SAND2015-1334PE

Comparison of CTH and miniAMR

Courtenay T. Vaughan

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Comparison with CTH

- Run on 128 cores
- CTH problem is a sphere that hits a block at an oblique angle and produces a shock wave
  - modeled in miniAMR as a deforming spheroid with an expanding hemisphere to represent the shock
- CTH averages 140.9 blocks/core over the run
  - average core has 16.3 messages per communication stage that average 261 KB
- miniAMR averages 141.9 blocks/core over the run
  - average core has 18.4 messages per communication stage that average 224 KB

Comparison to CTH continued

128 cores on Cielo	CTH	miniAMR
calculation	27.3%	35.4%
communication	61.5%	64.0%
refinement	11.2%	0.6%

In this comparison, communication includes just the communication of ghost cells between blocks. The calculation time includes calculation and communication necessary to calculate the timestep and other things using global reductions.

miniAMR Communication Patterns

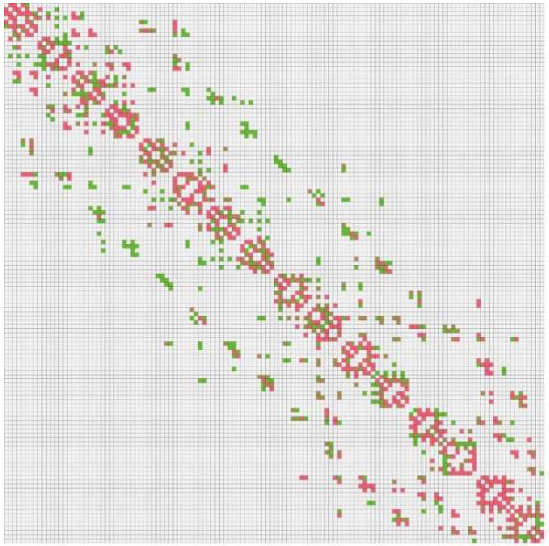

timestep 1748

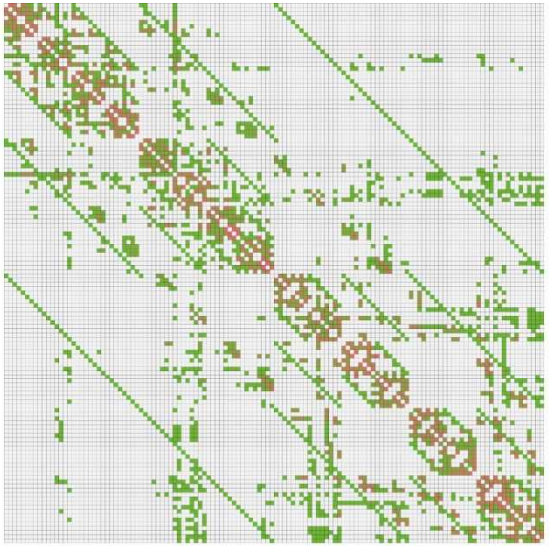

timestep 1749

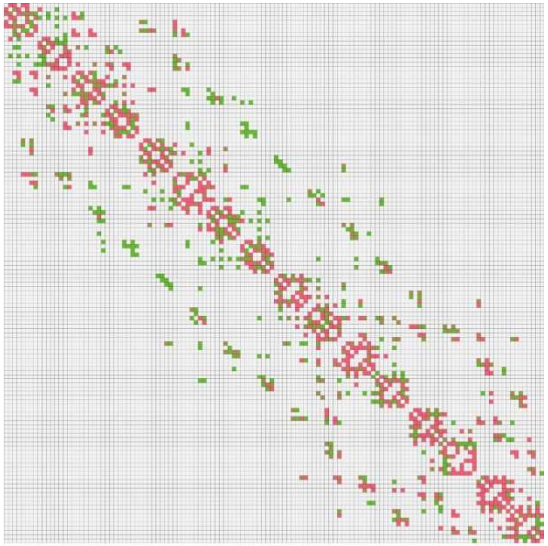

timestep 1750

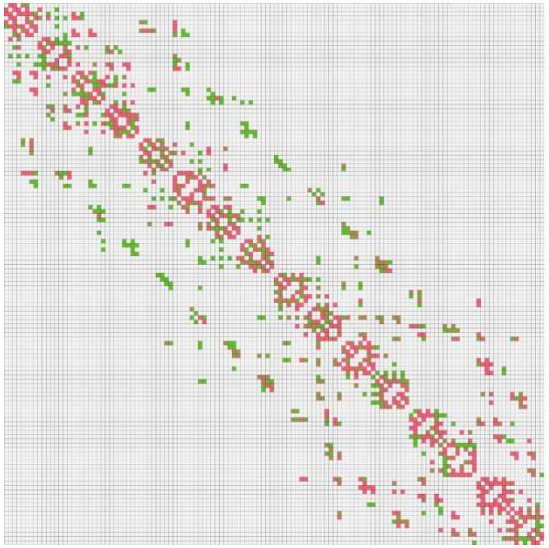

timestep 1749
ghost communication

timestep 1749
refine communicaton

626 messages
2.0 MB
average per core

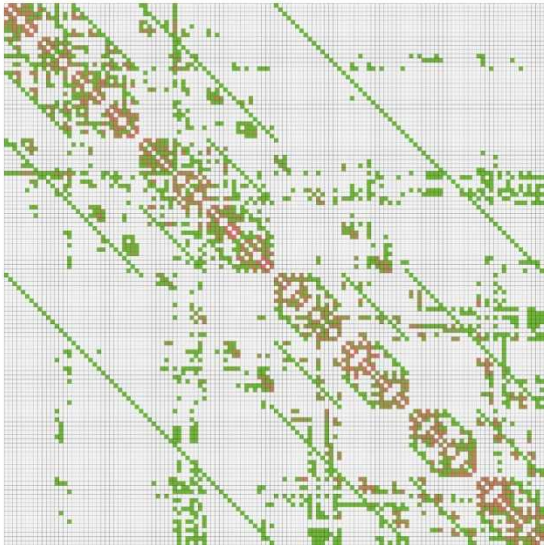

CTH Communication patterns

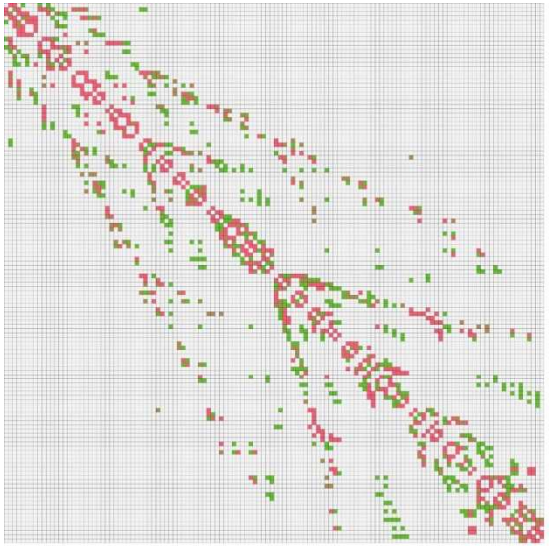

ts 3600 comm

ts 3601

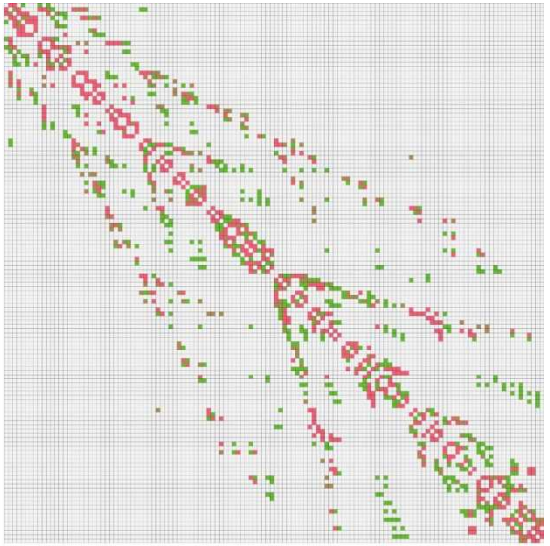

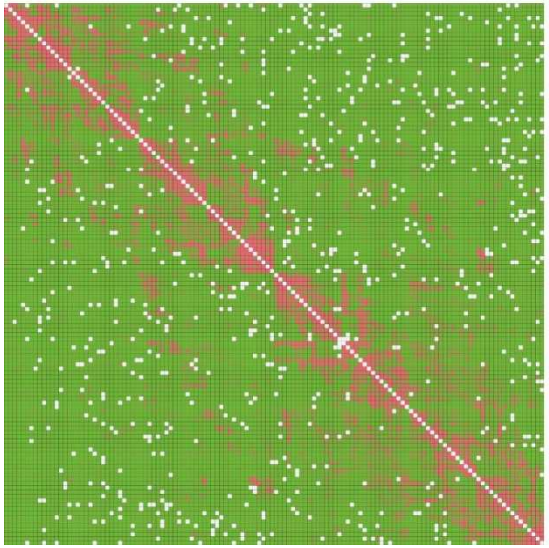

ts 3600 refine

21416 mesg
112 MB
per core

ts 3600

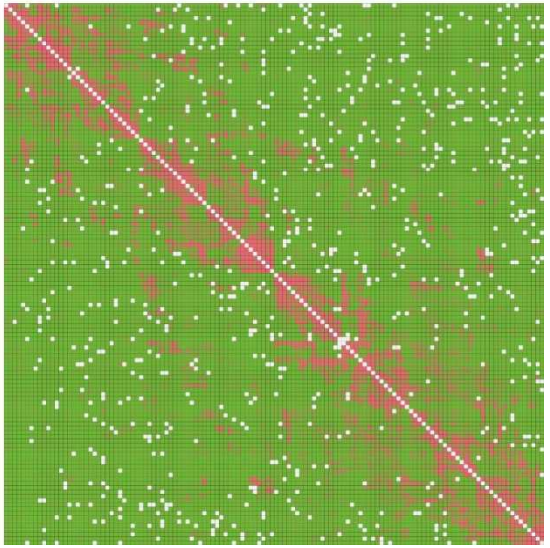

Communication Differences

- The CTH refinement step uses 34 times as many messages and communicates 56 times as much information as the miniAMR refinement step
- Both use RCB (recursive coordinate bisection) to do the load balancing, but ordering of the cuts is handled differently
- CTH also load balances parent blocks, which causes communication with more processors during the refinement step

Four Spheres Problem

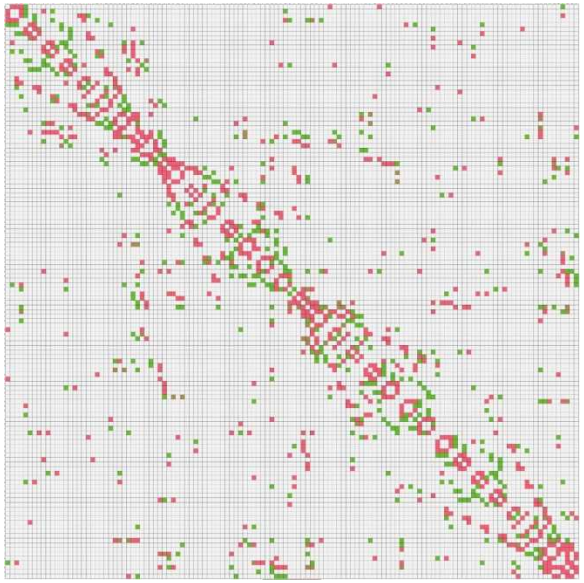

CTH
ts 1003

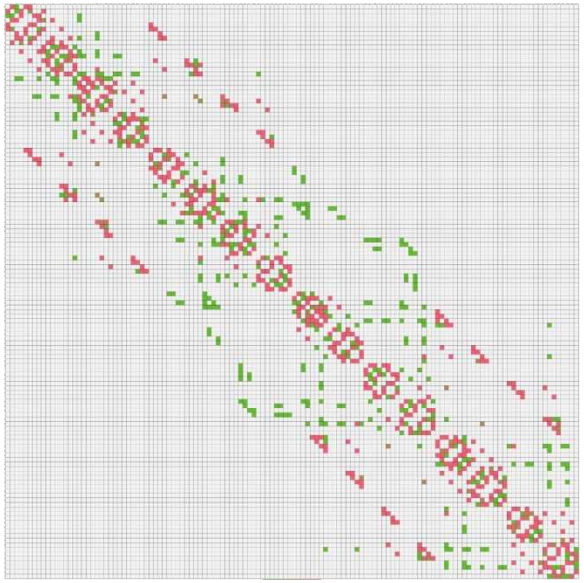

miniAMR (normal)
ts 122

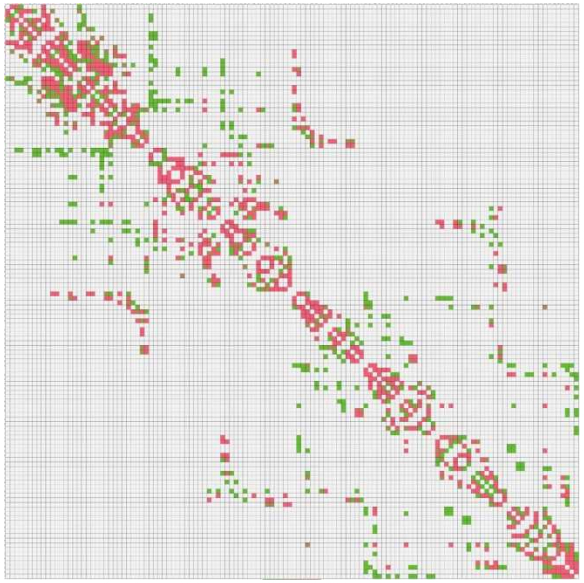

miniAMR (modified)
ts 122

Differences with CTH

- A CTH for a timestep involves calculations that are interspersed with ghost value communication and MPI_Allreduce calls
  - For the CTH problem mentioned, there are 16 communication steps in a timestep and 128 calls to MPI_Allreduce
- The calculation is much more regular in miniAMR than it is in CTH
- Communication has similar volumes, but is implemented differently
  - miniAMR maintains lists of faces and processors to communicate those with while CTH builds that information on the fly each communication

Differences with CTH - refinement

- miniAMR and CTH do the same basic communication operations: whether blocks are refining or unrefining with neighboring blocks (on and off core) and parents
- Both use RCB for load balancing
  - CTH also load balances parent blocks (they are complete blocks, where in miniAMR they do not contain array structures)

Calculation Comparison

- The following slides compare the calculation portion of the runs between CTH and miniAMR
- For CTH, we use one of the convection routines
  - Together these three routines use 20% to 25% of the computation time of the runs
  - These routines are fairly representative of the computations in CTH
- For miniAMR, use the stencil calculation

CrayPat counters - group 1

	CTH	miniAMR
PAPI_L1_DCM	5.132M/sec	13.559
PAPI_TLB_DM	0.072M/sec	0.378
PAPI_L1_DCA	550.807M/sec	1110.161
PAPI_FP_OPS	22.396M/sec	475.385
Computational intensity	0.01 ops/cycle	0.20
TLB utilization	7646.93 refs/miss	2938.87
D1 cache hit,miss ratios	99.1% hits	98.8%
D1 cache utilization	107.33 refs/miss	81.88

CrayPat Counters - Group 2

	CTH	miniAMR
PAPI_L1_DCM	5.355M/sec	13.719
PAPI_L1_DCA	558.250M/sec	1109.755
DATA_CACHE_REFILLS L2:	3.238M/sec	2.409
DATA_CACHE_REFILLS_SYSTEM:	2.388M/sec	22.093
D1 cache hit ratios	99.0% hits	98.8%
D2 cache hit ratio	57.6% hits	9.8%
D1+D2 cache hit ratio	99.6% hits	98.9%
D1+D2 cache utilization	245.63 refs/miss	89.71

CrayPat Counters - Groups 5 & 7

	CTH	miniAMR
PAPI_FML_INS	6.156M/sec	68.090
PAPI_FAD_INS	16.768M/sec	408.537
PAPI_FDV_INS	3.401M/sec	0.0
RETIRED_MMX_AND_FP_INSTRUCTIONS:		
PACKED_SSE2	79.518M/sec	553.360
DISPATCH_STALLS	408.896M/sec	1717.377
STALL_FOR_FPU_FULL	4.517M/sec	127.488
STALL_FOR_LS_FULL	150.160M/sec	1006.037

Sandia National Laboratories

CrayPat Counters - Groups 8 & 9

	CTH	miniAMR
PAPI_BR_TKN	165.535M/sec	76.875
PAPI_BR_MSP	3.162M/sec	1.477
PAPI_TOT_INS	1287.141M/sec	1755.939
PAPI_L2_ICM	0.110M/sec	0.038
PAPI_L1_ICA	384.626M/sec	358.463
INST_CACHE_MISSES	4.297M/sec	0.185
L1 Inst cache misses	4.196M/sec	0.155
I cache hit ratio	98.9% hits	100.0%