Benchmarking HEP code on various machines

March 23, 2004

Three pieces of code were run. First was MCFOCUS, the FOCUS Monte Carlo and reconstruction program. Each test was run with 50,000 generic (ccbar) events with the standard trigger applied. The time reported by the program (CPU time) was used to determine the performance. The second program was BTeVGeant which is the BTeV Monte Carlo program. This program read in $B^0 \to J/\psi K^0_S$ events and processed them through the detector simulation. No additional minimum bias interactions were added to keep it simple. Finally, the simple HEP benchmark program TINY was run.

Two versions of the MCFOCUS library are available for Linux. One version is compiled with gcc version 2.95.2 while the other is compiled with gcc version 3.3.1. Also, TINY was compiled both with gcc version 2.95.2 and 3.3.1. In all cases, code compiled with gcc version 3.3.1 is faster than code compiled with gcc version 2.95.2. The run-to-run variation of the MCFOCUS benchmark is less than 0.5%. The run-to-run variations of the TINY benchmark are around 1%.

The standard BTeVGeant executable (`ge1`) is run with two different sets of switches. The first run is based on muon studies performed some time ago. This requires tracing particles through the muon system. The second run is based on studies which do not require muon information. The command file for the runs look like:

Run 1 (with muon system)

```
int_crossing     0.0
beam_position  0. 0. 0.
beam_size   0.005 0.005 30.
TRIG   50
KINE 1 0.0 0.0 0.0 0.0 0.0 0.0
DEBUG 1 0 0
SWIT 1 2 0 1 0 15 0
```

Run 2 (without muon system)

```
int_crossing       0.0
beam_position    0. 0. 0.
beam_size        0.005 0.005 30.
stop_point  wall  ecal muon dipole
stop_output richdump
trace_secondaries true
geom_calor projective
mideventanalysis false
sizeFSilPlane 27.
sizePixXPlane 7.0
RICHLiquid On
beamPipeMaterial  Beryllium
disableArmZMinus -75.
TRIG  500
KINE 1 0. 0. 0. 0. 0. 0.
SWIT 1 0 0 0 0 0 0
ECUT 0.001 global
ECUT 0.1 ecal
ECUT 0.100 muon
HCUT 0.010 local
```

Unfortunately, the run-to-run variations of the both BTeVGeant benchmarks were quite large. From a study of about 15 identical runs, the RMS variation was about 4%. These variations are random in the sense that different runs give different results without any apparant difference in the execution. However, during the run, CPU times of individual events are printed out. In the muon case, 14 of the 50 times are printed out while in the non-muon case 23 of the 500 times are printed out. Each of these times is affected exactly the same as the final time. That is, if the time difference between two runs is 2%, each of the individual time differences is also 2% (within measurement errors). Therefore, there is obviously some non-random element at work. Note also that the elapsed time reported by the `time` command agrees very well with the times reported by the BTeVGeant program. Because I don't understand the source of the variation, I don't know how much to trust the BTeVGeant numbers. This is why I resticted myself to two significant figures for the BTeVGeant results.

1 Performance

Processor	Clock (MHz)	GCC Ver	MCFOCUS Evt/s	MCFOCUS Evt/s/GHz	MCFOCUS Norm	BTeVGeant (muon sys) Evt/s	BTeVGeant (muon sys) Evt/s/GHz	BTeVGeant (muon sys) Norm	BTeVGeant (no muon sys) Evt/s	BTeVGeant (no muon sys) Evt/s/GHz	BTeVGeant (no muon sys) Norm	Tiny mvup	SPEC 2000 INT	SPEC 2000 FP
PIII 500	501	2.95.2	7.0	14.0	0.49	0.0075	0.015	0.53	0.10	0.20	0.53	414	231	191
PIII 550	551	2.95.2	7.6	13.8	0.53	0.0081	0.015	0.56	0.11	0.19	0.57	456		
PIII 800	800	2.95.2*	11.9	14.8	0.83	-	-	-	-	-	-	670	380	269
PIII 800	800	3.3.1	12.9	16.1	0.90	-	-	-	-	-	-	740	380	269
PIII 1000	998	2.95.2	14.4	14.4	**1.0**	0.014	0.014	**1.0**	0.19	0.19	**1.0**	830	457	297
PIIIS 1400	1400	2.95.2*	20.3	14.5	1.41	-	-	-	-	-	-	1170	648	437
PIIIS 1400	1400	3.3.1	22.3	15.9	1.55	-	-	-	-	-	-	1290	648	437
AMP 1200	1195	2.95.2	20.4	17.0	1.42	0.016	0.013	1.1	0.19	0.16	1.0	1110	495	433
AMP 1600+	1400	2.95.2	23.9	17.1	1.66	0.017	0.012	1.2	0.20	0.14	1.1	1300	550	483
AMP 1800+	1533	2.95.2*	26.4	17.2	1.83	-	-	-	-	-	-	1420	587	504
AMP 1800+	1533	3.3.1	29.2	19.1	2.03	-	-	-	-	-	-	1530	587	504
AMP 2400+	2001	2.95.2	33.1	16.6	2.30	0.020	0.010	1.4	0.20	0.10	1.1	1850	782	641
AXP 2400+	2001	2.95.2	33.6	16.8	2.33	0.022	0.012	1.5	0.23	0.12	1.3	1850	782	641
AXP 2400+	2001	2.95.2*	35.0	17.5	2.43	0.022	0.011	1.6	0.24	0.12	1.3	1850	782	641
AXP 2400+	2001	3.3.1	38.2	19.1	2.65	-	-	-	-	-	-	1980	782	641
P4 1700	1695	2.95.2	17.2	10.2	1.19	-	-	-	-	-	-	1160	582	659
P4 2000	1994	2.95.2	21.2	10.6	1.48	-	-	-	-	-	-	1420	744	773
P4X 2000	1994	2.95.2	21.4	10.7	1.49	0.020	0.010	1.4	0.24	0.12	1.3	1420	744	773
P4 2400	2387	2.95.2	26.3	11.0	1.83	-	-	-	-	-	-	1640	865	813
P4H 2800	2806	2.95.2*	31.2	11.1	2.17	0.033	0.012	2.3	0.42	0.15	2.2	1930	1031	1031
P4H2 2800	2806	2.95.2*	35.4	12.6	2.46	0.045	0.016	3.1	0.52	0.19	2.7	~2800	1031	1031
P4H 2800	2806	3.3.1	35.3	12.6	2.45	-	-	-	-	-	-	1980	1031	1031
P4H2 2800	2806	3.3.1	38.8	13.8	2.69	-	-	-	-	-	-	~2900	1031	1031
Opt 240	1395	2.95.2*	26.9	19.3	1.87	0.031	0.022	2.2	0.46	0.33	2.4	1350	880	934
Opt 240	1395	3.3.1	31.7	22.7	2.20	-	-	-	-	-	-	1490	880	934
Opt 244	1793	2.95.2*	34.2	19.1	2.38	0.038	0.021	2.7	0.56	0.31	3.0	1660	1095	1122
Opt 244	1793	3.3.1	39.7	22.2	2.76	-	-	-	-	-	-	1780	1095	1122

Table 1: Results of running 50,000 MCFOCUS events or 50 (500) BTeVGeant events. Tiny and SPEC CPU2000 results are also presented. The GCC version indicates how the code was compiled. In general the code was run under the same platform. GCC 2.95.2 corresponds to Red Hat 7.2 while GCC 3.3.1 corresponds to Red Hat 9. The 2.95.2* runs were compiled with GCC 2.95.2 but run under Red Hat 9. The "Norm" column is referenced to the 1 GHz Pentium III score.

Description of processors. The three numbers in parantheses are the cpu family, model, and stepping as returned by /proc/cpuinfo.

- PIII 500 is a self assembed dual processor Pentium III at 500 MHz and 512 kB of 1/2 speed L2 cache (6,7,3).
- PIII 550 is a self assembed dual processor Pentium III at 550 MHz and 512 kB of 1/2 speed L2 cache (6,7,3).

- PIII 800 is a self assembed dual processor Pentium III at 800 MHz and 256 kB of L2 cache (6,8,3).
- PIII 1000 is a self assembed dual processor Pentium III at 1 GHz and 256 kB of L2 cache (6,8,6) in a VIA
- PIIIS 1400 is a self assembled dual processor Pentium III at 1.4 GHz and 512 kB of L2 cache (6,11,1–4).
- AMP 1800+ is self assembled dual processor Athlon at 1.533 GHz and 256 kB of L2 cache (6,6,2). motherboard with PC133 RAM.
- AMP 1200 is a self assembled dual processor Athlon MP at 1.2 GHz and 256 kB of L2 cache (6,6,1) with a Tyan motherboard and PC2100 RAM.
- AMP 1600+ is a self assembled dual processor Athlon MP at 1.4 GHz and 256 kB of L2 cache (6,6,2).
- AXP 2400+ is a self assembled single processor Athlon XP at 2.0 GHz and 256 kB of L2 cache (6,8,1).
- AMP 2400+ is a self assembled dual processor Athlon MP at 2.0 GHz and 256 kB of L2 cache (6,8,1).
- P4 1700 is Dell Precision Workstation 340 with a 1.7 GHz Pentium 4 processor and 256 kB of L2 cache (15,1,2).
- P4 2000 is Dell Precision Workstation 340 with a 2.0 GHz Pentium 4 processor and 512 kB of L2 cache (15,2,4).
- P4X 2000 is an IBM workstation with two Xeon processors at 2.0 GHz and 512 kB of L2 cache (15,2,7).
- P4 2400 is Dell Precision Workstation 340 with a 2.4 GHz Pentium 4 processor and 512 kB of L2 cache (15,2,7).
- P4H 2800 is a self assembled single processor Pentium 4 with Hyperthreading at 2.8 GHz and 512 kB of L2 cache (15,2,5).
- P4H2 2800 is a self assembled single processor Pentium 4 with Hyperthreading at 2.8 GHz and 512 kB of L2 cache (15,2,5) running two copies of the program.
- Opt 240 is an AMD Opteron dual processor system at 1.4 GHz and 1 MB of L2 cache (15,5,1).
- Opt 244 is an AMD Opteron dual processor system at 1.8 GHz and 1 MB of L2 cache (15,5,1).

2 Analysis

Generally, it appears that the AMD Athlon machines are best for running MCFOCUS. An Athlon MP 2800+ at $170 is probably as good as the top of the line dual processors Opterons and Pentium 4 processors at $800. For the BTeVGeant code, however, the newest Pentium 4 processors and the Opeteron systems seem to perform much better. The simultaneous multi threading (SMT) ability of the newer Pentium 4 processors (Hyperthreading) increases the performance by 10% for MCFOCUS and 25–30% for BTeVGeant when two jobs are run on a single processor system. Since the MCFOCUS and BTeVGeant memory footprints are quite small (less than 50 MB) this can easily be done. If these jobs were ever memory limited then one would require twice as much memory per processor compared to a non SMT machine.