Software Group | Compilation Technology

IBM pSeries Compiler Roadmap

Roch Archambault
IBM Toronto Laboratory
archie@ca.ibm.com

Agenda

- The pSeries Compiler Products
- Roadmaps
  - XL Fortran
  - VisualAge C++
  - Performance
- Multiple compiler installation
- Performance Results
- Q&A

The pSeries Compiler Products

- Latest Versions (All POWER4 enabled)
  - VisualAge C++ Version 6.0 for AIX
  - VisualAge C++ Version 6.0 for Linux on pSeries
  - XL C/C++ Advanced Edition Version 6.0 for Mac OS X (PPC970 enabled)
  - C for AIX V6.0
  - XL Fortran Version 8.1 for AIX (latest level is version 8.1.1)
  - XL Fortran Version 8.1 for Linux on pSeries (latest level is version 8.1.1)
  - XL Fortran Advanced Edition Version 8.1 for Mac OS X (PPC970 enabled)
- Upcoming Versions (All POWER4, POWER5 and PPC970 enabled)
  - XL C/C++ Enterprise Edition Version 7.0 for AIX, Linux
  - XL C/C++ Advanced Edition Version 7.0 for Mac OS X (PPC970 only)
  - XL Fortran Enterprise Edition Version 9.1 for AIX, Linux
  - XL Fortran Advanced Edition Version 9.1 for Mac OS X (PPC970 only)

XL Fortran Roadmap: Strategic Priorities

- Premium Customer Service
  - Continue to work closely with key ISVs and customers in scientific and technical computing industries
- Compliance to Language Standards and Industry Specifications
  - OpenMP Fortran API V2.0
  - Fortran 77, 90 and 95 standards
  - Emerging 2003 Standard
- Exploitation of Hardware
  - Committed to maximum performance on POWER4, PPC970, POWER5 and successors
  - Continue to work very closely with processor design teams

XL Fortran Version 8.1 for AIX

- **Fully compliant Fortran 77/90/95 compiler**
- **Partial Fortran 2003 support**
  - Allocatable components, IEEE module, INTENT of F90 pointers
- **Many industry extensions**
  - 128-bit float, 64-bit integer, Cray pointers, structure record, union map, BYTE, STATIC, AUTOMATIC, asynchronous I/O, SIZE intrinsic
- **32- and 64-bit support**
- **Optimized OpenMP Fortran V2.0 support**
- **Symbolic Debugging support**
  - TotalView, IBM Distributed Debugger (Tech Preview) and dbx/pdbx
  - Full support for debugging of OpenMP programs
  - Partial support for debugging of optimized code
- **Portfolio of Optimizing Transformations**
  - Comprehensive path length reduction
  - Whole program analysis
  - Loop optimization for parallelism, locality and instruction scheduling
  - Tuned support for all pSeries processors (including POWER4)
- **More info: www.ibm.com/software/ad/fortran**

Software Group | Compilation Technology

IBM

XL Fortran Version 8.1 for Linux on pSeries

- Based on XL Fortran Version 8.1 for AIX Product
  - Leverage proven industry leading performance compiler technology
- Fully compliant Fortran 77/90/95 compiler
- Partial Fortran 2003 support
  - Allocatable components, IEEE module, INTENT on F90 pointers
- Many industry extensions
  - 128-bit float, 64-bit integer, Cray pointers, structure record, union map, BYTE, STATIC, AUTOMATIC, asynchronous I/O, SIZE intrinsic
- 32- and 64-bit support
- Optimized OpenMP Fortran V2.0 support
- Symbolic Debugging support with gdb
- Portfolio of Optimizing Transformations
  - Comprehensive path length reduction
  - Whole program analysis
  - Loop optimization for parallelism, locality and instruction scheduling
  - Tune support for all pSeries processors (including POWER4)
- Supports SUSE Linux Enterprise Server 8 (SLES 8)
- Supports RedHat Enterprise Linux 3.0 (RHEL 3)
- Support for TurboLinux (TLES 8) and Conectiva

SciComp9 | IBM pSeries Compiler Roadmap

© 2004 IBM Corporation

XL Fortran V8.1.1 for AIX GA 06/03

- Fortran 2003
  - Stream I/O
  - VALUE and PROTECTED attributes and statements
- Performance
  - Customer benchmarks, SPECCPU2000 improvements, scalability for SPECOMP
  - 64-bit TPO enablement (compiler component)
  - Stream_unroll and Unroll_and_fuse directives
  - Unroll directive extended to outer loops
  - Intrinsic Performance improvements
- Customer Requirements
  - -qsuppress suboption
  - -qextname suboption
- Improved Debug support with TotalView
- Installation and operation on OS/400 PASE on iSeries

Beyond XL Fortran V8.1.1

- XL Fortran Enterprise Edition V9.1 for AIX (2004)
  XL Fortran Enterprise Edition V9.1 for Linux (2004)
  XL Fortran Advanced Edition V9.1 for Mac OS X (2004)
  - Continued rollout of Fortran 2003
  - POWER5 and PPC970 support
  - Improved performance
  - Improved compile time and reduce space usage
  - Improved debugging
  - New directives and intrinsics
  - New customer requirements

All information subject to change without notice

Tentative rollout of Fortran 2003 (XLF V9.1)

- XL Fortran V9.1 (2004)
  - BIND attribute and statement, BIND on derived types, common blocks, and procedures
  - PROCEDURE, IMPORT and FLUSH statements
  - PUBLIC/PRIVATE on components of a type
  - ISO_FORTRAN_ENV and ISO_C_BINDING intrinsic modules
  - ASSOCIATE Construct
  - Command Line Argument and NEW_LINE intrinsics
  - Intrinsic attribute on USE
  - IOMSG= specifier

All information subject to change without notice

Prioritized rollout of Fortran 2003 beyond 9.1

1. OO extensions
2. ENUM/ENUMERATOR
3. Derived type I/O
4. Asynchronous I/O additions
5. Allow character on MAX/MIN family of intrinsics
6. Procedure pointers
7. I/O specifiers (DECIMAL=, SIGN=, ROUND=, and PAD=)
8. Complex literal
9. Derived type parameters
10. Abstract interface
11. Enhance structure constructors
12. Enhanced array constructors
13. Pointer assignment enhancement (specify bounds)
14. Support for International Usage

All information subject to change without notice

C/C++ Roadmap: Strategic Priorities

- **Premium Customer Service**
- **Compliance to Language Standards and Industry Specifications**
  - ANSI / ISO C and C++ Standards
  - OpenMP C/C++ API V2.0
- **Exploitation of Hardware**
  - Committed to maximum performance on POWER4, PPC970, POWER5 and successors
  - Continue to work very closely with processor design teams
- **Exploitation of OS and Middleware**
  - Synergies with operating system and middleware ISVs (performance, specialized function)
  - Committed to AIX Linux affinity strategy and to Linux on pSeries
- **Reduced Emphasis on proprietary Tooling**
  - Affinity with GNU toolchain

VisualAge C++ Version 6.0 for AIX
C for AIX Version 6.0

- **Fully compliant ISO C 1999 and ISO C++ 1998 standards**
  - AIX 5.2 required for ISO C 1999 *complex* runtime support
- **Partial GNU C/C++ language and options compatibility**
- **32- and 64-bit support**
- **Interprocedural Analysis (IPA) including specialized C++ optimizations for templates, exception handling and virtual functions**
- **Automatic Parallelization for C/C++**
- **Optimized OpenMP C/C++ API V1.0 support**
- **Template instantiation improvements**
  - Reuse of template instantiations to reduce code size
  - Parsing templates at instantiation time
  - Improved automatic template instantiations
- **Symbolic Debugging Support**
  - TotalView, IBM Distributed Debugger and dbx/pdbx
  - Full support for debugging of OpenMP programs
  - Partial support for debugging of optimized code
  - Runtime memory debug support

VisualAge C++ Version 6.0 for AIX (continued)
C for AIX Version 6.0 (continued)

- Portfolio of optimizing transformations
  - Comprehensive path length reduction
  - Whole program analysis
  - Loop optimization for parallelism, locality and instruction scheduling
  - Tuned support for all pSeries processors (including POWER4)
- More info:
  - **www.ibm.com/software/ad/vacpp**
  - **www.ibm.com/software/ad/caix**

VisualAge C++ Version 6.0 for Linux on pSeries

- Based on VisualAge C++ Version 6.0 for AIX Product
  - Leverage proven industry leading performance compiler technology
- Fully compliant ISO C 1999 and ISO C++ 1998 standards
- Partial GNU C/C++ language and options compatibility
- 32- and 64-bit support
- Automatic Parallelization for C/C++
- Optimized OpenMP C/C++ API V1.0 support
- Portfolio of Optimizing Transformations
  - Comprehensive path length reduction
  - Whole program analysis
  - Loop optimization for parallelism, locality and instruction scheduling
  - Tune support for all pSeries processors (including POWER4)
- Supports SUSE Linux Enterprise Server 8 (SLES 8)
- Supports RedHat Enterprise Linux 3.0 (RHEL 3)
- Support for TurboLinux (TLES 3) and Conectiva

Beyond VisualAge C++ V6.0

- **XL C/C++ Enterprise Edition Version 7.0 for AIX (2004)**
  XL C/C++ Enterprise Edition Version 7.0 for Linux (2004)
  XL C/C++ Advanced Edition Version 7.0 for Mac OS X (2004)
  - Addenda to ISO C/C++ standards
  - Further GNU C/C++ compatibility features
  - OpenMP C/C++ API V2.0
  - POWER5 and PPC970 support
  - Improved performance
  - 64-bit TPO enablement (compiler component)
  - Boost and STLport libraries (as they become better accepted by the standard committee)
  - IOStream Performance improvements
  - New pragmas / options

All information subject to change without notice

Tentative GNU C/C++ Compatibility Enhancements

- **Labels as values / computed goto**
- **Nested functions**
- **Naming types**
- **Conditionals with omitted operands**
- **Complex**
- **Hex floats**
- **Zero length arrays**
- **Variable length arrays**
- **Compound literal**
  - Constructor expressions
- **Labeled elements (C only)**
- **Case ranges**
- **Cast to union (C only)**
- **Function Attributes**
  - Support
    - Noinline, always_inline, format, format_arg, section
  - Accept and ignore
    - used

All information subject to change without notice

Tentative GNU C/C++ Compatibility Enhancements (Continued)

- Character <esc> in char literal
- Variable Attributes
  - Support
    - Nocommon, transparent_union
- Type Attributes
  - Support
    - Aligned, packed
  - Accept and ignore
    - Transparent_union
- __extension__
- Incomplete enums
- Function names as strings
- Partial Asm support

All information subject to change without notice

Performance Improvements Delivered in 2003

- **Included in XLF V8.1.1 release**
- **OpenMP**
  - **Much faster barrier implementation (over 5x faster on 32-way p690)**
  - **Much faster uncontended atomic (over 4x faster on 32-way p690)**
  - **Improved parallel region startup (25%)**
- **POWER4**
  - **Expanded scope and precision of instruction scheduling**
  - **More precise loop unrolling for pipelining and parallelization of reductions**
  - **Use of dcbz to optimize store streams**
  - **Loop unrolling for prefetch stream utilization**
- **Loop Optimization**
  - **Aggressive loop fusion to exploit data reuse**
  - **Index set splitting and gather/scatter to move branches out of loops**
  - **Loop peeling for automatic parallelization**
  - **Improved loop interchange for automatic parallelization**
  - **Strip-mining of vectorized loops**
  - **Vector calls to specialized routines for POWER4 (vsqrt, vrsqrt, vtan, vlog, vexp)**

Expected Performance Improvements in 2004

- **POWER5**
  - Modified scheduling machine model
  - Usage of improved prefetch facilities
  - Usage of new instructions
- **PPC970**
  - Automatic generation of VMX code on Linux and Mac OS X (SIMD vectorization)
  - Interprocedural pointer alignment propagation
- **OpenMP**
  - Tuned support for 64-way SMP
  - Continued improvements in overhead
  - Cache-conscious loop scheduling
- **Intrinsic functions**
  - MATMUL, CSHIFT, RANDOM_NUMBER, TRANSFER, SHAPE, INDEX, SIZE, LBOUND, UBOUND, complex transcendentals

All information subject to change without notice

Software Group | Compilation Technology

IBM

Expected Performance Improvements in 2004

- **Loop Optimization**
  - Modulo scheduling of loops which contain branches
  - Further improvements to loop fusion for data reuse (e.g. loop alignment)
  - Perform vectorization on all platforms (including Linux and Mac OS X)
  - Enhancement of vectorization (additional functions, loop versioning, vector merging)
  - Tiling for BLAS-like and streaming loop nests
  - Predictive Commoning (common subexpression elimination across loop iterations)
  - Improved data dependence analysis
- **Tuning assists**
  - BLOCK_LOOP and LOOPID directives to specify which set of loops to tile, interchange or stripmine
  - NOVECTOR and NOSIMD directive to tell compiler not to vectorize or simdize loop
  - Builtin functions for generating software divides (full double precision on POWER5)
  - Thread binding (set via XLSMPOPTS env variable)
  - Environment variable to control number of threads used by MATMUL and RANDOM_NUMBER
  - View and manipulate information gathered by profile directed feedback (-qpdf1/-qpdf2) (showpdf and mergepdf tools)
  - New prefetch directives for POWER5

All information subject to change without notice

SciComp9 | IBM pSeries Compiler Roadmap

© 2004 IBM Corporation

Installation of Multiple Compiler Versions

- Installation of multiple compiler versions is supported
- The vacppndi and xlfndi scripts shipped with VisualAge C++ 6.0 and XL Fortran 8.1 allow the installation of a given compiler release or update into a non-default directory
- The configuration file can be used to direct compilation to a specific version of the compiler

  Example: xlf_v8r1 –c foo.f

  May direct compilation to use components in a non-default directory
- Care must be taken when multiple runtimes are installed on the same machine (details on next slide)

Coexistence of Multiple Compiler Runtimes

- **Backward compatibility**

  C, C++ and Fortran runtimes support backward compatibility.

  Executables generated by an earlier release of a compiler will work with a later version of the run-time environment.
- **Concurrent installation**

  Multiple versions of a compiler and run-time environment can be installed on the same machine

  Full support in xlfndi and vacppndi scripts is now available
- **Limited support for coexistence**

  LIBPATH must be used to ensure that a compatible runtime version is used with a given executable

  Only one runtime version can be used in a given process.

  Renaming a compiler library is not allowed.

  Take care in statically linking compiler libraries or in the use of *dlopen* or *load* .

  Details in the compiler FAQ (http://www.ibm.com/software/awdtools/fortran/xlfortran/support/)

Documentation

- An information center containing the documentation for the current versions of the AIX compilers (both Fortran and C/C++) is available at: **http://publib.boulder.ibm.com/infocenter/comphelp/index.jsp**
- This information center contains all the html documentation shipped with the compilers. It is completely searchable.
- We are planning to make similar information centers available for our Linux and Mac OS X compilers.
- Please send any comments or suggestions on this information center or about the existing C, C++ or Fortran documentation shipped with the products to compinfo@ca.ibm.com.

32-way EPCC results on AIX 5.2 p690 system (1.1 GHZ)

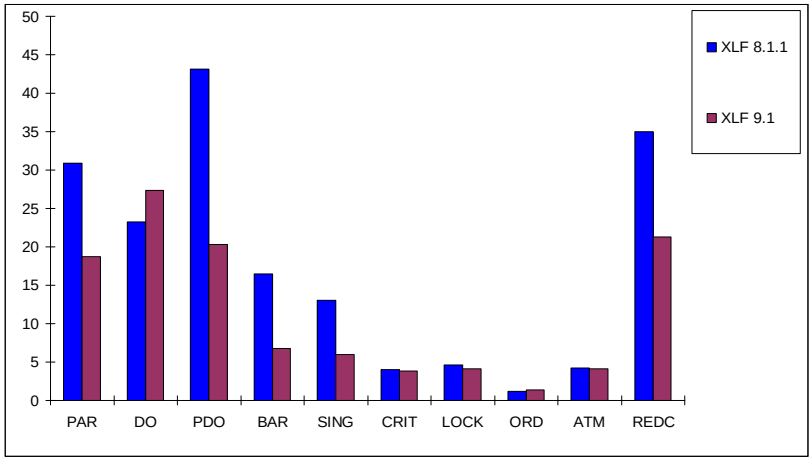

All information subject to change without notice

RANDOM_NUMBER Improvements (REAL*4)

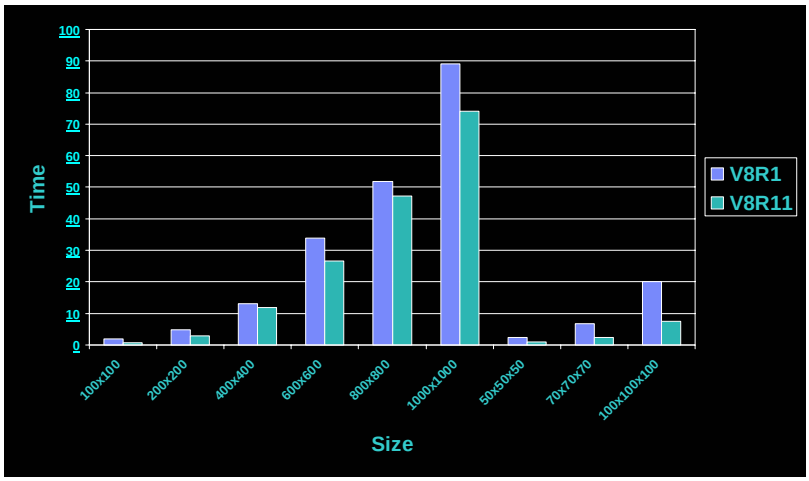

MATMUL Improvements (INT*4, REAL*4, COMPLEX*4)

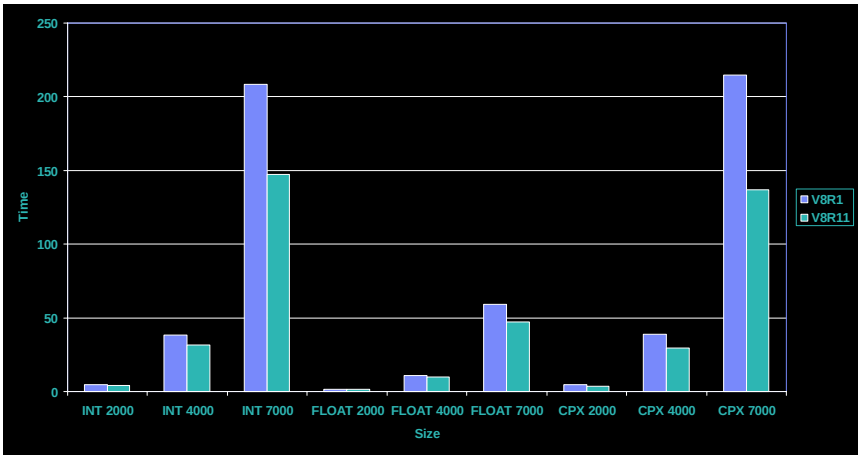

Gcc vs XLC Performance Comparison (POWER4)

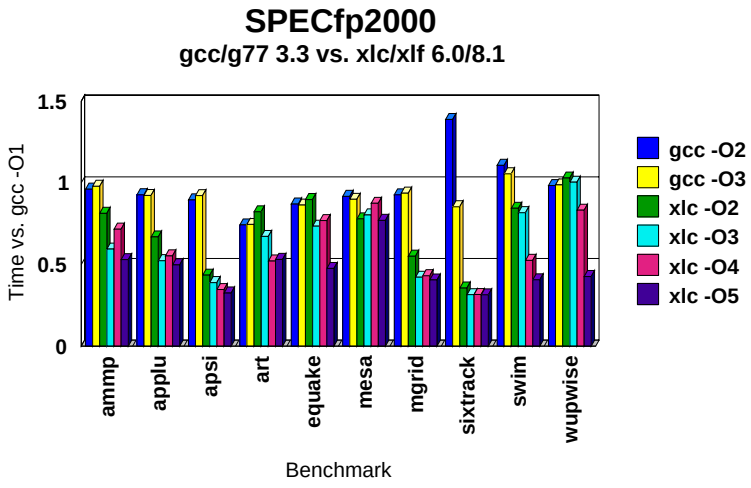

Gcc vs XLC Performance Comparison (POWER4)