Schultz's elimination of the need for disk file BLOCKD51.BIN), and suddenly the GIVE1 disk was found to overflow. By switching to the more efficient Version 2.04, the shortage was converted into a surplus. Although the compression utility itself is bigger (193 Kbytes vs. 131 Kbytes for the old), this loss of 62 Kbytes is overshadowed by the saving of 166 Kbytes on TPBIG: 925 Kbytes vs. 1091 Kbytes for the old. A favorable review of Version 2.04 can be found on page 40 in the *First Looks* section of *PC Magazine* dated 16 March 1993. Under the headline "*PKZIP Now Faster, More Efficient*" is found the following summary: "*PKZIP, Version 2.04, uses a new compression type called 'deflating.'* The results: ... compressed files were typically 5 to 15 percent smaller ... and the program performed up to 30 percent faster." The increased speed would seem to come from exploitation of newer hardware: "*PKZIP and PKUNZIP will use expanded and extended memory if available, make use of 286- and 386-specific code if they discover a compatible CPU, and run 32-bit code if they find a DPMI (DOS Protected Mode Interface) server.*"

LU6VRT is a new parameter of STARTUP that allows the user to delay the writing of LUNIT6 output to disk until the output first has accumulated sufficiently in RAM. This is yet another innovation from Robert A. Schultz of New York Power Authority (NYPA) in White Plains, as explained in a separate, later story.

Utilities to copy or zero vectors became available first among the NYPA contributions, however. These were a fringe benefit. Initially considering simpler parts of the Schultz revolution, your Editor decided to implement

these changes first. Modified on November 25th were segments MOVE, MOVE0, MOVER, MOVER0, and COPYI. As an illustration, consider MOVE that copies an integer vector. For length N3 words, the universal form would be:

DO 5873 J=1, N3

5873 INTB(J) = INTA(J)

Some 2 or 3 years ago, your Editor noted Salford library routine MOVE@ was faster, so he changed from such code to : N4 = 4 * N3

CALL MOVE@ (INTA, INTB, N4)

Now, from Mr. Schultz, we have:

CALL VEC_\$ICOPY (INTA, INTB, N3)

The difference is, this latest utility has been written by Mr. Schultz using Intel assembly language inside of a FORTRAN subroutine! Section 1.7 of the compiler manual says : "*FTN77 / 486 supports a CODE / EDOC facility for in-line Intel 80486 assembly instructions in 32-bit protected mode.*" Well, writing assembly language seems to be yet another of Mr. Schultz's many talents, and the speed of Salford EMTP simulation has increased as a result. In those cases where assembly language replaced FORTRAN DO-loops, the speed gain seems obvious. But how can the Schultz code beat Salford library utilities such as MOVE@ (see above)? This seems less logical, but apparently is no less true according to NYPA benchmarks. To conclude, a combination of DO-loops and calls to Salford library functions have been replaced by calls to Schultz-coded routines that correspond to HP 9000 Model 720 library functions (that is where names such as VEC_\$ICOPY originated).

Timing segment TGRUN1 has been simplified by Mr. Schultz. Your Editor had used the library routine CLOCK that returned seconds since midnight. This had the disadvantage of requiring special logic for cases that terminated during a later day, and was wrong for jobs lasting more than 24 hours. Mr. Schultz avoided such problems by using SECONDS_SINCE_1980@ . The improvement became available November 26th.

TIME_IN is the name of a new subroutine from Mr. Schultz that documents the start of execution on the screen only (not the .LIS file). An example follows :

ATP Started at 13:36:14 on Sunday, 19 December 1993

Parallel routine TIME_OUT produces a corresponding line as Salford EMTP halts ("ATP Finished at"). A second line then gives total elapsed time. For example:

Total Execution Time was 13.02 seconds

The standard 25 rows and 80 columns of old screen displays should **not** be assumed. This is the excellent suggestion contained in E-mail dated the 2nd of December from Harald Wehrend of the University of Hannover in Germany. "*I noticed ... that the screen always returns to standard text mode with 80 columns when graphics is finished This happens even when the screen was in the 132-column mode before. Wouldn't it be better to return to the text-mode that was selected before?*" Yes,

it certainly would! This is a great idea. The change was made the following day. Salford TEXT_MODE@ had been used, and all calls to this were replaced by calls to Robert Schultz's RESTORE_TEXT_SCREEN (this is the name of the routine that restores screen text following a return from graphic mode as explained at the bottom of column 1 on page 2 of the April issue). Here, we do not much care about the restoration of old text. But we do care about the mode, and Mr. Schultz did properly sense this using Salford GET_VIDEO_MODE@ and restore it using SET_VIDEO_MODE@. It would seem that the only thing Mr. Schultz missed was the desirability of removing all calls to TEXT_MODE@ from the program (his attention then was confined to CALCOMP PLOT graphics).

An asterisk or star "*" is accepted as an input parameter for Salford EMTP where the .LIS file name is to be specified. This is for users who want the output file to be parallel to the input. It is the batch-mode equivalent of <CR> or -R for interactive use. The following two lines are equivalent, beginning November 2nd:

```
CALL RUNTP DISK DC3. DC3. -R
CALL RUNTP DISK DC3. * -R
```

It is this new, second usage that can be found in file RUNDUBE.DAT (to run all 35 of the new MODELS examples) that was created by Laurent Dubé using the MS Windows editor on the output of **DIR *.DAT /B**.

NO{STRIP is a new UTPF ident (the serialization information of columns 1-8) that is used to prevent the stripping of in-line comments. Recall Adobe PostScript was added as an extra, auxiliary output for Salford EMTP plotting as mentioned in column 1 on page 11 of the previous issue. Well, the NYPA PostScript, which involves imbedded braces, requires some such exception to usual practice of translators. Whereas with TPPLOT we used "Protect open brace" on the extreme right (ending in column 80), here it seems easiest to use a special UTPF ident. Anyone else who might perform translations should take note of this change.

Variable LENREC of STARTUP is a new binary flag to control only C-like .PL4 files. Value zero means no change from the past whereas value one is taken as a request for extra empty space between the header and the initial time step. The suggestion for this extension came from Prof. Bruce Mork of Michigan Tech in Houghton. The initial communication was dated February 12th, 1992, although your Editor did nothing until just last month. In E-mail dated December 6th, Prof. Mork was very specific about the need: *"To my knowledge, standard fortran does not allow a record to be written (or read) if that record is longer than the RECL= that was specified in the OPEN statement. Therefore, it seems more logical to start the timestep data at an offset which is an integer multiple of the length of the data for one timestep. The fortran compiler I use (waterloo Watfor77) is a standard F77*

implementation. If you make the above-described implementation, the efficiency with which PL4 files are read could be greatly improved." So, the extension was added, and on December 10th, an illustrative DC3.PL4 was sent by E-mail to Prof. Mork for verification.

C-like .bin files no longer are used for disk storage of EMTP tables by Salford EMTP. The change will be noted in test cases such as DC-32 (changes were made to the archive ALLDAT.ZIP on January 10). Of course, the spectacular improvement of Robert Schultz's turbo table dumping (see separate story) forced the change. But it might have been made anyway. See the separate story about virtual LUNIT6 output, and those benchmark tests PROOF*.FTN. In retrospect, what was great 2 or 3 years ago --- before disk caching --- no longer seems to be worth special effort. Salford EMTP times have changed once again for the better --- but only for those who use disk caching (preferably DOS SMARTDrive if one can afford the associated, special memory).

Improvements to Salford TPPLOT

Parameters were added to TPPLOT following a December 7th telephone call from Scott Williams of Westinghouse in Pittsburgh, Pennsylvania. He agreed that one could use the RELAY subcommand of the EXPORT command to produce uniform columns of numbers that specify signals as functions of time. However, he wanted to do the .PL4 file conversion outside of TPPLOT, as an MS-DOS command. Well, with arguments, TPPLOT can be run from a command file so that it behaves like a separate, custom utility. The name of the disk file to be processed is the one parameter that must be passed into TPPLOT, for this example. A new command file RELAY.BAT illustrates this use beginning December 8th. To process DC3.PL4 the user would just send **RELAY DC3** at the MS-DOS prompt. This requires about 7 seconds on your Editor's 486/33.

Who needs interactive plotting? Not NYPA, as Robert Schultz explained in E-mail dated December 21st:

"Some day we will document the EMTP 'user culture' at NYPA, which I believe is very different from most other production environments. In essence, we use Apollo scrollable graphics, or the functionally equivalent Salford Display Manager, to produce high resolution CALCOMP CRT plots and we use the PostScript. We do not use TPPLOT during the engineering phases of EMTP use. TPPLOT is sometimes used for final report generation, although with the new generation PostScript plots, this is becoming even less of a requirement. The key to this style of EMTP usage is the availability of first-class CALCOMP graphics. I guess this contrasts sharply with other users around the world whose day-to-day routine revolves around TPPLOT!). This explains our lack of concern about the status of LECPLOT up to now."

Dr. Matt Donnelly of Pacific Northwest Laboratory (Battelle) suggests that new versions of MATLAB might be programmed to accept unformatted (even C-like) .PL4 files directly. This interesting possibility was contained in E-mail of the Fargo list server dated November 16th. Those wanting to correspond with Dr. Donnelly are advised that he uses address mk_donnelly@pnl.gov He wrote the following: *"Matlab 4.x on all platforms allows for more flexible data input than just the traditional 'load' command. One could set up an m-file to read even binary PL4 files with a little work. It is my opinion that this input data 'filtering' should occur in the secondary application program. Mathworks is not the first vendor to recognize the importance of allowing for flexible input filters. ATP currently supports a good mix of output file formats and ATP developers cannot be expected to modify output formats for an arbitrary mix of post-processing applications."* Amen, and thanks for the contribution!

News from Outside USA and Canada

The dominant EMTP news from overseas must be the end of LEC in Europe (see later, separate story).

A printed copy of the October newsletter was mailed by BPA to each of its primary EMTP contacts (and this includes LEC for the last time) on November 23rd --- one day after imminence of the operation was announced in E-mail of Prof. Bruce Mork's Fargo list server.

Iran established contact with the Can/Am user group by means of a letter dated October 12th from the Electric Power Research Center (EPRC) in Tehran. Managing Director S.M. Tabatabaee wrote in response to a copy of your Editor's 28-page memo about LEC dated September 10th. Unfortunately, the U.S. government and Iran have their differences, and higher BPA management decided that contact was to be discouraged. So, Iranians were referred back to Europe as announced in E-mail of the Fargo list server dated November 17th. Does Europe have a unified position concerning cooperation with Iran? What about Libya? If so, what is European policy?

One HP9000 Model 720 RISC workstation is used by New York Power Authority (NYPA) in White Plains for part of its support of EMTP. Along with a 3-page explanatory letter by NYPA's Robert J. Meredith dated November 3rd, a 4mm DDS cartridge tape containing all associated files was mailed to Taipei for interested users in Taiwan, Hong Kong, or Singapore. The initial request was from a Prof. Chang-Fa Yang of National Taiwan Institute of Technology, and he is expected to share the NYPA files with others who might have interest within the Republic of China EMTP User Group.

LEC's journal, EMTP News, has not been reprinted and mailed this year. First, the obligation was badly

delayed; then, since June, the possibility was suspended during writing about LEC irregularities. Now, with the year done, there is some sentiment for satisfying the obligation just one last time. But no December issue was received either by the Can/Am user group or BPA (which pays for a separate subscription), and the publisher (LEC) is out of business. Your Editor is not motivated to continue. Fortunately, no checks for \$15 should have been cashed, so there should be no money to return. If anyone knows otherwise, please notify the user group.

More about Electronic Mail (E-mail)

The October newsletter was made available on the Fargo server as announced on November 17 by list server mail from Prof. Bruce Mork. Available at the same time was **steer02.zip** which is the WordPerfect disk file of the 10-page "Part II" memorandum about LEC financial and other irregularities (pages 11-13 of the July issue).

Details of E-mail use by ATP developers in Portland can be found in a separate, later story. Included are both the user group's Agora and also BPA's recent extension to AT&T mail, which can exchange with Internet.

New York Power Authority (NYPA) in White Plains has gained access to E-mail. This is the superb news from Robert A. Schultz, who for the first time sent your Editor a message via Internet on November 22nd. No, NYPA does not yet provide the service. But Mr. Schultz seems to have found his own inexpensive local BBS (Bulletin Board System) that is named Execnet : **robert.schultz@execnet.com** \$56/year allows virtually unlimited access using a local telephone number (free calls) at 14.4 Kbps. Mr. Schultz pays another \$39/year for Internet, with FTP and TELNET coming in March.

Mr. Wehrend's E-mail address at school is changing, to the primitive **nhmchawe@rrzn-user.uni-hannover.de** In E-mail dated January 27th, he advises that usage of his old address be discontinued soon: *"This change is caused by the planned switch off for the old CYBER CDC system at the beginning of March. Up to this date, all mail addressed to my old address will be auto-forwarded to the new one."*

An inquiry about FIX SOURCE1 was broadcast by the Fargo list server on December 4th. Originating on the campus of Texas A&M University in College Station, this must have confused many subscribers until December 7th when your Editor clarified that the data in question applies not to ATP but rather to the DCG / EPRI EMTP! The author, Joydeep Mitra, explained in private E-mail that he did not know there was a difference in data. Apparently Mr. Mitra had not yet progressed to advanced control system modeling using Laurent Dubé's MODELS (which is not available in DCG / EPRI EMTP) !

Mid-night speed of the Fargo list server is noteworthy. Both MODELS author Laurent Dubé and this writer have experienced demonstrations of the phenomenon during recent months using PINE on Unix computers. Of course, one first sends his message to Fargo. Before leaving PINE on Agora, this writer may save his incoming mail, which takes a minute or two. Then, before he is able to save from the **Sent Mail** folder a copy of what had just been sent to Fargo, PINE issues a news bulletin indicating new mail! Sure enough, this is the return from Fargo, and it is saved instead of the outgoing copy (which then can be deleted).

Warsaw, Poland, was first heard from by E-mail on November 24th when Dr. Włodzimierz Kalat sent a message to the Fargo list server. Agora's Pine was used to store the address, and then send a reply the following day. Apparently because of Bitnet rather than Internet use in Warsaw, the message was sent indirectly through Fargo (note no indication of the ultimate destination in Europe) : <pwlok%plwatu21.bitnet@vm1.nodak.edu>

Ordinary E-mail can be used in place of FTP to obtain files from remote computers such as Prof. Bruce Mork's Fargo server. How does this work? An interested user merely sends his request in the form of E-mail to some cooperating site (e.g., Princeton) that has FTP. This site then acquires the file, converts it to E-mail, and sends it to the address of the original request. Neat, eh? Your Editor first heard of such service from Laurent Dubé, perhaps last summer. But we Oregonians were a little skeptical of speed (it might take hours or days), so we decided instead to pursue our own, real FTP using Agora (and later ORST for Mr. Dubé). But the idea resurfaced again at the end of last year in E-mail from Robert A. Schultz of New York Power Authority (NYPA), who had been using BITFTP of Princeton University. In E-mail dated December 30th, Mr. Schultz provided documentation about the service ("*The following is the BITFTP.TXT from BBS file _ennet.zip*"). Information of the following paragraph is entirely from this file.

BITFTP -- Princeton BITNET FTP Server

BITFTP provides a mail interface to the FTP portion of the IBM TCP/IP product ("FAL") running on the Princeton VM system, to allow BITNET/NetNorth/EARN users to ftp files from sites on the Internet. To use BITFTP, send mail containing your ftp commands to BITFTP@PUCC (or to BITFTP@PUCC.Princeton.edu). The first command to BITFTP must be "FTP", "FTPLIST", "HELP", or "VMS". Use "HELP" to request a current copy of this help file. The following is an example of a BITFTP request:

```
FTP f.ms.uky.edu NETDATA
USER anonymous
CD /pub/msdos/Games
DIR
BINARY
```

```
GET robotron.arc msdos.robotron
```

```
QUIT
```

.... BITFTP does not provide a PUT capability, and there is no intention to make it do so in the future. If BITFTP is unable to connect to the host you specify, it will send you mail after the first attempt, but will keep trying at intervals over three days. The load on BITFTP is often very heavy, and network backlogs are often so great that it may take several days for a file to get to you once BITFTP sends it, so please be patient

NYPA Optimizes EMTP Computation

Using compilers that optimize for parallel (vector) computation, Robert A. Schultz and Robert J. Meredith of New York Power Authority (NYPA) in White Plains have made important structural changes to EMTP. The remainder of this story summarizes changes that were extracted from a disk file NYPASRC.EXE dated November 10th.

The isolated global SAVE statements of segments GUTS2A and GUTS2B were replaced by selective statements that only preserve needed variables. Mr. Meredith wrote the following comment about the source of his intelligence: "*The following were identified by Apollo a88k compiler.*" Tools clearly have improved since your Editor was using Apollo (with no such detailed help, your Editor could not be sure what was needed so he decided to allow a global SAVE in these two input routines). APPEND was another module so improved -- temporarily, until SPY @5 execution using Salford EMTP demonstrated erroneous execution. Restoration of the original global SAVE then corrected the error. If the Apollo compiler really made the recommendation in APPEND as stated, it was wrong. February 10th E-mail from Robert Schultz promised a revision "*next week.*" It also reiterated the importance for NYPA : selective use of SAVE is "*mandatory for optimizer use on the HP.*"

SUBROUTINE MULT which multiplies a symmetric matrix and a vector has been replaced by 3 new routines MULT0, MULT1, and MULTN. Each of these new modules provides one of the 3 modes of original usage corresponding to the 5th and final argument IOPT that no longer is used. Unnecessary IF statements inside loops seem to be deadly for computers with vector capability. In the case of MULT, NYPA decided to triple the code to remove one argument and one such IF statement. According to Mr. Schultz, even Salford EMTP benefits some: "*Salford vector routines do literally exploit Intel vector instructions, which are extremely limited in scope compared to a CRAY, for instance.*"

In TACSUP, a new variable MTEST = N + J was defined to save one addition inside an inner loop. This is one of several such "*modifications to allow the Salford*

Optimizer to function. All ATP code now compiles perfectly with option /OPTIMISE" (not yet used at BPA).

GUTS2B illustrates use of a new vector routine that was added by NYPA . The old code was:

```
DO 1001 I=1, NJ
1001 ILNTAB(NDX1-I) = ILNTAB(NDX2-I)
Mr. Schultz has replaced this by reliance upon an all-new
routine: CALL vec_Sicopy_i ( ILNTAB(NDX2-1), -1,
1           ILNTAB(NDX1-1), -1, NJ )
```

OVER3 illustrates in-line vector initialization that has been modularized by NYPA. The old code had:

```
DO 8005 I=1, LBRNCH
  LITYPE(I) = 0
8005 IMODEL(I) = 0
```

Well, such code is no longer being tolerated. Instead, Mr. Schultz uses two calls to vector zeroing routines.

RENUMB illustrates simple sequential computations that have been combined into a single, higher-level statement. The new ILOC(I) = NORDER(ILOC(I)) replaces the old:

```
J = ILOC(I)
J = NORDER(J)
ILOC(I) = J
```

Documentation is complicated by such changes to higher-level statements, unfortunately. Whereas before 3 in-line comments explained the three operations step by step, now a separate comment card would be required.

Segments LU2WRT and LU2RED modularized the WRITE and READ statements of LUNIT2 table dumping. This was done 5 or 6 years ago for David Szymanski's C. Well, apparently your Editor was not careful with the FORTRAN. Since the length of the vector was an argument, the vector could be dimensioned with this value rather than unity. Then the WRITE or READ statement could involve the entire vector without subscripting. For example, in LU2WRT the original statement WRITE (LUNIT2) (KAR(J), J=1, NUM) has been replaced by WRITE (LUNIT2) KAR which Mr. Schultz knew was more efficient.

Alignment of SPY vectors FKAR1 and FKAR2 is another of Mr. Schultz's accomplishments. For years, compilers such as the one for DEC VAX / VMS have issued warnings about these. Of course, no one cared about optimization of interactive SPY, so the warnings were not important as long as consequences (unoptimized compilation) were confined to the subroutines of usage. But apparently the HP9000 720 compiler used by NYPA refused to optimize non-SPY subroutines as long as all variables were not aligned. Here *alignment* refers to the 25- or 30-year-old half-word problem of IBM /360 and /370 FORTRAN. For any variable in any COMMON block, the number of bytes preceding it in the block should be a multiple of the length of the variable. Intentionally, either FKAR1 or FKAR2

violated this rule. So, Mr. Schultz devised a way to remove FKAR1 and FKAR2 by using one new FUNCTION FKAR and one new SUBROUTINE REL_REAL8_STORE which have been placed at the end of the SPYTAC segment. Other affected segments were EXAMIN, DEPOSI, ANALYT, and SPYINK.

BLOCKD51.BIN is the disk file of program text that is connected using I/O channel LUNTEX (defined in the STARTUP file). The newer, faster way advocated by Mr. Schultz is chosen by a negative LUNTEX (the value -11 will be found on the GIVE1 disk beginning January 10th). In this case, no disk file is needed (and none is distributed by the user group) because the very same text has been built into the program using DATA statements. Program text is being stored in RAM so Salford paging will do the input of text rather than FORTRAN READS from a disk file. A separate, later story discusses details.

LUNIT6 output can be accumulated in virtual memory prior to flushing to disk. Mr. Schultz realized that this is faster than writing each line to disk as each line is produced. A later story provides details.

N15_TABLE is a new character mapping that was introduced into TFLUSH by NYPA to replace twelve successive character checks such as the following one:

```
IF ( VECEXT(K:K) .EQ. 'B' ) N15 = 12
```

Mr. Schultz's alternative is a single line that uses one standard intrinsic function (presumably in-line) and one reference to new array N15_TABLE :

```
N15 = N15_TABLE ( ICHAR ( VECEXT(K:K) ) )
```

Repeated, successive, identical diagnostic conditions such as IF (IPRSUP .GE. 7) of MATVEC have been replaced by a single IF - THEN - ENDIF construction. The saving may be so small as to be unmeasurable, but your Editor endorses Mr. Schultz's reform in any case.

Buffer for LUNIT6 Output to Disk

LU6VRT is a new parameter of STARTUP that allows the user to delay the writing of LUNIT6 output to disk until the output first has accumulated sufficiently in RAM . This is yet another innovation from Robert A. Schultz of New York Power Authority (NYPA) in White Plains, where the code in question had the name *overdrive* associated with it. This term is believed to come from the American automobile industry. Forty years ago, those big American cars had manual transmissions with only three forward gears. Some manufacturers (e.g., Ford) offered an optional, automatic, higher gear that was named *overdrive*. It was designed for high-speed cruising on level ground --- clearly what one wants for batch-mode EMTP simulation!

The initial implementation of virtual LUNIT6 output

at NYPA was limited to Salford EMTP for MS-DOS computers. Mr. Schultz's code was ingenious in that he only made a few changes in two places. But he used a virtual scratch file to provide the new buffering, and this is not universal. A second problem had to do with a compiler bug: although execution using Salford compiler and DOS extender version 2.71 seemed to be perfect, the output using DBOS 2.66 involved minor errors (typically one or 2 extra bytes at the end of some lines). Since the user group was not anxious to distribute a new Salford DOS extender DBOS with each new program version, this was a problem. Finally, your Editor wondered why I/O unit LUNIT6 should be used at all, if what one really wants is a RAM buffer. Why not just program the RAM buffer directly, and hope for improved efficiency (compared with Salford scratch files) as a fringe benefit? This is how your Editor spent the weekend of December 11th and 12th : programming the new RAM buffer.

Value zero of LU6VRT means no change from years past (i.e., direct output of each LUNIT6 line as it is created). On the other hand, positive LU6VRT is taken as the size in bytes of the RAM buffer that will be used for output to disk. For Salford EMTP, there is no practical limit, although $2^{**}16 = 65536$ has initially been imposed for compatibility with other systems (which may well have limits). Value $32768 = 2^{**}15$ will be found in STARTUP on the GIVE1 disk of Salford EMTP distribution.

Salford EMTP uses WRITEF@ to flush to disk the text that has accumulated in RAM. Since WRITEF@ is the same Salford library routine that is used to create those high-speed C-like .PL4 files, one must wonder whether part of the speed gain came from the use of WRITEF@ instead of **WRITE (LUNIT6)**. Surprisingly, the answer seems to be no. The problem with old LUNIT6 output lay not in the use of **WRITE (LUNIT6)**, but rather in the fact that records being written were so short (never more than 132 bytes). To demonstrate this, your Editor wrote little programs to create and dump 80-column card images 12500 times. On his AT&T 486/33 using 2 Mbytes of MS-DOS SMARTDrive disk cache, this required about 18 seconds for either output method. But a big gain resulted from the addition of a 1025000-byte holding vector. This then could be dumped once, at the end, in 6 and 7 seconds. As expected, creating the big vector in RAM took little time (output began after about 1/4 second). Anyone wanting the test cases for use with some other compiler should ask for PROOF*.FTN

Exactly LU6VRT bytes will be written to disk during each flush except the final one (which generally will be smaller). Mr. Suhrbier suggested such logic as a replacement for your Editor's original implementation that dumped all text through the end of the first line that exceeded the limit. As coded for DEC VAX/VMS, the transfer must always begin on a block boundary,

according to Mr. Suhrbier. By having ATP always dump a multiple of one disk block (2048 bytes), redundant rewriting of the remainder is avoided.

Note carefully that the RAM buffer for LUNIT6 output only applies to text that goes to disk. Output to the screen is never being buffered (unless by the operating system). Execution must begin with either DISK or BOTH for there to be any effect from positive LU6VRT. If output goes only to the screen, nothing has changed.

Two new lines will be seen at the start of each disk file of LUNIT6 output. Illustrating this for DC-1 :

EMTP begins. Send (SPY, file_name, DISK, HELP,..
--- 1323 cards of disk file read into card ...

Previously, these went to the screen rather than to disk when DISK was used. The new destination is logical enough, but it left the interactive user without a second or later prompt, and wondering what to do next. So, two new screen prompts have been added for the DISK and BOTH cases, respectively:

Ok, output goes to disk. Send input data file name:
Ok, output goes to both disk and screen. Send ...

More output that will be seen for the first time in .LIS output are the messages about \$INCLUDE and data sorting by class. This will be illustrated for DC-17 :

--- Pass 1. Card = 13. Ready to open \$INCLUDE ...
--- Pass 2. Card = 24. Ready to open \$INCLUDE ...
<====> Done with "/"-card sorting by data class. ...

This is more progress (such output really does belong with the solution). It was absent from the disk file before because it preceded the prompt for the disk file name. Now, since output is stored in RAM, there is no need for a file name until the program is ready to dump its RAM buffer.

Isolated blank lines are handled more efficiently than before. New, special code will add the line-termination character (for Unix) or characters (<CR> and <LF> for both MS-DOS and VAX / VMS) to the output buffer. Previously, service of each CALL BLANK6 was forced through the general output logic, and this resulted in the output of one blank byte. Not only is the code now faster, that former blank byte is saved.

DIAGNOSTIC output is incompatible with the new virtual output just as it always has been incompatible with Salford windows, or Szymanski's Unix windows (which were controlled using C-language). No, DIAGNOSTIC output is not lost, but unfortunately it is separated from the production output. Now, DIAGNOSTIC output is directed to the DEBUG.LIS disk file. That is the bad news. The good news is that all unintentional use of **WRITE (LUNIT6)** for production printout was similarly misdirected, so was obvious by its absence from the .LIS file. The new virtual output provided a perfect tool for locating all such shortcomings, which have been removed. Included were all output associated with Prof. Chen's

transformer model (USE SEATTLE XFORMER), and Dr. Tsu-huei Liu's LINE MODEL FREQUENCY SCAN (LMFS). So, program structure has been purified.

Execution has been sped noticeably when output goes only to disk. For example, on December 22nd the time to execute RUN.BAT (for all standard test cases) had dropped to 21:33 (3:27:23 - 3:05:50). This was using LUNTEX = -11 and LU6VRT = 32768 (compare with the 26:21 mentioned earlier).

DEC VAX/VMS is compatible with the new virtual LUNIT6 output thanks to coding by Randy Suhrbier of BPA, who supplied subroutines that mimic the operation of Salford library functions OPENRW@, WRITEF@, and CLOSEF@. Installation-dependent INCLUDE files TFLUS* will be found in VAXET.DAT and the new utilities themselves have been inserted near the end of VAXMODS.RUM . The change was made January 13th.

Mr. Schultz will be allowed the final word on this work that he inspired. Interestingly, he believes that we may not need it forever. In E-mail dated December 21st, he offered the following observation: *"In a more perfect world, ideally, the efficiency problem would be addressed by the compiler writers or runtime library or operating system. I would hope that your new incarnation will not be necessary in the future when the native performance improves to 'state-of-the-art'. This is already witnessed with the HP9000 Model 700. I will try to convince Salford to develop higher performance formatted I/O. The parts are in place. Some Fortran versions allow specification of a 'blocksize' (buffer) parameter which can go almost all the way in terms of improving performance."*

LEC Financial and Political Problems

This is a continuation of the story that began on page 8 of the preceding issue. It documents the end of LEC (the Leuven EMTP Center in Belgium).

Your Editor's 10-page response to LEC Chairman Van Dommelen (see the first paragraph on page 9 of the preceding issue) was mailed from Portland on Saturday, November 20th, to 160 LEC members plus 7 overseas user groups. Two more copies were mailed to Leuven : 1) Hans Claes, the Administrator of Leuven R&D; and 2) Prof. R. Govaerts, the Chairman of the Electrical Engineering Department. All 169 were sent the same way, with \$.95 in stamps (for 1 ounce) affixed to business-sized Air Mail envelopes.

LEC Manager Guido Empereur finally broke his long silence on November 18th with 3 E-mail messages from the Fargo list server. The big news contained in one of them was the following: *"LEC staff was fired on October*

1st, 1993, and workman's compensation became effective. We all agreed to continue serving our members at best (sic) up to December 31st, 1993, without accepting new jobs in that period."

Prof. Correia de Barros of IST in Lisbon did preside over that meeting of the Steering Committee (SC) at the Hotel Binnenhof in Leuven on Nov. 25th. A total of 14 LEC members were present, including 5 members of the Steering Committee. As should be detailed in the April issue, plans for a replacement user group were made.

That final LEC meeting the following day must have been rather anticlimactic. No LEC representative having any technical knowledge of EMTP was present. The usual sloppy accounting of money was presented, with just a few innovations. Whereas in years past the accounting was said to cover a full 12 months (even though income was recorded for only 8 of these, typically), this year the period was January 1st through October 31st, which left gaps at both ends. New this year is an unbelievable claim of a loss of 3.253.648 BF. More accurately characterized, this would seem to reflect separation payments for those fired employees --- paid using unreported membership fees of years past (about one third of the estimated total of 10 million francs)! To conclude, financial credibility of LEC politicians has not been enhanced by its final report.

Was this *accounting fiction* by LEC not challenged by anyone in attendance? This was your Editor's question of Prof. Correia de Barros during a telephone conversation immediately following the meeting. The answer was *no*. Apparently all decided to stop arguing about the past and move on to a replacement EMTP user group for Europe.

The numerous irregularities of LEC would seem to be approved by Department Chairman Govaerts, since he has refused to investigate the hidden income, denial of voting rights, etc. Quoting Chairman Govaerts from Annex 1 of the final report (an attachment to a letter from Leuven R&D Administrator Hans Claes dated December 13): *"The academic authorities of our University have decided that Prof. Van Dommelen should not attend this meeting. This decision has been taken after hearing about the aggressive and unrestrained language with which Prof. Van Dommelen has been confronted at the Steering Committee meeting in Lisbon. They have also investigated the written allegations made by Dr. Meyer in a 6 page letter to me in which he accuses K.U. Leuven R+D, Prof. Van Dommelen and his staff of, I quote 'If it were necessary for you to pick one English word with possible criminal implications to describe LEC's financial irregularities, my preference would be fraud.' Our University does not accept such an approach nor does it answer citations of anonymous critics. I have therefore been mandated to chair this meeting and to defend the interest of the University."*

Aggressive and unrestrained language? Such as

what? Prof. Van Dommelen objected to many things in the pages that he wrote in response to your Editor's 28-page memo dated September 10th, but this was not one of his many arguments. Well, maybe questions such as "*How much unreported income has LEC been hiding?*" seemed even more threatening to the Chairman of the Department in which that money was being spent! In the absence of evidence to the contrary, your Editor must suspect Prof. Govaerts' accusation about language to be nothing more than a weak, last-minute excuse to explain the otherwise-embarrassing absence of an LEC Chairman who did not want to face his unsympathetic members. If any participant of the Lisbon SC Meeting believes otherwise, please send specifics in writing to Portland for consideration.

Prof. Govaerts' mention of *anonymous critics*, and his objection to one sentence of this writer's private communication with him, are peculiarly and deliberately misleading. If this is his best excuse for not investigating LEC abuses, it seems unlikely that he will convince many readers, as the remainder of this story will explain.

First consider this *anonymous* nonsense. If Prof. Govaerts has not met LEC critics personally, it is not the fault of the critics. Many times your Editor encouraged him in writing to meet with SC members at the Binnenhof on November 24th. But he did not appear, and the acting Chairman, Prof. Correia de Barros, has reported that she never was contacted about the possibility. Prof. Govaerts also was given names and telephone numbers of several Europeans who should be able to corroborate different aspects of the complaints about LEC. On page 4 of his October 29th letter, your Editor wrote the following: "*If any particular detail of my writing concerns you, please mention it, and I will try to be more specific. I would be happy to telephone you with Mr. Dubé on the line, if you want to save time while you form your thoughts about LEC operation. ... Finally, either I or Co-Chairman Liu (my colleague here at BPA, Dr. Tsu-huei Liu) could be present in Leuven soon (next paragraph) if you wanted.*" Your Editor knows of no indication that Prof. Govaerts ever made any attempt to pursue any of these numerous avenues to learn specifically and personally (not at all anonymously) what was wrong with LEC. Finally, public E-mail of the Fargo list server dated February 7th attempted to locate "*anyone (not just these named persons) who might have been contacted by Prof. Govaerts for this purpose prior to the late-November meetings in Leuven.*" No one has reported such contact.

To understand just how misleading the quotation by Prof. Govaerts is, read more context from what your Editor wrote to Prof. Govaerts on October 29th: "*In your first letter, you used the word embezzlement. You do have a way of cutting through more polite language to the central issue! Interestingly, to the best of my recollection, you were the first person ever to have*

suggested this word for the case of LEC. In this country, 'embezzlement' is often used by newspapers in reporting about employees of banks who steal money secretly from their employers. My dictionary says 'to embezzle' is 'to take (money or property) for one's own use in violation of a trust.' Yes, there is an element of this in the LEC case. But the word fraud more commonly is seen in newspaper stories about criminal prosecution of so-called white - collar (i.e., non-violent) crime. Fraud is more general, and might apply to a broader range of LEC activity (e.g., those published K.U. Leuven accounts of LEC work as discussed in Section VI-B). My dictionary defines fraud as 'a deception deliberately practiced in order to secure unfair or unlawful gain; a piece of trickery; a swindle.' No, you have not misinterpreted the seriousness with which I and others regard the LEC matter --- even though a different English word seems to have occurred to you to summarize the problem with LEC. If it were necessary for you to pick one English word with possible criminal implications to describe LEC's financial irregularities, my preference would be 'fraud.' Yet, I point out to you that I did not do this in my 28-page memorandum. My writing merely describes, using less-judgmental language, what seems to have happened."

NYPA Eliminates BLOCKD51 . BIN

As mentioned in the story about NYPA optimization, program text file BLOCKD51.BIN has been eliminated by Robert A. Schultz of New York Power Authority as a way of gaining speed. While an excellent idea, there are some tradeoffs that should be understood.

Retention of the possibility of using BLOCKD51.BIN seemed to be desirable for several reasons. The switch is made easily enough. If STARTUP parameter LUNTEX is given a positive value (typically 11), operation will continue as in years past. One reason to allow this old alternative was to provide a means for unbiased timing at any time.

Speed of execution for your Editor's AT&T 486/33 is easily documented for the full set of standard test cases. Prior to the first, and after the last execution, MS-DOS TIME commands are issued. The same program was used, and condition of the disk was comparable. First one test (Schultz's scheme) was performed, and then the other (after erasing the minus sign from LUNTEX). In both cases, batch file RUN.BAT was used. Schultz's new scheme required 26:21 whereas BLOCKD51.BIN required 28:22 (nearly 8% longer). In both cases, 2 Mbytes of MS-DOS SMARTDrive disk cache were used.

Mr. Schultz gained more speed because he eliminated the subroutine READ11 that modularized the original READ for computer expert David Szymanski, who wanted to use C-language for faster I/O under Unix). Later, if

we decide to eliminate BLOCKD51.BIN entirely, it would be reasonable to eliminate READ11 at the same time. Changes would be confined to just two subroutines: TFLUSH and HELPER . Particularly for the faster Unix-based RISC workstations, there should be a speed gain from storing text in RAM . C-language coding might beat FORTRAN coding for input, but most likely it would not beat paging of the operating system. The days of modularized READ from LUNTEX probably are numbered. What seemed like a good idea in 1987 or 1988 no longer looks like a good idea here at the end of 1993. Computers have changed once again!

Multilinguality is adversely affected by moving text into the program, unfortunately. When text was only in a separate disk file, the program was nearly independent of language (e.g., the choice between English and Japanese). One could change the language by changing the disk file. Now that such text is linked into the program, one can still change the language, but that choice must be made earlier -- at linking rather than execution time. It is true that RAM could be loaded by reading from disk, but no such extension has yet been provided. The first person who wants to use some language other than English for program text is advised to check with your Editor.

Utility OVER51AI.FTN has in the past been used by developers to create BLOCKD51.BIN as identified near the end of the STARTUP file. Well, Mr. Schultz modified this utility to produce DATA statements. But since your Editor wanted to preserve both alternatives, he decided to combine the two, which in honor of NYPA has been named OVER51NY. Whether a user wants both of them or not, both BLOCKD51.BIN and also BLOCKD51.FTN automatically will be produced each time OVER51NY is executed.

Computers with minimal RAM might be adversely affected by lack of BLOCKD51.BIN which kept the text on disk rather than in precious RAM. Fortunately, this is not an issue for the average user, who has 4 or more Mbytes. If some reader wanted to run comparisons using 2 Mbytes or less of RAM, it might be informative.

BPA : a Government Corporation ?

"BPA Inc. must pursue social goals, not business only, states warn." This is the headline of a story on pages 7 through 9 of the January 17th issue of the *Clearing Up* newsletter. The present mention is a continuation of the story that began on page 9 of the preceding issue. About the NewsData Corp. newsletter, it seems that there are different issues for different regions of the country. The version that concerns itself with BPA is for *"Northwest Energy Markets"* (a subtitle). Content is supposed to be *"Energy and utility news for the US Pacific Northwest and Western Canada."*

So, the struggle for control of a reworked BPA is well under way, and suspicions of the proposed changes abound. Whereas the user group is concerned because of computer programs about possible loss of FOIA, social engineers share the concern, but for completely different reasons. Without existing FOIA, environmentalists, states, Indian tribes, and others no longer could see as clearly what BPA was doing, so might be adversely positioned to challenge some policies. There are relevant statements from the Washington State Energy Office and the Natural Resources Defense Council. While BPA customers seem to be supportive, generally, *"The Public Power Council joined the general chorus of concerns over BPA's wanting authority to keep appropriate information proprietary. 'We are concerned that BPA could abuse its market power by withholding information,' said PPC manager Bill Drummond."* The story began as follows: *"BPA held hearings last week in Seattle, Spokane and Missoula on draft legislation to make Bonneville a government corporation ... with utilities, industrial customers and business interests urging passage ... and environmentalists and state and tribal governments expressing dismay over BPA's attempts to redefine itself as a business without sufficient assurances that it will continue to pursue social goals."*

On the salmon-saving front (those fish and wildlife obligations), thoughtful readers may have wondered how BPA possibly could spend \$300 million a year. Salaries of employees would not explain much of this, obviously. If every BPA employee were fired, the saving of wages would not reach this staggering total, probably. So how was the money spent? In fact, about half of it may not have been spent at all. As explained in a short story on page B16 of the December 7th issue of *The Oregonian*, half was never earned: *"Because of the need to store water for fish flows, she [Mahar] said, the BPA wasn't able to sell as much power later in the year when it had extra. That meant a \$146 million drop in sales."* That's right, to generate electricity, water must flow through the turbines. There has been less of that in recent years, and this is why some persons jokingly refer to BPA as BFA (the F stands for **Fish**, of course).

E-mail in Portland : BPA and Agora

Agora continues to be used to receive EMTP-related E-mail for both the user group and BPA. Except for a week around the end of last year, when it was being rebuilt, Agora seems to have received mail flawlessly.

But the sending of mail from Agora became difficult around the end of last year. Unfortunately, the trouble was discovered in a message to the Fargo list server, which all recipients (including your Editor, using Agora) found to be hopelessly encoded. Eventually (as with Mohave SSR, it took 2 disasters before principles were

understood), it was concluded that the trouble occurred only for missives that pulled in (included) an external disk file using Ctrl-R within Pine. In E-mail dated January 10th, Prof. Bruce Mork reported the key clue: *"In the header of the files that are encoded appears the line Content-Transfer-Encoding: BASE64"* Responding to this on January 11th, owner/operator Alan Batie wrote: *"Pine seems to be using MIME to encode included sections, and it seems that a lot of remote sites aren't able to cope. I'll have to investigate, but it'll take me a little while."* Not wanting to wait, Laurent Dubé and your Editor abandoned Pine in favor of the older Elm during work together on January 17th. But did this solve the problem? John Coemans at ULB in Brussels reportedly received empty messages twice --- again, for cases where an external file had been included using **Ctrl-R**. However, as Agora's usefulness for sending files became questionable, an alternative appeared (next paragraph).

BPA connection to Internet is finally available. Your Editor first learned of this important development from BPA's James Hall on January 19th, who said that BPA's internal mail system has been connected to Internet. This is described in a 3-page internal BPA memorandum from Adrianne Martin, Route EL, dated December 1st. The news begins as follows: *"Mail can be sent to any place that has an X.400 E-mail Gateway. This includes such services as: MCI mail; ATT mail; Compuserve mail; Internet mail."* Upon seeing this, Dr. Tsu-huei Liu said our local advisor, Laura Young, earlier had mentioned the X.400 connection along with a qualification that FTP is not yet available. This is part of the problem: BPA is not yet directly connected to Internet. Instead, messages to Internet addresses are to be sent through AT&T mail using a burdensome addressing scheme that was illustrated for others in list server mail dated January 26th. Nonetheless, the new BPA connection is important because the use is so convenient. There is no need for a modem (telephone connection), and the windowed environment of MS Mail that runs under MS Windows is very convenient. Using any one of several editing tools (Notepad, Write, or Word for Windows), external files can be pasted into the text of a mail message that is being created.

The Agora menu offers an FTP entry, but this is restricted to Anonymous FTP use. A way to avoid this unfortunate limitation was discovered only on January 17th. During the first week of November, your Editor had attempted to send some files to Harald Wehrend's workstation at the University of Hannover in Germany. But, as announced using the Fargo list server on December 15th, the attempt to connect (FTP OPEN) failed. Mr. Wehrend or his E-mail experts recognized Agora's improper attempt to use Anonymous FTP, but they offered no cure. It was Laurent Dubé who had the inspiration to avoid the Agora menu by attempting to run FTP directly from an Agora Unix prompt. This worked!

Ordinary (secure) FTP transfers between OSU and Agora were successfully demonstrated January 17th during Laurent Dubé's work at BPA. This is important progress. Two weeks earlier, Mr. Dubé reported trying similar experiments with Prof. Bruce Mork's account on the Fargo computer, and had found it impossible to push (**put**) files directly from either OSU or Agora. He was able to log onto Fargo using Telnet, and then pull the files, however. But this involved security concerns: he had to reveal his own password (OSU or Agora) to the remote site (Fargo). There was no such trouble between OSU and Agora on January 17th. First, from the Unix prompt of OSU, Mr. Dubé sent **ftp**. Then, from the **ftp>** prompt, he sent the command **open agora.rain.com**. The response was as follows:

Connected to agora.rain.com.

220 agora.rain.com FTP server (Ver ...

Name (agora.rain.com:dube):

Mr. Dubé then responded to this prompt by sending **atp** which Agora acknowledged with a final prompt for a password:

331 Password required for atm.

Password:

Once this secret information had been sent, the connection was confirmed:

230 User atm logged in.

Subsequent **get startup** and **put dum.lis** transferred small ASCII files from and to Agora. Finally, **binary** was sent to cancel the default switch **ascii**, and a zipped text file was pulled to OSU by **get course01.zip**. It worked fine, and transfer rates are reasonable considering noontime overload. For STARTUP, the transmission commentary ended with a summary: **"3144 bytes received in 3.3 seconds (0.93 Kbytes/s)."** Finally, **quit** ended the connection to Agora. The following day, your Editor performed comparable operations from Agora, to **put over51ai.zip** (and also .exe and .ftn) on OSU. But .EXE and .ZIP files were received broken for some yet-unknown reason.

Lack of an intelligent mail-interfacing program to be run on our PC_s is a continuing problem of Agora. CompuServe provided CIM, recall. Unfortunately, no such comparable interface is known for the Unix computer of Agora. This is not to criticize Pine or Elm, which is the E-mail handlers of Agora. The problem is, these programs are remote rather than local. We want to compose outgoing mail on a PC before connecting to Agora. Also, we want to store incoming mail on the PC before it is read. For ordinary E-mail, we want to connect to Agora just long enough (typically half a minute or less) for the batch-mode exchange of all outgoing and incoming messages. This is what the **Send/Receive All Mail** entry of the CIM Mail menu did for us. Does any reader know of such an MS-DOS program that will interface with either Pine or Elm on a Unix computer?

The **<CR><LF>** that terminates each line of an MS-DOS file was an initial annoyance of manual use. If a standard MS-DOS file is uploaded to Agora, there will be trouble. The Unix computer expects only a **<LF>** to

terminate each line, and the extra <CR> sometimes (in some software) shows up as an unsightly ^M at the end of each line. At first, Laurent Dubé chose to remove all <CR> characters on the DOS computer before sending the message. But this involves an extra step, and creates a second file, which is messy. Robert Sarfi of the University of Waterloo in Ontario, Canada, suggested an alternative: perform a global substitution using the Unix editor **vi** on the Agora computer. Yet, your Editor never adopted either procedure. He decided to leave the ^M unmodified until subscribers of the Fargo list server complained (which never happened). The typical DOS subscriber may never have seen any ^M since by the time the message is back on a DOS computer, it no longer is abnormal. The same is true for VAX/VMS.

News about Laurent Dubé's MODELS

Salford EMTP object files were mailed from BPA to Massimo Ceraolo at the University of Pisa in Italy on December 3rd. This was at the request of MODELS author Laurent Dubé. The use of MODELS in Pisa is creative, and those in Pisa have contributed creative suggestions in the past.

Mr. Dubé has sped his Boca modem from the 9600 baud mentioned in the previous issue to 14.4 Kbaud. This connects with PROCOMM PLUS running at 57.6 Kbaud as recommended by Boca. So, if compression works well, this (57.6) is the actual transmission speed, then.

Mr. Dubé will be available near Toronto, Canada, on February 19th (see separate story about free workshop).

DCG / EPRI EMTP Information

Electrotek Concepts, Inc. of Knoxville, Tennessee, seems to be EPRI's agent for marketing the DCG / EPRI version of EMTP for MS Windows. A single, undated sheet of advertising for this product was received November 20th from Timothy Tibbals of Schweitzer Engineering Laboratories in Pullman, Washington. Recall that this was requested in the preceding issue (see the middle of column 1 on page 11). To summarize the table showing Electrotek prices of "EMTP-PC Windows Workstation," the boldface summary lines are:

EPRI/DCG EMTP for Windows ...	\$10,000
EMTP Documentation	2,000
User's Group Membership (1st year)	1,800
Shipping and Handling	750
Total Package Price	\$14,550

There is no explanation of whether this price is per computer or per site. The text immediately before the table mere says: "Purchase of the Windows package includes:" What a deal! Write out that check **now**,

before the price rises (remember the original EPRI proposed pricing of \$50K or more in the late '80s)!

Of course, those higher prices were proposed in 1984 or 1985, when DCG and EPRI still were operating under the arrogant delusion that they could monopolize EMTP (including the sale, as a sole source, of BPA work on EMTP). Recall that it took FOIA (the federal Freedom of Information Act) in the Spring of 1985 to demonstrate the folly of this belief. Then, for years afterward, the user group searched for the first customer who would admit to having purchased EMTP from EPRI at those absurd, advertised prices. Ten years later, the total still adds to zero. But what about Electrotek, which would seem to represent a new sales channel? Could any reader identify any customer who has paid Electrotek \$10K or more for EMTP? If so, please inform the user group.

Transients is the quarterly newsletter published by Electrotek Concepts and EPRI for users of the DCG/EPRI EMTP. The October issue is noteworthy by its lack of substance --- even by DCG / EPRI standards. It is worth summarizing this for the record. After removing the 4-page combined survey and resubscription form, one is left with just 4 pages. But half of the back page serves as an envelope for mailing, and half of page 3 has been left blank. Additional space is taken by pictures of Editor Grebe, Dr. Adapa, and one Dr. Tang who is advertized in a way that might be typical for consulting companies: "*If a problem can be solved or simulated using EMTP, Dr. Le Tang is the one who can do it.*" This is very useful information. Previously, your Editor would recommend different experts for different types of problems. But now we have one single person who seems to know everything! This is amazing (sarcasm)! Testimonials are interesting, too. The one from Sharma Kolluri of Entergy Services, Inc. summarizes how "*the Electromagnetic Transients Program*" has been used. But there is no explanation of **whose** program was used. Maybe a 10-year-old, public-domain version of EMTP from BPA? Or even ATP (yes, Entergy in New Orleans has been a user)? Lack of precision in EPRI writing never ceases to amaze your Editor. As for any possible new development of the EPRI-licensed version of EMTP, your Editor sees nothing. Has modeling changed in recent years? Or, instead, has DCG / EPRI money for new development been exhausted? One question of the survey seems to offer a clue: "*Would you consider providing funding for future developments to EMTP and EMTPOUT?*" Clearly, that project that 6 or 7 years ago was said to involve \$8 million worth of work has fallen on hard times! Finally, it would appear that abandonment of the original, inferior OS / 2 environment has not yet occurred: "*The OS / 2 workstation is presently being upgraded to version 2.1 of the operating system.*" Is this the normal EPRI snail's pace for such work? Or, has EPRI refused to pay for the work, having already been burned badly by performance of the original code from

Electrotek for OS / 2 ? Newsletter Editor Grebe could provide a lot of clarification in the blank spaces of his four pages, but that does not seem to be the DCG / EPRI way (Prof. Willis Long in Madison seemed to be no more informative during earlier years at a comparable job).

Miscellaneous NYPA Changes to UTPF

The SAVE declaration is never used for a local variable that is initialized by a DATA statement. This has been standard UTPF practice dating to the late '70s when Stonewall Jackson McMurray of Ebasco Services in New York City was advising your Editor about such matters from his perspective as a user of Burroughs computers, which were different in many ways. Well, in E-mail dated December 1st, Robert A. Schultz of NYPA stated the belief that SAVE is not implied in the ANSI FORTRAN 77 standard for DATA variables. Maybe not, but the practice has been proven using more than 30 distinct compilers. It is inconceivable that compiler writers would change now. So, this is one NYPA change that your Editor declines to implement until a practical need can be demonstrated.

Any \$LISTOFF that might remain in effect at the end of the phasor solution is cancelled at that point. This almost always is what the user intended. Following this NYPA reform at the top of OVER11, the request for output of the phasor solution (variable KSSOUT of the integer miscellaneous data card), takes precedence over \$LISTOFF. Previously, that was not the case.

Within DATIME, DATE and TIME were replaced by the newer DATE@ and TIME@. Apparently the performance is identical. Since other library routines end with "@", and a newer compiler manual shows these, the change seems to be appropriate.

The modal propagation table of LINE CONSTANTS now includes the actual velocity in addition to the velocity for a lossless line. The former is new, whereas the latter always has been shown. To illustrate the new output, consider the heading and first row of the table produced by DC-59. The following shows the first and the final three columns:

Mode	Lossless and actual velocity in [miles/sec]	Attenuation nepers / mile
1	1.200190E+05	1.168528E+05

Actually, Messrs. Schultz and/or Meredith had added to LCMODE one line of output for each mode, and the new output lines were interleaved with the rows of the modal table. This wasted vertical space, and was not pleasing to the eye. Your Editor knew some change had to be made, but could not envision it. Dr. Tsu-huei Liu was the one who noted that the old modal table would support one more column within the 131-byte limit (great idea). So, that is where the extra NYPA output was put.

Florida Short Course March 7 - 11

Prof. Dennis Carroll again will be offering his 4.5-day EMTP short course during spring break at the University of Florida : Monday, March 7th, through Friday, March 11th, 1993 .

The IEEE mailing list **will** be used this year. Prof. Carroll seems convinced that PES staff has corrected its unorthodox service of two years ago (see column 2 on page 9 of the January, 1992, issue). Yet, mailing to IEEE PES addresses will be confined to the Southeast.

The graphic preprocessor (EMTP data assembler) ATPDRAW should be in use and available at the short course. Recall that this BPA-supported software still is being developed at EFI in Trondheim, Norway, by author H. K. Hoidalen (see pages 18-20 of the July newsletter).

Minor changes to the 4th subcase of DC-22 occurred on February 1st as a result of Dr. Kurt Fehrle's preparation for the course in Gainesville. In consultation with Gabor Furst, the original source of this SVC (static var control) simulation, Dr. Fehrle balanced capacitances of the 7th harmonic filter. Now, all three phases have value 22.2 (previously, the final two had values 44.5). The MVAR value that is mentioned on the preceding comment card has been dropped from 20 to 9.95 at the same time. The resulting .LIS file is completely different when viewed as characters (FC usage), but differences seem small in an engineering sense.

Mohan Course: San Francisco, July 23-24

Prof. Ned Mohan of the University of Minnesota will be giving his portable EMTP short course immediately prior to the 1994 IEEE PES Summer Meeting in San Francisco. As with offerings of the past three years, only instructors will be using computers, which will use color projection for students to watch. As successfully tried in Vancouver during July of 1993, the course will be restricted to power electronics: "*Modeling Power Electronics in Power Systems using EMTP*" is the title. The price is \$690, and the subtitle of advertising has been changed to reflect the BPA-supported data assembler by Hoidalen of EFI in Trondheim: "*ATP PC Version using ATPDRAW graphical preprocessor for schematic capture.*"

The course is scheduled to begin at 8:00 Saturday morning, July 23rd, and end around 17:00 the following day (Sunday) at the Radisson Hotel near the San Francisco airport. The course brochure shows good news about cost: "*A block of rooms has been reserved for the course at the rate of \$70 plus tax*" Since most persons attending the Summer Meeting will pass through the airport, some may want to come a day or two early and stay at the

Radisson Airport before heading into the city for the meeting. This is the reason for weekend scheduling.

Free Waterloo ATP Workshop Feb. 19

The University of Waterloo (UW), which is located some 80 Km west of Toronto, Ontario, Canada, will be giving a one-day ATP workshop on February 19, 1994. This was first announced in a January 11th letter of invitation from Professors R. S. Ramshaw and M. M. A. Salama of the Department of Electrical and Computer Engineering. That one-page letter, plus 3 pages of attached printed matter, were mailed to some 69 Ontario and Quebec addresses that already were known to have ATP interest. Graduate student Robert Sarfi announced the meeting to others in the world using E-mail of the Fargo list server on January 28th. Any registered ATP user is invited, and the workshop is free. See the next issue for more about this precedent-setting initiative to which, on February 1st, the user group decided to pay Laurent Dubé's way as its representative. So, expect a Dubé-written summary of the event in the next issue.

Real - Time EMTP - like Simulation ?

"The Digital TNA has Arrived!" proclaims the bold headline of the July issue of *Centre Journal* from the Manitoba HVDC Research Centre in Winnipeg, Canada. Well, maybe. If not the digital TNA itself, at least its **salesmen**, have arrived! As your Editor views the issue, the problem with the digital TNA is similar to the problem with the analog TNA, or even just a conventional digital super-computer: due to the high cost of initial purchase and continuous maintenance, it may well be appropriate for only a handful of sites around the world. Yes, for simplified networks one might be able to simulate in real time. But for the average utility engineer having interest, this would only be after having made a reservation weeks or months in advance, purchased airline tickets, and made hotel reservations. In an age of interactive computing, when a \$1000 home computer will support much slower but more accurate simulation using EMTP, the idea of less-accurate, real-time solution at an expensive, remote site probably will not be attractive to many. Only those with a pressing need for real-time solutions should be interested, it would appear.

Yet, relay engineers are such people! They represent a very important minority that should not be discounted or forgotten. These are the same people who in recent years have built expensive test beds to drive actual relays using EMTP-generated signals. Recall that BPA's pioneering effort was cited erroneously nearly three years ago in a McGraw-Hill publication (see column 1 on page 8 of the July, 1991, newsletter). Well, such test beds are operated

off-line, in the batch-mode (using signals that have been prepared ahead of time), and this is a handicap. If it were possible for EMTP to simulate in real time, it could be interfaced with the test bed in real time, and testing could be done interactively. The concept is appealing, even if today's digital computers lack the horsepower for realistic, detailed modeling of large networks. Remember that figure of 1812 for the 33-MHz *"486 speed ratio to real time."* This was for the big hvdc simulation of BPA's Daniel Goldsworthy as described in column 1 on page 13 of the July, 1992, newsletter. Of course, the average data case would be smaller, there would be less switching (ac as opposed to dc), larger time steps might be tolerable, computers now are faster (or one could use a supercomputer), etc. Still, three orders of magnitude represent an enormous shortfall. Do not hold your breath waiting for detailed, conventional EMTP simulation in real time.

But Dr. Jose Marti of JMARTI SETUP fame has proposed an innovative, limited alternative to general EMTP simulation in a paper that he presented at the 1993 IEEE PES Summer Meeting in Vancouver. Anyone who has interest in speedy simulation is advised to study the concept carefully. Unfortunately, a thorough discussion would require more space than can be spared in this issue.

Turbo Table Dumping from NYPA

Robert A. Schultz of New York Power Authority (NYPA) in White Plains is the inventor of new code that has speed the dumping or restoring of tables as required for START AGAIN (or the preceding, related data case having MEMSAV = 1), Monte Carlo (STATISTICS) studies, and SYSTEMATIC simulations. The present writing provides a conclusion to the story that began on page 15 of the preceding (October, 1993) issue.

The size of compressed tables randomly varies a little as dimensions are changed. This will be illustrated using the DC-32 and DC-49 matched pair. Varying the multiplier of default dimensioning within LISTSIZE.DAT from 0.8 through 5.5, the following disk file sizes in Kbytes were observed using the DOS DIR command:

0.8	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5
40	42	36	36	37	36	37	37	37	32	38

All solutions were the same, as verified by using FC on the DC49.LIS files that resulted. Note the slight increase on the left, as dimensioning approaches zero. This would seem to be the result of minimum run length (Mr. Schultz ignores undefined intervals that are shorter than a certain minimum). For a larger data case, the variation would be relatively smaller. The 10 Kbytes of total variation illustrated above is small compared with the size of default dimensioning, which is about 256 Kbytes.

The efficiency of Mr. Schultz's compression is

summarized by one line of output upon completion of the operation. If output goes to disk (common for production studies), this new line will be seen on the screen, and can also be found in the DEBUG.LIS diagnostic file. Using DC-16 as an illustration, this line is:

End turbo TAPSAV dump. KNT = 1

99.19% compress 1669 words 0.33 sec

Variable KNT is the energization number, and the 1669 words are 0.81 % (100 - 99.19) of total size LTLABL = 205763 (this corresponds to the usual 3 times default dimensioning). Finally, Mr. Schultz's code did its work in one third of a second on your Editor's AT&T 486/33. Subsequent restoration is even faster, as the following first such output shows:

End turbo TAPSAV table restoration.

KNT = 2 Seconds = 0.11

Restoration for the third and later energizations is faster still (0.05 seconds) --- presumably because all pages by then have become firmly held in the disk cache. It can only be concluded that Mr. Schultz's code runs like greased lightning, and the name *turbo* is well-deserved.

Compression of each of the 108 individual COMMON blocks is being sent to DEBUG.LIS as a table of which the following is an example from DC-16 :

	Words	Bytes	% Full	Tot req	Tot wrt
1.	Compress D.	1500	16	0.13	12000 16
2.	Compress D.	7560	16	0.03	72480 32
< < Etc. (105 rows are missing) > >					
108.	Compress I.	300	40	3.33	823052 5560

Explaining this quickly, the "D" of the first two rows indicates double precision (REAL*8) whereas the "I" of the last one indicates INTEGER*4. Similarly, "C" is used for CHARACTER variables. The first column, labeled *Words*, is the number of cells to be dumped. The second column, *Bytes*, shows the actual number to be dumped (16 = 0.13% of 1500 * 8). The final two columns are cumulative totals, so they end at four times LTLABL and 1669, respectively. Later, this output probably will be put under ordinary diagnostic printout control; but for now, it always is produced.

EQUIVALENCE-d INTEGER and REAL vectors seem to represent the only known precarious aspect of Turbo TAPSAV. Because word structure and operating system protections depend on the computer being used, general advice is difficult. So, let's consider Salford EMTP for MS-DOS computers, which are of interest to most users. During table dumping, Salford DBOS might open its error window and issue a complaint such as the following: **Error: Invalid floating Point number.**

In TAPSAV_DWRITE at line 181.

This particular one was produced by simulation of DC-2 to which MEMSAV = 1 was added. That line 181 was traced to TACS storage SPTACS, of which the very first cell was corrupted. Why? Because the first 4 bytes are in fact one integer, and the second 4 are another! Together, they do not constitute a legitimate REAL*8

number. In E-mail dated December 21st, Mr. Schultz explained: *"The SPTACS utilization in Turbo TAPSAV required integer, rather than REAL*8 handling ... The floating point number representations include validation fields to account for NaN (not a number), denormalized, underflow, etc. conditions. When an 'equivalenced' integer is assigned an integer value, all the rules are broken governing the floating point protection mechanisms ... This is a fatal problem, which could be more properly solved by eliminating the equivalences in the first place. The work-around solution of Turbo TAPSAV was to treat the N REAL*8 array as a 2*N INTEGER*4 array. An INTEGER*4 datum can assume any bit pattern."*

As received from Mr. Schultz, only SPTACS was being treated as a double-length INTEGER to avoid the just-explained problem. This was originally introduced for TACS control system modeling. In fact, many (some 40 or so) vectors share this COMMON space. But other modeling of the program uses similar sharing, so might be expected to have the same trouble if treated as double-length REAL. This has been confirmed experimentally for two more COMMON blocks: 1) vector SPUM that stores all data of Prof. Hian Lauw's Universal Machine (U.M.); and 2) vector RRTACS that stores all non-CHARACTER data of Laurent Dubé's MODELS. The need was demonstrated using DC-35 for the U.M., and the second subcase of DC-33 for MODELS. So, these three vectors --- SPTACS, SPUM, and RRTACS --- now are dumped as double-length INTEGER vectors.

That ends the general, pervasive, fundamental problems with EQUIVALENCE for compressed tables. There also are exceptional, isolated EQUIVALENCEs, and some of these might require manual treatment. So far, four have been noted in Portland. First, the table dumping at the end of DC-32 failed. Why this executed correctly in White Plains is a mystery. In any case, the correction was simple enough (a change made near the end of OVER3 because of SUBR3 use of REAL*8 CRIT as INTEGER*4 ICRIT). Similarly, an attempt to dump tables of the second subcase of DC-22 failed until a correction was placed near the bottom of SOLVNL (use of REAL*8 GSLOPE for INTEGER*4 KINDEP). The second subcase of DC-7 pointed to the need for a second correction in SOLVNL (use of REAL*8 CCHAR for INTEGER*4 KSING). Finally, table dumping for DC-47 failed until corrections were made to both SMDATA and SOLVSM to eliminate local vector MASSEX that had been equivalenced to Type-59 S.M. vector HISTQ. Will there be other such discoveries later, by others? Any Salford EMTP user who might experience trouble with table dumping is advised to contact developers in Portland immediately. All standard test cases involving simulation have had tables dumped successfully, so it is hoped the average user will not experience any trouble.

JTURBO is a new parameter of STARTUP that allows the user either to accept (value unity) or reject (value zero) turbo table dumping. Of course, unless there is some exceptional trouble, use is recommended. The ordered pair of test cases DC-32 and DC-49 show why. Using the usual three times default dimensioning, disk file size, the time to dump LACOM, and the time to restore it, are as follows:

JBURBO	Size of file	32 dump	49 load
0	830890 bytes	6.5 sec	2.2 sec
1	25306 bytes	1.0 sec	0.2 sec

As usual, the experiment was performed using your Editor's AT&T 486/33 --- in recent months with 2 Mbytes of MS-DOS SMARTDrive disk caching rather than earlier use of Salford DBOS /DISK_CACHE .

Programming practices of both TACS (old control system modeling) and the Universal Machine (U.M.) are fully compatible with data compression. This will be illustrated for the U.M. First, no significant increase in compressed table sizes was observed as the multiple of default dimensioning was varied between unity and five. That was for the original data. Then the four numbers of the ABSOLUTE U.M. DIMENSIONS (AUMD) card were doubled, and the experiment was repeated. Again, no significant change in the compressed table size. The disk file sizes in Kbytes are:

Multiple of default :	1.0	2.0	3.0	4.0	5.0
Original DCN1 :	31	27	27	28	28
Doubled AUMD :	---	26	26	27	28

Unfortunately, MODELS (new control system modeling) was not as well conditioned initially when the associated compression was measured on December 22nd. Disk file size was observed to increase linearly with dimensioning as follows:

Multiple of default :	1.0	2.0	3.0	4.0	5.0
-----------------------	-----	-----	-----	-----	-----

DC33B (MODELS) : 56 76 101 127 152

When questioned about this, MODELS author Laurent Dubé explained that he initialized all available memory at the start in order to protect against the use of undefined storage. Of course, that was fine before Mr. Schultz's compression -- when we had lower expectations! But it no longer is acceptable. So, Mr. Dubé did some thinking, and decided that his error checking could be made more sophisticated. Following the adoption of many corrections that were delivered to BPA on January 17th, a repeat of the test revealed exemplary performance:

DC33B (MODELS) : 28 23 22 22 22

Undefined beginnings of COMMON blocks may occur for certain simple and/or degenerate data cases. An example is DC-33 which has no nodes of unknown voltage and hence no need to triangularize [Y]. Although the original NYPA compression logic handled the tables of DC-33 without error, it failed to compress the initial

undefined intervals of the two associated vectors YKM and KM. As a result, tables were substantially larger than necessary (e.g., for 3 times default dimensioning, 114 Kbytes were required rather than 22 Kbytes). It was found that the movement of one statement of the NYPA code (JPOINTER = J + LEN_RLE) one line downward (past the ENDIF so as to be outside the IF - THEN - ENDIF block) regained the wasted space, however.

Consistency of LISTSIZE.DAT is an important new consideration. Previously, the general rule was that tables could only be loaded back into the same program that created them. A different executable-program disk file generally would have different dimensions (determined at linking time), of course. The total size of the tables, LTLABL, was all that was checked in previous months. Now, with dynamic dimensioning, it is not the program that is the problem, since no one actually uses the sizes that are linked into the program (these limiting sizes are nearly infinite). Rather, it is the apparent sizes of program tables at the instant of usage that must be compatible. These come from LISTSIZE.DAT first, but perhaps NEW LIST SIZES later. Protection against inconsistency has been added to universal table-dumping SUBROUTINE TABLES. The user is familiar with only 30 list sizes, but VARDIM uses other dependent list sizes, too. All 100 dimensioned cells of LSTNEW are checked for consistency. If any is found to disagree, the program will automatically try to change its own sizes to match those of the tables. This decision will be noted in LUNIT6 output by the following message:

???? Current program list sizes are incompatible with those of the tables to be loaded. Attempt an automatic recovery by changing program list sizes, and restarting ...

Just as following NEW LIST SIZES (NLS), LUNIT6 output then will begin again. This assumes that the corrected list sizes yield the correct total size for all tables. If not (if more than just apparent dimensioning is incompatible), the attempt to restart will not be made. Instead, execution will be halted with a KILL = 201 error message that includes the incompatible totals. In the case of only apparent incompatibility, all 100 cells of both the program vector and the external table vector can be found on diagnostic file DEBUG.LIS . The 1018 format used for this output is compatible with LISTSIZE.DAT, so correction of this disk file is easy. Just remember to subtract 2 from list 1 (750 rather than 752 in the case of 3 times default dimensioning).

Text arrays KENFNT and JBEGOV were moved from the second partition of BLKCOM to the fourth. This eliminated the handling of just over 10 Kbytes every time tables are transferred to or from disk, and it saved the associated storage space on disk. In addition, there would seem to be improved consistency. It made no sense to transfer just text pointers without the associated text. When all was consistent, this was wasted effort: if ever

there were inconsistency, loading just pointers would not have helped matters. While studying TABLES for the turbo implementation, the idea for this unrelated reform (saving 10 Kbytes) occurred to your Editor. The reduced size of disk storage for DC-24 was confirmed January 9th when DC24AT40.BIN was observed to occupy only 13.7 Kbytes. Compare this with Robert Schultz's figures of 20 Kbytes and 28 Kbytes in column 1 on page 16 of the preceding issue.

WTURBO, RTURBO, and LTURBO are three new INCLUDE files of old SUBROUTINE TAPSAV which has undergone extensive modification to account for turbo table dumping. Unlike other INCLUDE files of ATP, these 3 new ones are not UTPF segments. Rather, they are constructed by the translator. Each has 108 lines --- one for each COMMON block of LACOM that must either be dumped (WTURBO), restored (RTURBO), or measured for the LTLABL computation (LTURBO). The exceptions for integer dumping of real vectors have been built right into the translator. It is important to remember this, in case changes are required. New structure has entered the translation process.

For ordinary START AGAIN use, a copy of program tables (e.g., DC32COPY.BIN for DC-32 and DC-49) no longer is used. Beginning January 9th, it will be seen within DC-49 that a scratch file is connected after tables have been loaded. This is progress, even if it does not extend to Monte Carlo use. That is, the matched pair of test cases DC-32 and DC-49 have been simplified whereas DC-24 and DC-40 have **not**. Do **not** try to connect a scratch file in DC-40 as was done in DC-49, since STATISTICS use will attempt to read tables from I/O channel number LUNIT2 = 22 .

One unnecessary layer of vector modularization has been eliminated in most cases, as a final step in the adoption of the NYPA code. From the beginning, EMTP has modularized simple operations such as the copying and zeroing of vectors. This involved the use of names such as COPYI, MOVE, MOVE0, MOVER, and MOVER0. These had been edited by Mr. Schultz to add, for long vectors (typically 18 or more cells), calls to his all-new vector modules that begin with VEC_\$: IINIT, DINIT, INIT, ICOPY, and DCOPY. Yet, Mr. Schultz's VEC_\$IINIT was functionally identical to the original COPYI, and comparable similarities existed for other pairings. To improve efficiency and uniformity, it was logical enough to imbed Mr. Schultz's code (Intel assembly language) directly into the original EMTP modules. All-new ICOPY_I and DGATHER had no comparable, old EMTP modules, however, so a different technique was used for these. It was found that each was only CALLED from one place, so installation-dependent INCLUDE modules (new GUTS2C and OVR11A) provided an easy way to manipulate the new code. In the process of the consolidation, the code of one

of Mr. Schultz's modules (INIT) was found to be totally unused, so of course was removed.

VAX/VMS translation was what inspired the final clean-up of the preceding paragraph. Existing translator logic only provided for 6-character names, so the removal or replacement of Mr. Schultz's long names posed more than a mechanical challenge. Not wanting to change each translator any more than necessary, the insertion of Mr. Schultz's code into EMTP modules provided the easiest avoidance. For non-Salford translations, there is still the need to destroy a few modules having Intel assembly language, but the rule is simple and effective: any subroutine having a match for the first 6 characters will be honored. So, VAX translator input VAX11ET.DAT will be seen to have a REMMOD (remove module) request for VEC_I, and this serves to destroy Salford VEC_ICOPY_I and any relatives. VAX / VMS code for this function now is supplied within GUTS2C .

Isolated Corrections from LEC

Guido Empereur, the former Manager of LEC, sent corrections to Portland via Laurent Dubé last July. The present story is a continuation of the one that spans pages 16 through 18 of the preceding issue. It provides details about two changes that were summarily dismissed three months ago due to lack of space for explanation.

Robert A. Schultz of New York Power Authority (NYPA) in White Plains was the first to explain what was wrong with LEC's published correction entitled "*Malfunctioning of K. C. Lee model.*" Recall that in column 2 on page 17 of the preceding issue, LEC's correction to OVER12 (which was explained on page 16 of the September, 1992, issue of *EMTP News*) was dismissed as follows: "*Your Editor rejects this. Robert Schultz of NYPA has better understanding (later).*" There was not space enough then, but your Editor is going to take the time, and make the space, to clarify this now.

Consider the trouble immediately below S.N. 3534 of OVER12 , which appears as follows :

3534 H2 = CIK(K)
H3 = H2 * DLTINV
I = H3

This code dates to Prof. Hermann W. Dommel, and it may be 30 years old. There was nothing wrong with it! One begins with H2 which is the modal travel time in seconds. This is divided by time-step size DELTAT to find the number of time steps that it contains. Finally, it is truncated to the integer number of time steps I (i.e., the fractional remainder is discarded). Of course, when the travel time is an exact multiple of DELTAT, there is an ambiguous situation because of numerical roundoff. If one wanted to protect against rounding downward (e.g., from 3.9999... to 3), he might consider adding a very

small number such as $FLZERO = 1.E-12$ (the default value for 64-bit computers). But why add 0.1 as LEC did? Your Editor is baffled by the LEC modification. Yes, Mr. Empereur wrote: "*results were excellent*" for the data cases he tested. The problem is, all data cases of his article were exceptional. Casual speculation by your Editor suggests that adding 0.1 will be wrong in 10% of the cases (those for which the remainder lies in the open interval between 0.9 and 1.0).

There is no indication that LEC ever understood what really was happening. The correction from Robert Schultz of NYPA forces the bad integer truncation $I = H3$ into a separate function named INT_CONVERT. Comments in this explain the reason. This "*forces correct integer conversion of 80-bit extended precision numbers which exist internally in the FPU by forcing a 64-bit conversion through this call mechanism.*" The error is not confined to Salford and/or Intel, either. In E-mail dated November 30th, Mr. Schultz explained: "*Back in 1992, we sent you fax outlining the machine language implementation of both the Salford and Apollo Fortran. Simply stated, the optimizer kept the 80 bit values on the floating point stack (it 'remembered' the impending usage) and these values were not formally rounded to 64 bits before integer truncation. That was the essence of the problem. It was not possible to fool the compiler, either. We tried every trick in the book.*" Once again, that extraordinary intellect in White Plains has triumphed. Can/Am congratulations to the NYPA experts!

Any reader who sees the following error message is asked to communicate the associated data to developers:

***** Trouble with integer truncation *****
OVER12 below 3534. K, I, IP =

The "I" of this message is Mr. Schultz's value, so this should always be right. ATP actually uses this. The IP is produced by : $IP = H3 * (1.0 + FLZERO)$. The message will be seen if and only if the two values differ. In fact, there is no trouble, since IP is not yet being used. If after a year or two of use on different computers no one has ever seen this trouble message, your Editor might be inclined to remove the function and use IP. This should have the effect of producing the same value on all computers when the travel time is a multiple of the time step. This seems to be a desirable goal.

START AGAIN use with TACS STAND ALONE is another LEC change that was not fully explained 3 months ago (see bottom of column 1 on page 17). The reason for rejection lies in the possibility of inconsistency. As conceived by your Editor a decade or two ago, the feature would reload tables and then transfer control to the time-step loop where execution was at the time tables were saved. Unfortunately, TACS STAND ALONE data cases perform their simulation within TACS2 of overlay 12 --- not within TACS3 of overlay 16 (the time-step loop). So, the use of START AGAIN for such a case

transfers control to new code that was not used for the preceding simulation. In effect, LEC changes converted the TACS STAND ALONE case to a TACS HYBRID data case. This is tricky, and anyone who looks closely will see differences (e.g., printout of the time-step loop is optimally encoded whereas that for TACS STAND ALONE is not). It is somewhat surprising that the idea works at all . LEC is to be congratulated for its innovation if not its prudence! Talking to TACS and MODELS author Laurent Dubé about the matter, it was agreed that a safer approach would be the one that has been talked about for MODELS STAND ALONE (which does not yet exist). The program could build a dummy electric network to make a hybrid problem. In the case of TACS STAND ALONE this would be needed only if MEMSAV = 1 (the request to save tables). On the other hand, there also has been thought about the possible destruction of TACS2 simulation code in order to minimize duplication, and ensure consistency. If this were to be done, then every TACS STAND ALONE data case would automatically be converted into an equivalent TACS HYBRID data case. Then the now-missing MODELS STAND ALONE would be added, and would be treated similarly.

Miscellaneous Intel PC Information

Are Japanese increasingly tolerant of IBM-standard PCs? This is important for ATP because Salford DBOS is incompatible with some Japanese computers that are produced for sale within Japan. Well, page B20 of the October 19th issue of *The Oregonian* contains a short story entitled "*Fujitsu plans to start selling IBM-compatible PCs in Japan.*" The second of three short paragraphs of this story from Tokyo reads : "*Fujitsu, Japan's largest computer maker was one of two holdouts that had refused to make IBM-compatible machines in Japan, preferring instead to sell its own systems that didn't work easily with other machines.*"

DEC , too, is hurting. IBM's economic troubles have been summarized in several previous issues because IBM's story is the biggest news. IBM is the country's largest computer maker, and it has suffered most as the industry moves from its large, once-profitable mainframe computers to single-user PCs and workstations. But the second largest American computer manufacturer, DEC, has been hit comparably by the declining fortunes of time-shared, multi-user computers. Page B1 of the October 7th issue of Portland's newspaper *The Oregonian* indicates this in a column by Mike Francis. Under the section heading entitled **Gallows Humor** can be found this joke: "*At Digital Equipment Corp., which has slashed 20 percent of its work force in the last two years, this was the riddle making the rounds of the company's computer screens: 'what's the difference between DEC and Jurassic Park?' Answer: 'One is a high-tech theme park*

full of dinosaurs; the other is a Steven Spielberg movie."

Apple Computer, too, has been a victim of the Intel-Microsoft offensive that is based on MS-DOS and MS Windows. *"Apple comes down off its high horse"* is the headline of a story on page B14 of *The Oregonian* dated November 16th. This is an important story that should worry any present user of Macintosh. *"Apple for years has charged a premium on its Macintosh personal computers But the company failed to bring prices down quickly when software made the rival design sold by IBM, Compaq, Dell and hundreds of others nearly as simple. That caused Apple's market presence to slip, and the company was driven this year to match price cuts of rivals. That spurred lower profits and the elimination of 2,500 jobs, a 15 percent cut. It also led to the departure of John Sculley, who had been Apple's CEO since 1984."* DEC hopes to save itself using Alpha. On what is Apple betting its future? Michael Spindler, new head of Apple, said at Comdex in Los Vegas that Apple's future is tied to the new RISC PowerPC chip from Motorola (see column 2 on page 17 of the July newsletter). But this will take time. *"Until the PowerPC models are available, Apple will produce a Macintosh that runs Microsoft Corp.'s Windows operating software."* Yes, Apple Computer does sound a whole lot more humble and accommodating than in years past.

"Challengers line up against Intel" is the title of an AP story by John Enders two days later on page E7 of *The Oregonian*. This explains why PowerPC is believed by some to be a serious challenger: *"Intel's Pentium has been criticized for running too hot, using too much energy and costing too much. The PowerPC is cooler, smaller and less expensive. Its makers say they can achieve greater economies of scale because more chips can be cut from 8-inch silicon wafers."* With IBM and Apple behind it, some believe PowerPC has a chance.

Miscellaneous Small Items

Randy Suhrbier, BPA's local DEC VAX / VMS expert, provided logic to mimic the Salford compiler's COMMAND_LINE routine. As a result, the Salford use of input parameters now is available for VMS users beginning November 4th. Yet, details are different and important. Whereas for Salford the arguments are used on batch file RUNTP, no such extra file is required or desired for VMS. Instead, a symbol to execute EMTP must be defined: **TP:==\$DISK22:[SCOT]EMTP**. Then program execution is trivial. Sending **TP** without any arguments provides execution as in years past. But if one wants arguments, these follow. For example, to simulate using the data of DC-3, sending the output to disk as in the preceding paragraph, send: **TP DISK DC3. *** VAX / VMS library routine LIB\$GET_FOREIGN then communicates the parameters to VAX / VMS EMTP.

User of LINE MODEL FREQUENCY SCAN (LMFS) are urged to run their data cases this way. For the illustration of DC-51, there was trouble using the former assignment of the input file to FOR003 (for command-file execution involving a variable input file name). There was some conflict involving unit LUNIT3 = 3 . Using parameters avoids any such problem.

OLDTACS is the utility that converts EMTP data from old (before the M39. update of July, 1984) TACS format to current TACS format. This was described 18 months ago in the article about hvdc (see column 2 on page 13 of the July, 1992, issue). The program works well, but has one problem: it creates a single disk file as output whereas the input may have come from many files if \$INCLUDE were used. Around the end of September, BPA's expert hvdc simulator, Daniel Goldsworthy, asked for another version that would preserve his existing data modularization. So, your Editor went to work on a new utility that has been named OLDTACS2. Testing was done on an MS-DOS PC first, using code that created .NEW output files in parallel with the input files. But since VAX / VMS as used by Mr. Goldsworthy was the real desired destination, output file naming then was changed to produce higher version numbers of the same input file names. As Randy Suhrbier, the local VMS expert observed, movement would be easy using the qualifier /SINCE=TODAY on either VMS DELETE or RENAME commands. The VMS version became available October 8th. It has been used for practical data cases by Mr. Goldsworthy beginning November 12.

Type-18 sources (ideal transformers) were limited in number to 9 by your Editor's automatic internal naming procedure that assumed 18TYP? (where the sixth and final byte here is a decimal digit) for blank BUS-X data field (the name of the internal node). In a letter dated October 7th, LEC Manager Guido Empereur provided two modifications --- one line in each of SUBR5 and OVER16 --- to extend the limit from 9 to 99. Rather than check for the first five bytes 18TYP in node names, the program now uses only the first four of these. It seems unlikely that any users will experience difficulty as a result of the change (other node names that begin with 18TY represent a conflict). So, your Editor likes the idea, whose time apparently has come (someone is using many ideal transformers). The change was made in Portland on October 17th. Yet, your Editor can not avoid wondering how long it will be before someone will propose reducing 18TY to 18T in order to expand the new limit of 99 to 999! Question : who is using large numbers of Type-18 sources, and for what?

Free-format branch cards immediately following a Marti frequency-dependent line led to a false error message prior to correction on November 23rd. The problem first was noted by Glenn Wrate, a graduate student of Prof. Bruce Mork at Michigan Tech in

Houghton. Mr. Wrate had followed his Marti line by a Type-51,52, etc. component that had been used with DCG/EPRI in years past. In fact, the cards seemed to be EMTP-generated using BCTRAN or TRELEG. This is old by ATP standards, which years ago switched to the more-practical wide alternative (\$VINTAGE, 1) of fixed-format. Well, correction was simple enough: NFRFLD = 1 was added to INDIST after the fixed-format Marti line was done using this variable.

New variable KOMENT of STARTUP is a binary flag that either accepts or discards comment cards during data input. Value zero means no change : comment cards will be seen in the LUNIT6 output. Switching this to unity means no comment cards should be seen. Prior to this addition on December 5th, value two of KOMPARE (see next paragraph) was used for this control of comments. Now, ATP has a separate control

Old variable KOMPARE of STARTUP has been given 2 new values. Previously, this was a binary switch for which value zero meant LUNIT6 output would be normal, and value unity meant that output would be simplified for easier comparisons with previous solutions of the same (or almost identical) problem. Salford EMTP for MS-DOS computers is the most widely used version, and Mike Albert's shareware utility FC is used by many to compare solutions. When KOMPARE = 1, 2 actions were taken: 1) comment cards were not serialized (so adding or subtracting a comment will not change all later comment-interpretation lines); and 2) case-summary statistics were created using the 80-column alternative to remove possible differences of dimensioning. This was good, but it was not enough. A new value two does more. First, it sets the date and time to that corresponding to the Armistice ending World War I (which began at 11:00 on 11 November 1918). Second, it zeros all component times of the case-summary statistics except the total time. This reduced the differences for all standard test cases from a file of about 250 Kbytes down to about 45 Kbytes, which saves considerable wear and tear on the eyeballs of developers who must check such files repeatedly. That total time is left so every data subcase will have at least one difference (this one line). Finally, KOMPARE = 3 will zero the total time, too, resulting in no differences at all. In this mode, the differences file is reduced to 20 Kbytes. Typical of the repetitive, minimal content is:

FC version 2.0 - Copyright (c) 1990 Mike Albert ...
Options are: -c1 -ds5 -t8

Compare files:

dc2.lis	150762	12-16-93	4:14a
dc2.sal	150609	12-11-93	9:29p

Lines in files are identical

The DEC VAX/VMS .LIS file that is produced for either DISK or BOTH usage had its attributes changed on December 16th by Randy Suhrbier of BPA. The added qualifier CARRIAGECONTROL='LIST' was added to

eliminate the carriage control character of column 1. A corresponding change to JCOLU1 of STARTUP was made, from the original unity (which caused one blank byte at the start of each line) to zero (meaning no such extra byte). So now, nearly 15 years after BPA first acquired a DEC VAX-11/780 for EMTP simulation, one no longer need be concerned about losing column 1 when the file is transferred (e.g., by Kermit from VAX/VMS to a PC) or displayed (e.g., the VAX / VMS TYPE command). Finally, VAX / VMS EMTP .LIS files behave the same way as editor files for such operations.

START AGAIN use involving TACS was slightly in error for many years. All dynamic history was not updated on the final time step prior to correction on December 31st. So, when the simulation was resurrected, one step had been partially skipped. Of course, for real hvdc simulations of BPA's Daniel Goldsworthy (see disk file HVDC.ZIP of Salford EMTP distribution), the loss had no engineering significance since the first half was merely drifting in the quasi-steady state, anyway. But when your Editor finally looked very closely as a check on the new Turbo table dumping by Robert A. Schultz of NYPA in White Plains, he was surprised to find a small error in variable FLUX of the first subcase of DC-33 when it was continued.

Monte Carlo (STATISTICS) simulations end with miniature printer plots of switch closing times T-close. For dependent switches, these plots were in error prior to January 28th when Dr Tsu-huei Liu corrected an error that first was recognized by BPA's Robert Hasibar. Random number testing is similarly affected. In fact, it was a case such as DC-50 that led Mr. Hasibar to conclude that the plots were wrong. Typically one would not be suspicious of actual simulations because the number of energizations would be too small (e.g., 100 or 200 shots) to produce smooth printer plots. But with testing only, the dice can be rolled thousands of times. A careful skeptic who wants to verify everything himself, this is what Mr. Hasibar did: 5000 rolls of the dice for his problem having dependent switches. The resulting printer plots did have smooth experimental as well as theoretical curves, but the two were clearly distinct! The mean agreed, but the standard deviation of the bell-shaped curve was distinctly smaller than for the theoretical curve. Dr. Liu studied the problem intently, and suddenly realized the error while at home during the evening of January 27th. Surprisingly, she discovered four lines of extraneous (i.e., extra, unwanted, and wrong) coding immediately preceding S.N. 1660 in DICTAB. Removal of these four lines not only cured Mr. Hasibar's real case, but also changed a little the miniature printer plots of DC-48 and DC-50 (standard test cases that involve one or more dependent switches). Nothing else changed, it should be emphasized. That is, switch closing times themselves never were in error; it was only their graphical display (those miniature printer plots) that was imperfect.