Parallelizing a micromagnetic program for use on multi-processor shared memory computers

Michael J. Donahue

National Institute of Standards and Technology
Gaithersburg, Maryland

5-May-2009

Abstract

Parallelization of a finite difference micromagnetic program on shared memory computer systems is studied. Efficiency is found to be limited by memory bandwidth, and techniques are introduced to reduce memory traffic. Computations are sped up by a factor of three with four processor cores, and a factor of over six with eight cores on some systems. This corresponds to a Karp-Flatt serial fraction of 3–7% for small core counts.

DISCLAIMER: The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology. All trademarks mentioned herein belong to their respective owners.

Parallel OOMMF

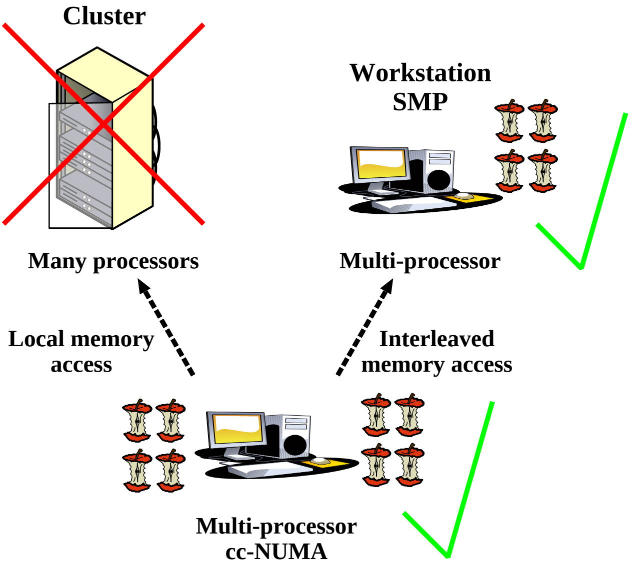

Platforms tested

- **2-way quad-core Intel Xeon X5365, 3 GHz**
  L1-cache I/D: 32 KB/32 KB 8-way set assoc.,
  L2-cache (2x): 4 MB 16-way set assoc.
  Memory access: Front-side bus
  Total cores: 8 Operating System: Mac OS X

- **4-way dual core AMD Opteron 885, 2.6 GHz**
  L1-cache I/D (2x): 64 KB/64 KB 2-way set assoc.,
  L2-cache (2x): 1 MB 16-way set assoc.
  Memory access: Point-to-Point
  Total cores: 8 Operating System: Linux

- **8-way dual core AMD Opteron 885, 2.6 GHz**
  L1-cache I/D (2x): 64 KB/64 KB 2-way set assoc.,
  L2-cache (2x): 1 MB 16-way set assoc.
  Memory access: Point-to-Point
  Total cores: 16 Operating System: Linux

Memory topologies

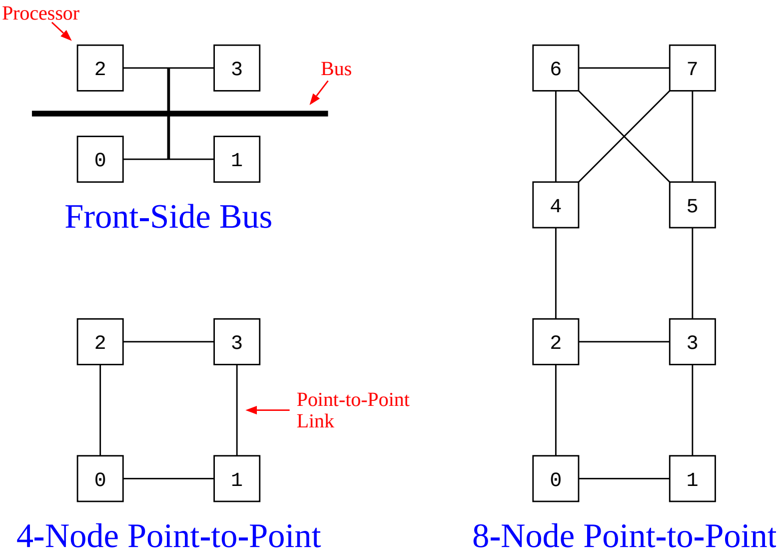

Computation components

- State-to-state transition
  (Stepper: LLG or energy minimization)

- Energy and field computation
  - "Local" energies (Zeeman, exchange, anisotropy)
  - Demagnetization

Local energies computation

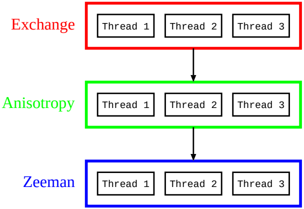

Sequential Computation
of Energy Terms

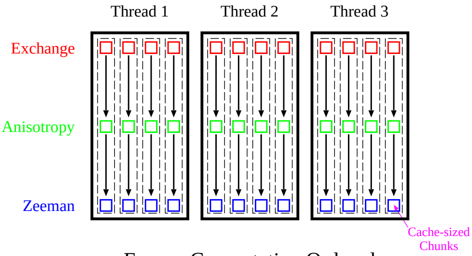

Energy Computation Ordered
by Cache-Sized Chunks
(reduces memory traffic)

Demag computation

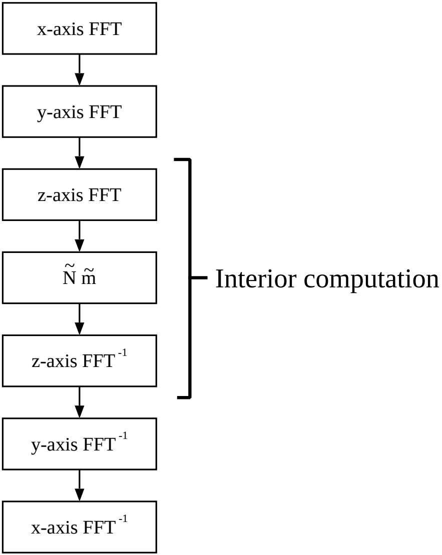

Zero padding

Naive cost: $8CN_xN_yN_z \log 8N_xN_yN_z$

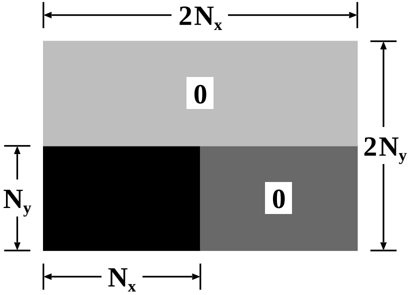

Skip zeros: relative savings ($N_x \geq N_y \geq N_z$):

$$1 - \frac{\log 2^7 N_x N_y^2 N_z^4}{\log 2^{12} N_x^4 N_y^4 N_z^4}$$

Savings is at least 41%, typically $> 50\%$.

Demag interior computation (a)

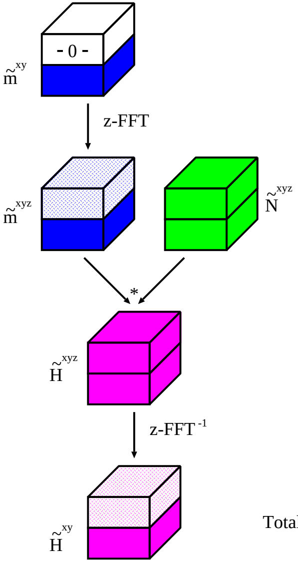

	Loads	Stores
	4	8
	8+4	8
	8	4+4
Totals:	24	24

Demag interior computation (b)

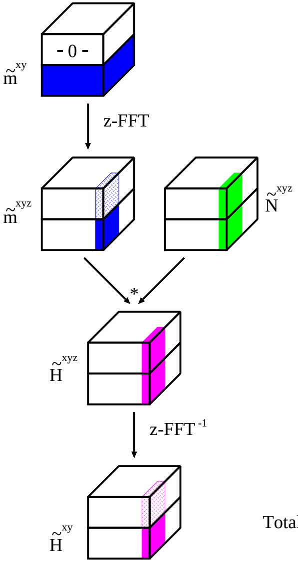

	Loads	Stores
z-FFT	4	~~8~~
*	~~8~~+4	~~8~~
z-FFT $^{-1}$	~~8~~	4+4
Totals:	~~24~~ 8	~~24~~ 8

Demag interior computation (c)

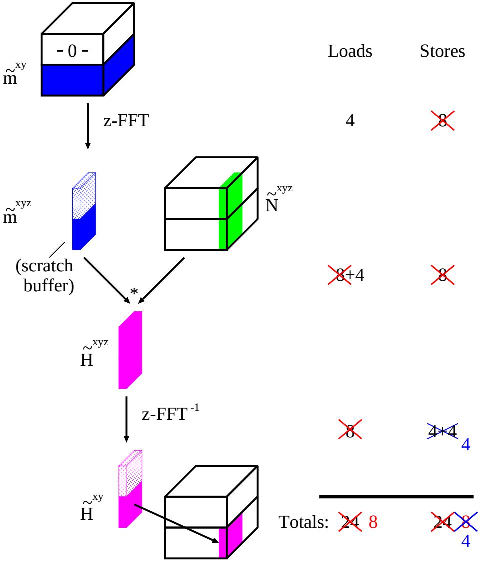

Parallelization results

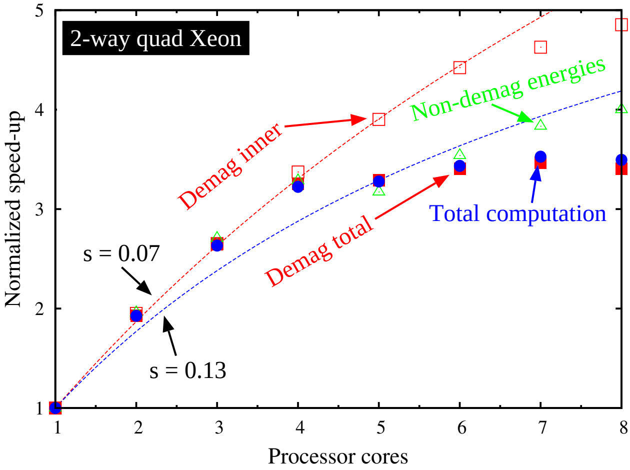

Dashed curves: $\frac{1}{s+(1-s)/n}$

Parallelization results

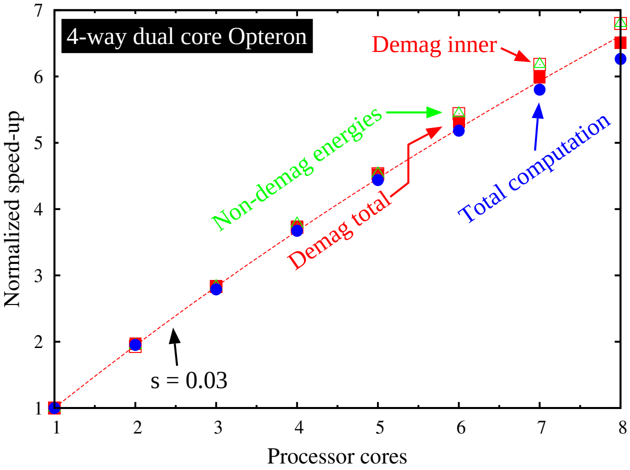

Dashed curves: $\frac{1}{s+(1-s)/n}$

Parallelization results

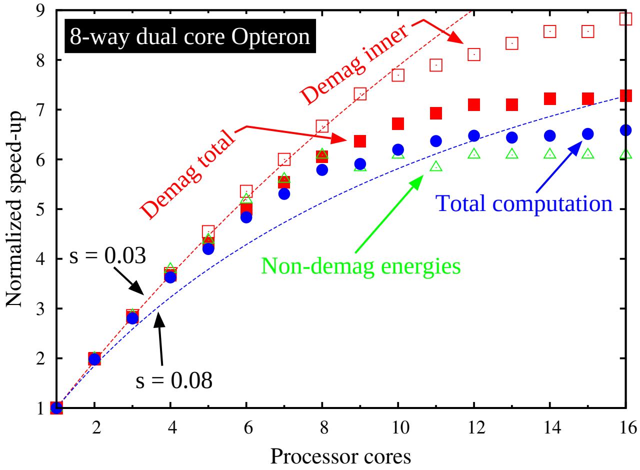

Dashed curves: $\frac{1}{s+(1-s)/n}$

Seconds/field eval for one million cell simulation

Intel Xeon X5365 (2-way quad core), 3 GHz				
Cores	Non-demag	Demag-inner	Demag Total	Step Total
1	0.092	0.199	0.579	0.790
2	0.047	0.102	0.300	0.410
3	0.034	0.075	0.219	0.300
4	0.028	0.059	0.178	0.245
5	0.029	0.051	0.176	0.241
6	0.026	0.045	0.170	0.230
7	0.024	0.043	0.167	0.224
8	0.023	0.041	0.170	0.226

Seconds/field eval for one million cell simulation

AMD Opteron 885 (4-way dual core), 2.6 GHz				
Cores	Non-demag	Demag-inner	Demag Total	Step Total
1	0.136	0.272	0.755	1.021
2	0.070	0.141	0.384	0.523
3	0.048	0.096	0.267	0.366
4	0.036	0.073	0.202	0.278
5	0.030	0.060	0.167	0.230
6	0.025	0.050	0.142	0.197
7	0.022	0.044	0.126	0.176
8	0.020	0.040	0.116	0.163

Seconds/field eval for one million cell simulation

AMD Opteron 885 (8-way dual core), 2.6 GHz				
Cores	Non-demag	Demag-inner	Demag Total	Step Total
1	0.140	0.300	0.859	1.146
2	0.070	0.151	0.429	0.579
3	0.049	0.105	0.304	0.409
4	0.037	0.081	0.234	0.316
5	0.032	0.066	0.199	0.273
6	0.027	0.056	0.172	0.237
7	0.025	0.050	0.155	0.216
8	0.023	0.045	0.142	0.198
9	0.024	0.041	0.135	0.194
10	0.023	0.039	0.128	0.185
11	0.024	0.038	0.124	0.180
12	0.023	0.037	0.121	0.177
13	0.023	0.036	0.121	0.178
14	0.023	0.035	0.119	0.177
15	0.023	0.035	0.119	0.176
16	0.023	0.034	0.118	0.174

Demagnetization computation details

Demag computation bandwidth requirements
(in "Real" units per spin node)

Component	Read	Write	R+W same	R+W diff
f-FFTx		2		8
f-FFTy		6	12	
Inner	12		24	
i-FFTy			24	
i-FFTx	10		8	

Demag results, one million spins ($500 \times 250 \times 8$)

Computed memory bandwidth time (secs/field eval)

Component	2-way Xeon	4-way Opteron	8-way Opteron
f-FFTx	0.018	0.014	0.015
f-FFTy	0.025	0.021	0.020
Inner	0.041	0.029	0.031
i-FFTy	0.025	0.021	0.020
i-FFTx	0.022	0.017	0.020
Total	0.131	0.102	0.106

Measured time (secs/field eval)

Component	2-way Xeon	4-way Opteron	8-way Opteron
f-FFTx	0.019	0.015	0.016
f-FFTy	0.035	0.022	0.022
Inner	0.041	0.040	0.034
i-FFTy	0.035	0.021	0.023
i-FFTx	0.023	0.019	0.021
Total	0.167	0.116	0.118

Measured memory bandwidth (GB/s)

Intel Xeon X5365 (2-way quad core), 3 GHz

Mode	Number of cores							
	1	2	3	4	5	6	7	8
Read only	3.6	5.4	5.1	5.7	5.6	5.6	5.7	5.7
Write only	2.2	3.7	3.4	3.8	3.8	3.6	3.8	3.8
R+W same	4.4	7.3	6.3	7.7	6.7	7.0	7.5	7.7
R+W different	2.8	4.4	3.9	3.5	4.4	4.5	4.5	4.6

Measured memory bandwidth (GB/s)

AMD Opteron 885 (4-way dual core), 2.6 GHz

Mode	Number of cores 1	2	3	4
— interleaved memory —				
Read only	3.4	4.6	5.9	7.6
Write only	1.9	2.7	3.6	4.6
R+W same	3.8	5.4	7.3	9.2
R+W different	2.5	3.8	4.5	6.1
— local memory —				
Read only	3.4	6.7	9.3	12.2
Write only	1.9	3.8	5.7	7.6
R+W same	3.8	7.6	11.4	15.1
R+W different	2.5	4.9	7.3	9.7

Measured memory bandwidth (GB/s)

AMD Opteron 885 (8-way dual core), 2.6 GHz

Mode	Number of cores							
	1	2	3	4	5	6	7	8
— interleaved memory —								
Read only	2.2	4.2	4.6	5.7	4.8	5.4	5.5	6.0
Write only	1.8	2.5	3.1	4.0	3.8	4.5	4.3	4.7
R+W same	3.6	5.0	6.3	8.2	7.7	9.1	8.7	9.5
R+W diff	2.2	3.4	3.7	5.0	4.4	5.1	5.0	5.5
— local memory —								
Read only	2.2	4.4	6.3	8.1	6.1	6.9	7.5	9.0
Write only	1.9	3.5	5.2	6.8	6.1	7.0	7.7	9.0
R+W same	3.6	6.9	10.3	13.8	12.1	13.7	15.1	17.8
R+W diff	2.2	4.4	6.2	8.1	6.1	6.9	7.6	9.0

Conclusions

- The 2-way (8 core) Xeon box achieves a speed-up of 3.2 on 4 cores (80% efficiency), and 3.5 on 7 cores (50% efficiency). The speed-up is limited by the bandwidth of the front-side bus.
- The multi-node point-to-point memory architecture of the 4-way (8 core) Opteron box allows it to scale well up to 8 cores, with a total speed-up of 6.25 (78% efficiency).
- On this test, the 8-way (16 core) Opteron box runs slower than the equivalent 4-way box, probably due to cache-coherency issues. There is little improvement beyond 12 cores, where the speed-up is 6.5 (54% efficiency).