

A large, detailed image of a microchip die is the central focus of the slide. The die is a complex grid of various colored regions, including red, orange, yellow, green, and blue, representing different functional blocks. The image is tilted at an angle. A large purple triangle is overlaid on the bottom left, and a green triangle is overlaid on the bottom right, both pointing towards the center.

ANALYSIS OF BLIS MULTITHREADED GEMM AND HPC WORKLOADS

PRADEEP RAO
AMD

AGENDA



- ▲ BLIS DGEMM Multi-threaded performance analysis on the AMD EPYC™ processor
- ▲ BLIS overheads in HPC workloads
- ▲ Q & A

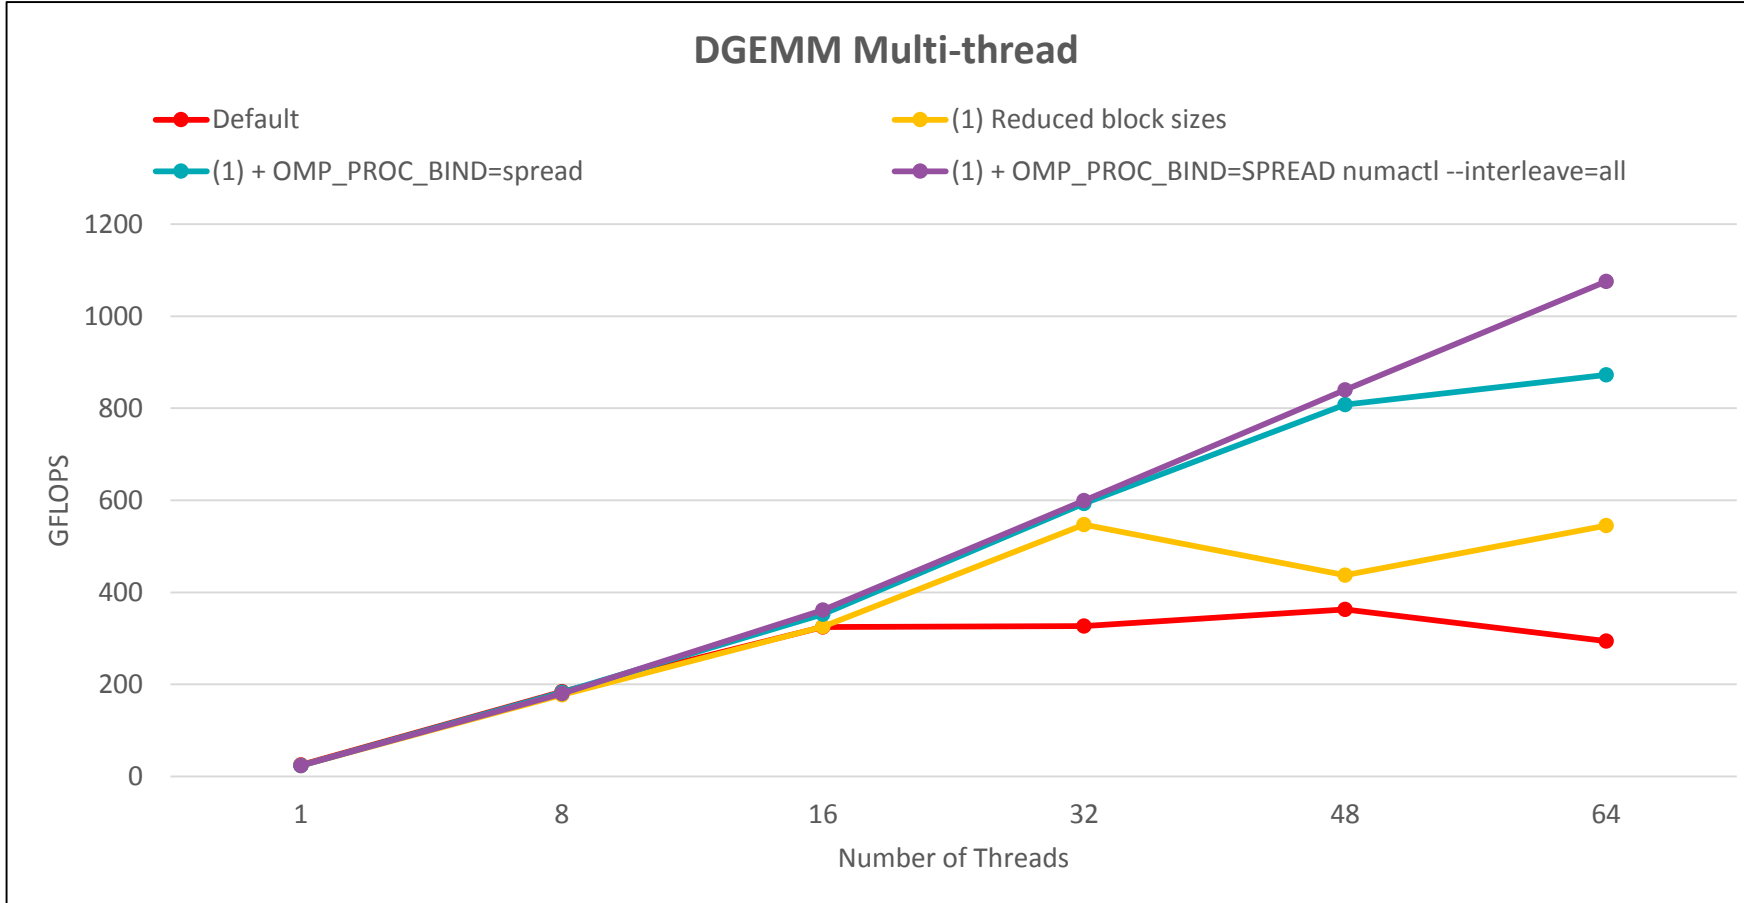
BLIS DGEMM Multi-threaded performance on the AMD EPYC™ processor

The background features a large teal triangle pointing downwards from the top left. Below it, there are several overlapping green and teal geometric shapes, including a large green trapezoid and a smaller teal trapezoid, creating a layered, abstract design.

TUNING PARALLEL DGEMM PERFORMANCE



MATRIX SIZE-5000X5000



1. Reduce block sizes
 - Each thread works on smaller blocks
2. OpenMP settings
 - OMP_PLACES = cores
 - OMP_PROC_BIND=spread
3. numactl
 - interleave

▲ Issue

- Memory contention increases as the number of threads increase

▲ Tuning strategy

- Reduce the block sizes
 - Each thread works on smaller blocks of data
 - Improved per thread cache
- Higher available bandwidth per thread achieved by dividing the threads over two sockets using `OMP_PROC_BIND=spread`
- Distribute the memory across all the nodes
- Reduced CPU stalls

BLIS overheads in HPC workloads



▲ CP2K

- Package to perform atomistic simulations of solid state, liquid, molecular, and biological systems

▲ GROMACS

- Molecular dynamics package mainly designed for simulations of proteins, lipids, and nucleic acids

▲ ScaLAPACK

- Library of high-performance linear algebra routines for parallel distributed memory machines

BLIS OVERHEADS IN HPC WORKLOADS



CP2K – 26%

- Input : 128 H₂O atoms
- DGEMM (6%), packing 5.7%
- Small matrix sizes: 1 x 1, 5 x 3, 7 x 1....128 x 1152
- HERK, TRMM, TRSM – invoked but negligible

GROMACS - 63%

- Test : Covariance Analysis, size 1167x1167
- DGEMM (29%), Dotxf, Saxpyf

ScaLAPACK – 49%

- Test : xzpbllt - Symmetric or Hermitian positive definite band matrix solve
- CGEMM, DGEMM, AXPY, ZAMAX
- Small matrix sizes : 1 x 1, 2 x 2,1023 x 1023

BLIS IN HPC WORKLOADS



- ▲ BLAS functionalities have considerable overhead in most HPC workloads
- ▲ Operations on small matrices dominate in many such applications such as CP2K and GROMACS
- ▲ There is scope to optimize the smaller matrices operations further

REFERENCES & SOURCE CODE



- ▲ <https://www.amd.com/en/products/epyc-server>
- ▲ <http://developer.amd.com/amd-cpu-libraries/>
- ▲ <https://github.com/amd/blis>
- ▲ <https://www.cp2k.org/>
- ▲ <http://www.gromacs.org/>
- ▲ <http://www.netlib.org/scalapack/>

Questions



DISCLAIMER & ATTRIBUTION



Pradeep Rao is a Senior Member of Technical Staff for CPU Libraries team at AMD. His postings are his own opinions and may not represent AMD's positions, strategies or opinions. Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of non infringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

AMD, the AMD Arrow logo, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

© 2018 Advanced Micro Devices, Inc. All rights reserved.

▲ Slide 4

- DGEMM performance for 5000 x 5000 matrices using AMD optimized BLIS. Testing were conducted as of 5th September 2018 on the test system comprising AMD EPYC™, 64 cores clocked at 2.5 GHz. The machine had DDR4 RAM of 256 GB, clocked at 1.2 GHz and 1 TB of Hard disk. The test system had Ubuntu operating system installed. PC manufacturers may vary configurations yielding different results

▲ Slide 8

- Testing were conducted as of 27th August 2018 on the test system comprising AMD EPYC™, 64 cores clocked at 2.5 GHz. The machine had DDR4 RAM of 256 GB, clocked at 1.2 GHz and 1 TB of Hard disk. The test system had Ubuntu operating system installed. PC manufacturers may vary configurations yielding different results