Intel® Math Kernel Library (Intel® MKL)

Software and Solutions Group

www.intel.com/software/products

intel Software Products

History - Intel Math Kernel Library

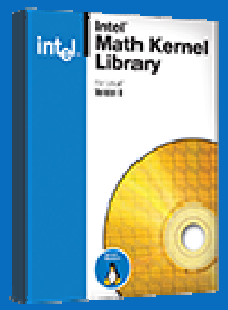

Library started ~1994 – called blas.lib

Mathematically skilled staff – PhD's from:

- Berkeley, Cornell, Cambridge, Russian Universities

Staff includes:

- Professors from NSU*
- Members of Russian Academy of Science
- Intact team from early 90's

Geographies

- Portland, Oregon
- Nizhniy Novgorod, Novosibirsk, Sarov, Russia

Always coded in Russia

*Novosibirsk State University

Goals for Intel MKL

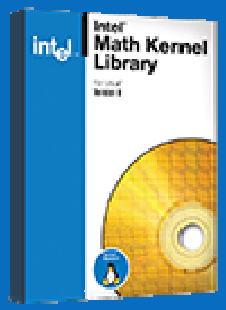

- Outstanding Performance on Intel processors
  - Intel Core® 2, Pentium® 4, Pentium III, Xeon® processors
  - Intel Itanium® processor
- Parity with best software on Intel compatible architectures
- Address computationally intensive SW
- Make software easy to use
  - Compiler independence
  - F95 Interfaces to BLAS, LAPACK
  - FFTW Interface to Intel MKL FFTs, and DFTi
- Parallelism everywhere it makes sense
  - Multi-core processors, distributed memory clusters
- Do not add software without performance!

Intel® Math Kernel Library 10.0

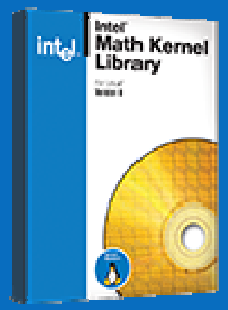

Description

- A highly optimized math library for scientific, engineering, financial and energy applications

Function Domains

- Linear Algebra: BLAS, LAPACK, ScaLAPACK
- Linear Algebra: Sparse Solvers
- Fast Fourier Transforms (SMP and DM)
- Vector Math Library
- Vector Statistical Library
  - Random Number Generators
- Poisson Solver
- ODE Solver (WhatIf)

- ☑ Windows
- ☑ Linux
- ☑ Mac OS

- ☑ Itanium® 2
- ☑ Xeon®
- ☑ Core® 2 Duo processors

Intel® Math Kernel Library 10.0

What's New in Version 10.0?

- Optimized for Penryn
- Mac OS X Leopard support
- Debian* and Ubuntu* Linux distributions support
- Cluster functionality added to base Intel MKL product
  - ScaLAPACK & Distributed Memory FFTs
  - Multi-core is driving developers to adopt parallelism, which increases interest in the already growing use of distributed (cluster) parallel computing.
- Re-architected for multiple compiler and interface support in single package.
  - Interface Layer
    - LP64 – Long Pointer 64
    - ILP64 - Integer Long Pointer 64
  - Threading layer
    - Intel Compiler, Microsoft Compiler, GCC compiler
    - Serial version
- PARDISO* Sparse Solver Out-Of-Core memory
  - Enables solving very large problems on large HW, or medium-large problems on desktop systems
  - More flexible solution than buying memory, cost-effective
- New Vector Math Enhanced Performance mode
  - 3 levels of accuracy to provide maximum performance for specific customer needs.
  - Compiler vectorization defaults to Low Accuracy.

"We achieved a 300-500% performance gain using the Intel Math Kernel Library PARDISO Solver with our NISA Finite Element Analysis tool during our trials on multi-core machines.

Ramdass Keshavamurthy
VP Development, Cranes Software

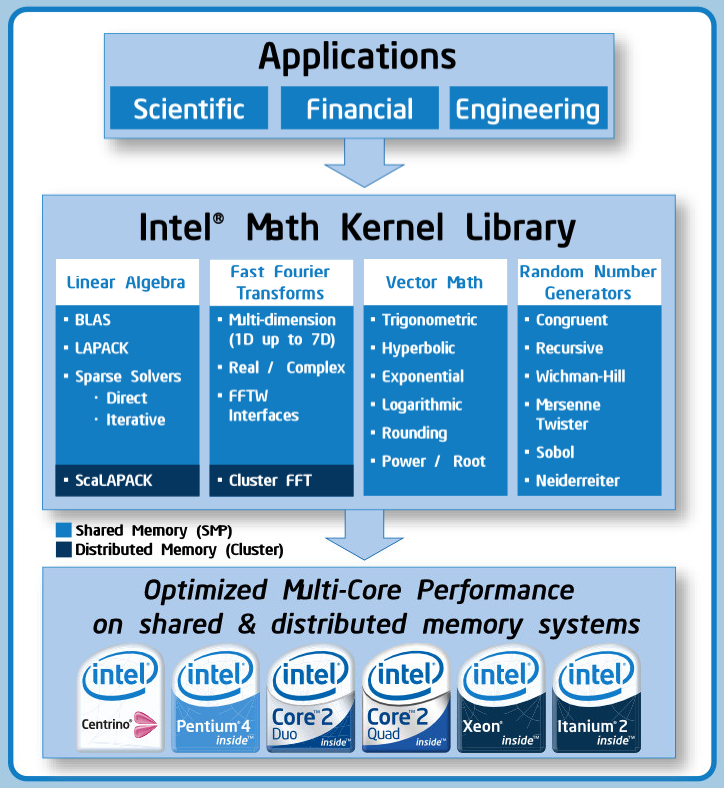

BLAS & LAPACK *Performance*

The industry standard for Linear Algebra Problems

BLAS/LAPACK
ScaLAPACK
Sparse Solver
FFTs
VML
VSL

Intel® Xeon® Processor

- Outstanding scaling performance for 2, 4, and 8 threads
- Intel MKL can deliver 100% or more performance than ATLAS*

Intel® Itanium® 2 Processor

- Outstanding scaling performance for 2, 4, and 8 threads
- Intel MKL can deliver tremendous performance benefits over ATLAS*

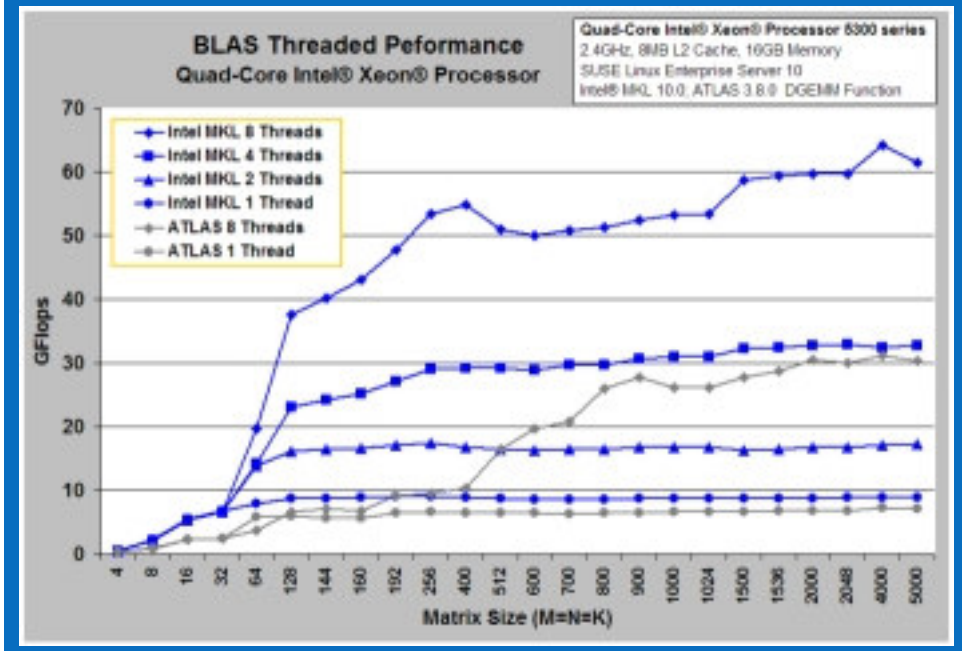

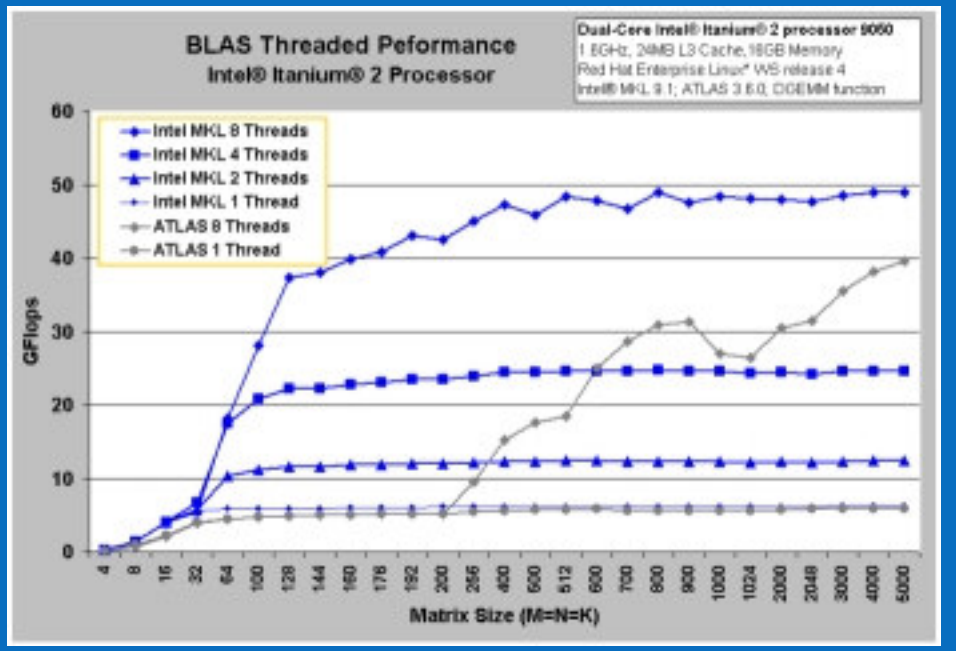

Excellent Multi-processor Performance

ScaLAPACK – *Performance*

Increasing problem size proportionally with increase in cluster size

Key Points

- Intel® MKL has excellent scaling
- Intel® MKL is significantly faster than Netlib
  - 40% faster than Netlib using Intel MKL BLAS
  - 100% faster than Netlib using ATLAS BLAS

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/software/products or call (U.S.) 1-800-628-8686 or 1-916-356-3104

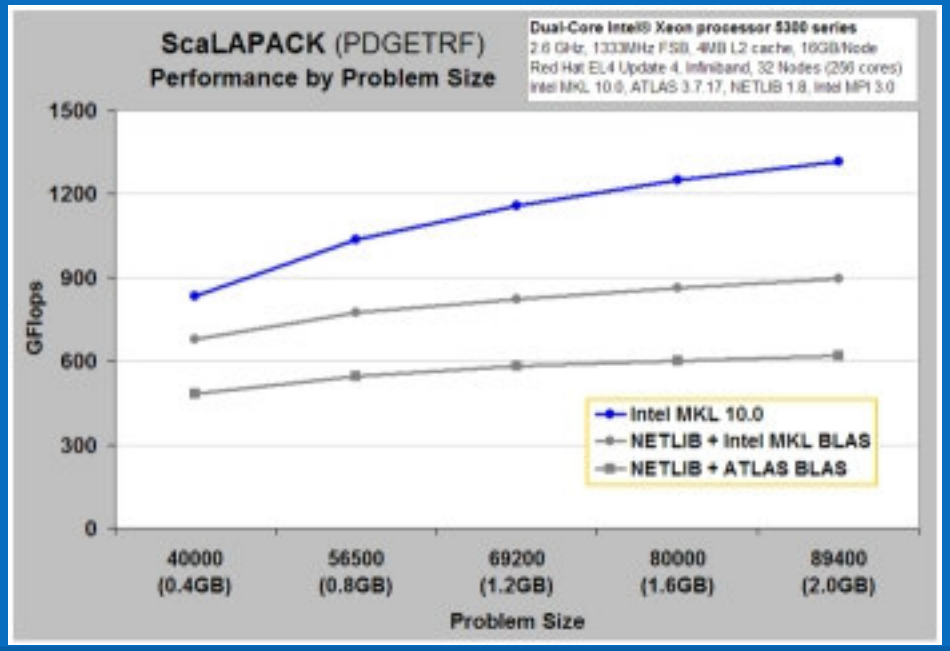

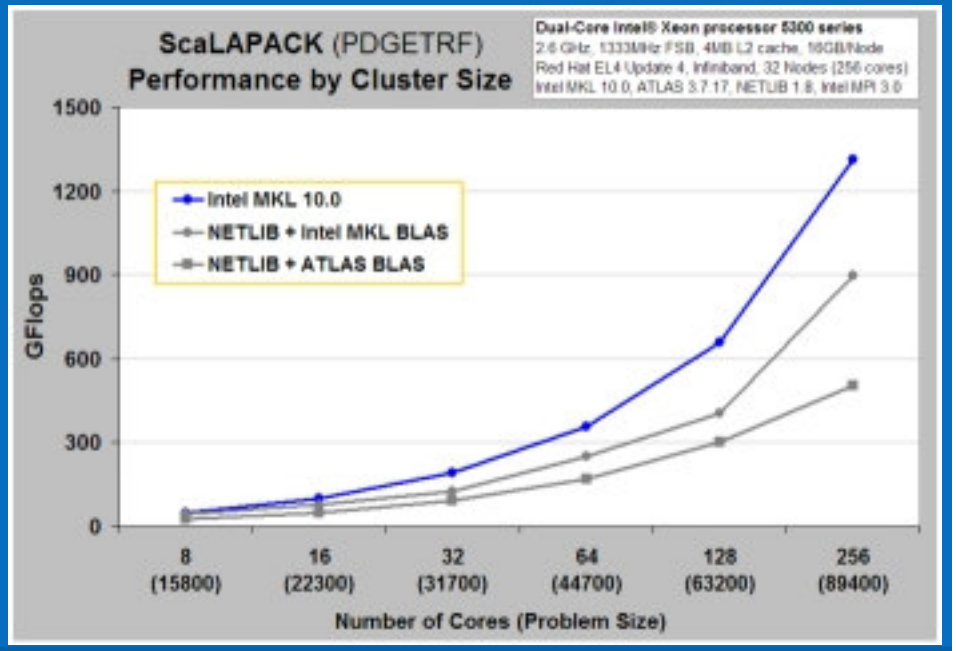

Fast Fourier Transform *Performance*

Threading Optimizations

- Intel MKL FFTs are extensively threaded to take advantage of multiprocessors

Intel® Xeon® Processor

- Intel MKL single thread performance is often greater than FFTW* 2 and 4 thread performance
- Intel MKL can deliver >100% gain

Intel® Itanium® 2 Processor

- Outstanding scaling performance for 2, 4, and 8 threads
- Intel MKL can deliver significant performance benefits over FFTW*

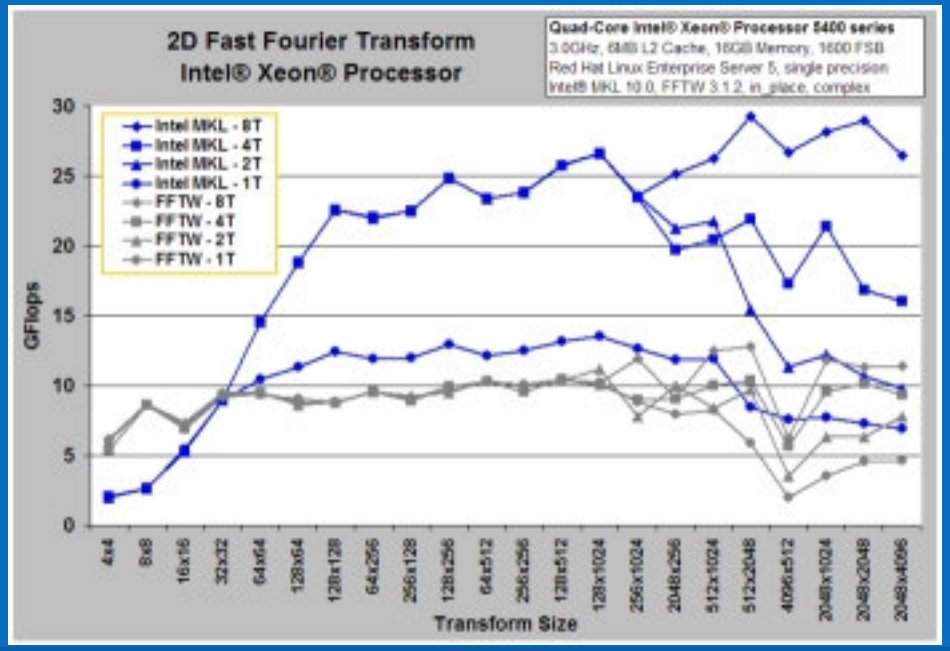

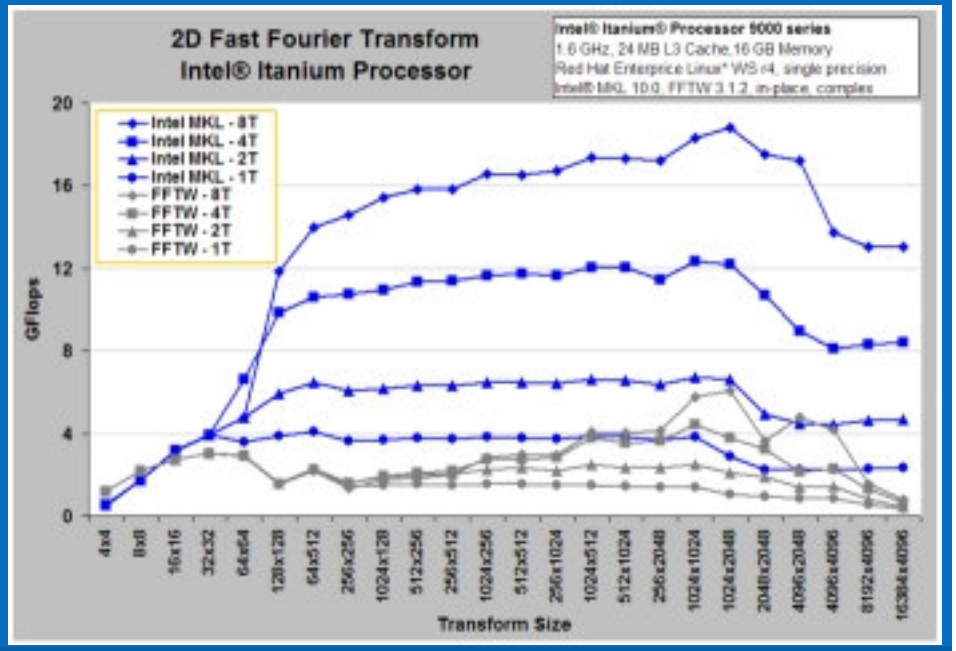

Vector Math Library (VML)

Highly optimized vector implementations of computationally intensive core mathematical functions

Arithmetic	Trigonometric	Hyperbolic	Power/Root	Rounding	Other
Add	Sin	Sinh	Pow	Floor	Ln
Sub	Cos	Cosh	Powx	Ceil	Log10
Div	SinCos	Tanh	Pow2o3	Round	Log1p
Sqr	CIS	Asinh	Pow3o2	Trunc	Exp
Mul	Tan	Acosh	Sqrt	Rint	Expm1
Conj	Asin	Atanh	Cbrt	NearbyInt	Erf
MulByConj	Acos		InvSqrt	Modf	Erfc
Abs	Atan		InvCbrt		ErInv
Inv	Atan2		Hypot		

- Real and Complex Data
- Single and Double Precision
- Threaded
- 3 levels of accuracy
  - (HA) High Accuracy* ~53 bits accurate
  - (LA) Low Accuracy ~51 bits accurate
  - (EP) Enhanced Performance ≥ 26 bits accurate

Performance

- Intel MKL Vector Math Library can provide tremendous speedup.
- No need to rely on compiler to vectorize
- Multiple Accuracy modes so developers can optimize performance per needed accuracy

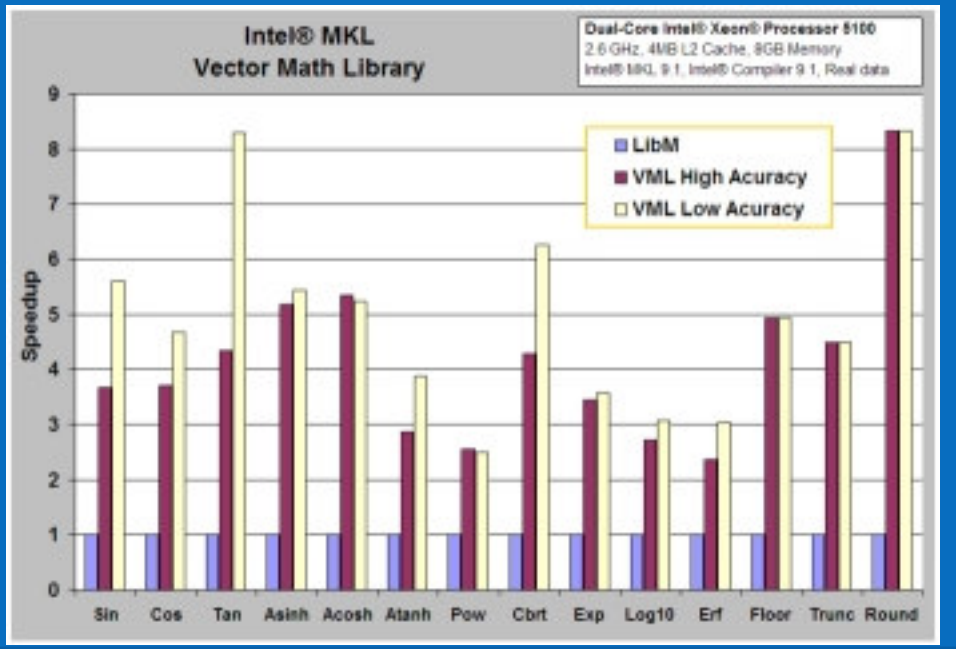

* 99% results correctly rounded to the last bit

Vector Math Library

Different Accuracy for Different Problems

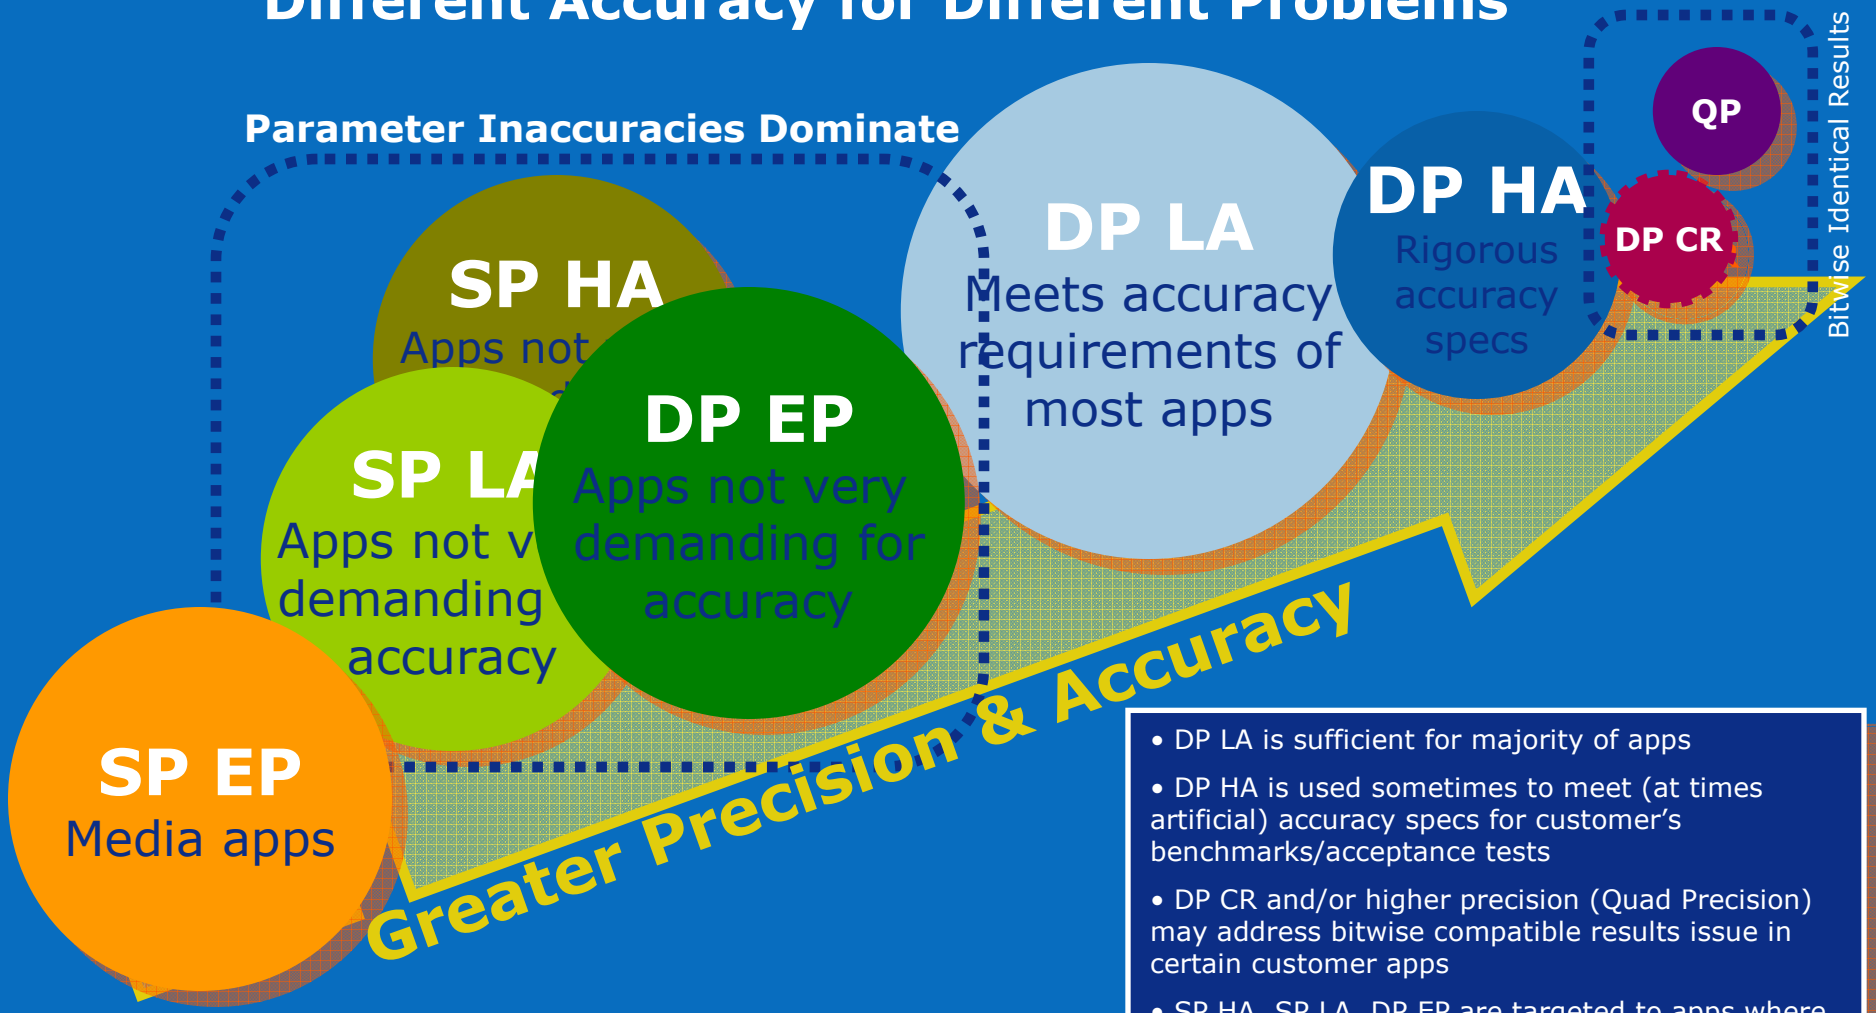

- DP LA is sufficient for majority of apps
- DP HA is used sometimes to meet (at times artificial) accuracy specs for customer's benchmarks/acceptance tests
- DP CR and/or higher precision (Quad Precision) may address bitwise compatible results issue in certain customer apps
- SP HA, SP LA, DP EP are targeted to apps where math function inaccuracies don't dominate parameter inaccuracies, e.g. Monte Carlo simulations
- SP EP is targeted to class of media/graphics apps

Vector Statistical Library (VSL)

BLAS/LAPACK
ScaLAPACK
Sparse Solver
FFTs
VML
VSL

A collection of Random Number Generators with a large set of probability distributions

Performance

- Thread Safe ⇨ Call in Parallel
- Excellent multi-core scaling

Random Number Generators
Pseudo-random
Multiplicative Congruential 59-bit
Multiplicative Congruential 31-bit
Multiple Recursive
Feedback Shift Register
Wichman-Hill
Mersenne Twister 2203
Mersenne Twister 19937
Quasi-random
Sobol
Niederreiter

Intel® Xeon® Processor 5300 series	Running Time (seconds)	Speedup vs. rand() (times)
Standard C rand() function	**40.52**	**1**
Intel® MKL VSL random number generator	**6.88**	**5.89**
OpenMP* version (8 threads)	**0.92**	**44.04**

Distribution Generators	
Continuous	Discrete
Uniform	Uniform
Gaussian	UniformBits
GaussianMV	Bernoulli
Exponential	Geometric
Laplace	Binomial
Weibull	Hypergeometric
Cauchy	Poisson
Rayleigh	PoissonV
Lognormal	Negative bionomial
Grumbel	
Gamma	
Beta	

Configuration Info:
- **Quad-Core Intel® Xeon® processor 5300 series**
- **2.4 GHz, 2x8MB L2 cache, 4 GB memory**
- **Windows Server* 2003 Enterprise x64 Edition**
- **Test run on a vector of 1000 elements**
- **Intel MKL 10.0 and Intel® C++ Compiler 10.1**

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approx performance of Intel products as measured by those tests. A difference in system hardware or software design or configur may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systen components they are considering purchasing. For more info on performance tests and on the performance of Intel produc reference www.intel.com/software/products or call (U.S.) 1-8 8686 or 1-916-356-3104

Vector Statistical Library (VSL)

Performance vs. NAG & VNI IMSL

- Intel MKL Random Number Generators can provide significant performance boost over alternatives.

Library	Random Number Generator	Option Value (Exact Value)		Absolute Error (Standard Error)		Time (seconds)	Speedup (times faster)
		Call	Put	Call	Put		
Intel MKL	MCG31m1	16.7306 (16.7341)	7.2177 (7.2179)	0.0036 (0.0019)	0.0002 (0.0009)	4.67	**8.78**
	MCG59	16.7364 (16.7341)	7.2162 (7.2179)	0.0023 (0.0019)	0.0017 (0.0009)	4.86	**8.44**
	MT19937	16.7349 (16.7341)	7.2164 (7.2179)	0.0007 (0.0019)	0.0015 (0.0009)	5.08	**8.08**
NAG	Original	16.7339 (16.7341)	7.2182 (7.2179)	0.0002 (0.0019)	0.0003 (0.0009)	11.45	**3.58**
IMSL	MT19937	16.7324 (16.7341)	7.2178 (7.2179)	0.0017 (0.0019)	0.0001 (0.0009)	35.71	**1.15**
	Minimal Standard	16.7343 (16.7341)	7.217 (7.2179)	0.0001 (0.0019)	0.0009 (0.0009)	41.03	**1**

Configuration Info:

- **Quad-Core Intel® Xeon® processor 5300 series**
- **2.4 GHz, 2x8MB L2 cache, 4 GB memory**
- **Windows Server* 2003 Enterprise x64 Edition**
- **Test run on a vector of 1000 elements**
- **Intel MKL 10.0, IMSL F90 6.0, NAG SMP (mark 21)**

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/software/products or call (U.S.) 1-800-628-8686 or 1-916-356-3104

Intel® MKL vs. ATLAS

Performance on Intel® Xeon™ processor

Small Matrices: Intel® MKL ~100% faster

Large Matrices: Intel® MKL ~25-50% faster

Performance on AMD* Opteron

Small Matrices: Intel® MKL faster than ATLAS!

Large Matrices: Intel® MKL within 20% of ATLAS

Intel may make changes to specification, product descriptions, and plans at any time, without notice. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/software/products or call (U.S.) 1-800-628-8686 or 1-916-356-3104

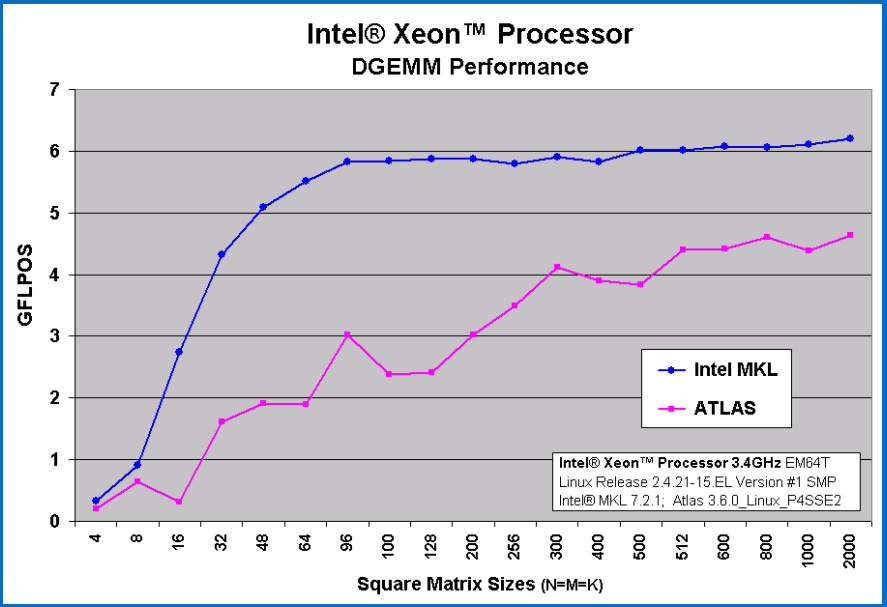

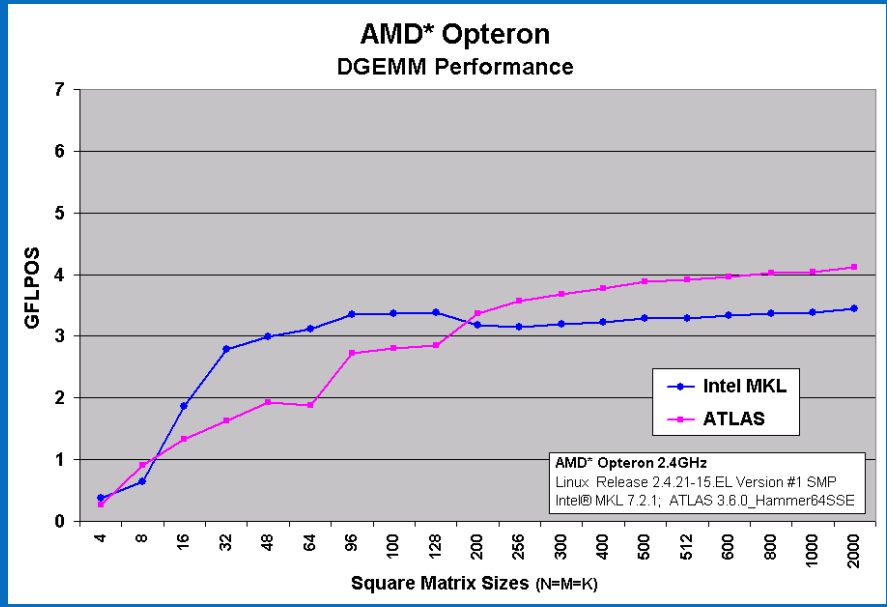

Intel® Math Kernel Library best choice for overall performance

Next Step … www.intel.com/software/products


intel®
Leap ahead™

Academgorodok

- 50 years of success
- First scientific city in the world
- 15,000 researchers
- 3,000 SW developers in the private sector

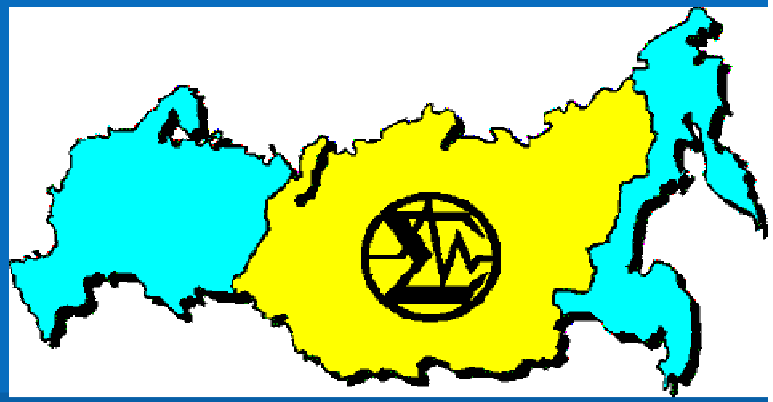

Siberian Branch of The Russian Academy of Sciences

Novosibirsk (Academgorodok) Russia