



Hewlett-Packard Research & Development Update May 2002

58X speedup in SAS/IML[®] software applications for HP platforms

Statement of Opportunity

SAS Institute's software offerings have emerged from many years of research and development, as well as input from thousands of SAS users. Today, SAS software covers a vast array of application areas including data warehousing, business and electronic intelligence, and Analytical Customer Relationship Management (ACRM); the list continues to grow. In addition to these integrated solution offerings, many SAS users still develop their own custom applications using SAS software as a state-of-the-art 4th generation development environment. Unlike pre-defined integrated SAS applications, there is no a clear way to predict the different ways that customers will utilize SAS software.

However, over the years we have been able to characterize many different types of customer-developed SAS applications and we have come to some basic conclusions. Whether the SAS applications are integrated, like data mining solutions, or the software system is custom developed, we find either application can be very demanding on computer resources. Two of the most important computer resources under pressure are:

- Numerical routines, i.e. SAS procedures (PROCs),
- File system input and output (I/O).

We began joint projects with SAS to understand in more detail the numerically intensive parts of the execution of SAS. We started with profiling real-world customer environments, and tracking the execution paths. Our findings led us to recommend changes to the compiler flags being used to build some of the algorithms. This work is discussed in an earlier HP/SAS R&D update, "Optimizing SAS Software Numerical Applications for HP Platforms", April 2001. This update discusses our initial investigation into SAS numerical routines, with a result of between 20% and 40% performance gain in SAS application execution. These improvements are included in the SAS 8.2 release (both 32 and 64 bit versions).

We proceeded with the investigations to examine further potential performance gains on HP servers. We asked the question: *What can we do to further improve the performance of numerical procedures in Version 9 SAS software on HP-UX servers and workstations?* In this report, we present the results of recent work done jointly by the researchers in SAS' Open Systems Research and Development (OSRD) department and HP High Performance Technical Computing Division to optimize performance for SAS computational procedures on HP systems.

Real World Applications

One of the fundamental data elements in SAS/IML is the matrix, a two-dimensional (row-by-column) array of numeric or character values.¹ The benchmark SAS provided utilizes the SAS/IML software, working with 300x300 matrices. HP's Mathematical Software Product, MLIB, contains BLAS (Basic Linear Algebra Subprograms) that directly benefits the performance of SAS/IML. More specifically, MLIB's Level 3 BLAS addresses matrix – matrix operations. MLIB's implementation of BLAS 3 routines is a combination approach of algorithms and assembly language. The algorithm manages cache behavior while the assembly code exploits the parallel instruction set offered by PA-RISC assembly language on HP machines. With help from SAS researchers, we were able to link MLIB's highly optimized version of BLAS 3 into the benchmark. The results were outstanding and discussed in the following section.

Results

Working together, the HP and SAS R&D teams have improved the performance of matrix multiply in SAS/IML significantly - over 58X!! The performance gains over the original SAS/IML on HP-UX are summarized in the following table:

SAS/IML Application Performance Gains

SAS/IML for HP-UX	HP Server N-Class
Original	1
Modified	58.55X

Gains of over 58X were made on the 8-way N-Class midrange server in SAS/IML testing conducted at the HP High Performance Technical Computing Division (HPTCD).

These new subroutines are currently being tested for inclusion into SAS Version 9 for HP-UX.

Next Steps

MLIB contains highly tuned mathematical subprograms to support HP-UX compilers, PA architectures, and Itanium[®] architecture. We have shown the benefits of SAS software developers building SAS software products by linking with MLIB. The use of MLIB will allow these developers to boost their software performance on HP platforms with little additional effort and leverage future HP systems based on both PA and Itanium architectures.

HP continues to work with SAS R&D on the Intel Processor Family (IPF) architecture, co-invented by HP and Intel. IPF offers unique new features that can be utilized to improve performance. These features include architecture resources, predication, register rotation, speculation, and larger integer and floating point units. MLIB has been implemented a several of SAS' VSUBS routines and the results are promising. We look forward to identifying other algorithms and working with SAS' R&D to continue bringing the best solutions to our customers.

To Learn More

For a free trial version of MLIB and additional information please visit <http://www.hp.com/go/mlib>. In addition to BLAS, there are other components of MLIB, like Sparse Linear Algebra, FFT's, Convolutions, LAPACK V3.0 and ScaLAPCK that may improve your mathematical programs. If you have any questions, comments or concerns please email the Mathematical Software Team at mllibcore@rsn.hp.com.

We invite you to contact us to learn more about this project and the other activities jointly being done by SAS and HP that can benefit you. We can be reached on the web at: <http://www.sas.com/partners/directory/hp>.

###

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. HP and all other HP product or service names are registered trademarks or trademarks of Hewlett-Packard Company.

Technical information contained in this document is subject to change without notice.

© Copyright Hewlett-Packard Company 2002.

All Rights Reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

ⁱ <http://www.sas.com/rnd/app/da/iml.html>