

IBM COBOL for MVS & VM Version 1 Release 2 Performance Tuning

January 26, 1996

R. J. Arellanes

IBM Corporation
Software Solutions
555 Bailey Avenue
San Jose, CA 95141

Disclaimer

The performance considerations contained in this paper were obtained by running sample programs in a particular hardware/software configuration using a selected set of tests and are presented as illustrations. Since performance varies with configuration, sample program characteristics, and other installation and environment factors, results obtained in other operating environments may vary. We recommend that you construct sample programs representative of your workload and run your own experiments with a configuration applicable to your environment.

IBM does not represent, warrant, or guarantee that a user will achieve the same or similar results in the user's environment as the experimental results reported in this paper.

Distribution Notice

Permission is granted to distribute this paper to IBM customers. IBM retains all other rights to this paper, including the right for IBM to distribute this paper to others.

| **Third Edition, January 1996**

| This edition applies to IBM COBOL for MVS & VM Version 1 Release 2 running with IBM Language Environment for MVS & VM Version 1 Release 5, and to all subsequent releases and modifications until otherwise indicated in new editions.

| This edition replaces all previous editions of this document. All changes made in the third edition are marked with change bars as indicated to the left of this paragraph.

Contents

Introduction	1
Referenced IBM Publications	1
Background	1
Tuning the Run-Time Environment	3
Compiler Options that Affect Run-Time Performance	3
AWO or NOAWO	3
CMPR2 or NOCMPR2	3
DATA(24) or DATA(31)	4
DYNAM or NODYNAM	4
FASTSRT or NOFASTSRT	4
NUMPROC - NOPFD, MIG, or PFD	5
OPTIMIZE(STD), OPTIMIZE(FULL), or NOOPTIMIZE	5
RENT or NORENT	6
RMODE - AUTO, 24, or ANY	6
SSRANGE or NOSSRANGE	6
TEST or NOTEST	7
TRUNC - BIN, STD, or OPT	7
Run-Time Options that Affect Run-Time Performance	8
AIXBLD	8
ALL31	8
CHECK	9
DEBUG	9
RPTOPTS	9
RPTSTG	9
RTEREUS	10
STORAGE	10
TEST	11
TRAP	11
VCTRSAVE	11
COBOL and LE Features that Affect Run-Time Performance	11
Storage Management Tuning	11
Calling IGZERRE	12
Using the CEEENTRY and CEETERM Macros	13
Using Preinitialization Services (CEEPIPI)	13
Using Library Routine Retention (LRR)	14
Tailoring COBPACKs	15
Library in the LPA/ELPA (MVS) or NSS (CMS)	15
Using CALLs	15
Other Product Related Factors that Affect Run-Time Performance	16
Mixing VS COBOL II or COBOL/370 Rel 1 with COBOL for MVS & VM Rel 2	16
Mixing OS/VS COBOL with COBOL/370 Rel 1 or COBOL for MVS & VM Rel 2	16
First Program Not COBOL	17
IMS	18
CICS	19
DB2	20
DFSORT	21
Efficient COBOL Coding Techniques	22
Data Files	22
QSAM Files	22
Variable-Length Files	22
VSAM Files	23
Data Types	23

BINARY (COMP or COMP-4)	23
Data Conversions	24
DISPLAY	24
PACKED-DECIMAL (COMP-3)	24
Comparing Data Types	25
Fixed-Point vs Floating-Point	25
Indexes vs Subscripts	25
OCCURS DEPENDING ON	26
Program Design	26
Algorithms	26
Data Structures and Data Types	26
Coding Style	27
Factoring Expressions	27
Symbolic Constants	27
Subscript Checking	27
Subscript Usage	27
Searching	28
Recent Performance Improvements	29
A Performance Checklist	30
Summary	31
Appendix A. Intrinsic Function Implementation Considerations	32
Appendix B. History of Prior Performance Improvements	34
Appendix C. Coding Examples	35
Using CEEENTRY	35
Using CEELRR	36
Using IGZERRE	38
Using CEEPIPI with Call_Sub	40
Using CEEPIPI with Call_Main	42
COBOL Example - COBSUB	44

Introduction

This paper identifies some of the factors that can affect the performance of COBOL programs when using IBM COBOL for MVS & VM Version 1 Release 2 running with IBM Language Environment for MVS & VM Version 1 Release 5 (LE). It also contains a summary of the performance experiences that we have had with COBOL for MVS & VM. The tuning methods discussed here are intended to assist users in selecting from the various options which are available for compiling COBOL programs with Version 1 Release 2 of COBOL for MVS & VM and executing them with Version 1 Release 5 of LE for MVS & VM. These methods provide the COBOL programmer with the opportunity to tune the LE run-time environment to potentially improve CPU time performance and the use of system resources.

First, we will look at some compiler and run-time options that affect the run-time performance of a COBOL application. Next, we will look at some efficient COBOL coding techniques that may improve the run-time performance. Finally, we will look at a check list of items to be aware of if a performance problem is encountered with COBOL for MVS & VM.

Referenced IBM Publications

Throughout this paper, all references to OS/VS COBOL refer to OS/VS COBOL Release 2.4, all references to MVS refer to MVS/ESA, and all references to CMS refer to CMS under VM/ESA, unless otherwise indicated. Additionally, several items have page number references after them, enclosed in parentheses. The manual abbreviations are in **bold face** followed by the page numbers in that manual. The abbreviations and manuals referenced in this paper are listed in the table below:

Abbrev	IBM Publication and Number
COB PG	<i>IBM COBOL for MVS & VM Programming Guide Release 2, SC26-4767-01</i>
COB MIG	<i>IBM COBOL for MVS & VM Release 2 Compiler and Run-Time Migration Guide, GC26-4764-03</i>
LE PG	<i>IBM Language Environment for MVS & VM Programming Guide Release 5, SC26-4818-05</i>
LE REF	<i>IBM Language Environment for MVS & VM Programming Reference Release 5, SC26-3312-02</i>
LE INST	<i>IBM Language Environment for MVS & VM Installation and Customization on MVS Release 5, SC26-4817-05</i>

These manuals contain additional information regarding the topics that are discussed in this paper, and it is strongly recommended that they be used in conjunction with this paper to receive increased benefit from the information contained herein.

Background

In general, we do most of our measurements on MVS batch and CMS, but sometimes we also measure the language products on CICS and IMS. In measuring the performance of the language products on MVS batch and CMS, we use a variety of different tools that we have developed specifically for this purpose. Our tools can collect data such as CPU time used by a program (Virtual and Total CPU Time on CMS; TCB and SRB Time on MVS), elapsed time, SIO or EXCP counts, virtual storage usage, and paging activity (CMS only). When measuring on CICS and IMS, we usually enlist the help of other departments that are more familiar with transaction environments in collecting and analyzing data from RMF and SMF.

In measuring COBOL, we have three different types of benchmarks: compile-only, compile and execute, and kernels. The compile-only set contains programs which are designed to be compiled but not executed. These programs measure the compiler performance as a function of program size and range from 200 statements to more than 7,000 statements. The compile and execute set contains programs which are either subsets of

"real" application programs or are entire application programs. These programs are somewhat representative of "real" application programs. Some are actual customer programs and the rest have been written by IBM. The kernel set contains programs which are a collection of specific language statements enclosed in a loop of 1,000 to 100,000 times. Each kernel program consists of several of these loops, with each loop having a different variation of that language statement. Each loop is surrounded by a call to a timing routine so that we can measure the performance of the individual statements.

| The performance considerations reported in this paper were made on a 9121-440 running VM/ESA 2.2
| Service Level 9401 with CMS Level 11 Service Level 401 (using SET MACHINE 370 or SET MACHINE
| ESA, depending on the individual test) or on a 9021-720 running MVS/SP 5.2.2 PTF Level 9507. The
| programs used were batch-type (non-interactive) applications. **Unless otherwise indicated, all performance
| comparisons made in this paper are referencing CPU time performance and not elapsed time performance.**

Tuning the Run-Time Environment

This section focuses on some of the options that are available for tuning an application, as well as the overall LE run-time environment. This in itself may not produce high performing code since both the coding style and the data types can have a significant impact on the performance of the application. In fact, the coding style and data types usually have a far greater impact on the performance of an application than that of tuning the application via external means (e.g., compiler options, run-time options, space management tuning, setting up the COBPACKs, and placing the library routines in shared storage). First, we will look at each of these external options in a little more detail.

Compiler Options that Affect Run-Time Performance

Compiler options, by far, have been the cause of a vast majority of the performance problems reported by customers. Many customers are not aware of the performance implications that many of the compiler options have at run time, especially the AWO, CMPR2, DYNAM, FASTSRT, NUMPROC, OPTIMIZE, RENT, SSRANGE, TEST, and TRUNC options. Let's look at each of these options to see how they might affect the performance of an application program.

| (COB PG: p 247)

AWO or NOAWO

The AWO compiler option causes the APPLY WRITE-ONLY clause to be in effect for all physical sequential, variable-length, blocked files, even if the APPLY WRITE-ONLY clause is not specified in the program. With APPLY WRITE-ONLY in effect, the file buffer is written to the output device when there is not enough space in the buffer for the next record. Without APPLY WRITE-ONLY, the file buffer is written to the output device when there is not enough space in the buffer for the maximum size record. If the application has a large variation in the size of the records to be written, using APPLY WRITE-ONLY can result in a performance savings since this will generally result in fewer calls to Data Management Services to handle the I/Os.

Performance considerations using AWO:

One program using variable-length files and AWO was 10% faster than NOAWO.

| (COB PG: pp 22, 249, 428)

CMPR2 or NOCMPR2

The CMPR2 compiler option generates code that is compatible with code generated by VS COBOL II Release 2 (ANSI 1974 Standard) and does not allow the use of many of the new features available with VS COBOL II Release 3 and later, COBOL/370, or COBOL for MVS & VM (ANSI 1985 Standard). NOCMPR2 allows the full use of all supported 1985 Standard COBOL language features. Since additional features are available when using NOCMPR2 than when using CMPR2, there may be times when CMPR2 is faster than NOCMPR2.

Note: CMPR2 is provided as a migration aid to allow you to gradually convert your applications to the 1985 Standard support provided by NOCMPR2. All future enhancements to COBOL for MVS & VM will be provided only under NOCMPR2. Therefore, it is recommended that all applications be converted to NOCMPR2 as soon as possible.

Performance considerations using CMPR2:

| On the average, CMPR2 was equivalent to NOCMPR2, with a range of equivalent to 2% slower.

| (COB PG: p 250)

| **DATA(24) or DATA(31)**

| Using DATA(31) with your RENT program will help to relieve some below the line virtual storage constraint problems. When you use DATA(31) with your RENT programs, most QSAM file buffers can be allocated above the 16MB line. When you use DATA(31) with HEAP(,ANYWHERE), most of the WORKING-STORAGE and the FD record areas can be allocated above the 16MB line.

| With DATA(24), the WORKING-STORAGE and FD record areas will be allocated below the 16 MB line.

| **Note:** For NORENT programs, see the RMODE option to determine where WORKING-STORAGE will be allocated.

| Performance data using DATA is not available at this time.

| (COB PG: pp 130, 252-253, 428)

| **DYNAM or NODYNAM**

The DYNAM compiler option specifies that all subprograms invoked through the CALL literal statement will be loaded dynamically at run time. This allows you to share common subprograms among several different applications, allowing for easier maintenance of these subprograms since the application will not have to be re-link-edited if the subprogram is changed. DYNAM also allows you to control the use of virtual storage by giving you the ability to use a CANCEL statement to free the virtual storage used by a subprogram when the subprogram is no longer needed. However, when using the DYNAM option, you pay a performance penalty since the call must go through a library routine, whereas with the NODYNAM option, the call goes directly to the subprogram. Hence, the path length is longer with DYNAM than with NODYNAM.

Performance considerations using DYNAM with CALL literal (measuring CALL overhead only):

| On the average, for a CALL intensive application, the overhead associated with the CALL using DYNAM ranged from 80% slower to 350% slower than NODYNAM.

| **Note:** This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms is not degraded this much.

| (COB PG: pp 255, 343-347, 352-353, 428-429)

| **FASTSRT or NOFASTSRT**

For eligible sorts, the FASTSRT compiler option specifies that the SORT product will handle all of the I/O and that COBOL does not need to do it. This eliminates all of the overhead of returning control to COBOL after each record is read in or after processing each record that COBOL returns to sort. The use of FASTSRT is recommended when direct access devices are used for the sort work files since the compiler will then determine which sorts are eligible for this option and generate the proper code. If the sort is not eligible for this option, the compiler will still generate the same code as if the NOFASTSRT option were in effect. A list of requirements for using the FASTSRT option is in the COBOL programming guide.

Performance data using FASTSRT is not available at this time.

| (COB PG: pp 173-176, 256, 429)

NUMPROC - NOPFD, MIG, or PFD

Using the NUMPROC(PFD) compiler option generates significantly more efficient code for numeric comparisons. It also avoids the generation of extra code that NUMPROC(NOPFD) or NUMPROC(MIG) generates for most references to COMP-3 and DISPLAY numeric data items to ensure a correct sign is being used. With NUMPROC(NOPFD), sign fix-up processing is done for all references to these numeric data items. With NUMPROC(MIG), sign fix-up processing is done only for receiving fields (and not for sending fields) of arithmetic and MOVE statements. With NUMPROC(PFD), the compiler assumes that the data has the correct sign and bypasses this sign fix-up processing. NUMPROC(MIG) generates code that is similar to that of OS/VSE COBOL. Using NUMPROC(NOPFD) or NUMPROC(MIG) may also inhibit some other types of optimization. However, not all external data files contain the proper sign for COMP-3 or DISPLAY signed numeric data, and hence, using NUMPROC(PFD) may not be applicable for all application programs. For performance sensitive applications, NUMPROC(PFD) is recommended when possible.

Performance considerations using NUMPROC:

On the average, NUMPROC(PFD) was 1% faster than NUMPROC(NOPFD), with a range of 12% faster to equivalent.

On the average, NUMPROC(PFD) was 1% faster than NUMPROC(MIG), with a range of 8% faster to equivalent.

On the average, NUMPROC(MIG) was equivalent to NUMPROC(NOPFD), with a range of 11% faster to equivalent.

| (COB PG: pp 65-66, 266-267, 429-430)

OPTIMIZE(STD), OPTIMIZE(FULL), or NOOPTIMIZE

To assist in the optimization of the code, you should use the OPTIMIZE compiler option. With the OPTIMIZE(STD) or OPTIMIZE(FULL) options in effect, you may receive optimizations that include:

- eliminating unnecessary branches
- simplifying inefficient branches
- simplifying the code for the out-of-line PERFORM statement, moving the performed paragraphs in-line, where possible
- simplifying the code for a CALL to a contained (nested) program, moving the called statements in-line, where possible
- eliminating duplicate computations
- eliminating constant computations
- aggregating moves of contiguous, equal-sized items into a single move
- deleting unreachable code

| Additionally, with the OPTIMIZE(FULL) option in effect, you may also receive these optimizations:

- deleting unreferenced data items and the associated code to initialize their VALUE clauses

Many of these optimizations are not available with OS/VSE COBOL, but are available with COBOL for MVS & VM. NOOPTIMIZE is generally used while a program is being developed when frequent compiles are necessary. NOOPTIMIZE also makes it easier to debug a program since code is not moved; NOOPTIMIZE is required when using the TEST compiler option with a value other than TEST(NONE). OPTIMIZE requires more CPU time for compiles than NOOPTIMIZE, but generally produces more efficient run-time code. For production runs, OPTIMIZE is recommended.

Performance considerations using OPTIMIZE:

On the average, OPTIMIZE(STD) was 4% faster than NOOPTIMIZE, with a range of 17% faster to equivalent.

On the average, OPTIMIZE(FULL) was equivalent to OPTIMIZE(STD).

One RENT program calling a RENT subprogram with 500 unreferenced data items with VALUE clauses was 3% faster with OPTIMIZE(STD) and 5% faster with OPTIMIZE(FULL).

The same RENT program calling a RENT subprogram with 500 unreferenced data items with VALUE clauses using PROGRAM IS INITIAL was 6% faster with OPTIMIZE(STD) and 80% faster with OPTIMIZE(FULL).

Note: The two RENT program tests measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(COB PG: pp 268-269, 425-427, 430)

RENT or NORENT

Using the RENT compiler option causes the compiler to generate some additional code to ensure that the program is reentrant. Reentrant programs can be placed in shared storage like the Link Pack Area (LPA) or the Extended Link Pack Area (ELPA) on MVS or the Named Save Segment (NSS) on CMS. Also, the RENT option will allow the program to run above the 16 MB line on MVS or CMS. Producing reentrant code may increase the execution time path length slightly.

Note: The RMODE(ANY) option can be used to run NORENT programs above the 16 MB line.

Performance considerations using RENT:

On the average, RENT was equivalent to NORENT, with a range of equivalent to 1% slower.

(COB PG: pp 272, 369, 430)

RMODE - AUTO, 24, or ANY

The RMODE compiler option determines the RMODE setting for the COBOL program. When using RMODE(AUTO), the RMODE setting depends on the use of RENT or NORENT. For RENT, the program will have RMODE ANY. For NORENT, the program will have RMODE 24. When using RMODE(24), the program will always have RMODE 24. When using RMODE(ANY), the program will always have RMODE ANY.

Note: When using NORENT, the RMODE option controls where the WORKING-STORAGE will reside. With RMODE(24), the WORKING-STORAGE will be below the 16 MB line. With RMODE(ANY), the WORKING-STORAGE can be above the 16 MB line.

Performance data using RMODE is not available at this time.

(COB PG: pp 273)

SSRANGE or NOSSRANGE

Using SSRANGE generates additional code to verify that all subscripts, indexes, and reference modification expressions are within the proper bounds. This in-line code occurs at every reference to a subscripted or variable-length data item, as well as at every reference modification expression, and it can result in some degradation at run time. In general, if you need to verify the subscripts only a few times in the application instead of at every reference, coding your own checks may be faster than using the SSRANGE option. For performance sensitive applications, NOSSRANGE is recommended.

Performance considerations using SSRANGE with CHECK(ON):

On the average, SSRANGE was 4% slower than NOSSRANGE, with a range of equivalent to 31% slower.

(COB PG: pp 275-276, 389, 430-431)

TEST or NOTEST

The TEST compiler option generates additional code so that the program can be run under Debug Tool, the debugging component of AD/Cycle CODE/370. It also allows you to request that symbolic variables be included in the formatted dump produced by LE.

Using TEST with a value other than NONE can cause a significant performance degradation when used in a production environment since this additional code occurs at *each* COBOL statement. The TEST option with a value other than NONE should be used only when debugging an application. Additionally, when TEST is used this way, the OPTIMIZE option is disabled. For production runs, NOTEST or TEST(NONE) is recommended.

Performance considerations using TEST:

On the average, TEST(ALL,SYM) was 14% slower than NOTEST, with a range of equivalent to 59% slower.

On the average, TEST(NONE,SYM) was equivalent to NOTEST.

(COB PG: pp 276-277, 431)

TRUNC - BIN, STD, or OPT

When using the TRUNC(BIN) compiler option, all binary (COMP) sending fields are treated as either halfword, fullword, or doubleword values, depending on the PICTURE clause, and code is generated to truncate all binary receiving fields to the corresponding halfword, fullword, or doubleword boundary (base 2 truncation). The full content of the field is significant. This can add a significant amount of degradation since typically some data conversions must be done, which may require the use of some library subroutines. BIN is usually the slowest of the three sub options for TRUNC.

When using the TRUNC(STD) compiler option, the final intermediate result of an arithmetic expression, or the sending field in the MOVE statement, is truncated to the number of digits in the PICTURE clause of the binary (COMP) receiving field (base 10 truncation). This can add a significant amount of degradation since typically the number is divided by some power of ten (depending on the number of digits in the PICTURE clause) and the remainder is used; a divide instruction is one of the more expensive instructions. TRUNC(STD) behaves the same as TRUNC in OS/VSE COBOL.

However, with TRUNC(OPT), the compiler assumes that the data conforms to the PICTURE and USAGE specifications and manipulates the result based on the size of the field in storage (halfword, fullword or doubleword). TRUNC(OPT) behaves the same as NOTRUNC in OS/VSE COBOL.

TRUNC(STD) conforms to the ANSI and SAA standards, whereas TRUNC(BIN) and TRUNC(OPT) do not. TRUNC(OPT) is provided as a performance tuning option and should be used only when the data in the application program conforms to the PICTURE and USAGE specifications. For performance sensitive applications, the use of TRUNC(OPT) is recommended when possible.

Performance considerations using TRUNC:

On the average, TRUNC(OPT) was 27% faster than TRUNC(BIN), with a range of 92% faster to equivalent.

On the average, TRUNC(STD) was 26% faster than TRUNC(BIN), with a range of 88% faster to equivalent.

| On the average, TRUNC(OPT) was 5% faster than TRUNC(STD), with a range of 49% faster to
| equivalent.

| (COB PG: pp 278-280, 431, 464)

Run-Time Options that Affect Run-Time Performance

Selecting the proper run-time options is another factor that affects the performance of a COBOL application. Therefore, it is important for the system programmer responsible for installing and setting up the LE environment to work with the application programmers so that the proper run-time options are set up correctly for your installation. Let's look at some of the options that can help to improve the performance of the individual application, as well as the overall LE run-time environment.

| (LE PG: pp 116-126; LE INST: pp 112-118)

AIXBLD

The AIXBLD option allows alternate indexes to be built at run time. However, this may adversely affect the run-time performance of the application. It is much more efficient to use Access Method Services to build the alternate indexes before running the COBOL application than using the NOAIXBLD run-time option. Note that AIXBLD is not supported when VSAM datasets are accessed in RLS mode.

| Performance considerations using AIXBLD:

| One VSAM program was 8% slower when using AIXBLD compared to using NOAIXBLD.

| (COB PG: p 420; LE INST: pp 154-155; LE REF: p 11)

ALL31

The ALL31 option allows LE to take advantage of knowing that there are *no* AMODE(24) routines in the application. It specifies that the entire application will run in AMODE(31). This can help to improve the performance for an all AMODE(31) application because LE can minimize the amount of mode switching across calls to common run-time library routines. ALL31(OFF) is required for OS/VS COBOL programs that are not running under CICS, VS COBOL II NORES programs, and all other AMODE(24) programs. Note that when using ALL31(OFF), you must also use STACK(,BELOW).

| Performance considerations using ALL31 (measuring CALL overhead only):

| On the average, ALL31(ON) was 1% faster than ALL31(OFF), with a range of 4% faster to equivalent.

| One program with many library routine calls was 10% faster when using ALL31(ON).

Note: This test measured only the overhead of the CALL for a RENT program (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms will have different results, depending on the number of calls that are made to LE common run-time routines.

| (LE INST: pp 155-157; LE REF: pp 11-12; COB MIG: p 42)

CHECK

The CHECK option activates the additional code generated by the SSRANGE compiler option, which requires more CPU time resources for the verification of the subscripts, indexes, and reference modification expressions. Using the CHECK(OFF) run-time option deactivates this code but still requires some additional CPU time resources at every use of a subscript, index, or reference modification expression to determine that this check is not desired during the particular run of the program. This option has an effect only on a program that has been compiled with the SSRANGE compiler option.

Performance considerations using CHECK:

| On the average, CHECK(ON) with SSRANGE was 2% slower than CHECK(OFF) with SSRANGE,
| with a range of equivalent to 20% slower.

| (LE INST: pp 162-163; LE REF: p 15)

DEBUG

The DEBUG option activates the COBOL batch debugging features specified by the USE FOR DEBUGGING declarative. This may add some additional overhead to process the debugging statements. This option has an effect only on a program that has the USE FOR DEBUGGING declarative.

| Performance considerations using DEBUG:

| The eleven programs measured ranged from equivalent to 2080% slower when using DEBUG compared
| to using NODEBUG.

| **Note:** The programs measured in this test were modified to use WITH DEBUGGING MODE on the
| SOURCE-COMPUTER paragraph and to contain a USE FOR DEBUGGING ON ALL PROCEDURES
| declarative that did a DISPLAY DEBUG-ITEM. Since the debugging code in these cases is generated only
| for paragraph and section labels, other programs may have significantly different results.

| (LE INST: pp 164-165; LE REF: p 16)

RPTOPTS

The RPTOPTS option allows you to get a report of the run-time options that were in use during the execution of an application. This report is produced after the application has terminated. Thus, if the application abends, the report will not be generated. Generating the report can result in some additional overhead. Specifying RPTOPTS(OFF) will eliminate this overhead.

Performance considerations using RPTOPTS:

| On the average, RPTOPTS(ON) was equivalent to RPTOPTS(OFF), with a range of equivalent to 2%
| slower.

| (LE INST: pp 188-189; LE REF: pp 31-32)

RPTSTG

The RPTSTG option allows you to get a report on the storage that was used by an application. This report is produced after the application has terminated. Thus, if the application abends, the report will not be generated. The data from this report can help you fine tune the storage parameters for the application, reducing the number of times that the LE storage manager must make system requests to acquire or free storage. Collecting the data and generating the report can result in some additional overhead. Specifying RPTSTG(OFF) will eliminate this overhead.

Performance considerations using RPTSTG:

| On the average, RPTSTG(ON) was 5% slower than RPTSTG(OFF), with a range of equivalent to
| 37% slower. Note that when using call intensive applications, the degradation can be 160% slower or
| more.

| (LE INST: pp 191-192; LE REF: pp 33-34)

RTEREUS

The RTEREUS option causes the LE run-time environment to be initialized for reusability when the first COBOL program is invoked. The LE run-time environment remains initialized (all COBOL programs and their work areas are kept in storage) in addition to keeping the library routines initialized and in storage. This means that, for subsequent invocations of COBOL programs, most of the run-time environment initialization will be bypassed. Most of the run-time termination will also be bypassed, unless a STOP RUN is executed or unless an explicit call to terminate the environment is made (**Note:** using STOP RUN results in control being returned to the caller of the routine that invoked the first COBOL program, terminating the reusable run-time environment).

Because of the effect that the STOP RUN statement has on the run-time environment, you should change all STOP RUN statements to GOBACK statements in order to get the benefit of RTEREUS. The most noticeable impact will be on the performance of a non-COBOL driver repeatedly calling a COBOL subprogram (for example, an IMS environment or an assembler driver that repeatedly calls COBOL applications). The RTEREUS option helps in this case. However, using the RTEREUS option does affect the semantics of the COBOL application: each COBOL program will now be considered to be a subprogram and will be entered in its last-used state on subsequent invocations (if you want the program to be entered in its initial state, you can use the INITIAL clause on the PROGRAM-ID statement). **WARNING: This means that storage that is acquired during the execution of the application will not be freed.** Therefore, RTEREUS may not be applicable to all environments.

Performance considerations using RTEREUS (measuring CALL overhead only):

One program (Assembler calling COBOL) using RTEREUS was 99% faster than using NORTEREUS.

Note: This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

| (LE INST: pp 194-195; LE REF: pp 34-35; COB MIG: pp 41, 67)

STORAGE

This option clears all external data records acquired by a program to binary zeros when the storage for the external data is allocated. Additionally, the working storage acquired by a RENT program is cleared to binary zeros (unless a VALUE clause is used on the data item) when the program is first called or, for dynamic calls, when the program is canceled and then called again. Storage is not cleared on subsequent calls to the program. This can result in some overhead at run time depending on the number of external data records in the program and the size of the working storage section.

Performance considerations using STORAGE:

| On the average, STORAGE(00,00,00) was 7% slower than STORAGE(NONE,NONE,NONE), with a
| range of equivalent to 57% slower. Note that when using call intensive applications, the degradation can
| be 160% slower or more.

| (LE INST: pp 198-201; LE REF: pp 37-38; COB MIG: pp 42-44)

TEST

The TEST option specifies the conditions under which Debug Tool (the debugging component of AD/Cycle CODE/370) assumes control when the user application is invoked. Since this may result in Debug Tool being initialized and invoked, there may be some additional overhead when using TEST. Specifying NOTEST will eliminate this overhead.

Performance data using TEST is not available at this time.

| (LE INST: pp 203-205; LE REF: pp 39-40)

TRAP

The TRAP option allows LE to intercept an abnormal termination (abend), provide the abend information, and then terminate the LE run-time environment. TRAP(OFF) prevents LE from intercepting the abend. In general, there will not be any significant impact on the performance of a COBOL application when using TRAP(ON).

Performance considerations using TRAP:

On the average, TRAP(ON) was equivalent to TRAP(OFF).

| (LE INST: pp 208-211; LE REF: pp 42-43)

VCTRSAVE

The VCTRSAVE option specifies whether any language in the application uses the vector facility when the user-provided condition handlers are called. When the condition handlers use the vector facility, the entire vector environment has to be saved on every condition and restored upon return to the application code. Unless you need the function provided by VCTRSAVE(ON), you should run with VCTRSAVE(OFF) to avoid this overhead.

Performance data using VCTRSAVE is not available at this time.

| (LE INST: pp 212-213; LE REF: p 44)

COBOL and LE Features that Affect Run-Time Performance

COBOL for MVS & VM and LE for MVS & VM have several installation and environment tuning features that can enhance the performance of your application. We will now look at some additional factors that should be considered for the application.

Storage Management Tuning

Storage management tuning can reduce the overhead involved in getting and freeing storage for the application program. With proper tuning, several GETMAIN and FREEMAIN calls can be eliminated. To better understand the need for storage tuning, let's look at how storage management works.

First of all, storage management was designed to keep a block of storage only as long as necessary. This means that during the execution of a COBOL program, if any block of storage becomes empty, it will be freed. This can be beneficial in a transaction environment (or any environment) where you want storage to be freed as soon as possible so that other transactions (or applications) can make efficient use of the storage. However, it can also be detrimental if the last block of storage does not contain enough free space to satisfy a storage request by a library routine. For example, suppose that a library routine needs 2K of storage but

there is only 1K of storage available in the last block of storage. The library routine will call storage management to request 2K of storage. Storage management will determine that there is not enough storage in the last block and issue a GETMAIN to acquire this storage (this GETMAINED size can also be tuned). The library routine will use it and then, when it is done, call storage management to indicate that it no longer needs this 2K of storage. Storage management, seeing that this block of storage is now empty, will issue a FREEMAIN to release the storage back to the operating system. Now, if this library routine or any other library routine that needs more than 1K of storage is called often, a significant amount of CPU time degradation can result because of the amount of GETMAIN and FREEMAIN activity.

Fortunately, there is a way to compensate for this with LE; it is called storage management tuning. The RPTSTG(ON) run-time option can help you in determining the values to use for any specific application program. You use the value returned by the RPTSTG(ON) option as the size of the initial block of storage for the HEAP, ANYHEAP, BELOWHEAP, STACK, and LIBSTACK run-time options. This will prevent the above from happening in an all COBOL/370 or COBOL for MVS & VM application. However, if the application also contains OS/VSE COBOL programs that are being called frequently, the RPTSTG(ON) option may not indicate a need for additional storage. Increasing these initial values can also eliminate some storage management activity in this mixed environment.

The IBM supplied default storage options for batch applications are listed below:

```
ANYHEAP(16K,8K,ANYWHERE,FREE)
BELOWHEAP(8K,4K,FREE)
HEAP(32K,32K,ANYWHERE,KEEP,8K,4K)
LIBSTACK(8K,4K,FREE)
STACK(128K,128K,BELOW,KEEP)
```

If you are running only COBOL applications, you can do some further storage tuning as indicated below:

```
STACK(64K,64K,BELOW,KEEP)
```

Additionally, if all of your applications are AMODE(31), you can use ALL31(ON) and STACK(.,ANYWHERE).

Overall below the line storage requirements have been substantially reduced by reducing the default storage options and by moving some of the library routines above the line. Here is a comparison of storage usage for a minimal GOBACK COBOL program, using the IBM supplied default run-time options:

	Below	Above
LE/370 Release 3	1156K	156K
LE/370 Release 4	552K	316K
LE for MVS & VM Release 5	268K	928K

As you can see, by moving to LE for MVS & VM Release 5, you can reduce the amount of below the line storage used by your applications.

Performance data using Storage Management Tuning is not available at this time.

(LE INST: pp 157-160, 170-172, 174-175, 196-198; LE REF: pp 12-14, 20-23, 35-37, 254-256; LE PG: pp 168-175; COB MIG: p 42)

Calling IGZERRE

Another way to set up a reusable run-time environment for COBOL is by calling IGZERRE. This module can be invoked to explicitly initialize and terminate COBOL's portion of the LE run-time environment. It allows a non-COBOL, non-LE-conforming program to initialize the LE run-time environment, thereby effectively establishing itself as the main COBOL program. As a result, the use of STOP RUN will cause control to be returned to the caller of the routine that invoked the IGZERRE initialization. IGZERRE is an enhanced version of ILBOSTP0 (for OS/VSE COBOL) and accomplishes the same results as ILBOSTP0,

except that IGZERRE has been designed for COBOL for MVS & VM, COBOL/370, and VS COBOL II applications only. Using IGZERRE has the added benefits of supporting applications running above the 16 MB line, allowing the application to terminate the COBOL portion of the LE run-time environment and improving the performance of the application. Using ILBOSTP0 in an LE environment will set up both the OS/VS COBOL and the COBOL portion of the LE environments whereas using IGZERRE will set up only the COBOL portion of the LE environment. The semantic changes and performance benefits of using this method are the same as when using the RTEREUS run-time option. See “Using IGZERRE” on page 38 for an example of using IGZERRE.

Performance considerations using IGZERRE (measuring CALL overhead only):

One program (Assembler calling COBOL) using IGZERRE was 99% faster than not using IGZERRE.

See “First Program Not COBOL” on page 17 for additional performance considerations comparing calling IGZERRE with other environment initialization techniques.

Note: This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(COB MIG: p 66)

Using the CEEENTRY and CEETERM Macros

To improve the performance of Assembler calling COBOL, you can make the Assembler program LE-conforming. This can be done using the CEEENTRY and CEETERM macros provided with LE. See “Using CEEENTRY” on page 35 for an example of using the CEEENTRY and CEETERM macros.

Performance considerations using the CEEENTRY and CEETERM macros (measuring CALL overhead only):

One program (Assembler calling COBOL) using the CEEENTRY and CEETERM macros was 99% faster than not using them.

See “First Program Not COBOL” on page 17 for additional performance considerations comparing using CEEENTRY and CEETERM with other environment initialization techniques.

Note: This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(LE PG: pp 454-465)

Using Preinitialization Services (CEEPIPI)

LE preinitialization services (CEEPIPI) can also be used to improve the performance of Assembler calling COBOL. LE preinitialization services lets an application initialize the LE environment once, execute multiple LE-conforming programs, then explicitly terminate the LE environment. This substantially reduces the use of system resources that would have been required to initialize and terminate the LE environment for each program of the application. See “Using CEEPIPI with Call_Sub” on page 40 for an example of using CEEPIPI to call a COBOL subprogram and “Using CEEPIPI with Call_Main” on page 42 for an example of using CEEPIPI to call a COBOL main program.

Performance considerations using CEEPIPI (measuring CALL overhead only):

One program (Assembler calling COBOL) using CEEPIPI to invoke the COBOL program as a subprogram was 99% faster than not using CEEPIPI.

The same program using CEEPIPI to invoke the COBOL program as a main program was 95% faster than not using CEEPIPI.

See “First Program Not COBOL” on page 17 for additional performance considerations comparing using CEEPIPI with other environment initialization techniques.

Note: This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(COB MIG: p 77; LE PG: pp 470-497)

Using Library Routine Retention (LRR)

LRR is a function that provides a performance improvement for those applications or subsystems running on MVS with the following attributes:

- the application or subsystem invokes programs that require LE
- the application or subsystem is not LE-conforming (i.e., LE is not already initialized when the application or subsystem invokes programs that require LE)
- the application or subsystem repeatedly invokes programs that require LE running under the same MVS task
- the application or subsystem is not using LE preinitialization services

LRR is useful for assembler drivers that repeatedly call LE-conforming languages and for IMS/DC regions. LRR is supported only under MVS and not under VM. Also, LRR is not supported under CICS.

When LRR has been initialized, LE keeps a subset of its resources in memory after the environment terminates. As a result, subsequent invocations of programs in the same MVS task that caused LE to be initialized are faster because the resources can be reused without having to be reacquired and reinitialized. The resources that LE keeps in memory upon LE termination are:

- LE run-time load modules
- storage associated with these load modules
- storage for LE startup control blocks

When LRR is terminated, these resources are released from memory.

LE preinitialization services and LRR can be used simultaneously. However, there is no additional benefit by using LRR when LE preinitialization services are being used. Essentially, when LRR is active and a non-LE-conforming application uses preinitialization services, LE remains preinitialized between repeated invocations of LE-conforming programs and does not terminate. Upon return to the non-LE-conforming application, preinitialization services can be called to terminate the LE environment, in which case LRR will be back in effect. See “Using CEELRR” on page 36 for an example of using LRR.

Performance considerations using LRR:

One program (Assembler calling COBOL) using LRR was 96% faster than using not using LRR.

See “First Program Not COBOL” on page 17 for additional performance considerations comparing using LRR with other environment initialization techniques.

Note: This test measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(LE PG: pp 450-453; COB MIG: p 76)

Tailoring COBPACKs

Tailoring the COBPACKs can also affect the performance by reducing the directory search and program fetch time to load the COBOL component of the LE library. The ideal COBPACKs for a particular application are those that have only those library routines in them that are needed by the application. This will reduce the amount of overhead required to load the COBPACKs as well as reduce the amount of virtual storage required by not having unused library routines loaded. When properly tuned, the number of COBOL-specific library routines loaded can be minimized. However, this may not be practical since the COBPACKs will generally be shared among several different applications and cannot be tuned for one specific application. In this case, the COBPACKs should have, as a minimum, all library routines that are common to all application programs. Others can be added as virtual storage is available.

Performance data using tailored COBPACKs is not available at this time.

| (LE INST: pp 32-33, 128-131, 229-234)

Library in the LPA/ELPA (MVS) or NSS (CMS)

Placing the COBOL and the LE library routines in the Link Pack Area (LPA) or Extended Link Pack Area (ELPA) on an MVS system or the Named Save Segment (NSS) on a CMS system can also help to improve total system performance. This will reduce the real storage requirements for the entire system for COBOL/370, COBOL for MVS & VM, VS COBOL II RES, or OS/VS COBOL RES applications since the library routines can be shared by all applications instead of each application having its own copy of the library routines. For the COBOL COBPACKs, if you want them to go above the 16 MB line, you must be careful to include only those library routines that have AMODE 31 and RMODE ANY in the COBPACKs. A list of the eligible library routines with their AMODE and RMODE attributes can be found in the Installation and Customization manual.

Placing the library routines in a shared area will also reduce the I/O activity since they are loaded only once when the system is started and not for each application program.

Performance data using COBPACKs in the LPA/ELPA/NSS is not available at this time.

| (LE INST: pp 31-32, 126-128, 266-267, 269-271)

Using CALLs

When using CALLs, be sure to consider using nested programs when possible. The performance of a CALL to a nested program is faster than an external static CALL; external dynamic calls are the slowest. CALL identifier is slower than dynamic CALL literal. Additionally, you should consider space management tuning (mentioned earlier in this paper) for all CALL intensive applications.

With static CALLs, all programs are link-edited together, and hence, are always in storage, even if you do not call them. However, there is only one copy of the bootstrapping library routines link-edited with the application.

With dynamic CALLs, each subprogram is link-edited separately from the others. They are brought into storage only if they are needed. However, each subprogram has its own copy of the bootstrapping library routines link-edited with it, bringing multiple copies of these routines in storage as the application is executing.

Performance considerations for using CALLs (measuring CALL overhead only):

| CALL to nested programs was 50% to 60% faster than static CALL.

| Static CALL literal was 45% to 55% faster than dynamic CALL literal.

| Static CALL literal was 60% to 65% faster than dynamic CALL identifier.

Dynamic CALL literal was 15% to 25% faster than dynamic CALL identifier.

Note: These tests measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(COB PG: pp 342-353)

Other Product Related Factors that Affect Run-Time Performance

It is important to understand COBOL's interaction with other products in order to enhance the performance of your application. We will now look at some product related factors that should be considered for the application.

Mixing VS COBOL II or COBOL/370 Rel 1 with COBOL for MVS & VM Rel 2

If the COBOL for MVS & VM Release 2 application program statically calls a COBOL/370 Release 1 program or a VS COBOL II program, you must ensure that the proper bootstrap routines are linked with the application. If you are linking only object decks (output from the compiler), the normal link-edit process will do this. However, if you have previously link-edited load modules that you are including in the application, you must do the following:

- if the VS COBOL II RES program was link-edited with the VS COBOL II Release 4 run-time library without APAR PN74000, you must include a REPLACE IGZEBST control statement in the linkage editor input.
- if the VS COBOL II RES program was link-edited with the LE/370 Release 2 or 3 run-time library without APAR PN74011, you must include a REPLACE IGZEBST control statement in the linkage editor input.

Additionally, to have the best performance, you should also do the following:

- if the COBOL/370 Release 1 program was link-edited with the LE/370 Release 2, 3, or 4 run-time library without APAR PN74011, you should include a REPLACE IGZCBSN control statement in the linkage editor input.
- if the VS COBOL II NORES program was link-edited with the LE/370 Release 2, 3, or 4 run-time library without APAR PN74011, you should include a REPLACE IGZENRI control statement in the linkage editor input.
- if the VS COBOL II RES program was link-edited with the LE/370 Release 4 run-time library without APAR PN74011, you must include a REPLACE IGZEBST control statement in the linkage editor input.

(COB MIG: pp 58-78, 154-155)

Mixing OS/VS COBOL with COBOL/370 Rel 1 or COBOL for MVS & VM Rel 2

If the application program is an OS/VS COBOL program using the LE library or if the application program has a mixture of OS/VS COBOL and COBOL/370 or COBOL for MVS & VM, there will be some degradation at run time since both the OS/VS COBOL environment and the LE environment must be initialized and cleaned up. Converting the entire application to COBOL/370 or COBOL for MVS & VM will eliminate the need for setting up both environments.

Performance data using mixed OS/VS COBOL and COBOL/370 or COBOL for MVS & VM programs is not available at this time.

(COB MIG: pp 48-57, 154)

First Program Not COBOL

If the first program in the application is not COBOL or it is not LE-conforming, there can be a significant degradation if COBOL is repeatedly called since the COBOL environment must be initialized and terminated each time a COBOL main program is invoked. This overhead can be reduced by doing one of the following (listed in order of most improvement to least improvement):

- Use the CEEENTRY and CEETERM macros in the first program of the application to make it an LE-conforming program
- Call the first program of the application from a COBOL stub program (a program that just has a call statement to the original first program)
- Call CEEPIPI from the first program of the application to initialize the LE environment, invoke the COBOL program, and then terminate the LE environment when the application is complete
- Use the run-time option RTEREUS to initialize the run-time environment for reusability, making all COBOL main programs become subprograms
- Call IGZERRE from the first program of the application to make it appear as the COBOL main program and to initialize the run-time environment for reusability
- Call ILBOSTP0 (when using OS/VS COBOL) from the first program of the application to make it appear as the COBOL main program. This is provided for compatibility with OS/VS COBOL; calling IGZERRE is preferred over calling ILBOSTP0.
- Use the Library Routine Retention (LRR) function (similar to the function provided by the LIBKEEP run-time option in VS COBOL II)
- Place the LE library routines in the LPA, ELPA, or NSS. The list of routines to put in the LPA, EPLA, or NSS is release dependent and is the same routines listed under the IMS preload list considerations on page 19. Additionally, if the application is a VS COBOL II application using the LE library, include the following library routines: IGZCTCO, IGZEINI, IGZEPLF, and IGZEPCL. An alternative method (when not using RTEREUS, IGZERRE, or CEEPIPI) is to load these routines using the LOAD macro prior to the first time COBOL is called, or to use the NUCXLOAD command on CMS to load them as a nucleus extension (CEEBINIT cannot be loaded as a nucleus extension).

Make sure that you fully understand the implications of each one before you use them.

The performance considerations for Assembler calling COBOL will be presented in two different ways. The first will compare each of the tuning possibilities with using a standard, non-tuned assembler program to repeatedly call COBOL (the slowest method). The second will show the same comparison using an LE-conforming assembler program using the CEEENTRY and CEETERM macros (the fastest method).

Performance considerations for Assembler calling COBOL (measuring CALL overhead only) compared to a non-tuned assembler program, in order of most improvement to least improvement:

CALL overhead was 99% faster when using the CEEENTRY and CEETERM macros (LE-conforming assembler)

CALL overhead was 99% faster when calling the Assembler program from a COBOL stub.

CALL overhead was 99% faster when calling CEEPIPI to invoke the COBOL program as a subprogram.

CALL overhead was 99% faster when using the RTEREUS run-time option.

CALL overhead was 99% faster when calling IGZERRE before calling COBOL.

CALL overhead was 99% faster when calling ILBOSTP0 before calling COBOL.

CALL overhead was 97% to 95% faster when calling CEEPIPI to invoke the COBOL program as a main program.

CALL overhead was 96% faster when using LRR (MVS only).

CALL overhead was 88% to 72% faster when using the LOAD SVC to load the LE library routines before calling the COBOL program.

CALL overhead was 75% faster when NUCXLOADing (VM only) the LE library routines (except CEEBINIT) before calling the COBOL program.

Performance considerations for placing the library routines in the LPA, ELPA, or NSS are not available at this time.

Performance considerations for Assembler calling COBOL (measuring CALL overhead only) compared to an LE-conforming assembler program, in order of least degradation to most degradation:

CALL overhead was 2% slower when calling the Assembler program from a COBOL stub.

CALL overhead was 45% slower when calling CEEPIPI to invoke the COBOL program as a subprogram.

CALL overhead was 735% slower when using the RTEREUS run-time option.

CALL overhead was 745% slower when calling IGZERRE before calling COBOL.

CALL overhead was 750% slower when calling ILBOSTP0 before calling COBOL.

CALL overhead was 4,790% slower when calling CEEPIPI to invoke the COBOL program as a main program.

CALL overhead was 5,250% slower when using LRR (MVS only).

CALL overhead was 19,600% slower when using the LOAD SVC to load the LE library routines before calling the COBOL program.

CALL overhead was 39,400% slower when NUCXLOADing (VM only) the LE library routines (except CEEBINIT) before calling the COBOL program.

CALL overhead was 166,000% slower when not using any of the above methods (non-tuned assembler)

Note: These tests measured only the overhead of the CALL (i.e., the subprogram did only a GOBACK); thus, a full application that does more work in the subprograms may have different results.

(**COB MIG:** pp 65-67, 76-77, 192-193)

IMS

If the application is running under IMS, preloading the application program and the library routines can help to reduce the load/search overhead, as well as reduce the I/O activity. This is especially true for the library routines since they are used by every COBOL program. When the application program is preloaded, subsequent requests for the program are handled faster because it does not have to be fetched from external storage. The RENT compiler option is required for preloaded applications. If you are using a release of IMS/VS prior to Version 3 Release 1, the data for the IMS application program must reside below the 16 MB line, and hence, the DATA(24) compiler option is also required if you are using IMS services.

(**COB MIG:** p 257)

Using the Library Routine Retention (LRR) function can significantly improve the performance of COBOL transactions running under IMS/DC. LRR provides function similar to that of the VS COBOL II LIBKEEP run-time option. It keeps the LE environment initialized and retains in memory any loaded LE library routines, storage associated with these library routines, and storage for LE startup control blocks. To use LRR in an IMS dependent region, you must do the following:

- In your startup JCL or procedure to bring up the IMS dependent region, specify the PREINIT=xx parameter (xx is the 2-character suffix of the DFSINTxx member in your IMS PROCLIB dataset)
- Include the name CEELRRIN in the DFSINTxx member of your IMS PROCLIB dataset
- Bring up your IMS dependent region

You can also create your own load module to initialize the LRR function by modifying the CEELRRIN sample source in the SCEESAMP dataset. If you do this, use your module name in place of CEELRRIN above.

(LE INST: pp 134-135; LE PG: pp 450-453)

WARNING: If the RTEREUS run-time option is used, the top level COBOL programs of all applications must be preloaded. Note that using RTEREUS will keep the LE environment up until the region goes down or until a STOP RUN is issued by a COBOL program. This means that every program and its working storage (from the time the first COBOL program was initialized) is kept in the region. Although this is very fast, you may find that the region may soon fill to overflowing, especially if there are many different COBOL programs that are invoked.

When not using RTEREUS or LRR, it is recommended that you preload the following library routines:

- For LE/370 Release 1:
CEEBINIT, IGZCPAC, IGZCPCO, CEEEV005, and CEEVHEAP
- For LE/370 Release 2:
CEEBINIT, IGZCPAC, IGZCPCO, CEEEV005, CEEVHEAP, and CEEPLPKA
- For LE/370 Release 3:
CEEBINIT, IGZCPAC, IGZCPCO, CEEEV005, and CEEPLPKA
- For LE/370 Release 4:
CEEBINIT, IGZCPAC, IGZCPCO, CEEEV005, CEEPLPKA, IGZETRM, IGZEINI, IGZCLNK (for COBOL/370 Release 1 and VS COBOL II), and IGZCLNC (for OS/VS COBOL)
- For LE for MVS & VM Release 5:
CEEBINIT, IGZCPAC, IGZCPCO, CEEEV005, CEEPLPKA, IGZETRM, IGZEINI, IGZCFCC (for COBOL for MVS & VM Release 2), IGZCLNK (for COBOL/370 Release 1 and VS COBOL II), and IGZCLNC (for OS/VS COBOL)
- If the application contains VS COBOL II programs:
IGZCTCO, IGZEINI, IGZEPLF, and IGZEPCL
- If the application contains OS/VS COBOL programs:
IGZCTCO, IGZEINI, IGZEPLF, and IGZEPCL
any ILBO library routines that you previously preloaded.
If ILBOSTT0 is in the preload list, make sure that it is in there twice.
You should, at a minimum, include all heavily used ILBO library routines in your preload list.

Preloading should reduce the amount of I/O activity associated with loading and deleting these routines for each transaction.

The two COBPACKs, IGZCPAC and IGZCPCO, can be tailored to eliminate any library routines that you do not need. Additionally, all AMODE 24, RMODE 24 routines can be removed from the COBPACKs to allow the COBPACKs to reside above the 16 MB line. In this case, you should also preload any of the below the line routines that you need.

Additionally, heavily used application programs can be compiled with the RENT compiler option and preloaded to reduce the amount of I/O activity associated with loading them.

Performance data using COBOL for MVS & VM with IMS is not available at this time.

(COB PG: pp 432, 473-476)

CICS

In order to minimize the amount of below the line storage used by LE under CICS, you should run with ALL31(ON) and STACK(,ANYWHERE) as much as possible. In order to do this, you have to identify all of your VS COBOL II, COBOL/370, and COBOL for MVS & VM AMODE(24) programs. Then you can either make the necessary coding changes to make them AMODE(31) or you can link-edit a CEEUOPT with ALL31(OFF) and STACK(,BELOW) as necessary for those run units that need it. You can find out

how much storage a particular transaction is using by looking at the auxiliary trace data for that transaction. You do not need to be concerned about the OS/VS COBOL programs since the LE run-time options do not affect OS/VS COBOL programs running under CICS. Also, if the transaction is defined with TASKDATALOC(ANY) and ALL31(ON) is being used and the programs are compiled with DATA(31), then LE does not use any below the line storage for the transaction under CICS, resulting in some additional below the line storage savings.

The RENT compiler option is required for an application running under CICS. Additionally, if the program is run through the CICS translator (i.e., it has EXEC CICS commands in it), it must also use the NODYNAM compiler option. CICS Version 3 Release 2.1 or later is required for COBOL for MVS & VM.

(**COB PG:** pp 462-464; **COB MIG:** p 134)

COBOL for MVS & VM and COBOL/370 support static and dynamic calls to COBOL for MVS and VM, COBOL/370, and VS COBOL II subprograms containing CICS commands or dependencies. Static calls are done with the CALL literal statement and dynamic calls are done with the CALL identifier statement. Converting EXEC CICS LINKs to COBOL CALLs can improve transaction response time and reduce virtual storage usage. Neither COBOL for MVS & VM nor COBOL/370 support calls to or from OS/VS COBOL programs in a CICS environment. In this case, EXEC CICS LINK must be used.

(**LE PG:** pp 419-420; **COB MIG:** pp 146-147)

If you are using the COBOL CALL statement to call a program that has been translated with the CICS translator, you must pass DFHEIBLK and DFHCOMMAREA as the first two parameters on the CALL statement. However, if you are calling a program that has not been translated, you should not pass DFHEIBLK and DFHCOMMAREA on the CALL statement. Additionally, if your called subprogram does not use any of the EXEC CICS condition handling commands, you can use the run-time option CBLPSHPOP(OFF) to eliminate the overhead of doing an EXEC CICS PUSH HANDLE and an EXEC CICS POP HANDLE that is done for each call by the LE run-time.

(**LE INST:** p 161; **LE REF:** p 14; **COB PG:** pp 465-470; **COB MIG:** pp 44, 141-143)

As long as your usage of all binary (COMP) data items in the application conforms to the PICTURE and USAGE specifications, you can use TRUNC(OPT) to improve transaction response time. This is recommended in performance sensitive CICS applications. If your usage of any binary data item does not conform to the PICTURE and USAGE specifications, you should use TRUNC(BIN). Note that the CICS translator does not generate code that will cause truncation. If you were using NOTRUNC with your OS/VS COBOL programs without problems, TRUNC(OPT) on COBOL for MVS & VM will behave the same way. For additional information on the TRUNC option, please refer to the compiler options section of this paper. **Note:** This is the most up-to-date information for using TRUNC with CICS applications and may be different than what is recommended in the Programming Guides.

Performance data using COBOL for MVS & VM with CICS is not available at this time.

(**COB PG:** pp 432, 463-464; **LE PG** pp 410-412, 416-417; **COB MIG:** p 135)

DB2

As long as your usage of all binary (COMP) data items in the application conforms to the PICTURE and USAGE specifications, you can use TRUNC(OPT) to improve performance under DB2. This is recommended in performance sensitive DB2 applications. If your usage of any binary data item does not conform to the PICTURE and USAGE specifications, you should use TRUNC(BIN). If you were using NOTRUNC with your OS/VS COBOL programs without problems, TRUNC(OPT) on COBOL for MVS & VM will behave the same way. For additional information on the TRUNC option, please refer to the compiler options section of this paper. **Note:** This is the most up-to-date information for using TRUNC with DB2 applications and may be different than what is recommended in the Programming Guides.

| Performance data using COBOL for MVS & VM with DB2 is not available at this time.

| **DFSORT**

| Use the FASTSRT compiler option to improve the performance of most sort operations. With FASTSRT, the DFSORT product performs the I/O on input and/or output files named in either or both of the SORT ... USING or SORT ... GIVING statements. If you have input or output procedures for your sort files, you cannot use the FASTSRT option. The complete list of requirements is contained in the Programming Guides.

| Performance data using COBOL for MVS & VM with DFSORT is not available at this time.

| (**COB PG:** pp 173-179; **LE PG** pp 542-543)

Efficient COBOL Coding Techniques

This section focuses on how the source code can be modified to tune a program for better performance. Coding style, as well as data types, can have a significant impact on the performance of an application. Producing higher performing code usually has a far greater impact than that of tuning the application via compiler and run-time options, but producing such code may not be a viable option for many existing applications since it does require modifying the source code and at times may even require an extensive knowledge of how the program works. However, these techniques should be considered for all new applications.

| (COB PG: pp 417-433)

Data Files

Planning how the files will be created and used is an important factor in determining efficient file characteristics for the application. Some of the characteristics that affect the performance of file processing are: file organization, access method, record format, and blocksize. Some of these are discussed in more detail below.

QSAM Files

When using QSAM files, use large block sizes whenever possible by using the BLOCK CONTAINS clause on your file definitions (the default with COBOL is to use unblocked files). If you are using DFP Version 3 Release 1 or later, you can have the system determine the optimal blocksize for you by specifying the BLOCK CONTAINS 0 clause for any new files that you are creating and omitting the BLKSIZE parameter in your JCL for these files. This should significantly improve the file processing time (both in CPU time and elapsed time).

Additionally, increasing the number of I/O buffers for heavy I/O jobs can improve both the CPU and elapsed time performance, at the expense of using more storage. This can be accomplished by using the BUFNO subparameter of the DCB parameter in the JCL or by using the RESERVE clause of the SELECT statement in the FILE-CONTROL paragraph. Note that if you do not use either the BUFNO subparameter or the RESERVE clause, the system default will be used.

| QSAM buffers can be allocated above the 16 MB line if all of the following are true:

- | • the programs are compiled with VS COBOL II Release 3.0 or higher, COBOL/370 Release 1.0 or higher, or IBM COBOL for MVS & VM Release 2.0 or higher
- | • the programs are running with LE/370 Release 3.0 or higher, or IBM Language Environment for MVS & VM Release 5.0 or higher
- | • the programs are compiled with RENT and DATA(31)
- | • the ALL31(ON) run-time option is used (for EXTERNAL files)

| (COB PG: pp 119-120, 130, 253)

Variable-Length Files

When writing to variable-length blocked sequential files, use the APPLY WRITE-ONLY clause for the file or use the AWO compiler option. This reduces the number of calls to Data Management Services to handle the I/Os.

| (COB PG: pp 22, 249, 428)

VSAM Files

When using VSAM files, increase the number of data buffers (BUFND) for sequential access or index buffers (BUFNI) for random access. Also, select a control interval size (CISZ) that is appropriate for the application. A smaller CISZ results in faster retrieval for random processing at the expense of inserts, whereas a larger CISZ is more efficient for sequential processing. In general, using large CI and buffer space VSAM parameters may help to improve the performance of the application.

In general, sequential access is the most efficient, dynamic access the next, and random access is the least efficient. However, for relative record VSAM (ORGANIZATION IS RELATIVE), using ACCESS IS DYNAMIC when reading each record in a random order can be slower than using ACCESS IS RANDOM, since VSAM may prefetch multiple tracks of data when using ACCESS IS DYNAMIC. ACCESS IS DYNAMIC is optimal when reading one record in a random order and then reading several subsequent records sequentially.

Random access results in an increase in I/O activity because VSAM must access the index for each request.

If you use alternate indexes, it is more efficient to use the Access Method Services to build them than to use the AIXBLD run-time option. Avoid using multiple alternate indexes when possible since updates will have to be applied through the primary paths and reflected through the multiple alternate paths.

VSAM buffers can be allocated above the 16 MB line if all of the following are true:

- the programs are compiled with VS COBOL II Release 3.0 or higher, COBOL/370 Release 1.0 or higher, or IBM COBOL for MVS & VM Release 2.0 or higher
- the programs are running with LE/370 Release 3.0 or higher, or IBM Language Environment for MVS & VM Release 5.0 or higher

(COB PG: pp 161-162, 420)

Data Types

Using the proper data types is also an important factor in determining the performance characteristics of an application. Some of these are discussed below.

(COB PG: pp 419-425)

BINARY (COMP or COMP-4)

When using binary (COMP) data items, the use of the SYNCHRONIZED clause specifies that the binary data items will be properly aligned on halfword, fullword, or doubleword boundaries. This may enhance the performance of certain operations on some machines. Additionally, using signed data items with eight or fewer digits produces the best code for binary items. The following shows the performance considerations (from most efficient to least efficient) for the number of digits of precision for signed binary data items (using PICTURE S9(n) COMP):

- n is from 1 to 8
 - for n from 1 to 4, arithmetic is done in halfword instructions where possible
 - for n from 5 to 8, arithmetic is done in fullword instructions where possible
- n is from 10 to 17
 - arithmetic is done in doubleword format
- n is 9

fullword values are converted to doubleword format and then doubleword arithmetic is used (**this is SLOWER than any of the above**)

- n is 18

doubleword values are converted to a higher precision format and then arithmetic is done using this higher precision (**this is the SLOWEST of all for binary data items**)

Note: Using 9 digits is slower than using 10 digits.

| Performance considerations for BINARY:

| using 1 to 8 digits is the fastest

| using 10 to 17 digits is 20% to 30% slower than using 1 to 8 digits.

| using 9 digits is 50% slower than using 1 to 8 digits.

| using 18 digits is 1150% slower than using 1 to 8 digits.

| (COB PG: pp 419-420)

Data Conversions

Conversion to a common format is necessary for certain types of numeric operations when mixed data types are involved in the computation. This results in additional processing time and storage for these conversions. In order to minimize this overhead, it is recommended that the guidelines discussed below be followed.

| (COB PG: pp 64-65, 420)

DISPLAY

Avoid using USAGE DISPLAY data items for computations (especially in areas that are heavily used for computations). When a USAGE DISPLAY data item is used, additional overhead is required to convert the data item to the proper type both before and after the computation. In some cases, this conversion is done by a call to a library routine, which can be expensive compared to using the proper data type that does not require any conversion.

| Performance considerations for DISPLAY:

| using 1 to 5 digits is the fastest

| using 6 to 9 digits is 15% slower than using 1 to 5 digits

| using 10 to 13 digits is 50% slower than using 1 to 5 digits

| using 14 to 16 digits is 65% slower than using 1 to 5 digits

| using 17 to 18 digits is 120% slower than using 1 to 5 digits

| (COB PG: p 420)

PACKED-DECIMAL (COMP-3)

When using PACKED-DECIMAL (COMP-3) data items in computations, use 15 or fewer digits in the PICTURE specification to avoid the use of library routines for multiplication and division. A call to the library routine is very expensive when compared to doing the calculation in-line. Additionally, using a signed data item with an odd number of digits produces more efficient code since this uses an integral multiple of bytes in storage for the data item.

| Performance considerations for PACKED-DECIMAL:

| using an odd number of digits is 6% faster than using the next lower even number of digits

using the fewest odd number of digits as possible may result in an additional 5% to 15% savings compared to using the next larger number of odd digits

(COB PG: pp 419-420)

Comparing Data Types

When selecting your data types, it is important to understand the performance characteristics of them before you use them. Shown below are some performance considerations of doing several ADDs and SUBTRACTs on the various data types of the specified precision.

Performance considerations for comparing data types:

Packed decimal (COMP-3) compared to binary (COMP or COMP-4)

using 1 to 5 digits: packed decimal is 150% to 210% slower than binary

using 6 to 8 digits: packed decimal is 220% to 260% slower than binary

using 9 to 17 digits: packed decimal is 100% to 240% slower than binary

using 18 digits: packed decimal is 65% faster than binary

DISPLAY compared to packed decimal (COMP-3)

using 1 to 9 digits: DISPLAY is 60% to 95% slower than packed decimal

using 10 to 18 digits: DISPLAY is 100% to 180% slower than packed decimal

DISPLAY compared to binary (COMP or COMP-4)

using 1 to 5 digits: DISPLAY is 350% to 480% slower than binary

using 6 to 8 digits: DISPLAY is 500% to 570% slower than binary

using 9 to 17 digits: DISPLAY is 300% to 725% slower than binary

using 18 digits: DISPLAY is 10% faster than binary

Fixed-Point vs Floating-Point

Plan the use of fixed-point and floating-point data types. You can enhance the performance of an application by carefully determining when to use fixed-point and floating-point data. When conversions are necessary, binary (COMP) and packed decimal (COMP-3) data with nine or fewer digits require the least amount of overhead when being converted to or from floating-point (COMP-1 or COMP-2) data. Also, when using fixed-point exponentiations with large exponents, the calculation can be done more efficiently by using operands that force the exponentiation to be evaluated in floating point.

Performance considerations for fixed-point vs floating-point:

forcing an exponentiation to be done in floating point is 98% faster than doing it in fixed point

(COB PG: pp 76-77, 421)

Indexes vs Subscripts

Using indexes to address a table is more efficient than using subscripts since the index already contains the displacement from the start of the table and does not have to be calculated at run time. Subscripts, on the other hand, contain an occurrence number that must be converted to a displacement value at run time before it can be used. When using subscripts to address a table, use a binary (COMP) signed data item with eight or fewer digits (for example, using PICTURE S9(8) COMP for the data item). This will allow fullword arithmetic to be used during the calculations. Additionally, in some cases, using four or fewer digits for the data item may also offer some added reduction in CPU time since halfword arithmetic can be used.

- | Performance considerations for indexes vs subscripts (PIC S9(8)):
 - | using binary data items (COMP) to address a table is 56% slower than using indexes
 - | using decimal data items (COMP-3) to address a table is 426% slower than using indexes
 - | using DISPLAY data items to address a table is 680% slower than using indexes

| (COB PG: pp 81-84, 98, 421, 424)

OCCURS DEPENDING ON

When using OCCURS DEPENDING ON (ODO) data items, ensure that the ODO objects are binary (COMP) to avoid unnecessary conversions each time the variable-length items are referenced. Some performance degradation is expected when using ODO data items since special code must be executed every time a variable-length data item is referenced. This code determines the current size of the item every time the item is referenced. It also determines the location of variably-located data items. Because this special code is out-of-line, it may inhibit some optimizations. Furthermore, code to manipulate variable-length data items is substantially less efficient than that for fixed-length data items. For example, the code to compare or move a variable-length data item may involve calling a library routine and is significantly slower than the equivalent code for fixed-length data items. If you do use variable-length data items, copying them into fixed-length data items prior to a period of high-frequency use can reduce some of this overhead.

| (COB PG: pp 89-93, 421, 424-425)

Program Design

| Using the appropriate program design is another important factor in determining the performance characteristics of an application. Some of these are discussed below.

Algorithms

Examine the underlying algorithms that have been selected before looking at the COBOL specifics. Improving the algorithms usually has a much greater impact on the performance than does improving the detailed implementation of the algorithm. As an example, consider two search algorithms: a sequential search and a binary search. Clearly, both of them will produce the same results and may, in fact, have almost the same performance for small tables. However, as the table size increases, the binary search will be much faster than the sequential search. As in this case of the two searches, you may have to do some additional coding to maintain a sorted table for the binary search, but the additional effort spent here is more than saved during the execution of the program.

| (COB PG: p 417)

Data Structures and Data Types

After deciding on the algorithm, look at the data structures and data types. Ensure that both are appropriate for the selected algorithm. The algorithm may in general be a fast one, but if the wrong data structures or types are used, the performance can degrade significantly. As an example, consider two PERFORM VARYING loops, one using a USAGE DISPLAY data item for the loop variable and the other using a COMPUTATIONAL data item. In the case of DISPLAY data item, data conversion must be done for each iteration of the loop, whereas in the COMPUTATIONAL data item, binary fullword arithmetic can be used. Once again, they will both produce the same results, but the loop using the COMPUTATIONAL data item will be much faster than the loop using the DISPLAY data item.

| Performance considerations for loop control variables (PIC S9(8)):

| using a decimal (COMP-3) is 320% slower than using binary (COMP)

| using a DISPLAY is 690% slower than using binary (COMP)

| (COB PG: pp 419-421)

Coding Style

Examine the coding style. Ensure that the program is well structured, utilizing the structured coding constructs that are available with COBOL for MVS & VM. Avoid using the GO TO statement (in particular, the altered GO TO statement) and avoid using PERFORMed procedures that involve irregular control flow (for example, a PERFORMed procedure that cannot reach the end of the procedure). The optimizer can optimize the code better and over larger blocks of code if the programs are well structured and don't have a "spaghetti-like" control flow. Additionally, the programs will be easier to maintain because of the structured logic flow.

| (COB PG: pp 417-418)

Factoring Expressions

Factor expressions where possible, especially in loops. The optimizer does not do the factoring for you. For evaluating arithmetic expressions, the compiler is bound by the left-to-right evaluation rules for COBOL. In order for the optimizer to recognize constant computations (that can be done at compile time) or duplicate computations (common subexpressions), move all constants and duplicate expressions to the left end of the expression or group them in parentheses.

| (COB PG: pp 418-419)

Symbolic Constants

If you want the optimizer to recognize a data item as a constant throughout the program, initialize it with a VALUE clause and don't modify it anywhere in the program (if a data item is passed BY REFERENCE to a subprogram, the optimizer considers it to be modified not only at this CALL statement, but also at all CALL statements).

| (COB PG: p 418)

Subscript Checking

When using tables, evaluate the need to verify subscripts. Using the SSRANGE option to catch the errors causes the compiler to generate special code at each subscript reference to determine if the subscript is out of bounds. However, if subscripts need to be checked in only a few places in the application to ensure that they are valid, then coding your own checks can improve the performance when compared to using the SSRANGE compiler option.

| (COB PG: pp 98, 389)

Subscript Usage

Additionally, try to use the tables so that the rightmost subscript varies the most often for references that occur close to each other in the program. The optimizer can then eliminate some of the subscript calculations because of common subexpression optimization.

| Performance considerations for table reference patterns (PIC S9(8)):

| when referencing tables sequentially, having the leftmost subscript vary the most often can be 7% slower
| than having the rightmost subscript vary the most often

| (COB PG: pp 422-424)

Searching

When using the SEARCH statement, place the most often used data near the beginning of the table for more efficient sequential searching. For better performance, especially when searching large tables, sort the data in the table and use the SEARCH ALL statement. This results in a binary search on the table.

| Performance considerations for search example:

| using a binary search (SEARCH ALL) can be 18% faster than using a sequential search (SEARCH)

| (COB PG: pp 93-96, 98)

Recent Performance Improvements

COBOL has made performance improvements in some specific areas of the compiled code and library routines. Here is a brief summary of the latest improvements:

| IBM COBOL for MVS & VM Version 1 Release 2 running with LE for MVS & VM Release 5 provides improved performance over COBOL/370 Version 1 Release 1 running with LE/370 Release 3:

- | • Dynamic and Static CALLs
- | • Non-COBOL to COBOL CALLs
- | • Reduction in below the line storage (LE improvement)

| COBOL/370 Version 1 Release 1 with all current maintenance (Release 1.1) running with LE/370 Release 3 provides improved performance over COBOL/370 Version 1 Release 1 running with LE/370 Release 2:

- | • Eliminating storage and initialization code for unreferenced data items with the OPTIMIZE(FULL) option
- | • RMODE option added to support NORENT programs above the 16MB line
- | • Static initialization of WORKING-STORAGE variables at compile time for NORENT programs
- | • Optimized parameter list generated code for the USING phrase of the CALL statement
- | • Variable-length MOVEs for LINKAGE SECTION data items
- | • Optimized code generated for the main entry point
- | • Dynamic CALL literal
- | • Library Routine Retention (LRR) (available with LE/370 Version 1 Release 3 and later)

A Performance Checklist

The following list of questions will help to isolate the problem area if a performance-related problem arises with COBOL for MVS & VM. Most of the performance-related problems that have been reported in the past have fallen under one or more of the categories below. Each category does not always apply to all operating environments. Most of these have been discussed earlier in this paper, so additional detail will not be given here. This will just serve as a checklist of things to consider when investigating a performance problem.

1. What are all of the compiler options that were used? Is OPTIMIZE being used? Is TRUNC(OPT) being used? Make sure you understand the performance implications of the options you are using.
2. What run-time options were used? Make sure you understand the performance implications of the options you are using.
3. How are the COBPACKs defined? As a minimum, do they contain all of the heavily used library routines?
4. Is the program compiled with OS/VS COBOL or VS COBOL II and run with the LE library?
5. Does the application have a mixture of OS/VS COBOL or VS COBOL II with COBOL/370 or COBOL for MVS & VM programs?
6. Is the application called by any other non-COBOL or non-LE-conforming programs? For a non-COBOL driver repeatedly calling COBOL, is RTEREUS or IGZERRE or ILBOSTP0 being used? Is LRR or CEEPIPI being used? Are the CEEENTRY and CEETERM macros being used?
7. Does the application have any calls to any other programs? If so, what languages are involved? What is the approximate number of calls and the depth of the calls? Are the calls static or dynamic? How many unique programs are called?
8. What Space Management Tuning is being used? The RPTSTG(ON) run-time option can help you to determine the correct values to use.
9. If the problem is on MVS, what are the JOBLIB and STEPLIB datasets and where is the LE library in the search order?
10. For IMS, is the application and/or library preloaded? Is LRR being used?
11. Do you have an execution profiler? If so, have you used it to try to identify the "hot spots" of where time is spent in the application? This information can be very useful in identifying and solving performance problems.
12. What are the release levels of all COBOL products being used? Has all current maintenance been applied? If the most current release of COBOL for MVS & VM is not being used, you should try it before reporting the problem to IBM since the problem may have already been addressed.
13. What are the release levels of the operating and subsystems being used (IMS, CICS, CMS under VM/ESA, MVS/ESA)? Has all current maintenance been applied?
14. If using SORT, what release of DFSORT is being used? Has all current maintenance been applied? Is the FASTSRT compiler option being used?
15. In case you need to seek assistance from IBM in solving the performance problem, what other information can you tell us to help us understand the overall program structure (e.g., heavy use of a particular COBOL verb, the application program alters the save area in a non-standard way, subscripts that are not binary, data types used (USAGE DISPLAY, INDEX, COMP-n), etc.)?

Summary

This paper has identified some of the factors for tuning the performance of a COBOL application through the use of compiler options, run-time options, and efficient program coding techniques. Additionally, it has identified some factors for tuning the overall LE run-time environment. A variety of different tuning tips was provided for each of the above. The primary focus was on tuning the application using the compiler and run-time options with a secondary focus, for more in depth fine tuning, on examining the program design, algorithms, and data structures.

For the type of tuning suggested here, the costs and the skill level are relatively low, because in many cases, the program itself is not changed (except for, perhaps, data types). Hence, introducing errors is a low risk. The performance gains from this type of tuning may be sufficient to delay or eliminate the need for algorithmic changes, program structure changes, or further data type considerations.

In summary, there are many opportunities for the COBOL programmer to tune the COBOL application program and run-time environment for better CPU time performance and better use of system resources. The COBOL programmer has many compiler options, run-time options, data types, and language features from which to select, and the proper choice may lead to significantly better performance. Conversely, making the wrong choice can lead to significantly degraded performance. The goal of this paper is to make you aware of the various options that are available so that you -- both the system programmer installing the product as well as the COBOL programmer responsible for the application -- can choose the right ones for your application program that will lead to the best performance for your environment.

Appendix A. Intrinsic Function Implementation Considerations

The COBOL intrinsic functions are implemented either by using LE callable services, library routines, in-line code, or a combination of these. The following table shows how each of the intrinsic functions are implemented:

Table 1 (Page 1 of 2). Intrinsic Function Implementation			
Function Name	LE Service	Library Routine	In-line Code
ACOS	X		
ANNUITY			X
ASIN	X		
ATAN	X		
CHAR			X
COS	X		
CURRENT-DATE	X		
DATE-OF-INTEGER	X		
DAY-OF-INTEGER	X		X
FACTORIAL			X
INTEGER			X
INTEGER-OF-DATE	X		
INTEGER-OF-DAY	X		X
INTEGER-PART			X
LENGTH			X
LOG	X		
LOG10	X		
LOWER-CASE		X	
MAX			X
MEAN			X
MEDIAN		X	
MIDRANGE			X
MIN			X
MOD			X
NUMVAL		X	
NUMVAL-C		X	
ORD			X
ORD-MAX			X
ORD-MIN			X
PRESENT-VALUE		X	
RANDOM	X		X
RANGE			X

Table 1 (Page 2 of 2). Intrinsic Function Implementation			
Function Name	LE Service	Library Routine	In-line Code
REM (fixed point)			X
REM (floating point)	X		
REVERSE		X	
SIN	X		
SQRT	X		
STANDARD-DEVIATION		X	X
SUM			X
TAN	X		
UPPER-CASE		X	
VARIANCE		X	X
WHEN-COMPILED			X ¹

¹ WHEN-COMPILED is a literal that is used whenever it is needed.

Appendix B. History of Prior Performance Improvements

COBOL has made performance improvements in some specific areas of the compiled code and library routines. Here is a brief history of the improvements made for prior COBOL releases, version, or products:

COBOL/370 Version 1 Release 1 running with LE/370 Version 1 Release 1 provides improved performance over VS COBOL II Release 3.2 and 4.0 for:

- Inter-language call with Assembler and C
- Dynamic CALL identifier
- UNSTRING (some additional cases are done in-line)

VS COBOL II Release 3.2 provides improved performance over VS COBOL II Release 3.1 for:

- EVALUATE (when using EVALUATE TRUE or EVALUATE FALSE)
- Passing parameters of large data structures to subprograms

VS COBOL II Release 3.1 provides improved performance over VS COBOL II Release 3.0 for:

- INITIALIZE (for tables with many subordinate items)
- INSPECT, STRING, and UNSTRING (some additional cases are done in-line)

VS COBOL II Release 3.0 provides improved performance over VS COBOL II Release 2 for:

- SSRANGE processing (code is now done in-line instead of through a library routine call, significantly reducing the overhead of subscript range checking)
- Decimal divide
- INSPECT, STRING, and UNSTRING (some of the simple cases of these statements are now done in-line instead of through a library routine call)
- CALL (by reducing some of the overhead)
- INDEX (when comparing INDEXes with constants and when using the SEARCH statement)

Appendix C. Coding Examples

Using CEEENTRY

```
* =====
*   Bring up the LE environment
* =====
CEE2COB  CEEENTRY PPA=MAINPPA,AUTO=WORKSIZE
        USING WORKAREA,13
* =====
*   Set up the parameter list for the COBOL program
* =====
        LA   5,OP1           Get address of 1st parameter
        ST   5,PARM1         and store it in parm list
        LA   5,OP2           Get address of 2nd parameter
        ST   5,PARM2         and store it in parm list
        LA   1,PARMLIST      Load addr of parm list in Reg 1
* =====
*   Call the COBOL program
* =====
        L    15,COBPGM       Get the address of the COBOL program
        BALR 14,15           and branch to it
* =====
*   Terminate the LE environment
* =====
        CEETERM RC=0        Terminate with return code zero
*
* =====
*   Data Constants
* =====
COBPGM  DC   V(COBSUB)       Address of COBOL program
OP1     DC   X'00100C'       1st parameter for COBOL program
OP2     DC   X'00200C'       2nd parameter for COBOL program
*
MAINPPA CEEPPA ,            Constants describing the code block
* =====
*   Workarea
* =====
WORKAREA DSECT
        ORG   *+CEEDSASZ     Leave space for the DSA fixed part
*
PARMLIST DS   0F             Parameter list for COBOL program
PARM1    DS   A             Address of 1st parameter
PARM2    DS   A             Address of 2nd parameter
*
        DS    0D
WORKSIZE EQU  *-WORKAREA
        CEEDSA ,            Mapping of the Dynamic Save Area
        CEECAA ,            Mapping of the Common Anchor Area
*
        END   CEE2COB
```

Using CEELRR

```
LRR2COB CSECT
LRR2COB AMODE 31
LRR2COB RMODE ANY
*
      EXTRN COBSUB
*
* =====
* Save callers regs and chain save areas
* =====
*
      STM 14,12,12(13)      Store incoming registers
      LR  12,15             Base LRR2COB on Register 12
      USING LRR2COB,12
*
      LA  15,SAVEAREA      Get this program's save area
      ST  13,4(,15)        Save caller's save area pointer
      ST  15,8(,13)        Save this program's save area pointer
      LR  13,15            Load standard save area Register 13
*
* =====
* Initialize Library Routine Retention (LRR)
* =====
*
      CEELRR ACTION=INIT
*
* =====
* Set up the parameter list for the COBOL program
* =====
*
      LA  5,OP1             Get address of 1st parameter
      ST  5,PARM1           and store it in parm list
      LA  5,OP2             Get address of 2nd parameter
      ST  5,PARM2           and store it in parm list
      LA  1,PARMLIST        Load addr of parm list in Reg 1
*
* =====
* Call the COBOL program
* =====
*
      L   15,COBPGM         Get the address of the COBOL program
      BALR 14,15            and branch to it
*
* =====
* Terminate Library Routine Retention (LRR)
* =====
*
      CEELRR ACTION=TERM
*
* =====
* Return to our caller
* =====
*
      L   13,4(,13)         Point to incoming registers
      LM  14,12,12(13)      Restore caller's registers
      SR  15,15             Return code 0
      BR  14                Return to caller
*
```



```

* =====
*   Data Constants and Parameter Lists
*   =====
COBPGM  DC    V(COBSUB)           Address of COBOL program
OP1     DC    X'00100C'          1st parameter for COBOL program
OP2     DC    X'00200C'          2nd parameter for COBOL program
*
PARMLIST DS    0F                Parameter list for COBOL program
PARM1   DS    A                  Address of 1st parameter
PARM2   DS    A                  Address of 2nd parameter
*
SAVEAREA DS    18F              Standard Save Area
*
          END    LRR2COB

```

Using IGZERRE

```
*****
*
* The IGZERRE module can be invoked to set up the COBOL Run-time
* Environment running under LE before the first COBOL program is
* called. Invoking IGZERRE will explicitly drive the COBOL
* initialization and termination functions of LE.
*
* LOAD/DELETE of IGZERRE must be done by the user. This load module
* must remain loaded until after the LE run-time environment has
* been terminated (either by IGZERRE termination or a STOP RUN).
*
* When a reusable run-time environment has been created via IGZERRE
* initialization, subsequent use of STOP RUN will result in control
* being returned to the caller of the routine that invoked IGZERRE
* initialization.
*
*****
RRE2COB CSECT
RRE2COB AMODE 31          This routine is 31 bit addressable
RRE2COB RMODE ANY       And can reside above or below the
*                          line
* =====
* Save callers regs and chain save areas
* =====
          STM  14,12,12(13)  Store incoming registers
          LR   12,15         Base RRE2COB on Register 12
          USING RRE2COB,12
*
          LA   15,SAVEAREA   Get this program's save area
          ST   13,4(,15)     Save caller's save area pointer
          ST   15,8(,13)    Save this program's save area pointer
          LR   13,15        Load standard save area Register 13
*
          LOAD EP=IGZERRE   Issue LOAD for Reusable Run-time
*                          Environment INIT/TERM Routine
          ST   0,IGZERREA   Save address for termination
*
*****
* IGZERRE is AMODE(31), RMODE(ANY). The routine that invokes IGZERRE
* must do so via BASSM 14,15, if running on MVS/ESA in 24-bit mode.
* IGZERRE will always return via BSM 0,14, if running on MVS/ESA.
*****
          LA   1,1          Function code for init is "1"
          LTR  15,0         Get address of IGZERRE
          BM  ALTBRCH1     High bit on, so above the line
          BALR 14,15       Go initialize the COBOL environment
          B    TOCHECK     Check return code
ALTBRCH1 BASSM 14,15      Go initialize the COBOL environment
*
*****
* At this point, the user may wish to check the return codes:
* 0 - Function completed correctly
* 4 - LE already initialized (initialization only)
* 8 - Invalid function code (not 1 or 2)
* 16 - LE not initialized (termination only)
*****
```

```

*
TOCHECK LA 14,4 "4" or less is OK
        CR 14,15 Test the return Register 15
        BL ULTIMATE Leave if return higher than "4"
* =====
* Set up the parameter list for the COBOL program
* =====
        LA 5,OP1 Get address of 1st parameter
        ST 5,PARM1 and store it in parm list
        LA 5,OP2 Get address of 2nd parameter
        ST 5,PARM2 and store it in parm list
        LA 1,PARMLIST Load addr of parm list in Reg 1
* =====
* Call the COBOL program
* =====
        L 15,COBPGM Get the address of the COBOL program
        BALR 14,15 and branch to it
* =====
* Terminate the reusable environment
* =====
        L 15,IGZERREA Address of IGZERRE was saved here
        LA 1,2 Function code for term is "2"
        LTR 15,15 IGZERRE might be above the line
        BM PENULTMT If so, go issue BASSM, else
        BALR 14,15 Go terminate the COBOL environment
        B ULTIMATE COBOL environment cleaned up
* (unless return code non-zero)
PENULTMT BASSM 14,15 Go terminate the COBOL environment
*
* =====
* Ready to return to our caller
* =====
*
ULTIMATE DELETE EP=IGZERRE Delete IGZERRE
        L 13,4(,13) Point to incoming registers
        RETURN (14,12),RC=(15) Return to caller
*
* =====
* Data Constants and Parameter Lists
* =====
COBPGM DC V(COBSUB) Address of COBOL program
OP1 DC X'00100C' 1st parameter for COBOL program
OP2 DC X'00200C' 2nd parameter for COBOL program
*
PARMLIST DS 0F Parameter list for COBOL program
PARM1 DS A Address of 1st parameter
PARM2 DS A Address of 2nd parameter
*
SAVEAREA DS 18F Standard Save Area
IGZERREA DC A(0) Address of IGZERRE
*
        END RRE2COB

```

Using CEEPIPI with Call_Sub

```

PIPI2COB CSECT
*****
* Since CEEPIPI must run as AMODE ANY / RMODE 24, this assembler      *
* routine must either:                                               *
* *                                                                     *
* 1 - have the same AMODE / RMODE as CEEPIPI (i.e., ANY / 24), or   *
* *                                                                     *
* 2 - the CALL macro must be removed and coded manually, using the  *
*   appropriate mode switching instructions.                          *
* *                                                                     *
*****
PIPI2COB AMODE ANY
PIPI2COB RMODE 24
*
*       EXTRN COBSUB           COBOL program is external
*
*       STM  14,12,12(13)      Save caller's registers
*       LR   12,15             Get base address
*       USING PIPI2COB,12      Identify base register
*
*       LA   15,SAVEAREA       Get this program's save area
*       ST   13,4(,15)         Save caller's save area pointer
*       ST   15,8(,13)         Save this program's save area pointer
*       LR   13,15             Load standard save area Register 13
*
*       COMPSWT ON             Set flag to load modules (CMS only)
*       LOAD  EP=CEEPIPI       Load CEEPIPI rtn dynamically
*       ST   0,CEEPIPIA        Save the addr of CEEPIPI rtn
*       COMPSWT OFF            Set flag to not load modules (CMS)
*
* =====
* Set up the parameter list for the COBOL program
* =====
*       LA   5,OP1             Get address of 1st parameter
*       ST   5,PARM1           and store it in parm list
*       LA   5,OP2             Get address of 2nd parameter
*       ST   5,PARM2           and store it in parm list
*
* =====
* Initialize a new LE environment
* =====
*       L    15,CEEPIPIA       Get address of CEEPIPI routine
*       CALL (15),(INITSUB,@CEXPTBL,@SRVRTNS,RUNTMOPT,TOKEN)
*
*                               User may want to check return code
*
* =====
* Invoke the COBOL subprogram which is statically linked
* =====
*       L    15,CEEPIPIA       Get address of CEEPIPI routine
*       CALL (15),(CALLSUB,PTBINDEXTOKEN,PARMPTR,SUBRETC,SUBRSNC, X
*         SUBFBC)              Invoke CEEPIPI
*
* =====
* Terminate the LE environment
* =====
*       L    15,CEEPIPIA       Get address of CEEPIPI routine
*       CALL (15),(TERM,TOKEN,ENV_RC) Invoke CEEPIPI
*       DELETE EP=CEEPIPI      Delete CEEPIPI
*

```

```

* =====
* Standard exit code
* =====
SYSRET  L    13,4(,13)      Point to caller's save area
        RETURN (14,12),RC=(15)  Restore regs and return to caller
*
* =====
* Data Areas and Constants
* =====
OP1     DC    X'00100C'
OP2     DC    X'00200C'
*
PARMLIST DS    0F           Parameter list
PARM1   DS    A
PARM2   DS    A
*
SAVEAREA DS    18F
CEEPIPIA DS    A           Save the address of CEEPIPI routine
*
* =====
* Parameter list passed to a CEEPIPI(INIT-SUB)
* =====
INITSUB DC    F'3'         Function code for init subprogram
@CEXPBTL DC    A(PPTBL)    Address of PIPI Table
@SRVRTNS DC    A(0)        No service routines
RUNTMOPT DC    CL255' '    No run-time options
TOKEN   DS    F           Unique value returned
*
* =====
* Parameter list passed to a CEEPIPI(CALL-SUB)
* =====
CALLSUB DC    F'4'         Function code for calling subpgm
PTBINDEXT DC    F'0'       The row number of PIPI Table entry
PARMPTR  DC    A(PARMLIST) Pointer to parameter list
SUBRETC  DS    F           Subroutine return code
SUBRSNC  DS    F           Subroutine reason code
SUBFBC   DS    3F         Subroutine feedback token
*
* =====
* Parameter list passed to a CEEPIPI(TERM)
* =====
TERM     DC    F'5'         Function code for term subprogram
ENV_RC   DS    F           Environment return code
*
* =====
* PIPI Table
* =====
PPTBL    CEEXPIT           PIPI Table with index
        CEEXPITY COBSUB,COBSUB  Statically linked REENTRANT routine
        CEEXPITS
*
        END    PIP12COB

```

Using CEEPIPI with Call_Main

```
PIPM2COB CSECT
*****
* Since CEEPIPI must run as AMODE ANY / RMODE 24, this assembler *
* routine must either: *
* *
* 1 - have the same AMODE / RMODE as CEEPIPI (i.e., ANY / 24), or *
* *
* 2 - the CALL macro must be removed and coded manually, using the *
* appropriate mode switching instructions. *
* *
*****
PIPM2COB AMODE ANY
PIPM2COB RMODE 24
*
*          EXTRN COBSUB          COBOL program is external
*
*          STM  14,12,12(13)     Save caller's registers
*          LR   12,15            Get base address
*          USING PIPM2COB,12     Identify base register
*
*          LA   15,SAVEAREA      Get this program's save area
*          ST   13,4(,15)        Save caller's save area pointer
*          ST   15,8(,13)        Save this program's save area pointer
*          LR   13,15            Load standard save area Register 13
*
*          COMPSWT ON            Set flag to load modules (CMS only)
*          LOAD EP=CEEPIPI       Load CEEPIPI rtn dynamically
*          ST   0,CEEPIPIA       Save the addr of CEEPIPI rtn
*          COMPSWT OFF           Set flag to not load modules (CMS)
*
* =====
* Set up the parameter list for the COBOL program
* =====
*          LA   5,OP1            Get address of 1st parameter
*          ST   5,PARM1          and store it in parm list
*          LA   5,OP2            Get address of 2nd parameter
*          ST   5,PARM2          and store it in parm list
*
* =====
* Initialize a new LE environment
* =====
*          L    15,CEEPIPIA      Get address of CEEPIPI routine
*          CALL (15),(INITMAIN,@CEXPTBL,@SRVRTNS,TOKEN)
*
*          User may want to check return code
*
* =====
* Invoke the COBOL subprogram which is statically linked
* =====
*          L    15,CEEPIPIA      Get address of CEEPIPI routine
*          CALL (15),(CALLMAIN,PTBINDEXTOKEN,RUNTMOPT,PARMPTR,SUBRETC, X
*          SUBRSNC,SUBFBC)      Invoke CEEPIPI
*
* =====
* Terminate the LE environment
* =====
*          L    15,CEEPIPIA      Get address of CEEPIPI routine
*          CALL (15),(TERM,TOKEN,ENV_RC) Invoke CEEPIPI
*          DELETE EP=CEEPIPI     Delete CEEPIPI
*
```

```

* =====
* Standard exit code
* =====
SYSRET  L    13,4(,13)      Point to caller's save area
        RETURN (14,12),RC=(15)  Restore regs and return to caller
*
* =====
* Data Areas and Constants
* =====
OP1     DC    X'00100C'
OP2     DC    X'00200C'
*
PARMLIST DS    0F           Parameter list
PARM1   DS    A
PARM2   DS    A
*
SAVEAREA DS    18F
CEEPIPIA DS    A           Save the address of CEEPIPI routine
*
* =====
* Parameter list passed to a CEEPIPI(INIT-MAIN)
* =====
INITMAIN DC    F'1'         Function code for init main program
@CEXPBTL DC    A(PPTBL)     Address of PIPI Table
@SRVRTNS DC    A(0)         No service routines
RUNTMOPT DC    CL255' '     No run-time options
TOKEN   DS    F            Unique value returned
*
* =====
* Parameter list passed to a CEEPIPI(CALL-MAIN)
* =====
CALLMAIN DC    F'2'         Function code for calling main pgm
PTBINDEXT DC    F'0'         The row number of PIPI Table entry
PARMPTR  DC    A(PARMLIST)  Pointer to parameter list
SUBRETC  DS    F            Subroutine return code
SUBRSNC  DS    F            Subroutine reason code
SUBFBC   DS    3F           Subroutine feedback token
*
* =====
* Parameter list passed to a CEEPIPI(TERM)
* =====
TERM     DC    F'5'         Function code for term subprogram
ENV_RC   DS    F            Environment return code
*
* =====
* PIPI Table
* =====
PPTBL    CEEXPIT           PIPI Table with index
         CEEXPITY COBSUB,COBSUB  Statically linked REENTRANT routine
         CEEXPITS
*
        END    PIPM2COB

```

COBOL Example - COBSUB

```
      CBL  RENT
000100 IDENTIFICATION DIVISION.
000200   PROGRAM-ID. COBSUB.
000300*
000400 ENVIRONMENT DIVISION.
000500*
000600 DATA DIVISION.
000700   WORKING-STORAGE SECTION.
000800*
000900 LINKAGE SECTION.
001000   01 X PIC S9(5) COMP-3.
001100   01 Y PIC S9(5) COMP-3.
001200*
001300 PROCEDURE DIVISION USING X Y.
001400   COMPUTE X = Y + 1.
001500*
001600   GOBACK.
```