

ParMA Tools and Real-World Applications

Bettina Kramer – UVSQ

Benedetto Risio, Alexander Berreth – RECOM

Wilfried Schäfer, Karin Lukaszewicz – MAGMA

Koutaiba Kassem, René Menzel - GNS

ParMA

Parallel Programming for Multi-core Architectures
www.parma-itea2.org

ISC'10 BoF, Hamburg, 31 May 2010

Content

ParMA

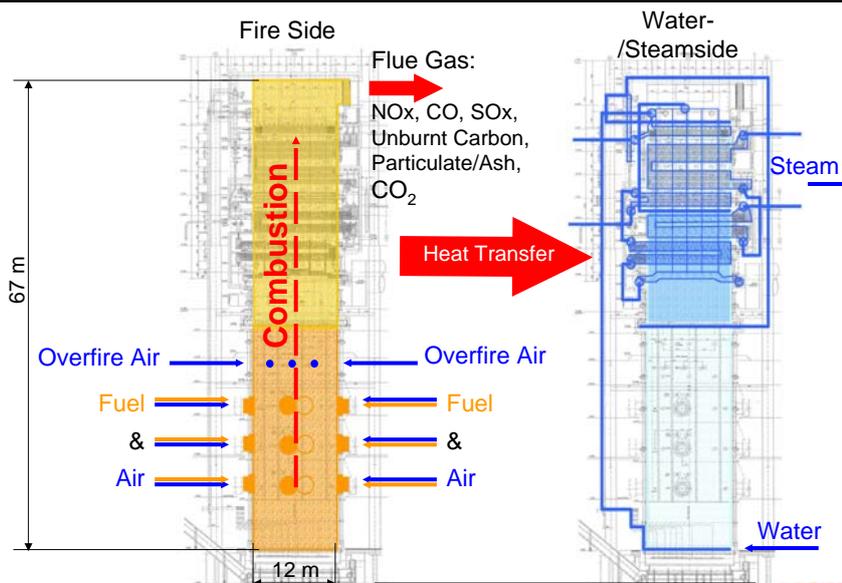
- **RECOM – combustion simulation**
- MAGMA – casting simulation
- GNS – metal forming simulation

Computer-Aided combustion optimisation in fossil-fueled power stations



Dr.-Ing. Benedetto Risio – RECOM Services GmbH

Schematic of power generation in fossil fueled power stations



Dr.-Ing. Benedetto Risio – RECOM Services GmbH

Compute-intensive simulation!

VAMPIR / SCALASCA

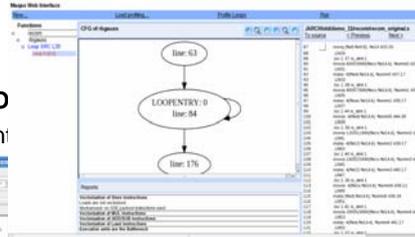
Analysis of MPI-parallel execution & communication



10% savings in computing time on 128 Cores (MPI_Waitall vs MPI_Wait)

MAQAO

Analysis of memory access and execution structure for memory intense subroutines



38 % savings in computing time on subroutine level, and 13.6 % savings in computing time on 16 cores for the entire application

(modified data structures to avoid stride 2 access)
Data type conversion error identified and removed

➤ **STEP:**

Automatic conversion of OpenMP-parallelized software into a hybrid parallel (OpenMP and MPI) version.

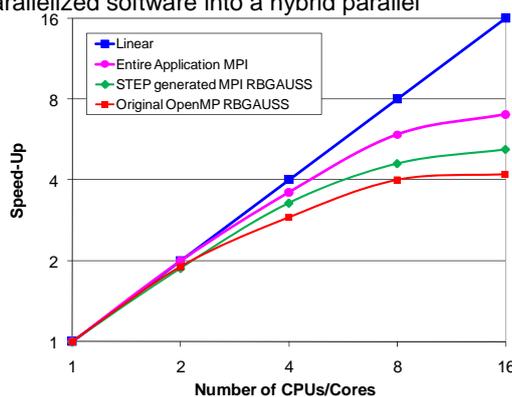
➤ **Target of the experimentation:**

Demonstrate the potential of STEP for the complex RECOM-AIOLOS Software package by automatically converting the RBGAUSS solver

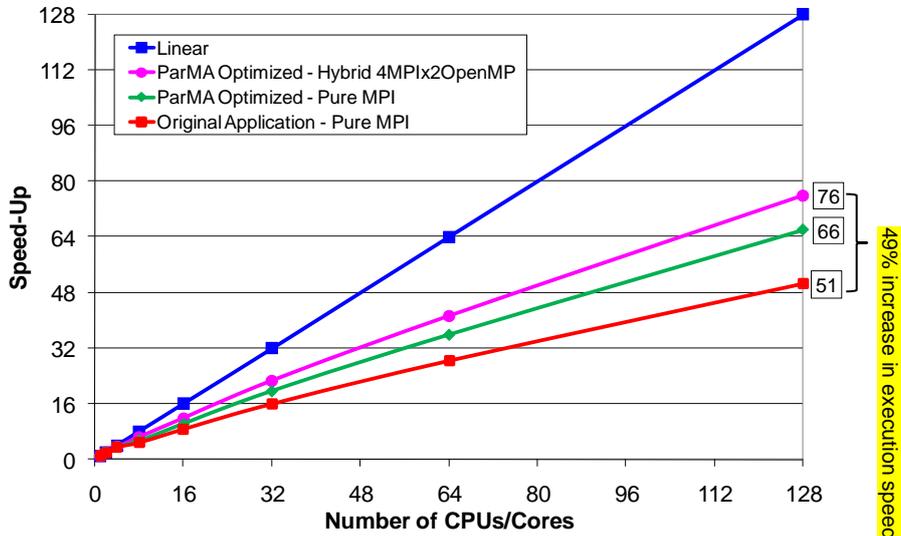
➤ **Results of the experimentation:**

STEP successfully converted an OpenMP-parallelized solver into an MPI-parallel version and achieved comparable parallel performance.

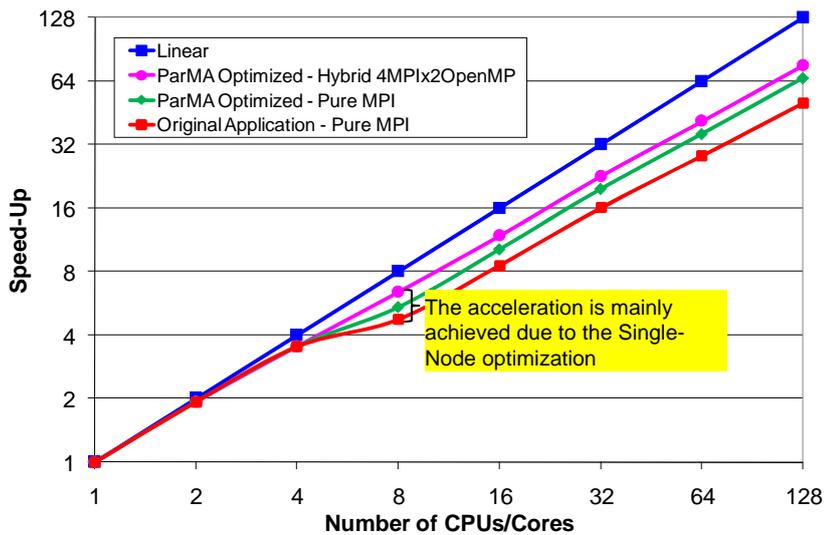
The STEP concept represents a useful approach to decrease development time and costs for MPI-programming.



Performance data on Nehalem

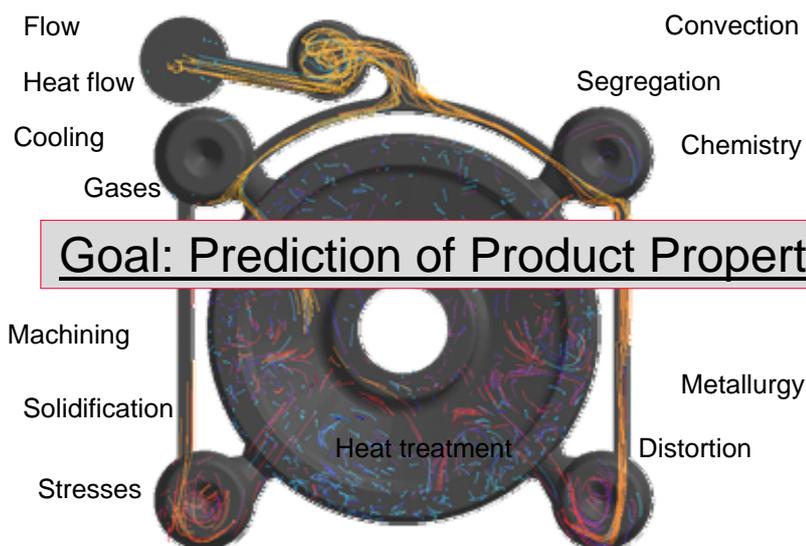


Performance data on Nehalem



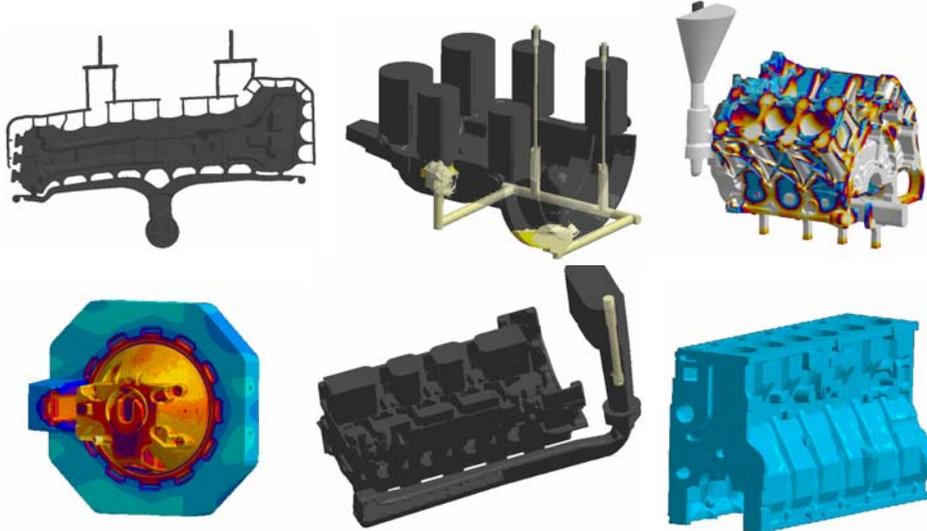
- RECOM – combustion simulation
- **MAGMA – casting simulation**
- GNS – metal forming simulation

MAGMA - Casting the most complex manufacturing process



Casting process simulation a core innovation

ParMA



Wilfried Schäfer, Karin Lukaszewicz – MAGMA GmbH

MAGMA 11

MAGMASOFT® Optimisation of different modules

ParMA

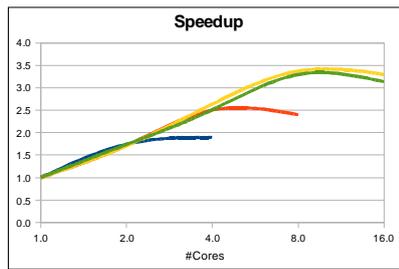
- **MAGMAfill**
 - Optimisation of the MPI communication with Scalasca
 - Replace collective by asynchronous communication
 - Overall MPI runtime cut by 3200 seconds (almost by 50%)
 - Overall runtime of the complete simulation reduced by 23%.
- **MAGMASolid**
 - **Integration of a new linear equation solver and optimisation with Scalasca**
 - **Optimisation of the CG-Solver with MAQAO**
- **MAGMAstress**
 - Optimisation of the communication with Scalasca
 - MPI time approx. reduced by 50%

Wilfried Schäfer, Karin Lukaszewicz – MAGMA GmbH

MAGMA 12

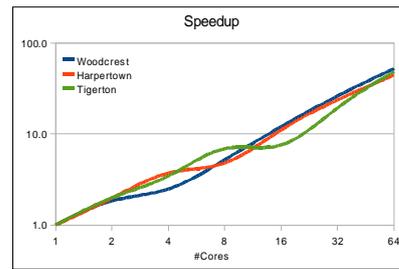
SMP-Version

The solver (ADI, Thomas Algorithm) parallelised with OpenMP does not scale well on a SMP machine and is not suitable for a cluster



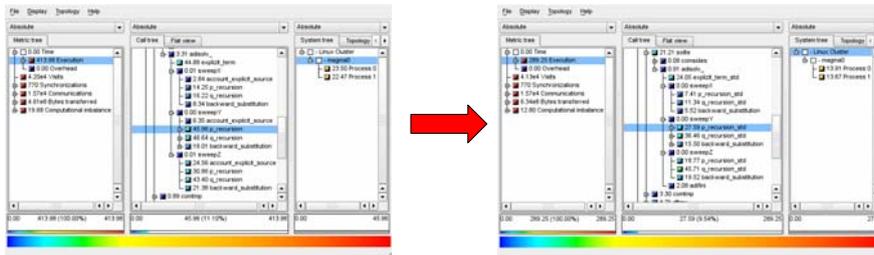
MPI-Version

The MPI solver (CGSolver) does scale quite well, but for less than 8 CPU's this solver is much slower than the SMP Version



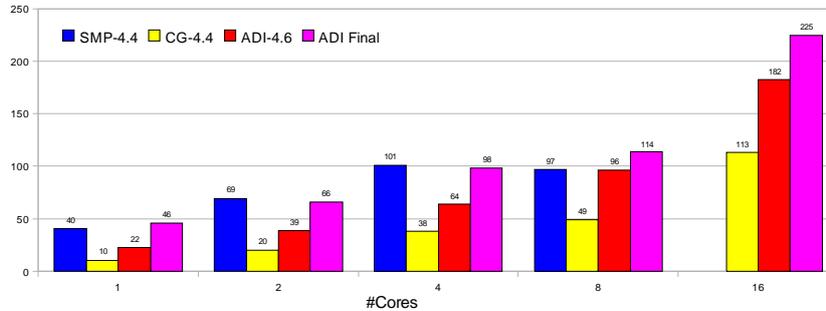
Starting point: The 2 code versions have to be maintained to make customers happy

- Integration of a parallel tridiagonal equation solver
- Optimisation with Scalasca



- Optimisation for 1, 2 and 4 cores
 - Using the sequential solver for one core in all three directions, for two cores in two directions and for four cores in one direction

Solidmarks



- New solver (ADI final) competitive with or better than old solver (SMP)
- CG solver sometimes still needed for stability reasons

Analysed Functions

- scalprod (Vector product of two float vectors)
- scalpnorm (Calculation of the L2 Norm and the MAX Norm)
- saxpy2 (vec1 = vec1 + const * vect2)
- matvec (Matrix Vector product)

Original Code

```

erg=0D0
do k = anf3,end3
do j = anf2, end2
do = anf1, end1
erg = erg +
DPROD(vec1(i,j,k),vec2(i,j,k))
end do
end do
end do
    
```

Optimised Code

```

end1_1=(anf1+(end1-anf1+1)/16)*16-1
erg = 0D0
do k = anf3, end3
do j = anf2, end2
do i = anf1, end1_1,16
call MM_PREFETCH (vect1(i+256,j,k), 3)
call MM_PREFETCH (vect2(i+256,j,k), 3)
erg = (((((((((((erg +
DPROD(vec1(i+ 0,j,k), vec2(i+
DPROD(vec1(i+ 1,j,k), vec2(i+ 1,j,k)))) +
...
DPROD(vec1(i+15,j,k), vec2(i+15,j,k))
end do
do i = end1_1+1, end1
erg = erg + DPROD(vec1(i,j,k), vec2(i,j,k))
end do
end do
end do
    
```

Prefetching
Loop unrolling (16)
Vector alignment

MAGMASolid matvec MAQAO Analysis (original loop)



Maqao Web Interface

CFG of matvec_

Loop 13 line: 821

Loop 13 line: 916

Loop 13 line: 877

Loop highly vectorised
Unrolled by 2 (assembly level)
Lots of movups instructions (unaligned vector load/store)

Wilfried

JARCHweigsdv2.s

```

948 movps -8(%r12,%rax), %xmm11 #1
949 movps -4(%r12,%rax), %xmm1 #19
950 movps 0(%r12,%r9), %xmm2 #197.4
951 movps -4(%r12,%r10), %xmm3 #19
952 movps -4(%r12,%rdi), %xmm2 #19
953 movps -4(%r12,%rax), %xmm4 #19
954 movps -4(%r12,%r11), %xmm4 #19
955 movps -4(%r12,%rdi), %xmm7 #20
956 movps -4(%r12,%rdi), %xmm6 #20
957 movps -4(%r12,%rax), %xmm9 #201
958 movps -4(%r12,%r14), %xmm8 #20
959 movps -4(%r12,%rdi), %xmm10 #21
960 movps 0(%r12,%rax), %xmm14 #1
961 movps 0(%r12,%rdi), %xmm15 #1
962 movps 0(%r12,%rdi), %xmm15 #1
963 movps 0(%r12,%r10), %xmm3 #198.40
964 movps 0(%r12,%rdi), %xmm2 #198.40
965 movps 0(%r12,%rdi), %xmm2 #198.40
966 movps 0(%r12,%rdi), %xmm2 #198.40
967 movps 0(%r12,%rdi), %xmm2 #198.40
968 movps 0(%r12,%rdi), %xmm2 #198.40
969 movps 0(%r12,%rdi), %xmm2 #198.40
970 movps 0(%r12,%rdi), %xmm2 #198.40
971 movps 0(%r12,%rdi), %xmm2 #198.40
972 movps 0(%r12,%rdi), %xmm2 #198.40
973 movps 0(%r12,%rdi), %xmm2 #198.40
974 movps 0(%r12,%rdi), %xmm2 #198.40
975 movps 0(%r12,%rdi), %xmm2 #198.40
976 movps 0(%r12,%rdi), %xmm2 #198.40
977 movps 0(%r12,%rdi), %xmm2 #198.40
978 movps 0(%r12,%rdi), %xmm2 #198.40
979 movps 0(%r12,%rdi), %xmm2 #198.40
980 movps 0(%r12,%rdi), %xmm2 #198.40
981 movps 0(%r12,%rdi), %xmm2 #198.40
982 movps 0(%r12,%rdi), %xmm2 #198.40
983 movps 0(%r12,%rdi), %xmm2 #198.40
984 movps 0(%r12,%rdi), %xmm2 #198.40
985 movps 0(%r12,%rdi), %xmm2 #198.40
986 movps 0(%r12,%rdi), %xmm2 #198.40
987 movps 0(%r12,%rdi), %xmm2 #198.40
988 movps 0(%r12,%rdi), %xmm2 #198.40
989 movps 0(%r12,%rdi), %xmm2 #198.40
990 movps 0(%r12,%rdi), %xmm2 #198.40
991 movps 0(%r12,%rdi), %xmm2 #198.40
992 movps 0(%r12,%rdi), %xmm2 #198.40
993 movps 0(%r12,%rdi), %xmm2 #198.40
994 movps 0(%r12,%rdi), %xmm2 #198.40
995 movps 0(%r12,%rdi), %xmm2 #198.40
996 movps 0(%r12,%rdi), %xmm2 #198.40
997 movps 0(%r12,%rdi), %xmm2 #198.40
998 movps 0(%r12,%rdi), %xmm2 #198.40
    
```

MAGMASolid matvec DECAN analysis



Impact of load/store instructions removal

matvec_omp_loop2_ORIGINAL

matvec_omp_loop2_stores

matvec_omp_loop2_stores_loads

matvec_omp_loop2_st_0x403e5a

matvec_omp_loop2_loads

matvec_omp_loop2_id_0x403e50

matvec_omp_loop2_id_0x403e36

matvec_omp_loop2_id_0x403e2c

matvec_omp_loop2_id_0x403e26

matvec_omp_loop2_id_0x403e1c

matvec_omp_loop2_id_0x403e12

matvec_omp_loop2_id_0x403e08

matvec_omp_loop2_id_0x403e02

matvec_omp_loop2_id_0x403df8

matvec_omp_loop2_id_0x403dee

matvec_omp_loop2_id_0x403de7

matvec_omp_loop2_id_0x403de1

matvec_omp_loop2_id_0x403dda

matvec_omp_loop2_id_0x403dd3

removing store instructions in the loop cuts down execution time by a factor of 2

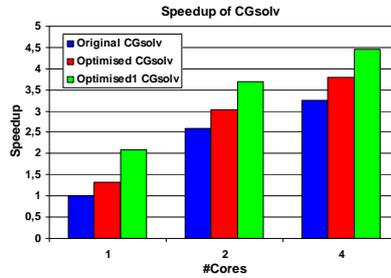
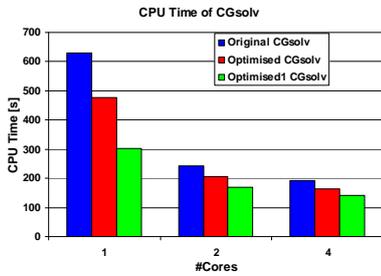
Reason: 4K-aliasing load-store conflict between the array to be stored and 2 other arrays

Solution: using different offsets for these three arrays.

Wilfried Schäfer, Karin Lukaszewicz – MAGMA GmbH

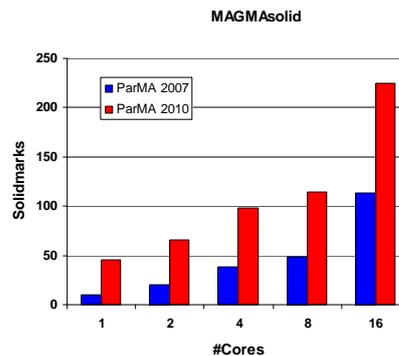
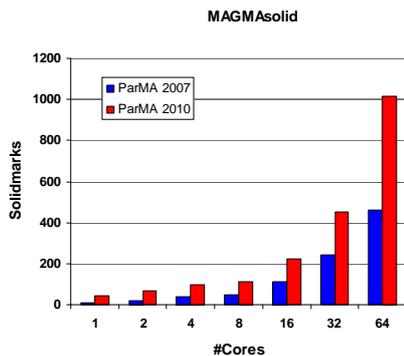
MAGMA

MAGMASolid Analysis of the CG-Solver with MAQAO



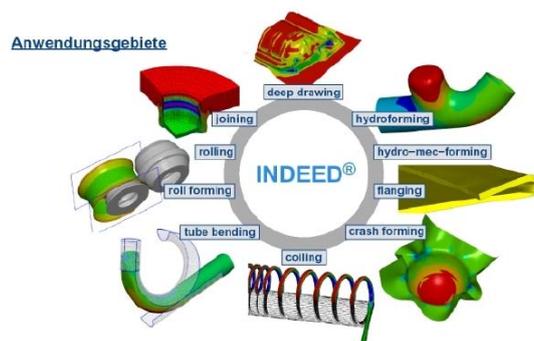
Full optimisation does not meet customers' requirement of having identical results

MAGMASolid Conclusion

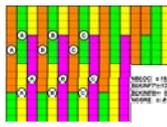
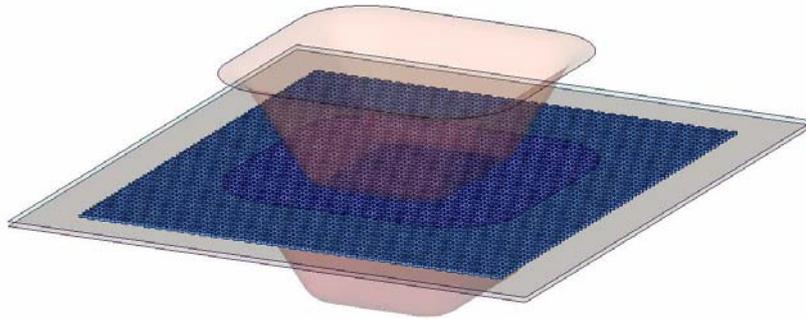


- achieved speedup for MAGMASolid more than a factor of 2
- switch to a single solver strategy by using the new MPI solver on all platforms → speed up further development.

- RECOM – combustion simulation
- MAGMA – casting simulation
- **GNS – metal forming simulation**



- Implicit Finite Element program for the simulation of metal forming processes.
- Vectorized data structures (Cray, NEC).
- Shared memory parallelism with OpenMP pragmas (since 2000).
- Distributed memory parallelism using MPI standard (started in 2005).

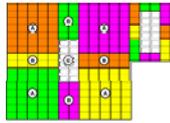


Analysis with Scalasca and VampirTrace showed many sequential amounts in the parallel Stiffness matrix assembly phase !

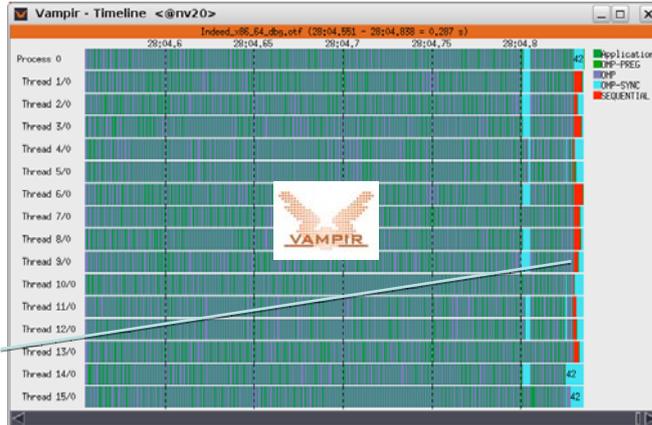
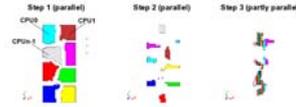


OpenMP Version: Improvements (2)

ParMA



- ① Core
- ② Realtime-Aufwandsindikator Subsystem
- ③ Realtime-Aufwandsindikator Zeitschleife
- ④ Realtime-Aufwandsindikator sich ändernde Zeitschleife



Minimized sequential amounts

K.Kassem, R.Menzel



25

OpenMP Version: Improvements (3)

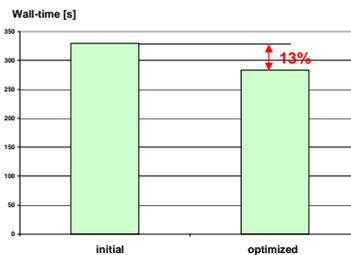
ParMA

MAQAO analysis: loop within compute intensive routine vectorized insufficiently

```

DO 480 J = 1, 6
  cDEC$ VECTOR ALIGNED
  DO 490 I = 1, NVELE -> 64
    VMAT(I,1,J) = -2.D+0*XK32 *
      &      ( QN(I,1)*NTX0(I,1) * XMAT(I,1,J)
      &      + QN(I,4)*NTX0(I,4) * XMAT(I,2,J)
      &      + QN(I,5)*NTX0(I,5) * XMAT(I,3,J)
      &      + (QN(I,4)*NTX0(I,1) + QN(I,1)*NTX0(I,4)) * XMAT(I,4,J)
      &      + (QN(I,5)*NTX0(I,1) + QN(I,1)*NTX0(I,5)) * XMAT(I,5,J)
      &      + (QN(I,5)*NTX0(I,4) + QN(I,4)*NTX0(I,5)) * XMAT(I,6,J) )
    !
    VMAT(I,2,J) = -2.D+0*XK32 *
      &      ( QN(I,7)*NTX0(I,7) * XMAT(I,1,J)
      &      + QN(I,2)*NTX0(I,2) * XMAT(I,2,J)
      &      + QN(I,6)*NTX0(I,6) * XMAT(I,3,J)
      &      + (QN(I,7)*NTX0(I,2) + QN(I,2)*NTX0(I,7)) * XMAT(I,4,J)
      &      + (QN(I,7)*NTX0(I,6) + QN(I,6)*NTX0(I,7)) * XMAT(I,5,J)
      &      + (QN(I,6)*NTX0(I,2) + QN(I,2)*NTX0(I,6)) * XMAT(I,6,J) )
    ...
  
```

Enforce vector alignment



K.Kassem, R.Menzel

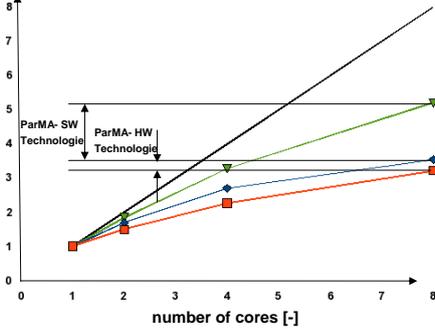


26

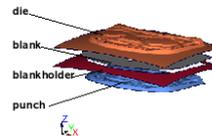
OpenMP Version: performance gain



Speed-up*[-]



Benchmark2



- Initial version on Itanium
- Initial version on Nehalem
- Optimized version on Nehalem

After restructuring the code runs up to **30% faster !!**

K.Kassem, R.Menzel

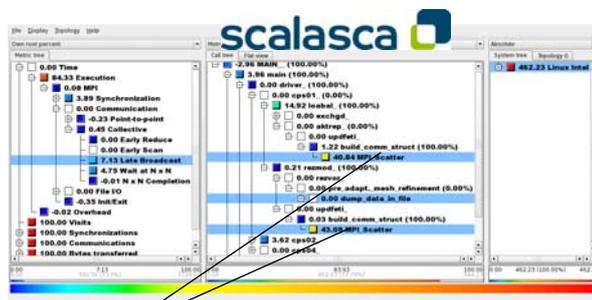


27

MPI Version: improvements



- ParMA tools showed no severe performance issues despite long computation times → inefficient algorithm.



- Integration of **different solver** leads to **significant stability + performance improvement**

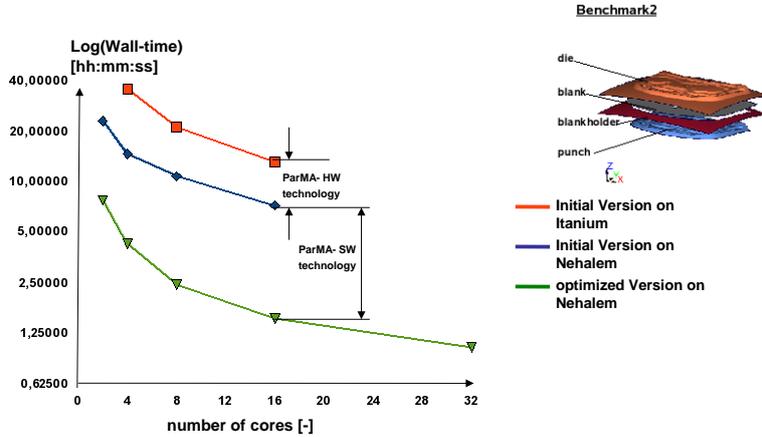
- Restructuring of rezoning & load balancing phase leads to **further performance improvement**

Rezoning & load balancing phases: Sequential amounts showed up as bottlenecks

K.Kassem, R.Menzel



28



Due to the modifications the code scales now up to 64 cores.
 On 16 cores now only **1.5h** are needed instead of **13h !!**

K.Kassem, R.Menzel



Thanks for your attention

Special thanks to all members of the ParMA project for the excellent work done in the past 3 years...

...and coming soon: our follow-up project
 H4H (Hybrid for Heterogeneous)

- **ParMA website:** <http://www.parma-itea2.org>
 - ParCo 2009, Minisymposium “Parallel Programming Tools for Multi-core Architectures”
 - Several papers presenting results with tools and applications
 - http://www.parma-itea2.org/index.php?option=com_content&task=view&id=67&Itemid=35
- **UNITE website:** <http://apps.fz-juelich.de/unite/>