

November 24, 1969

MPL-43

To: Multics Performance Log

From: J. H. Saltzer

Subject: Comparison of Compile Time, Run Time and Size of a Small Program Using BASIC, FORTRAN, and PL/I.

A short FORTRAN subroutine which reads a number, ℓ , and then computes and prints the $(\ell + 2)$ 'nd prime was borrowed from the Multics standard certifier script, and recoded in BASIC and PL/I. The three versions were then compared on several points; the results are presented here.

I. Source Programs

The two new language versions were coded to be as similar as possible in algorithm to the original FORTRAN routine. The lack of a modulo function in BASIC was bypassed by a direct computation using the integer function; this change being in the innermost loop may have affected execution time performance described in Section II.

The three source programs are listed for comparison in Figure I.

II. Execution Time

Each program was executed several times each, with input values of 5, 20, 40, 90, and 175. The input lines to both the command interpreter and the program itself were queued. In order to distinguish page fault time from pure execution time, each command was queued two or more times, in an attempt to drive the number of page-faults to zero on second and later executions. Each experiment was repeated several times, and the smallest cpu time observed was recorded, in an attempt to minimize the effect of interrupts whose execution time is currently charged to the executing process. The results are shown in Figure III. In general, this graph suggests that the execution time of the three object programs is in the ratio 1:2:3 for FORTRAN, PL/I, and BASIC, respectively. Note that with input value 175, the required execution time of all three programs solidly swamps out the end effects of program starting and input/output statements, and even of BASIC compilation time.

The difference in execution time of FORTRAN and PL/I appeared worth further study, so the object programs were compared in detail. Figure IV and V exhibit the two object programs. In terms of the physical program

I. Comparison of source program listings:

fpm.fortran

11/24/69

```
subroutine fpm
read(5,70) 1
70 format(i3)
m = 1
do 10 i = 3,100000
k = i-1
do 20 j = 2,k
if(mod(i,j)) 20,10,20
20 continue
m = m+1
if(m-1) 10,40,40
10 continue
40 write(6,60) m,i
60 format(7h Prime ,i4,3h is,i6)
return
end
```

pm.pl1

```
pm:      procedure;
declare (m,i,k,j,l) fixed binary;
call read_list_(1);
m = 1;
do i = 3 to 100000;
k = i - 1;
do j = 2 to k;
if mod(i,j) = 0 then go to 1190;
1160:   end;
m = m+1;
if m >= 1 then go to 1200;
1190:   end;
1200:   call ioa_("prime %d is %d", m, i);
end;
```

pm.basic

```
100 input 1
110 let m = 1
120 for i = 3 to 100000
130 let k = i - 1
140 for j = 2 to k
150 if (i = int(i/j)*j) then 190
160 next j
170 let m = m + 1
180 if m >= 1 then 200
190 next i
200 print "prime";m;"is";i
210 end
```

II. Queuing of command input to obtain zero-page-fault case:

basic pm
20
basic pm
20
basic pm

Compile time in ms. = 368, Page waits = 21
11/21/69 23:07

? PRIME 20 IS 71


r 2307 1.178 67

Compile time in ms. = 193, Page waits = 0
11/21/69 23:07

? PRIME 20 IS 71

r 2307 .611 0

III. Run time of Prime Program with 0 p.f.
 11/21/69
 jhs



sizes, the reason for the difference in execution time is clear. The inner and outer program loops measure out as follows:

| | FORTRAN | PL/I |
|------------|-----------------|---------------------------------|
| Inner loop | 11 instructions | 14 inst. + 5 in subroutine = 19 |
| Outer loop | 22 instructions | 39 inst. (including subroutine |

Upon inspecting the compiled code, the difference seems to be primarily that FORTRAN is very good on register optimization, while PL/I largely ignores the subject. Compare also the FORTRAN compiled "AOS" on line 152, compared with the PL/I sequence starting at 73, both for the statement "m = m + 1". It appears that FORTRAN misses few tricks.

III. Compile Time

The compile times were compared, using a similar command stacking technique to get the compiler "in core" and minimize the effect of missing page faults. It was found to be impossible to bring the number of page faults to zero when using PL/I as the compiler apparently does not fit into the available ($\sim 160K$) core. Results were as follows, in both cases after linking had been accomplished by an earlier command.

| | Compile Time/sec. | Page Faults |
|---------|----------------------|----------------|
| FORTRAN | 1.720 | 0 |
| PL/I | 4.698 | 170 |
| BASIC | .500(est.) | 0 |

The BASIC compile time is estimated from the lower asymptote of its compile-and-execute curve in Figure III.

IV. Working Set Size

Each program was compiled following a "flush" command, and then executed (with input value 5) following a "flush" command. The experiment was repeated several times to insure that other users had not distorted the result. The following table indicates the number of missing-page faults observed.

V. Compilation of source lines 4-12 by PL/I:

11/24/69

```

000036 aa 777745 2360 04 liq -27,ic
000037 aa 6 00124 7561 09 stq spl84

000040 aa 777747 2360 04 liq -25,ic
000041 aa 6 00125 7561 09 stq spl85
000042 aa 6 00125 2361 09 liq spl85
000043 aa 777745 1160 04 cmpq -27,ic
000044 aa 000002 6000 04 tze 2,ic
000045 aa 000041 6050 04 tpl 33,ic

000046 aa 6 00125 2361 09 liq spl85
000047 aa 777734 1760 04 sbq -36,ic
000050 aa 6 00126 7561 09 stq spl86

```

STATEMENT 1 ON LINE 4
000003 = 000000000001

STATEMENT 1 ON LINE 5
000007 = 000000000003

000010 = 000000303240
000046
000106

STATEMENT 1 ON LINE 6

000003 = 000000000001

```

000051 aa 6 00126 2361 09 liq spl86
000052 aa 6 00131 7561 09 stq spl89
000053 aa 777736 2360 04 liq -34,ic
000054 aa 6 00127 7561 09 stq spl87
000055 aa 6 00127 2361 09 liq spl87
000056 aa 6 00131 1161 09 cmpq spl89
000057 aa 000002 6000 04 tze 2,ic
000050 aa 000013 6050 04 tpl 11,ic

000061 aa 6 00125 2361 09 liq spl85
000062 aa 6 00127 3521 09 eapbp spl87
000063 aa 0 00704 6701 09 tsbip apl452
000064 aa 777716 1160 04 cmpq -50,ic
000065 aa 000002 6010 04 tnz 2,ic
000066 aa 000014 7100 04 tra 12,ic

```

STATEMENT 1 ON LINE 7

000011 = 000000000002

000061
000073

STATEMENT 1 ON LINE 8

mod fx1 *5 instruction subroutine*

000002 = 000000000000
000067
000102

STATEMENT 1 ON LINE 9

000003 = 000000000001

000055
STATEMENT 1 ON LINE 10

000003 = 000000000001

STATEMENT 1 ON LINE 11

000102
000106

STATEMENT 1 ON LINE 12

000003 = 000000000001

000042

000000000000

000000000000

```

000076 aa 6 00124 2361 09 liq spl84
000077 aa 6 00130 1161 09 cmpq spl88
000100 aa 000002 6040 04 tni 2,ic
000101 aa 000005 7100 04 tra 5,ic

000102 aa 6 00125 2361 09 liq spl85
000103 aa 777700 0760 04 aiq -64,ic
000104 aa 6 00125 7561 09 stq spl85
000105 aa 777735 7100 04 tra -35,ic

```

| | Compile | Run |
|----------|---------|-----|
| BASIC | 74 | |
| FORTTRAN | 105 | 15 |
| PL/I | 225 | 27 |

Note that the PL/I compiler taken only a few more page faults when it follows "flush" than when it follows itself. The unusually large number of pages required to execute the PL/I object program suggested further analysis. Figure VI shows a page trace of the PL/I program. Most of the trouble clearly arose from the use of subroutine "read_list_" for input, since that subroutine call resulted in 12 distinct page faults. It also appears that in system 4.7f, segment "p/l-operators" is not yet wired-down.

A second fact uncovered by the page trace is that segment 110, the Teletype DIM (bound-tty-active) is organized so that all six of its pages are touched by a write and read call. In earlier systems, this segment was "optimally bound" such that only three pages were touched. A review of the contents of bound-tty-active suggests that such a reordering of its components could again reduce the number of pages touched to three.

VI. Page fault trace of PL/I object program execution following a "flush" command:

19/23/69

meter_start;flush
measurement started
pm 5
5
meter_stop
r 819 2.564 293

Ready for 1 value
prime 5 is 11
r 819 .322 27

measurement stopped
r 819 .098 6

print_pages 35
18696668625 0 000124
18696643303 0 phcs_.link
18696508150 10 meter_start
18696489710 31 !BBBHwmPcGZQDDx

18696402098 5 !BBBHwmPcGZQDDx
18696359160 3 read_list_
18696328938 0 free_
18696286736 0 bin_oct
18696240114 0 read_list_
18696192060 4 read_list_
18696169407 2 read_list_
18696031631 6 write_out
18696006734 11 !BBBHwmPcGZQDDx
18695970285 2 stack_01
18695959267 1 stack_01
18695923550 3 stack_01
18695839472 1 read_list_
18695803045 1 pll_operators
18695770729 0 pll_operators

18695631200 0 pm
18695598870 32 !BBBHwmPcGZQDDx
18695567624 6 000110
18695546720 0 000110
18695516733 3 000110
18695475981 1 000053
18695432839 2 000110
18695384450 0 000112
18695345641 5 000110
18695319802 4 000110
18695288055 5 000016
18695266610 1 listen_

18695217248 1 hcs_.link
18695180876 0 000120
18695155433 0 process_info
18695135919 0 hcs_.link
r 820 3.227 63

↑
Time
in
msec.

Page #
Segment # or name if known

meter_stop command

12 pages touched by read_list_

27 pages touched

not wired yet

TTM DIM

tail end of flush command