



Математические библиотеки Intel

Сергей Гололобов

НГУ/Intelligent Computing

Что нового в высокопроизводительных библиотеках от компании Intel®?

Intel® Math Kernel Library – библиотека быстрых математических примитивов

Intel® Threading Building Blocks – Библиотека управления потоками для C++

Intel® Data Analytics Acceleration Library – Библиотека примитивов для анализа данных для C++/Python

Intel® Integrated Performance Primitives – Обработка изображений, сигналов и данных

Optimization Notice

Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



Параллельная STL: Ускоренные алгоритмы C++ STL

Intel® Threading Building Blocks версии 2018

Параллельная библиотека стандартных шаблонов

- Поддержка черновика стандарта C++ 2017
- Методы для распараллеливания и векторизации с помощью компилятора и минимальными изменениями кода
- Выбор способа исполнения работы
- Способы могут выбираться отдельно для каждого задания
- Полностью использует многопоточность из библиотеки Intel® TBB
- Работает с компилятором Intel® C/C++ для векторизации

Пример

```
for_each(exec policy, begin, end, function)
```

	Exec Policy	Режим
Стандарт	<i>seq</i>	Последовательный
	<i>par</i>	Параллельный
	<i>par_vec</i>	Векторизованный
Расширение	<i>fpga_policy_PSG</i>	Для PSG FPGA (ПЛИС)
	<i>my_policy_target</i>	Другое оборудование

Поддержка новейшего оборудования, улучшенная документация и прочее

Intel® Data Analytics Acceleration Library версии 2018

- Улучшена страница с [документацией](#) для быстрого поиска нужной информации
- Добавлены некоторые модификации интерфейса для упрощения использования библиотеки и придания ему большей логичности
- Добавлена поддержка алгоритма «случайного леса» с вычислением индекса Джини.
- Добавлен переменный параметр скорости обучения в алгоритме стохастического градиентного спуска
- Улучшена производительность отдельных частей библиотеки

Поддержка новейшего оборудования, ускоренное сжатие и прочее

Intel® Integrated Performance Primitives версии 2018

Оборудование

- Поддержка следующего поколения процессоров серии Intel® Xeon Phi™

Криптография

- Ликвидирована зависимость от основного пакета Intel® IPP и его версии

Сжатие

- Оптимизация алгоритма LZO (сжатие без потерь) с использованием инструкций SSE4.2 и AVX2
- Улучшенная производительность

Intel® Math Kernel Library



Energy



Science & Research



Engineering Design



Financial Analytics




Signal Processing



Digital Content Creation

- Ускоряет математические вычисления в задачах машинного обучения, научных, финансовых и инженерных приложений
- Содержит функции для линейной алгебры (BLAS, LAPACK, PARDISO), БПФ, векторной математики, статистики и др.
- Использует имеющиеся на рынке стандартные интерфейсы для лёгкой смены библиотек
- Оптимизированные, параллелизованные и векторизованные программы для максимальной производительности



Компоненты Intel® MKL 2018

Линейная алгебра

- BLAS
- LAPACK
- ScaLAPACK
- Sparse BLAS
- Sparse Solvers
 - Iterative
 - PARDISO*
- Cluster Sparse Solver

Быстрые преобразования Фурье

- Многомерные
- FFTW интерфейсы
- Cluster FFT

Векторные функции

- Тригонометрические
- Гиперболические
- Экспоненты
- Логарифмы
- Степенные
- Корни
- Генераторы случайных чисел

Статистика

- Эксцесс
- Коэффициенты вариации
- Статистика рангов
- Min/max
- Дисперсия, ковариация

А также

- Сплаины
- Интерполяция
- Trust Region
- Fast Poisson Solver
- Свёртки
- Нормализации
- И др. 😊

Как Intel® Math Kernel Library (Intel® MKL) работает?

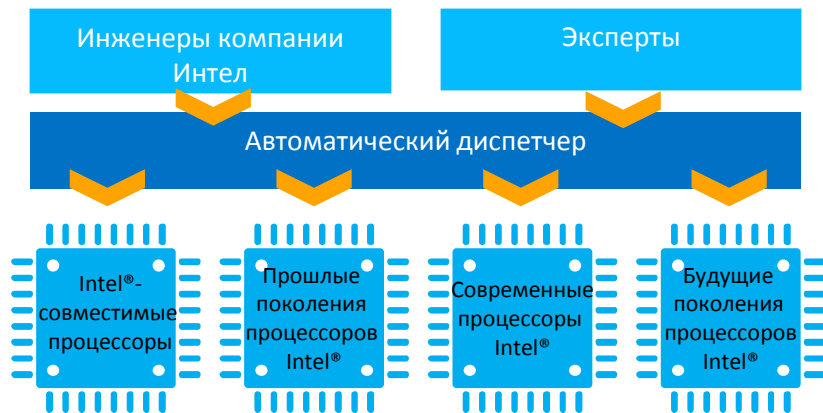
Оптимизированные математические примитивы

- Производительность – наше всё 😊!

Самая популярная библиотека математических примитивов в мире

- Используется везде 😊 (например, в Matlab, CADax)

Заточена под существующие и будущие процессоры компании Интел, НО работает на всех x86 совместимых процессорах



EDC North America
Development Survey 2016,
Volume I


Большинство пользователей
математических библиотек используют
MKL

**Знает всё о многоядерных
процессорах и системах**

- Поддерживает все платформы
- Векторизована, параллельна и работает с кластерами

Intel® MKL: оцените производительность на Зеонах!

DGEMM Performance On Intel® Xeon® Processor E5-2699 v4



Configuration Info - Versions: Intel® Math Kernel Library (Intel® MKL) 2017; Hardware: Intel® Xeon® Processor E5-2699 v4, 2 Twenty-two-core CPU (55MB smart cache, 2.2GHz), 64GB of RAM; Operating System: RHEL 7.2 GA x86_64;

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804.

Optimization Notice


Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Последняя версия Intel® MKL даёт высочайшую
производительность на всех процессорах компании Интел



Intel® MKL: оцените производительность на Зеонах!

Faster 1,000 Small Matrix Multiplications using Intel® MKL DGEMM_BATCH DGEMM_BATCH vs DGEMM, Intel® Xeon® Processor E5-2699 v4 (44 threads)



Configuration Info - Versions: Intel® Math Kernel Library (Intel® MKL) 2017; Hardware: Intel® Xeon® Processor E5-2699 v4, 2 Twenty-two-core CPU (55MB smart cache, 2.2GHz), 64GB of RAM; Operating System: RHEL 7.2 GA x86_64;

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804.

Optimization Notice


Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Последняя версия Intel® MKL даёт высочайшую
производительность на всех процессорах компании Интел



Intel® MKL: оцените производительность на Xeon Phi 😊!

DGEMM Performance On Intel® Xeon Phi™ Processor 7250



Configuration Info - Versions: Intel® Math Kernel Library (Intel® MKL) 2017; Hardware: Intel® Xeon Phi™ Processor 7250, 68 cores (34MB L2 cache, 1.4GHz), 96 GB of DDR4 RAM and 16 GB MCDRAM; Operating System: RHEL 7.2 GA x86_64;

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>.

Optimization Notice


Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Последняя версия Intel® MKL даёт высочайшую
производительность на всех процессорах компании Интел



Intel® MKL: оцените производительность на Xeon Phi😊!

Faster 1,000 Small Matrix Multiplications using Intel® MKL DGEMM_BATCH DGEMM_BATCH vs DGEMM, Intel® Xeon Phi™ Processor 7250 (68 threads)



Configuration Info - Versions: Intel® Math Kernel Library (Intel® MKL) 2017; Hardware: Intel® Xeon Phi™ Processor 7250, 68 cores (34MB L2 cache, 1.4GHz), 96 GB of DDR4 RAM and 16 GB MCDRAM; Operating System: RHEL 7.2 GA x86_64;

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804.

Optimization Notice


Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Последняя версия Intel® MKL даёт высочайшую
производительность на всех процессорах компании Интел




Новые и дополненные математические функции Intel® Math Kernel Library версии 2018


BLAS



Новые функции
целочисленного
умножения



Улучшенная
производительность
SGEMM




Оптимизированные
свёртки и
внутренние
произведения


BLAS Групповой интерфейс



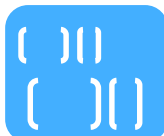
Эффективность и
производительность



Новый групповой
интерфейс для решения
треугольных матриц

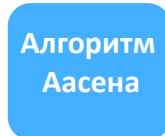


Ускоренный
GEMM_BATCH



Подходит для
маленьких и
средних
матриц

LAPACK



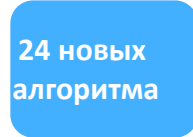
Алгоритм
Аасена

Функции
факторизации
и решения
систем



Улучшенная
производительность

Новые векторные функции



24 новых
алгоритма

Богатый выбор
оптимизированных
функций

Optimization Notice

Copyright © 2016, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Библиотеки компании Интел доступны БЕСПЛАТНО для любых
целей и любых организаций

ПРОБУЙТЕ И НАСЛАЖДАЙТЕСЬ ПРОИЗВОДИТЕЛЬНОСТЬЮ!

<https://software.intel.com/en-us/performance-libraries>

ВОПРОСЫ

Legal Disclaimer and Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Copyright © 2016, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

Optimization Notice

Intel’s compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

