45th Aerospace Sciences Meeting and Exhibit, January 2007, Reno, NV

TIMINGS OF AN UNSTRUCTURED-GRID CFD CODE ON COMMON HARDWARE PLATFORMS AND COMPILERS

Fernando Camelli, Rainald Löhner and Eric L. Mestreau

CFD Center, Dept. of Computational and Data Science
M.S. 6A2, College of Science, George Mason University
Fairfax, VA 22030-4444, USA

Advanced Technology Group, SAIC, McLean, VA, USA

Timings have been conducted for several benchmark testcases using the same code on a variety of common hardware platforms and compilers. The results indicate surprisingly little variation in the performance given the considerable number of vendors, architectures and compilers. The largest differences amounted to less than 30% of run-times. For the single processor runs, an increase in cache size reduced run-times, though not dramatically (approximately 10%). Going from 32 bits to 64 bits (with the same clockspeed, cache size, memory and compiler) in most cases produced a gain of 10%, although in some cases no gain was recorded. Overall, the chip with the slowest clocktime (Intel It II at 1.50 GHz) achieved the best performance. In some cases, it beat the 64-bit, 3.40 GHz P4/Xeon machines by a considerable margin. For shared memory parallelism, the best scaling was achieved by the SGI Altix. This should not come as a surprise, as SGI's CCNUMA technology has matured over the last decade. For the AMD Opteron, the SUN compiler exhibited the best scaling.

I. INTRODUCTION

The continuous advance and update of computer hardware (microchips, memory, motherboards) and software (compilers, field solvers) make it almost impossible to compare the large variety of options available in the market. It is often difficult for a user of Computational Fluid Dynamics (CFD) solvers to assess the relative merit of these different offerings from published BLAS[2], LAPACK[5], SPEC[17] or NAS Parallel Benchmarks[14] figures. In order to gain some insight into the state of hardware and software options for scientific computing at the end of 2006, a series of CFD benchmarks were conducted on commonly used hardware and compilers.

II. CFD CODE

The CFD code used for the benchmarks was FEFLO. FEFLO was conceived as a general-purpose CFD code based on the following principles:

- Use of unstructured grids (automatic grid generation and mesh refinement);
- Finite element discretization of space;
- Separate flow modules for compressible and incompressible flows;
- ALE formulation for moving grids;

- Embedded or immersed formulation for dirty CAD/ cracks/ shock- structure interaction;
- Edge-based data structures for speed;
- Optimal data structures for different architectures;
- Bottom-up coding from the subroutine level to assure an open-ended, expandable architecture;
- Standard Fortran-77 for maximal portability.

The code has had a long history of relevant development and applications[1,6–13,15,16]. As far as timing is concerned, the CPU-intensive loops are over **edges** (fluxes, Riemann solvers, artificial viscosity, etc.). The code renumbers points, elements and edges to minimize cache-misses while avoiding memory contention for pipelining and shared-memory parallel execution.[6]

III. BENCHMARKS

In order to cover a variety of solver techniques and the associated algorithms, the following benchmarks were selected:

3.1 NACA0012 Wing This is an incompressible Euler case. The geometry, as well as a typical solution, are shown in Figures 1a,b. The angle of attack was set to $\alpha = 5^o$. The grids used had approximately 0.37 Mels and 1.44 Mels.

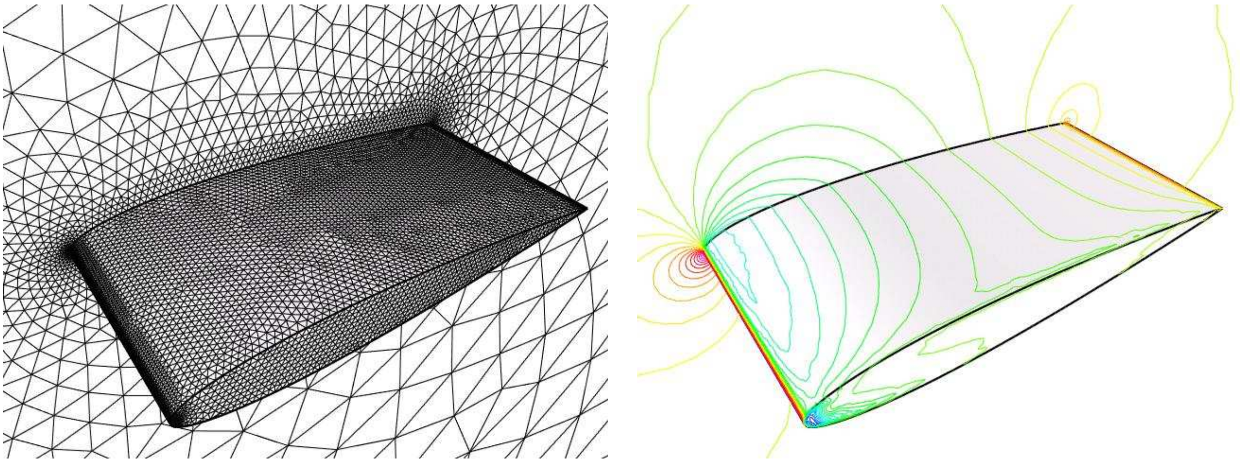

Figures 1a,b Surface Mesh and Pressure

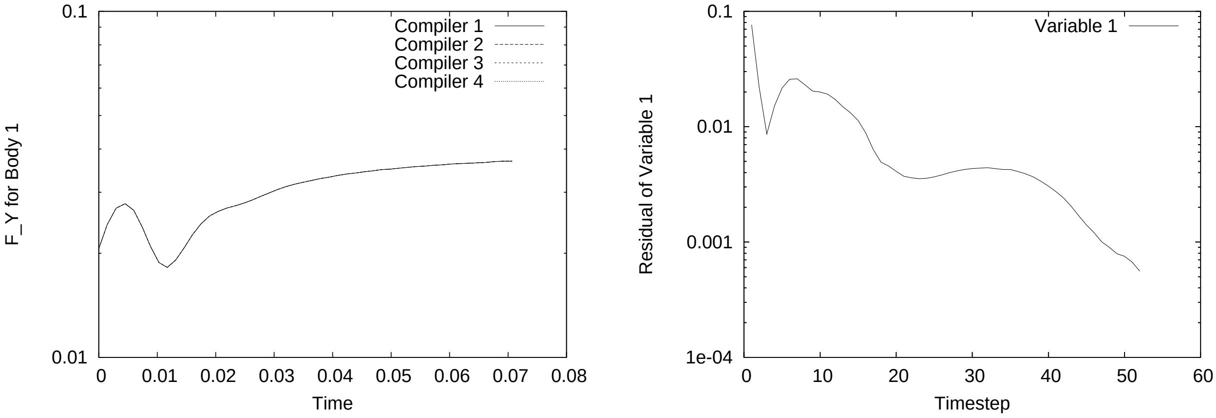

Figures 1c,d Lift Force and Residual History

The projection scheme[7] with three explicit stages per timestep[10,12] for the advective velocity prediction was run for 50 timesteps. At this point, for the coarser mesh, the lift force is converged to better than 2% as seen from Figure 1c, and the residual for pressures had dropped by more than 3 orders of magnitude. The main CPU intensive operations are the limiting of advective fluxes and the Poisson solver for the pressure.

3.2 Generic Missile This is a compressible Euler case. The geometry, as well as a typical solution, are shown in Figures 2a,b. The grids used had approximately 0.22 Mels and 2.53 Mels. The LU-SGS scheme[8, 9, 13, 16] with local timesteps, a Courant-nr. of $C = 10$ and 10 inner sweeps was run for 100 timesteps. At this point, the residual for the coarser grid has been lowered by five orders of magnitude (see Figure 2c). The lift force converges much quicker as seen from Figure 2d. The main CPU intensive operations are the limiting of fluxes, the HLLC-solver and the LU-SGS sweeps.

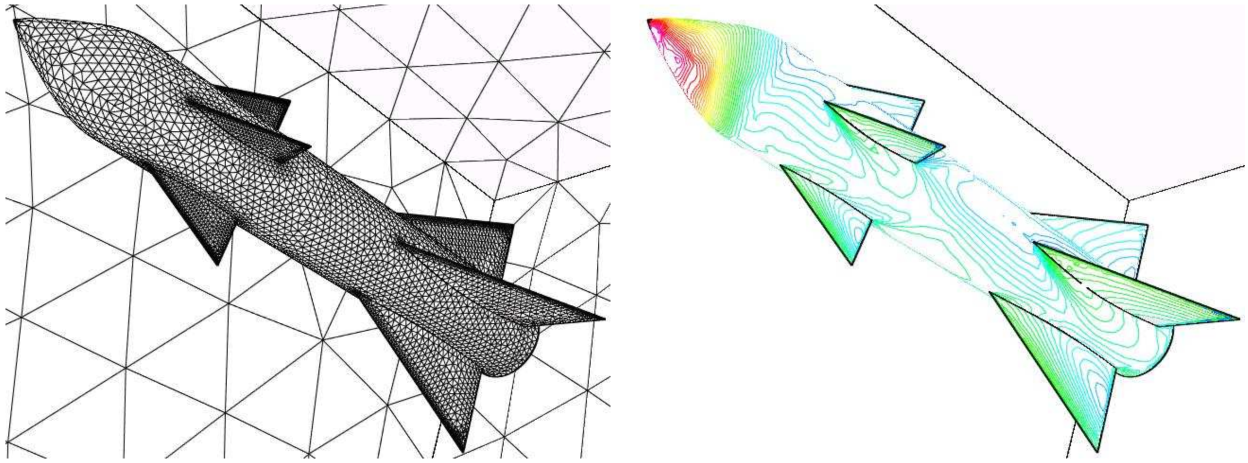

Figures 2a,b Surface Mesh and Pressure

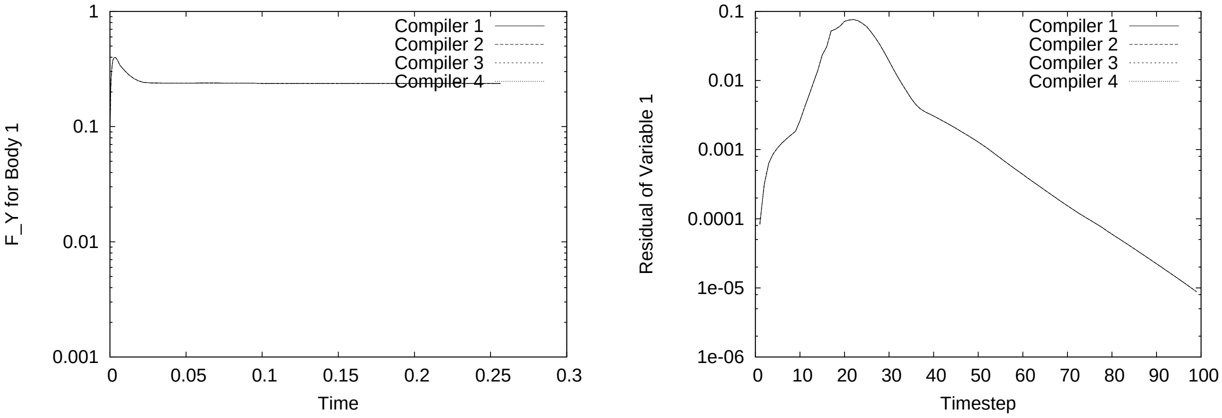

Figures 2c,d Lift and Residual History

3.3 LPD 17 Ship This is an incompressible Navier-Stokes case. The geometry, as well as a typical solution, are shown in Figures 3a,b. The wind is blowing from the front. The mesh used has approximately 16.6 Mels[3]. The projection scheme[7] with three explicit stages per timestep[10, 12] for the advective-diffusive velocity prediction was run for 50 timesteps. As with the first case, the main CPU intensive operations are the limiting of advective fluxes and the Poisson solver for the pressure.

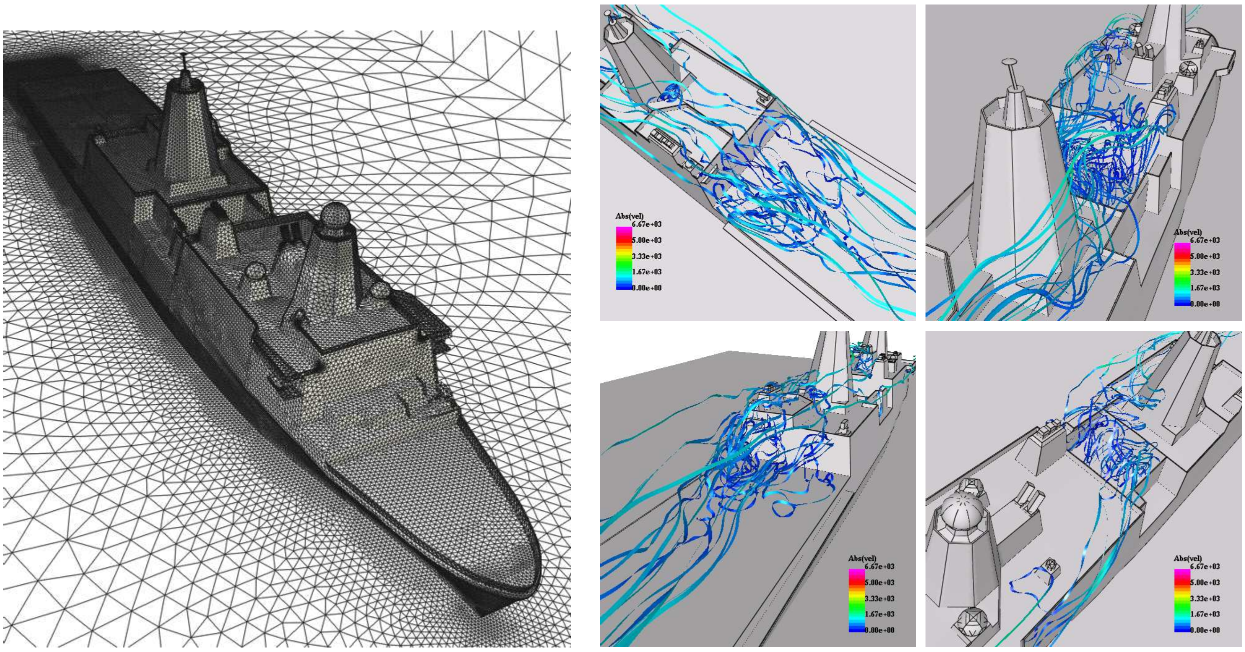

Figures 3a,b Surface Mesh and Stream Ribbons

IV. MACHINES

The machines that were used for the benchmarks, together with their characteristics, have been compiled in Table 1. Note that all the Intel P4/Xeon CPUs have hyperthreading, and that the AMD Opterons are dual core. We tried running in shared-memory mode (OpenMP) on the hyperthreaded machines, but the gains in speed were only of the order of 15%. However, on the same machines, running two separate jobs at the same time produced negligible degradation for each, i.e. the gain in throughput was almost 100%. The dual core AMD Opterons scaled very well in shared-memory mode (OpenMP).

V. COMPILERS

The following compilers and compiler flags were used for the benchmarks:

- Intel:

  fortran: ifort (build 20051201)

  cc: icc (build 20051201)

  flags (P4/X, AMD): `-tpp7 -132 -i4 -r8 -W1 -w95 -Vaxlib -zero -fpp -O3 -openmp`

  flags (Itanium II): `-ftz -w -extend_source -size_lp64 -i4 -r8 -align -WB -openmp -ivdep_parallel -tpp2 -fpp -O2`

- Portland Group: version 6.1-5 64-bit target on x86-64 Linux

  fortran: pgf90

  cc: pgcc

  flags: `-Mextend -i4 -r8 -tp k8-64 -fast -O3 -mp -Minfo=mp`

- SUN: sunstudio-20060314.linux.tar (alpha version)

  fortran: Sun Fortran 95 8.3 Build18.0 2006/03/13

Table 1. Machines used

Name	Vendor	CPU/Model	Clock (GHz)	Cache (MB)	Memory (GB)	# Cores	Distro	Version	kernel
quantz	DELL	Intel P4-32 Hyper Threading	3.4	1	4	1	Suse	9.3 x86	2.6.11.4-21.11-bigsmp
monteverdi	DELL	Intel P4-64 Hyper Threading	3.2	2	4	1	Suse	9.3 x86_64	2.6.11.4-21.11-smp
byrd	DELL	Intel Xeon Hyper Threading	3.4	2	4	1	Red Hat	ES4 x86_64	2.6.9-5.ELsmp
lassus	SGI	Intel Itanium II Altix 350	1.5	4	40	8	Suse	ES 9.0 ia64	2.6.5-7.252-sn2
bach	SGI	Intel Itanium II Altix 3700 Bx2	1.5	4	128	64	Suse	ES 9.0 ia64	2.6.5-7.252-sn2
nuke	SGI	Intel Itanium II Altix 4700	1.6	9	512	128x2	Suse	ES 10.0 ia64	2.6.16.21-0.25-default
handel	SGI	Intel Xeon 5160 Altix XE 210	3.0	4	16	2x2	Suse	ES 9.0 x86_64	2.6.5-7.244-smp
lully	Western Scientific	AMD Opteron 280 Dual Core	2.4	1	16	2	Suse	10.0 x86_64	2.6.13-15.8-smp
8-way	Western Scientific	AMD Opteron 875 Dual Core	2.2	1	32	8x2	CentOS	4.3 x86_64	2.6.13-15.8-smp

cc: Sun C 5.9 Build18.0 2006/03/13

flags: `-e -O3 -openmp`

- Gnu Fortran 95

  flags: `-fsloppy-char -i4 -r8 -fzero -O3`

All the multiprocessor cases were run in shared-memory (OpenMP) mode. This implied that the user did not have to go through the usual distributed-memory steps of splitting the domain, running in parallel and assembling the results at the end. The benchmarks shown are of fixed mesh topology, so they could have been run in distributed-memory mode (the CFD code has such options). However, many of the applications usually run with the CFD code used in the benchmarks require adaptive mesh refinement, remeshing, and regions of the flow where CPU requirements change drastically (e.g. particles or chemically reacting flows). For these cases, the distributed-memory option becomes rather involved, and the shared-memory option has been the preferred choice of users.

VI. TIMINGS

We have decided to show the timings in tabular form, as this is the only way to have all details present simultaneously. The list starts with runs carried out on one processor using the different hardware configurations and compilers, and then proceeds, in ascending number of processors, to the OpenMP shared-memory runs. These same tables have also been compiled in a different format at the end of the paper.

3.1 NACA0012 Wing (Incompressible, Euler)

We note that for this coarse mesh the largest differences in recorded run-times are due to compilers, cache size (10% gain) and 32 bits vs. 64 bits (another 10% gain). The SUN compiler, while worst on the uni-processor P4 machine, performs very well on the AMD Opteron, and also exhibits the best scaling for the OpenMP runs.

We note that for this fine mesh the results mirror, to a large extent, the runs on the coarse mesh. The 10% gain due to cache size increase is again present. The switch from 32 bits to 64 bits does not produce a performance increase. The reason is that the code has all real variables declared as `real*8`. As before, the SUN compiler, while worst on the uni-processor P4 machine, performs admirably on the AMD Opteron, and also exhibits the best scaling for the OpenMP runs. The SGI Altix, while powered by the (slow) Itanium II chip, achieves the best performance when using more than 2 processors.

3.2 Generic Missile (Compressible, Euler)

For this coarse mesh the largest differences in recorded run-times are due to compilers, cache size (10% gain) and 32 bits vs. 64 bits (another 10% gain). As in the first bechmark, the SUN compiler, while worst on the uni-processor P4 machine, performs very well on the AMD Opteron, and also exhibits the best scaling for the OpenMP runs.

For this fine mesh the results mirror, to a large extent, the runs on the coarse mesh. The 10% gain due to cache size increase, as well as the 10% gain from the 32/64 bits switch, are again present. As before, the SUN compiler, while worst on the uni-processor P4 machine, performs remarkably well on the AMD Opteron, and also exhibits the best scaling for the OpenMP runs. The SGI Altix, while powered by the (slow) Itanium II chip, achieves the best performance for any number of processors.

3.3 LPD 17 Ship

This large case did not fit into the machines with smaller memory. Therefore, we only report the results for the machines with larger memory.

For this fine mesh the Intel compiler performed best on the AMD Opteron, even for the multiprocessor runs. The larger memory SGI Altix while powered by the (slow) Itanium II chip, achieves the best performance, no matter how many processors are used. We also remark that this case clearly shows the (obvious) importance of large memory for large problems.

Table 2. Timings for NACA0012 (Incompressible, Euler, 370 Kels)

Vendor	CPU	Clock(GHz)	Cache(MB)	Mem(GB)	# Cores	Compiler	Time(sec)
Dell	Intel P4-32	3.4	1	4	1	ITL	60
Dell	Intel P4-32	3.4	1	4	1	PGI	65
Dell	Intel P4-32	3.4	1	4	1	SUN	73
Dell	Intel P4-32	3.4	1	4	1	G95	64
Dell	Intel P4-32	3.6	2	4	1	ITL	54
Dell	Intel P4-64	3.2	2	4	1	ITL	48
Dell	Intel Xeon	3.4	2	4	1	ITL	49
SGI	Intel It II	1.5	4	8	1	ITL	51
SGI	Intel It II	1.5	4	128	1	ITL	51
SGI	Intel It II	1.6	9	512	1	ITL	42
SGI	Intel Xeon 5160	3.0	4	16	1	ITL	39
Western	AMD Opteron 280	2.4	1	16	1	ITL	56
Western	AMD Opteron 280	2.4	1	16	1	PGI	51
Western	AMD Opteron 280	2.4	1	16	1	SUN	51
Western	AMD Opteron 875	2.2	1	32	1	ITL	87
SGI	Intel It II	1.5	4	8	2	ITL	29
SGI	Intel It II	1.5	4	128	2	ITL	29
SGI	Intel It II	1.6	9	512	2	ITL	25
SGI	Intel Xeon 5160	3.0	4	16	2	ITL	20
Western	AMD Opteron 280	2.4	1	16	2	ITL	32
Western	AMD Opteron 280	2.4	1	16	2	PGI	29
Western	AMD Opteron 280	2.4	1	16	2	SUN	27
Western	AMD Opteron 875	2.2	1	32	2	ITL	48
SGI	Intel It II	1.5	4	8	4	ITL	20
SGI	Intel It II	1.5	4	128	4	ITL	29
SGI	Intel It II	1.6	9	512	4	ITL	15
SGI	Intel Xeon 5160	3.0	4	16	4	ITL	17
Western	AMD Opteron 280	2.4	1	16	4	ITL	19
Western	AMD Opteron 280	2.4	1	16	4	PGI	18
Western	AMD Opteron 280	2.4	1	16	4	SUN	16
Western	AMD Opteron 875	2.2	1	32	4	ITL	24
SGI	Intel It II	1.5	4	8	8	ITL	11
SGI	Intel It II	1.6	9	512	8	ITL	10
Western	AMD Opteron 875	2.2	1	32	8	ITL	15
SGI	Intel It II	1.6	9	512	16	ITL	9
Western	AMD Opteron 875	2.2	1	32	16	ITL	12

Table 3. Timings for NACA0012 (Incompressible, Euler, 1.44 Mels)

Vendor	CPU	Clock(GHz)	Cache(MB)	Mem(GB)	# Cores	Compiler	Time(sec)
Dell	Intel P4-32	3.40	1	4	1	ITL	275
Dell	Intel P4-32	3.40	1	4	1	PGI	296
Dell	Intel P4-32	3.40	1	4	1	SUN	324
Dell	Intel P4-32	3.40	1	4	1	G95	295
Dell	Intel P4-32	3.60	2	4	1	ITL	244
Dell	Intel P4-64	3.20	2	4	1	ITL	243
Dell	Intel Xeon	3.40	2	4	1	ITL	245
SGI	Intel It II	1.50	4	8	1	ITL	247
SGI	Intel It II	1.50	4	128	1	ITL	250
SGI	Intel It II	1.60	9	512	1	ITL	197
SGI	Intel Xeon 5160	3.00	4	16	1	ITL	199
Western	AMD Opteron 875	2.20	1	32	1	ITL	417
Western	AMD Opteron 280	2.40	1	16	1	ITL	264
Western	AMD Opteron 280	2.40	1	16	1	PGI	241
Western	AMD Opteron 280	2.40	1	16	1	SUN	288
SGI	Intel It II	1.50	4	8	2	ITL	155
SGI	Intel It II	1.50	4	128	2	ITL	146
SGI	Intel It II	1.60	9	512	2	ITL	120
SGI	Intel Xeon 5160	3.00	4	16	2	ITL	141
Western	AMD Opteron 875	2.20	1	32	2	ITL	243
Western	AMD Opteron 280	2.40	1	16	2	ITL	192
Western	AMD Opteron 280	2.40	1	16	2	PGI	179
Western	AMD Opteron 280	2.40	1	16	2	SUN	151
SGI	Intel It II	1.50	4	8	4	ITL	108
SGI	Intel It II	1.50	4	128	4	ITL	111
SGI	Intel It II	1.60	9	512	4	ITL	72
SGI	Intel Xeon 5160	3.00	4	16	4	ITL	84
Western	AMD Opteron 875	2.20	1	32	4	ITL	138
Western	AMD Opteron 280	2.40	1	16	4	ITL	148
Western	AMD Opteron 280	2.40	1	16	4	PGI	143
Western	AMD Opteron 280	2.40	1	16	4	SUN	132
Western	AMD Opteron 875	2.20	1	32	8	ITL	66
SGI	Intel It II	1.50	4	8	8	ITL	60
SGI	Intel It II	1.60	9	512	8	ITL	42
Western	AMD Opteron 875	2.20	1	32	16	ITL	61
SGI	Intel It II	1.60	9	512	16	ITL	32
SGI	Intel It II	1.60	9	512	32	ITL	29

Table 4. Timings for Generic Missile (Compressible, Euler, 220 Kels)

Vendor	CPU	Clock(GHz)	Cache(MB)	Mem(GB)	# Cores	Compiler	Time(sec)
Dell	Intel P4-32	3.4	1	4	1	ITL	183
Dell	Intel P4-32	3.4	1	4	1	PGI	185
Dell	Intel P4-32	3.4	1	4	1	SUN	215
Dell	Intel P4-32	3.4	1	4	1	G95	209
Dell	Intel P4-32	3.6	2	4	1	ITL	158
Dell	Intel P4-64	3.2	2	4	1	ITL	123
Dell	Intel Xeon	3.4	2	4	1	ITL	128
SGI	Intel It II	1.5	4	8	1	ITL	98
SGI	Intel It II	1.5	4	128	1	ITL	100
SGI	Intel It II	1.6	9	512	1	ITL	79
SGI	Intel Xeon 5160	3.0	4	16	1	ITL	85
Western	AMD Opteron 875	2.2	1	32	1	ITL	215
Western	AMD Opteron 280	2.4	1	16	1	ITL	138
Western	AMD Opteron 280	2.4	1	16	1	PGI	136
Western	AMD Opteron 280	2.4	1	16	1	SUN	137
SGI	Intel It II	1.5	4	8	2	ITL	57
SGI	Intel It II	1.5	4	128	2	ITL	57
SGI	Intel It II	1.6	9	512	2	ITL	47
SGI	Intel Xeon 5160	3.0	4	16	2	ITL	64
Western	AMD Opteron 875	2.2	1	32	2	ITL	132
Western	AMD Opteron 280	2.4	1	16	2	ITL	83
Western	AMD Opteron 280	2.4	1	16	2	PGI	75
Western	AMD Opteron 280	2.4	1	16	2	SUN	84
SGI	Intel It II	1.5	4	8	4	ITL	58
SGI	Intel It II	1.5	4	128	4	ITL	65
SGI	Intel It II	1.6	9	512	4	ITL	41
SGI	Intel Xeon 5160	3.0	4	16	4	ITL	34
Western	AMD Opteron 875	2.2	1	32	4	ITL	86
Western	AMD Opteron 280	2.4	1	16	4	ITL	59
Western	AMD Opteron 280	2.4	1	16	4	PGI	N/A
Western	AMD Opteron 280	2.4	1	16	4	SUN	52
Western	AMD Opteron 875	2.2	1	32	8	ITL	65
Western	AMD Opteron 875	2.2	1	32	16	ITL	53

Table 5. Timings for Generic Missile (Compressible, Euler, 2.53 Mels)

Vendor	CPU	Clock(GHz)	Cache(MB)	Mem(GB)	# Cores	Compiler	Time(sec)
Dell	Intel P4-32	3.4	1	4	1	ITL	2294
Dell	Intel P4-32	3.4	1	4	1	PGI	2323
Dell	Intel P4-32	3.4	1	4	1	SUN	2758
Dell	Intel P4-32	3.4	1	4	1	G95	2637
Dell	Intel P4-32	3.6	2	4	1	ITL	2257
Dell	Intel P4-64	3.2	2	4	1	ITL	1859
Dell	Intel Xeon	3.4	2	4	1	ITL	1920
SGI	Intel It II	1.5	4	8	1	ITL	1627
SGI	Intel It II	1.5	4	128	1	ITL	1633
SGI	Intel It II	1.6	9	512	1	ITL	1449
SGI	Intel Xeon 5160	3.0	4	16	1	ITL	1489
Western	AMD Opteron 875	2.2	1	32	1	ITL	2919
Western	AMD Opteron 280	2.4	1	16	1	ITL	2187
Western	AMD Opteron 280	2.4	1	16	1	PGI	2337
Western	AMD Opteron 280	2.4	1	16	1	SUN	1821
SGI	Intel It II	1.5	4	8	2	ITL	1398
SGI	Intel It II	1.5	4	128	2	ITL	1065
SGI	Intel It II	1.6	9	512	2	ITL	848
SGI	Intel Xeon 5160	3.0	4	16	2	ITL	824
Western	AMD Opteron 875	2.2	1	32	2	ITL	1624
Western	AMD Opteron 280	2.4	1	16	2	ITL	1320
Western	AMD Opteron 280	2.4	1	16	2	PGI	1364
Western	AMD Opteron 280	2.4	1	16	2	SUN	1350
SGI	Intel It II	1.5	4	8	4	ITL	773
SGI	Intel It II	1.5	4	128	4	ITL	865
SGI	Intel It II	1.6	9	512	4	ITL	547
SGI	Intel Xeon 5160	3.0	4	16	4	ITL	601
Western	AMD Opteron 875	2.2	1	32	4	ITL	914
Western	AMD Opteron 280	2.4	1	16	4	ITL	1050
Western	AMD Opteron 280	2.4	1	16	4	PGI	1040
Western	AMD Opteron 280	2.4	1	16	4	SUN	939

Table 6. Timings for LPD 17 (Incompressible, Euler, 16.6 Mels)

Vendor	CPU	Clock(GHz)	Cache(MB)	Mem(GB)	# Cores	Compiler	Time(sec)
SGI	Intel It II	1.5	4	8	1	ITL	1543
SGI	Intel It II	1.5	4	128	1	ITL	1100
SGI	Intel It II	1.6	9	512	1	ITL	819
SGI	Intel Xeon 5160	3.0	4	16	1	ITL	771
Western	AMD Opteron 875	2.2	1	32	1	ITL	2212
Western	AMD Opteron 280	2.4	1	16	1	ITL	1260
Western	AMD Opteron 280	2.4	1	16	1	PGI	1222
Western	AMD Opteron 280	2.4	1	16	1	SUN	1335
SGI	Intel It II	1.5	4	8	2	ITL	909
SGI	Intel It II	1.5	4	128	2	ITL	642
SGI	Intel It II	1.6	9	512	2	ITL	472
SGI	Intel Xeon 5160	3.0	4	16	2	ITL	594
Western	AMD Opteron 875	2.2	1	32	2	ITL	1127
Western	AMD Opteron 280	2.4	1	16	2	ITL	712
Western	AMD Opteron 280	2.4	1	16	2	PGI	801
Western	AMD Opteron 280	2.4	1	16	2	SUN	842
SGI	Intel It II	1.5	4	8	4	ITL	473
SGI	Intel It II	1.5	4	128	4	ITL	406
SGI	Intel It II	1.6	9	512	4	ITL	259
SGI	Intel Xeon 5160	3.0	4	16	4	ITL	376
Western	AMD Opteron 875	2.2	1	32	4	ITL	632
Western	AMD Opteron 280	2.4	1	16	4	ITL	467
Western	AMD Opteron 280	2.4	1	16	4	PGI	586
Western	AMD Opteron 280	2.4	1	16	4	SUN	585
SGI	Intel It II	1.6	9	512	8	ITL	142
Western	AMD Opteron 875	2.2	1	32	8	ITL	336
SGI	Intel It II	1.6	9	512	16	ITL	79
Western	AMD Opteron 875	2.2	1	32	16	ITL	270
SGI	Intel It II	1.6	9	512	32	ITL	49
SGI	Intel It II	1.6	9	512	64	ITL	34
SGI	Intel It II	1.6	9	512	96	ITL	24

VII. CONCLUSIONS

Typical compressible and incompressible solvers/schemes were tested on a number of platforms using a variety of compilers.

As a general result/conclusion from the present benchmarks, the results indicate surprisingly little variation in the performance achieved using a variety of vendors, architectures and compilers. The largest differences amounted to less than 30% run-time variation. The low dependency of performance on the compiler used indicates that most of these take advantage of the hardware features offered, and that most of them have achieved near-optimal performance on these architectures. For the single processor runs, an increase in cache size reduced run-times, though not dramatically (approximately 10%). Similar findings were reported by Guler and Ali[4]. Going from 32 bits to 64 bits (with the same clockspeed, cache size, memory and compiler) in most cases produced a gain of 10%, although in some cases no gain was recorded.

Overall, the chip with the slowest clocktime (Intel It II at 1.50 GHz) achieved the best performance. In some cases, it beat the 64-bit, 3.40 GHz P4/Xeon machines by a considerable margin.

For shared memory parallelism, the best scaling was achieved by the SGI Altix. This should not come as a surprise, as SGI's CCNUMA technology has matured over the last decade. For the AMD Opteron, the SUN compiler exhibited the best scaling.

The dual core AMD Opterons scaled very well in shared-memory mode (OpenMP), even for the cases that required considerable memory (LPD-17 Ship).

VIII. AKNOWLEGMENTS

It is a pleasure to acknowledge the help of Stan Posey and John Gorski from SGI, as well as A.J. Jackson from Western Scientific, Inc., not only for providing computer resources, but also for guidance in setting compiler directives and system/ run-time environment variables which are sometimes less than obvious to the novice user.

References

[1] J.D. Baum, E. Mestreau, H. Luo, R. Löhner, D. Pelessone and Ch. Charman - Modeling Structural Response to Blast Loading Using a Coupled CFD/CSD Methodology; *Proc. Des. An. Prot. Struct. Impact/ Impulsive/ Shock Loads (DAPSIL)*, Tokyo, Japan, December (2003).

[2] `www.netlib.org/blas`

[3] F. Camelli, O. Soto, R. Löhner, W. Sandberg and R. Ramamurti - Topside LPD17 Flow and Temperature Study with an Implicit monolithic Scheme; *AIAA*-03-0969 (2003).

[4] B. Guler and R. Ali - Understanding the Behaviour of Computational Fluid Dynamics Applications on Dell PowerEdge 3250 Servers; *High Performance Computing*, 73-76 (2004), see also `www.dell.com/downloads`

[5] `www.netlib.org/lapack`

[6] R. Löhner - Renumbering Strategies for Unstructured-Grid Solvers Operating on Shared-Memory, Cache-Based Parallel Machines; *Comp. Meth. Appl. Mech. Eng.* 163, 95-109 (1998).

[7] R. Löhner, Chi Yang, E. Oñate and S. Idelssohn - An Unstructured Grid-Based, Parallel Free Surface Solver; *Appl. Num. Math.* 31, 271-293 (1999).

[8] R. Löhner, Chi Yang, J. Cebral, O. Soto, F. Camelli, J.D. Baum, H. Luo, E. Mestreau, D. Sharov, R. Ramamurti, W. Sandberg and Ch. Oh - Advances in FEFLO; *AIAA*-01-0592 (2001).

[9] R. Löhner, Chi Yang, J. Cebral, O. Soto, F. Camelli, J.D. Baum, H. Luo, E. Mestreau and D. Sharov - Advances in FEFLO; *AIAA*-02-1024 (2002).

[10] R. Löhner - Multistage Explicit Advective Prediction for Projection-Type Incompressible Flow Solvers; *J. Comp. Phys.* 195, 143-152 (2004).

[11] R. Löhner, J.D. Baum, E. Mestreau, D. Sharov, C. Charman and D. Pelessone - Adaptive Embedded Unstructured Grid Methods; *Int. J. Num. Meth. Eng.* 60, 641-660 (2004).

[12] R. Löhner, Chi Yang, J.R. Cebral, F. Camelli, O. Soto and J. Waltz - Improving the Speed and Accuracy of Projection-Type Incompressible Flow Solvers; *Comp. Meth. Appl. Mech. Eng.* 195, 23-24, 3087-3109 (2006).

[13] H. Luo, J.D. Baum and R. Löhner - A Fast, Matrix-Free Implicit Method for Compressible Flows on Unstructured Grids; *J. Comp. Phys.* 146, 664-690 (1998).

[14] `http://www.nas.nasa.gov/Software/NPB/`

[15] R. Ramamurti, W. Sandberg and R. Löhner - Simulation of Flow About Flapping Airfoils Using a Finite Element Incompressible Flow Solver; *AIAA*-99-0652 (1999).

[16] D. Sharov, H. Luo, J.D. Baum and R. Löhner - Implementation of Unstructured Grid GMRES+LU-SGS Method on Shared-Memory, Cache-Based Parallel Computers; *AIAA*-00-0927 (2000).

[17] `http://www.spec.org/`

Table 7. Timings for NACA0012 (Incompressible, Euler, 370 Kels) - 64 bits

CPUs	nuke It II 1.6 GHz 9MB/512GB		bach It II 1.5 GHz 4MB/128GB		lassus It II 1.5 GHz 4MB/8GB		lully Opteron 280 2.4 GHZ 1MB/16GB						8-way Opteron 875 2.2 GHz 1MB/32GB		handel Xeon 3.0 GHz 4MB/16GB	
	ITL		ITL		ITL		ITL		PGI		SUN		ITL		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	42	1.0	51	1.0	51	1.0	56	1.0	51	1.0	51	1.0	87	1.0	39	1.0
2	25	1.7	29	1.8	29	1.8	32	1.7	29	1.8	27	1.9	48	1.8	20	2.0
4	15	2.8	20	2.5	18	2.8	19	2.9	18	2.8	16	3.2	24	3.6	17	2.3
8	10	4.2	-	-	-	-	-	-	-	-	-	-	15	5.8	-	-
16	9	4.7	-	-	-	-	-	-	-	-	-	-	12	7.3	-	-

Table 8. Timings for NACA0012 (Incompressible, Euler, 370 Kels) - Hyper Threading

CPUs	monteverdi P4 3.2 GHz 2MB/4GB		quantz P4 3.4 GHz 1MB/4GB								byrd Xeon 3.4 GHz 2MB/4GB	
	ITL		ITL		PGI		SUN		G95		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	48	1.0	61	1.0	66	1.0	73	1.0	64	1.0	50	1.0
2	-	-	53	1.1	-	-	-	-	-	-	51	1.0

Table 9. Timings for NACA0012 (Incompressible, Euler, 1.44 Mels) - 64 bits

CPUs	nuke It II 1.6 GHz 9MB/512GB		bach It II 1.5 GHz 4MB/128GB		lassus It II 1.5 GHz 4MB/8GB		lully Opteron 280 2.4 GHZ 1MB/16GB						8-way Opteron 875 2.2 GHz 1MB/32GB		handel Xeon 3.0 GHz 4MB/16GB	
	ITL		ITL		ITL		ITL		PGI		SUN		ITL		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	197	1.0	250	1.0	247	1.0	264	1.0	241	1.0	288	1.0	417	1.0	199	1.0
2	120	1.6	146	1.7	155	1.6	192	1.4	179	1.3	151	1.9	243	1.7	141	1.4
4	72	2.7	111	2.3	108	2.3	148	1.8	143	1.7	132	2.2	138	3.0	84	2.4
8	42	4.7	60	4.1	-	-	-	-	-	-	-	-	66	6.3	-	-
16	32	6.2	-	-	-	-	-	-	-	-	-	-	61	6.8	-	-
32	29	6.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 10. Timings for NACA0012 (Incompressible, Euler, 1.44 Mels) - Hyper Threading

CPUs	monteverdi P4 3.2 GHz 2MB/4GB		quantz P4 3.4 GHz 1MB/4GB								byrd Xeon 3.4 GHz 2MB/4GB	
	ITL		ITL		PGI		SUN		G95		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	243	1.0	275	1.0	297	1.0	325	1.0	296	1.0	244	1.0
2	-	-	241	1.1	-	-	-	-	-	-	239	1.0

Table 11. **Timings for Generic Missile (Compressible, Euler, 220 Kels) - 64 bits**

CPUs	nuke It II 1.6 GHz 9MB/512GB		bach It II 1.5 GHz 4MB/128GB		lassus It II 1.5 GHz 4MB/8GB		lully Opteron 280 2.4 GHZ 1MB/16GB						8-way Opteron 875 2.2 GHz 1MB/32GB		handel Xeon 3.0 GHz 4MB/16GB	
	ITL		ITL		ITL		ITL		PGI		SUN		ITL		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	79	1.0	100	1.0	98	1.0	138	1.0	136	1.0	137	1.0	215	1.0	85	1.0
2	47	1.7	57	1.8	57	1.7	83	1.7	75	1.8	84	1.6	132	1.6	64	1.3
4	41	1.9	65	1.5	58	1.7	59	2.3	-	-	52	2.6	86	2.5	34	2.5
8	-	-	-	-	-	-	-	-	-	-	-	-	65	3.3	-	-
16	-	-	-	-	-	-	-	-	-	-	-	-	53	4.1	-	-

Table 12. Timings for Generic Missile (Compressible, Euler, 220 Kels) - Hyper Threading

CPUs	monteverdi P4 3.2 GHz 2MB/4GB		quantz P4 3.4 GHz 1MB/4GB								byrd Xeon 3.4 GHz 2MB/4GB	
	ITL		ITL		PGI		SUN		G95		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	183	1.0	187	1.0	181	1.0	215	1.0	210	1.0	131	1.0
2	-	-	161	1.2	-	-	-	-	-	-	140	0.9

Table 13. Timings for Generic Missile (Compressible, Euler, 2.53 Mels) - 64 bits

CPUs	nuke It II 1.6 GHz 9MB/512GB		bach It II 1.5 GHz 4MB/128GB		lassus It II 1.5 GHz 4MB/8GB		lully Opteron 280 2.4 GHZ 1MB/16GB						8-way Opteron 875 2.2 GHz 1MB/32GB		handel Xeon 3.0 GHz 4MB/16GB	
	ITL		ITL		ITL		ITL		PGI		SUN		ITL		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	1449	1.0	1633	1.0	1627	1.0	2187	1.0	2337	1.0	1821	1.0	2919	1.0	1489	1.0
2	848	1.7	1065	1.5	1398	1.2	1320	1.7	1350	1.7	1364	1.3	1624	1.8	824	1.8
4	547	2.6	865	1.9	773	2.1	1050	2.1	1040	2.2	939	1.9	914	3.2	601	2.5

Table 14. Timings for Generic Missile (Compressible, Euler, 2.53 Mels) - Hyper Threading

CPUs	monteverdi P4 3.2 GHz 2MB/4GB		quantz P4 3.4 GHz 1MB/4GB								byrd Xeon 3.4 GHz 2MB/4GB	
	ITL		ITL		PGI		SUN		G95		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	1859	1.0	2410	1.0	2329	1.0	2765	1.0	2637	1.0	1962	1.0
2	-	-	-	-	-	-	-	-	-	-	1775	1.1

Table 15. Timings for LPD17 (Incompressible, Euler, 16.6 Mels) - 64 bits

CPUs	nuke It II 1.6 GHz 9MB/512GB		bach It II 1.5 GHz 4MB/128GB		lassus It II 1.5 GHz 4MB/8GB		lully Opteron 280 2.4 GHZ 1MB/16GB						8-way Opteron 875 2.2 GHz 1MB/32GB		handel Xeon 3.0 GHz 4MB/16GB	
	ITL		ITL		ITL		ITL		PGI		SUN		ITL		ITL	
	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S	T(s)	S
1	819	1.0	1100	1.0	1534	1.0	1260	1.0	1222	1.0	1335	1.0	2212	1.0	771	1.0
2	472	1.7	642	1.7	909	1.7	712	1.8	801	1.5	824	1.6	1127	2.0	594	1.3
4	259	3.2	406	2.7	473	3.2	467	2.7	586	2.1	585	2.3	632	3.5	376	2.1
8	142	5.8	-	-	-	-	-	-	-	-	-	-	336	6.6	-	-
16	79	10.4	-	-	-	-	-	-	-	-	-	-	270	8.2	-	-
32	49	16.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-
64	34	24.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-
96	24	34.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-