

Redesigning Dense Linear Algebra Software for ManyCore and Beyond

Jack Dongarra
University of Tennessee
and
Oak Ridge National Laboratory

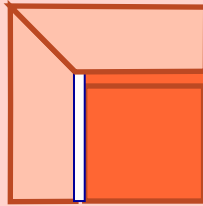
Major Changes to Software

- **Must rethink the design of our software**
 - **Another disruptive technology**
 - Similar to what happened with cluster computing and message passing
 - **Rethink and rewrite the applications, algorithms, and software**
- **Numerical libraries for example will change**
 - **For example, both LAPACK and ScaLAPACK will undergo major changes to accommodate this**

A New Generation of Software:

Software/Algorithms follow hardware evolution in time

LINPACK (70's)
(Vector operations)

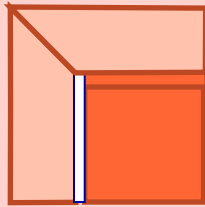


Rely on
- Level-1 BLAS
operations

A New Generation of Software:

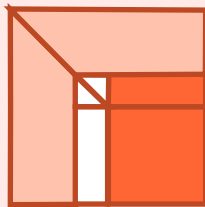
Software/Algorithms follow hardware evolution in time

LINPACK (70's)
(Vector operations)



Rely on
- Level-1 BLAS
operations

LAPACK (80's)
(Blocking, cache
friendly)

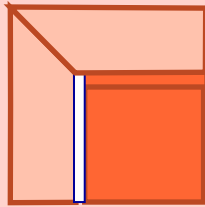


Rely on
- Level-3 BLAS
operations

A New Generation of Software:

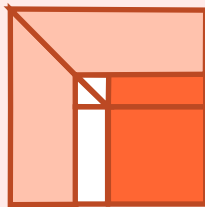
Software/Algorithms follow hardware evolution in time

LINPACK (70's)
(Vector operations)



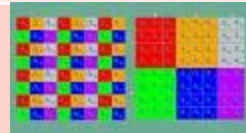
Rely on
- Level-1 BLAS operations

LAPACK (80's)
(Blocking, cache friendly)



Rely on
- Level-3 BLAS operations

ScaLAPACK (90's)
(Distributed Memory)



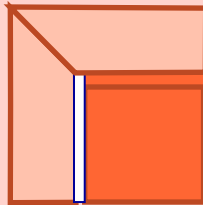
Rely on
- PBLAS Mess Passing

A New Generation of Software:

Parallel Linear Algebra Software for Multicore Architectures (PLASMA)

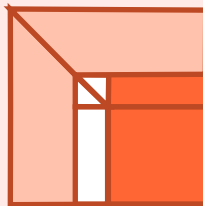
Software/Algorithms follow hardware evolution in time

LINPACK (70's)
(Vector operations)



Rely on
- Level-1 BLAS operations

LAPACK (80's)
(Blocking, cache friendly)



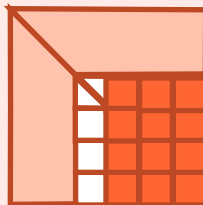
Rely on
- Level-3 BLAS operations

ScaLAPACK (90's)
(Distributed Memory)



Rely on
- PBLAS Mess Passing

PLASMA (00's)
New Algorithms
(many-core friendly)



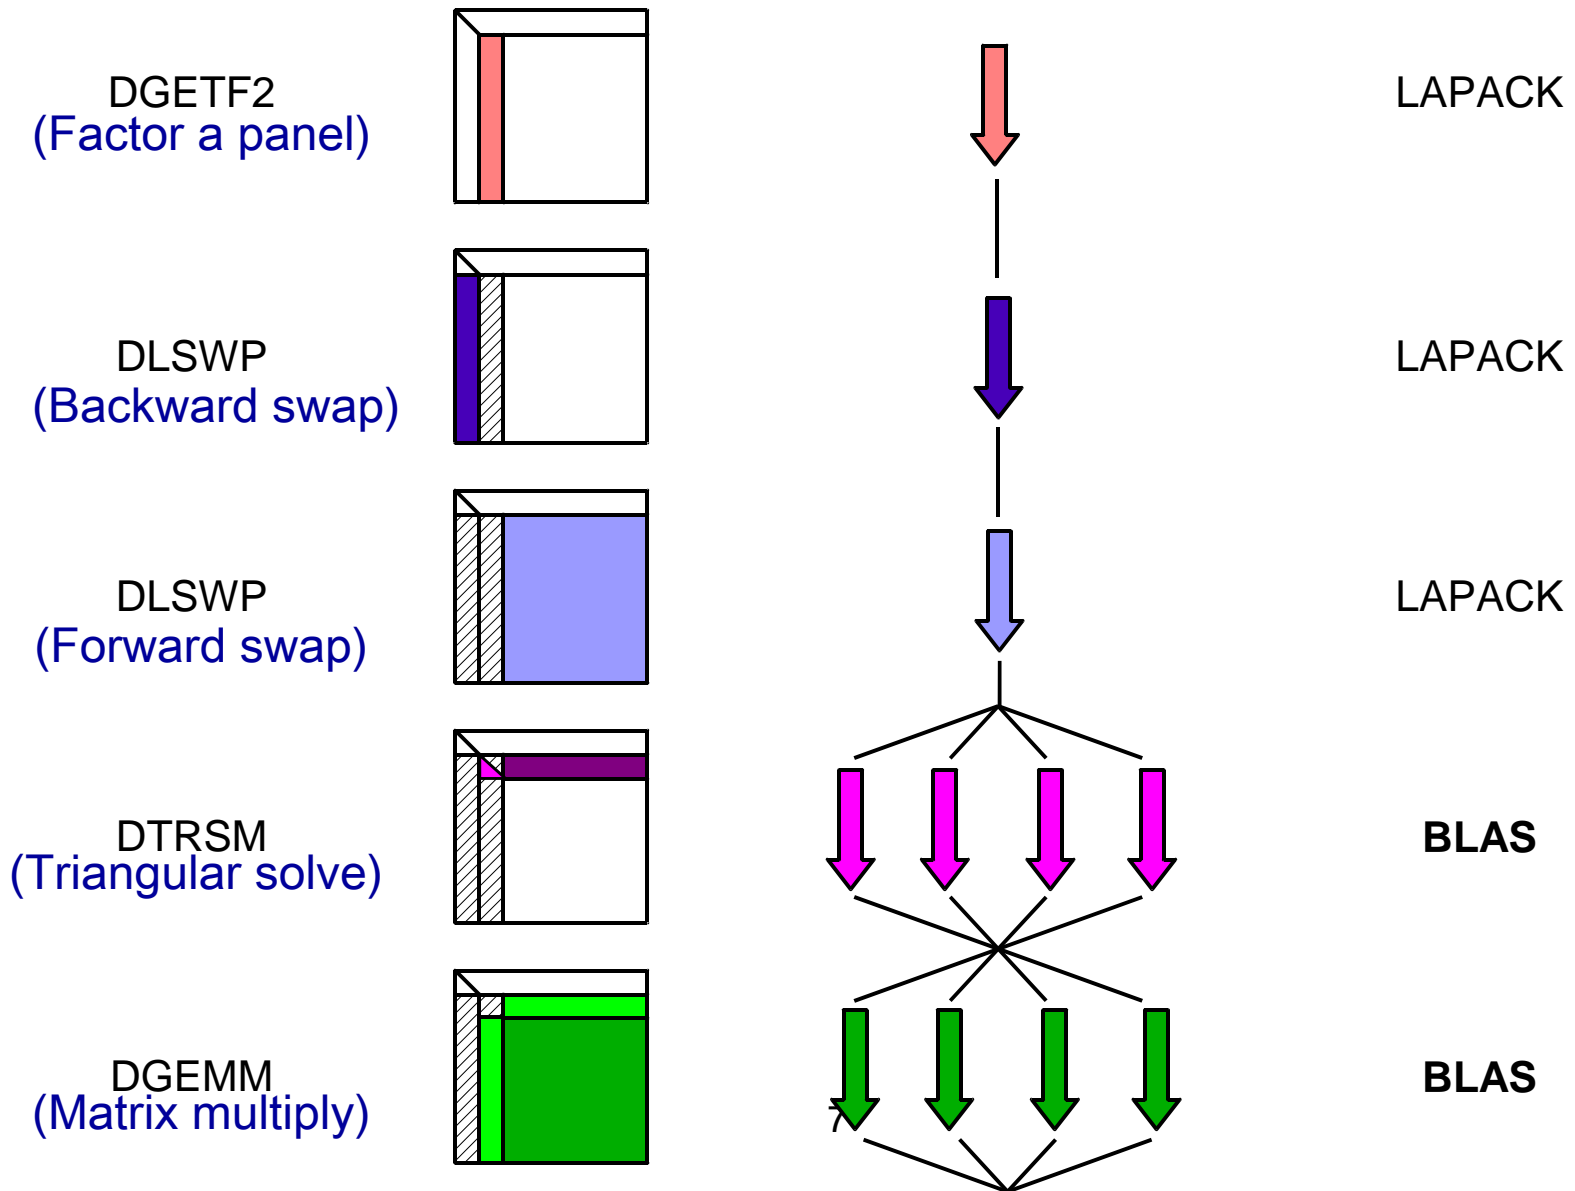
Rely on
- a DAG/scheduler
- block data layout
- some extra kernels

Those new algorithms

- have a very **low granularity**, they scale very well (multicore, petascale computing, ...)
- **removes a lots of dependencies** among the tasks, (multicore, distributed computing)
- **avoid latency** (distributed computing, out-of-core)
- **rely on fast kernels**

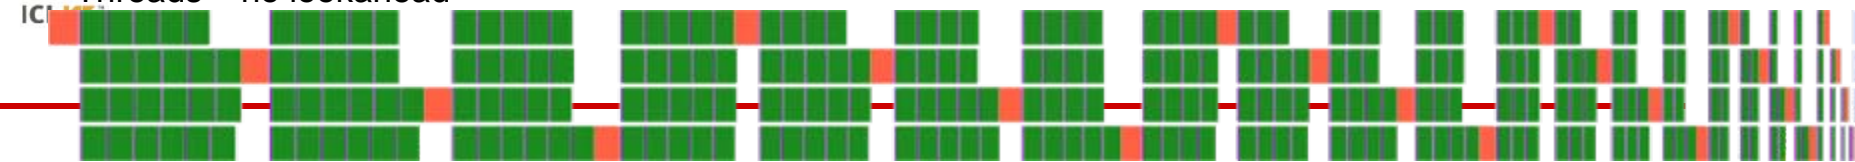
Those new algorithms need new kernels and rely on efficient scheduling algorithms.

Steps in the LAPACK LU

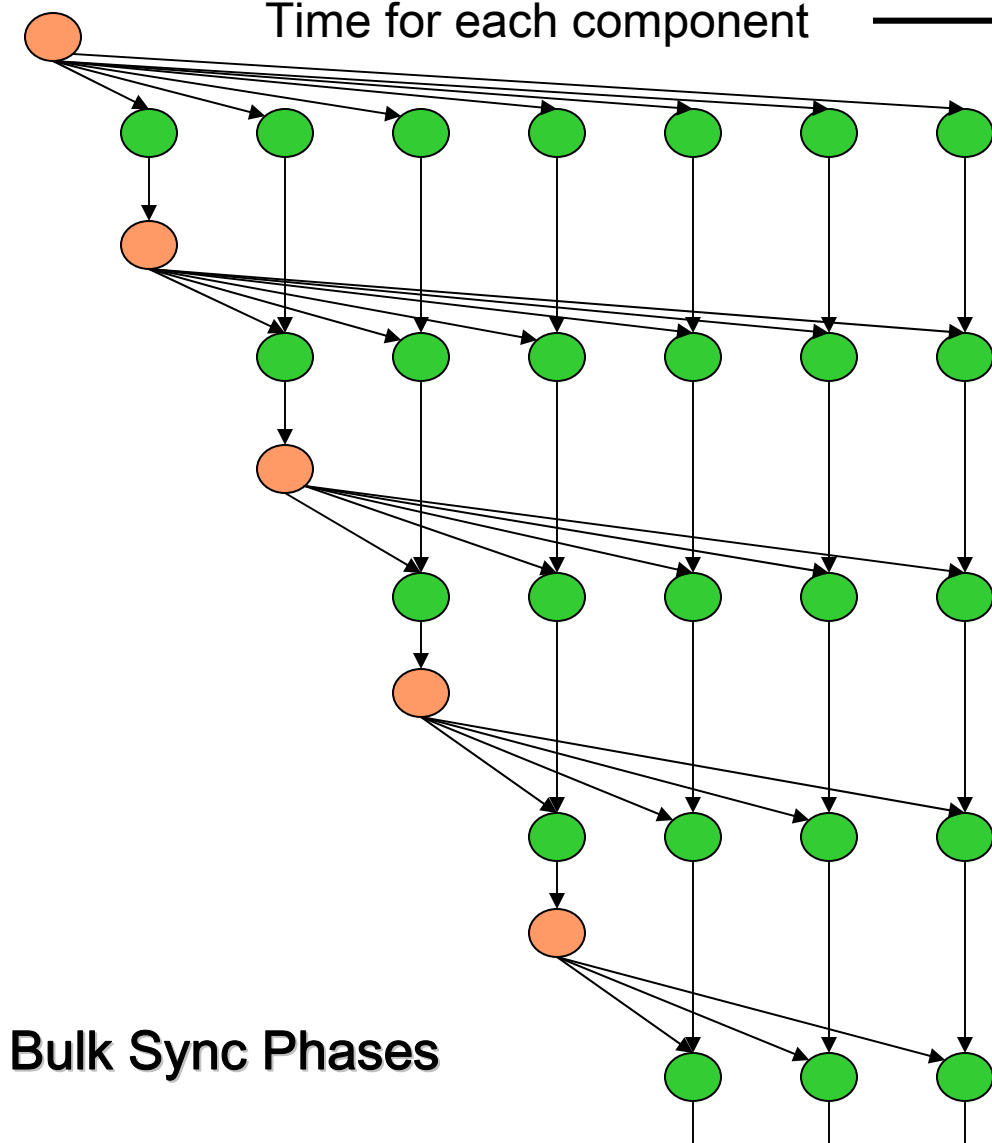


LU Timing Profile (4 processor system)

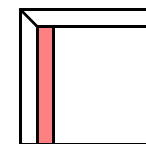
Threads – no lookahead



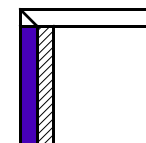
Time for each component →



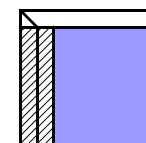
DGETF2



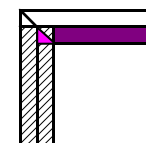
DLASWP



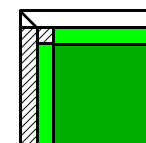
DLASWP



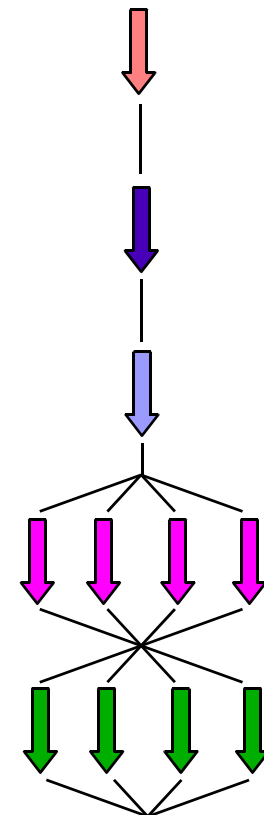
DTRSM



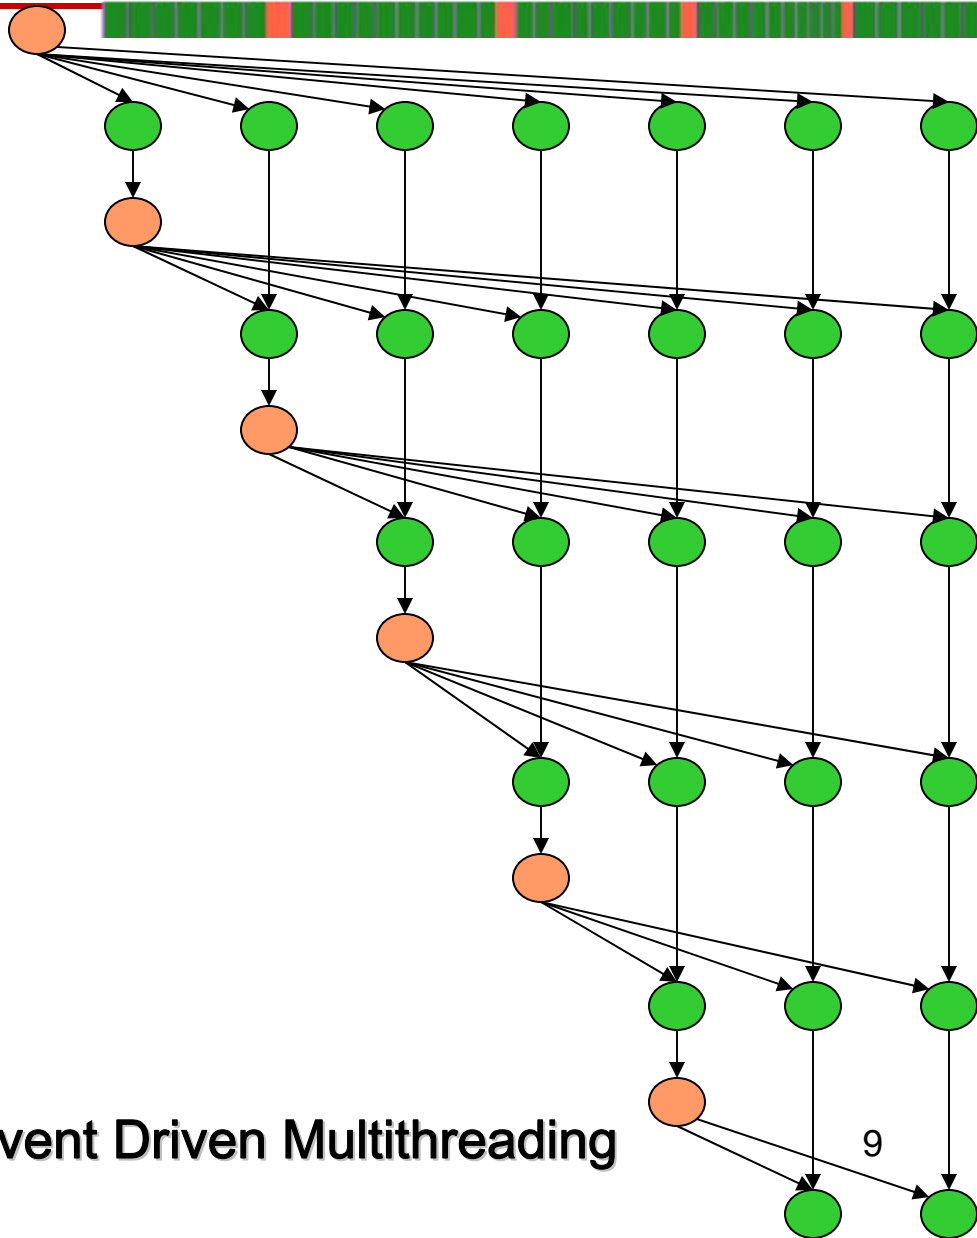
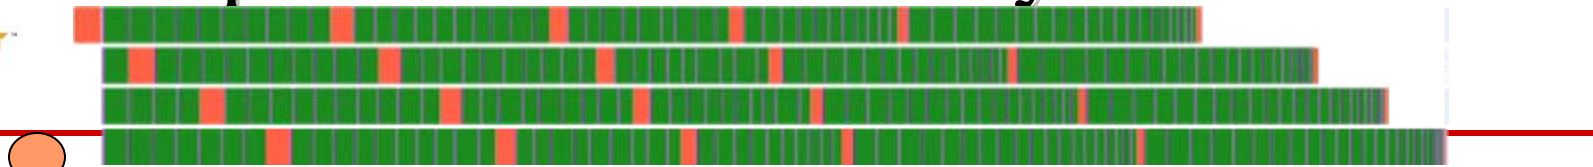
DGEMM



- DGETF2
- DLASWP(L)
- DLASWP(R)
- DTRSM
- DGEMM



Adaptive Lookahead - Dynamic



Event Driven Multithreading

9

```

while(1)
  fetch_task();
  switch(task.type) {
    case PANEL:
      dgetf2();
      update_progress();
    case COLUMN:
      dlaswp();
      dtrsm();
      dgemm();
      update_progress();
    case END:
      for()
        dlaswp();
      return;
  }
}

```

Reorganizing
algorithms to use
this approach

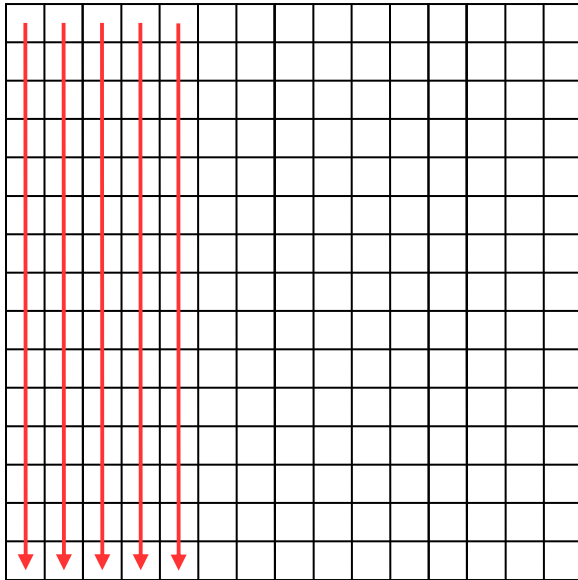
Redesign

- ◆ **Asynchronicity**
 - **Avoid fork-join (Bulk sync design)**
- ◆ **Dynamic Scheduling**
 - **Out of order execution**
- ◆ **Fine Granularity**
 - **Independent block operations**
- ◆ **Locality of Reference**
 - **Data storage - Block Data Layout**

Achieving Fine Granularity

Fine granularity may require novel data formats to overcome the limitations of BLAS on small chunks of data.

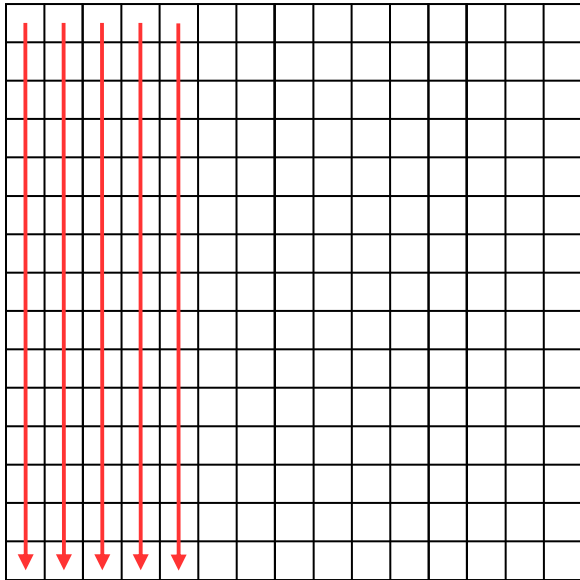
Column-Major



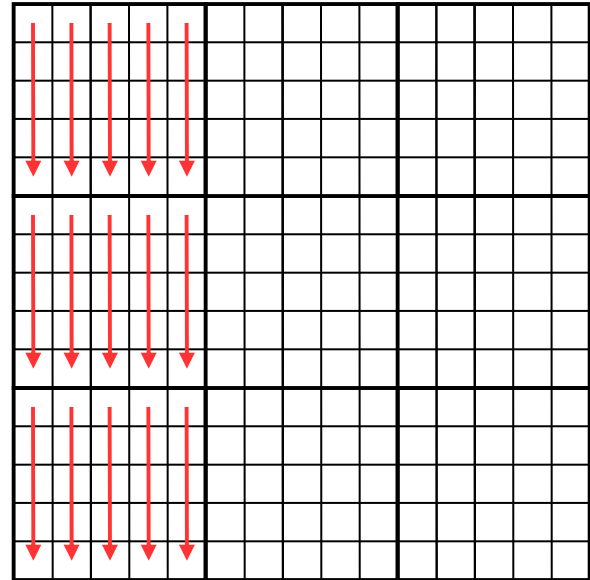
Achieving Fine Granularity

Fine granularity may require novel data formats to overcome the limitations of BLAS on small chunks of data.

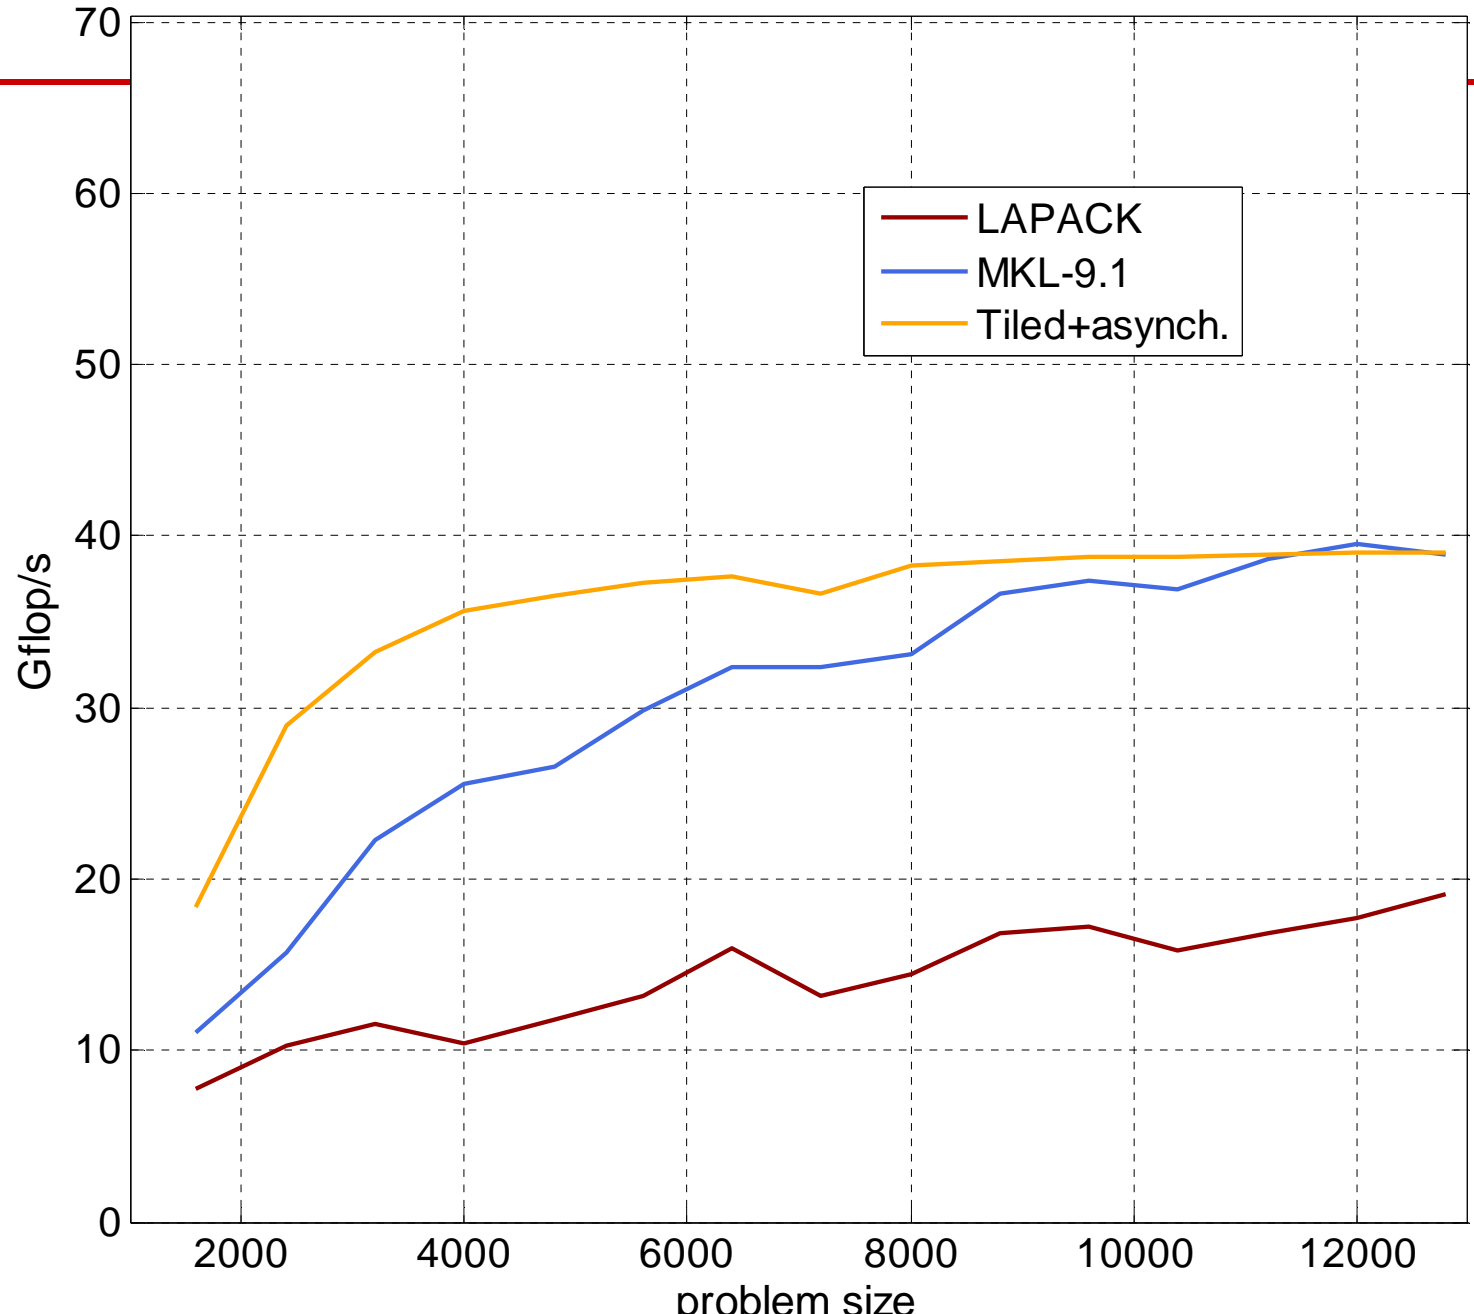
Column-Major



Block data layout

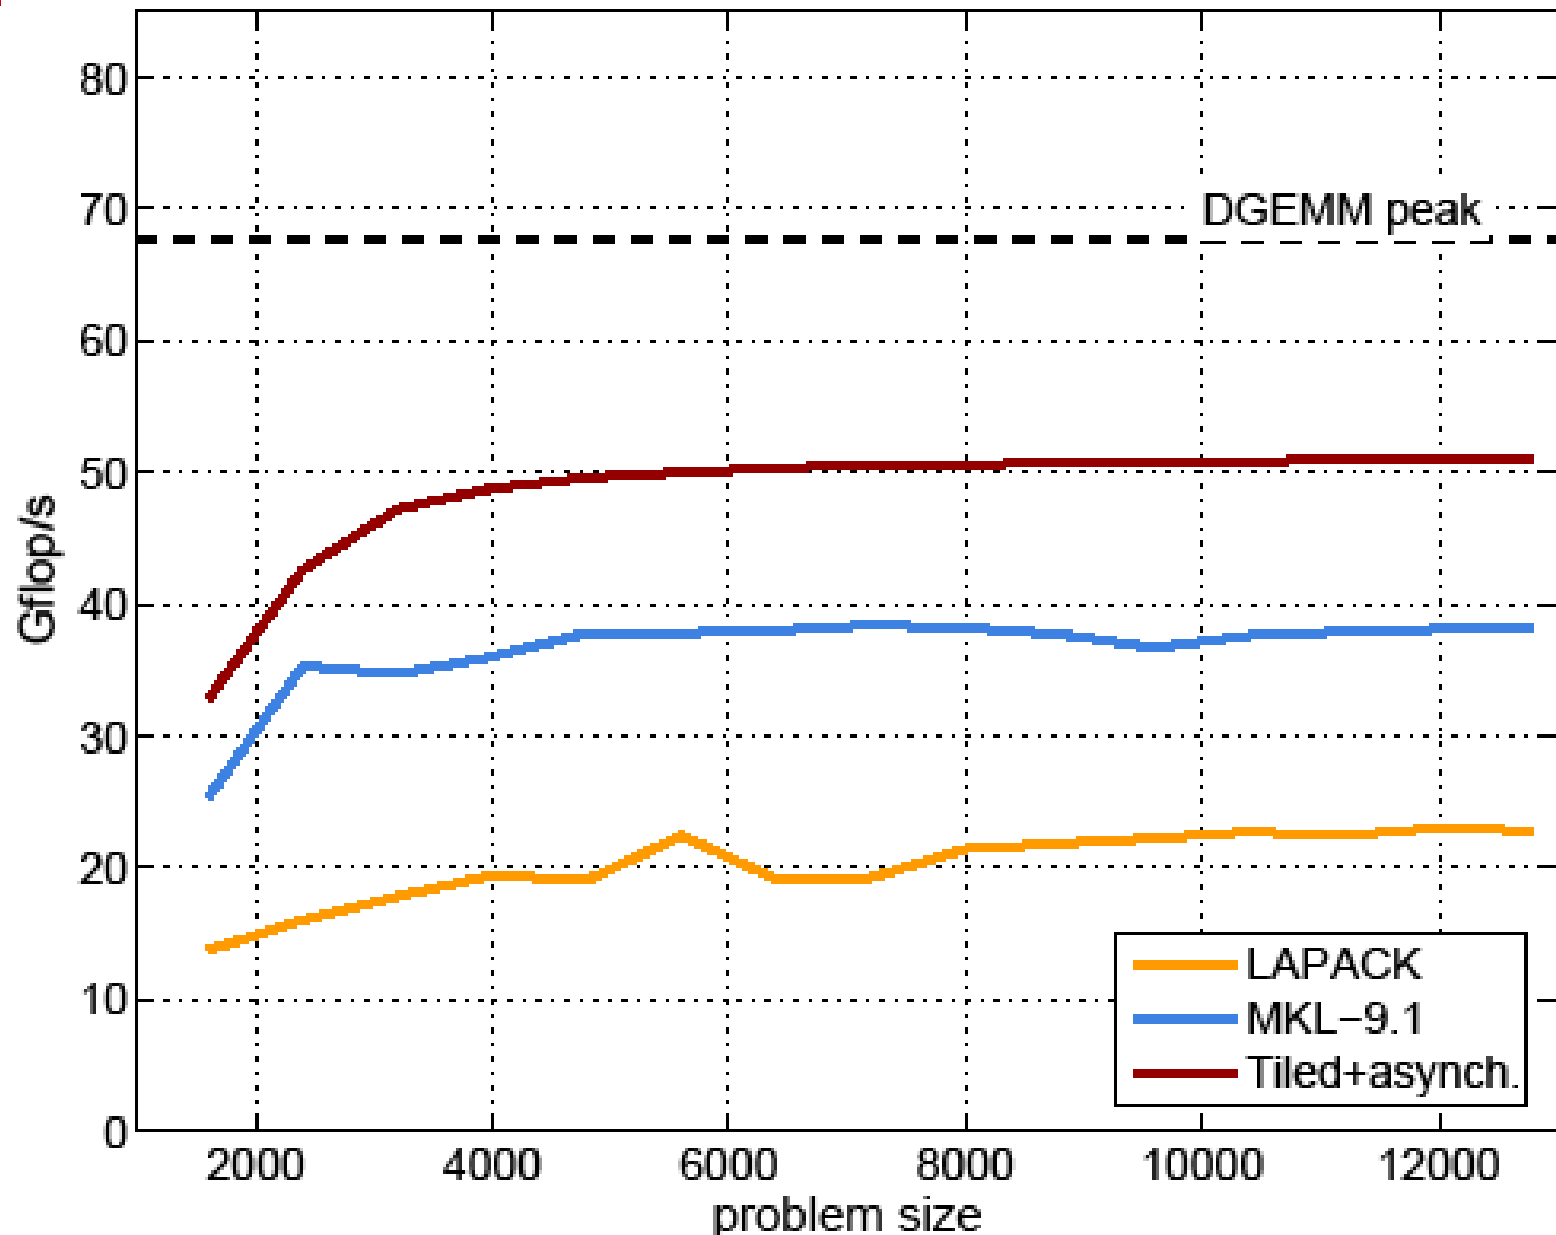


LU -- 8-way dual Opteron -- MKL-9.1

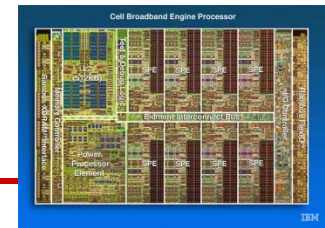




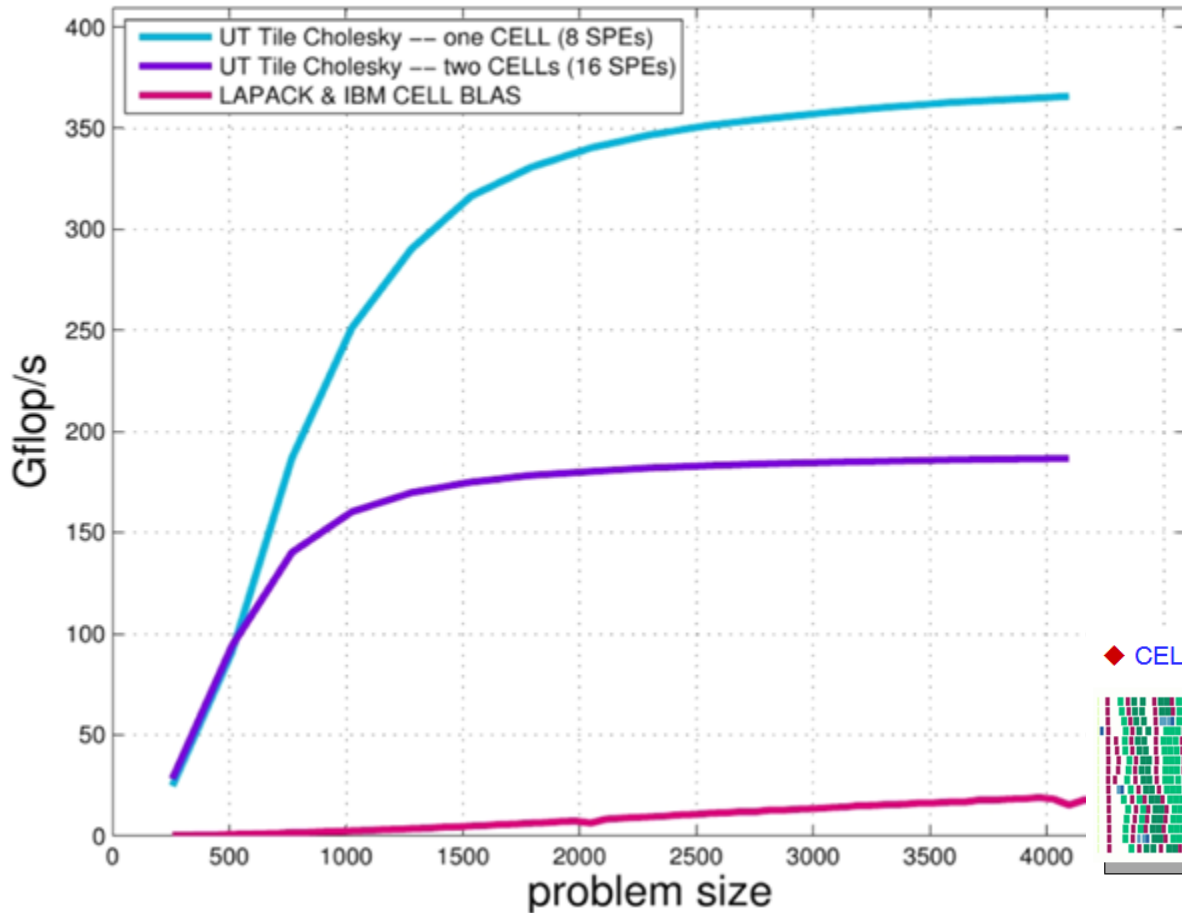
QR -- 2-way Quad Clovertown



Cholesky on the CELL



Cholesky -- CELL Processor



1 CELL (8 SPEs)

- 186 Gflop/s
- 91 % peak
- 97 % SGEMM peak

2 CELLS (16 SPEs)

- 365 Gflop/s
- 89 % peak
- 95 % SGEMM peak

◆ CELL Cholesky - 16 cores



◆ CELL Cholesky - 8 cores



Single precision results on the Cell

Next Steps

- ◆ Looking at implementations of the algorithms using:
 - Cilk, SMP SuperScalar, Intel Threading Building Blocks, UPC,...
- ◆ Distributed memory implementations will follow

Collaborators / Support

- ◆ U Tennessee, Knoxville
 - Julien Langou, Julie Langou, Piotr Luszczek, Jakub Kurzak, Stan Tomov, Remi Delmas, Peng Du
- ◆ UC Berkeley:
 - Jim Demmel, Ming Gu, W. Kahan, Beresford Parlett, Xiaoye Li, Osni Marques, Christof Voemel, David Bindel, Yozo Hida, Jason Riedy, Jianlin Xia, Jiang Zhu

Microsoft



GoogleTM
English

Web [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#)^{New!} [more »](#)

dongarra

Google Search

I'm Feeling Lucky

[Advanced Search](#)
[Preferences](#)
[Language Tools](#)

[Advertising Programs](#) - [About Google](#) - [Go to Google.com](#)

[Make Google Your Homepage!](#)

©2005 Google - Searching 8,058,044,651 web pages