Gaussian G09 Scaling Benchmarks

Jemmy Hu
SHARCNET
June-July, 2009

Systems:

Name	CPUs/node	RAM/node	OS	Interconnect
Saw	8 (2 quad-core), Xeon @ 3.0 GHz	16.0 GB	HP Linux XC 4	InfiniBand
Narwhal	4 (2 dual-core), Opteron @ 2.2 GHz	8.0 GB	HP Linux XC 3.1	Myrinet 2g (gm)
Silky	SGI, Itanium2 @ 1.6 GHz	256GB	SUSE Enterprise 10	SMP, NUMA
Hound	16 (4 quad) Xeon@2.4GHz, 32 (8 quad) Opteron@2.2 GHz	128 GB	Centos 5	InfiniBand, NFS storage file system

Molecules and Methods/Models:

	Molecule\Module	B3LYP	MP2	CISD	CCSD
		Opt + Freq	Opt + Freq	Opt + Freq	Opt +Freq
I	**C4H14Cl2P2Pd** (test job 445)	BS on card	BS on card		
III	**CH3OH** (test job 58)			6-311g(2df,p)	
IV	**CH3CH2** (test job 684)				6-311g*, 6-311g(2df,p)

Gaussian versions:

Gaussian versions		
G09-A.01	Binary versions from Gaussian Inc	
G09-A.02	Compiled from source on Silky	Binaries for others
G03-E.01	Binary versions from Gaussian Inc	

Target goals:

[1] Scaling results for typical models/methods in Gaussian 09
[2] Scaling on different systems: clusters (saw, narwhal, hound) vs. SMP (silky)
[3] G03 vs G09

General conclusions:

1. Gaussian 09 scales quite good for shared memory jobs.

   Silky (SMP machine): DFT type of methods scale very good to 16 processors
   (small speedup from 16 to 32 processors)
   MP2 type of methods scale very good to 8 CPUs
   (small speedup from 8 to 16 processors)

   Saw (8-cpu nodes): DFT scale good to 8 processors
   MP2 scales to 4 processors (small speedup for 8 processors)

2. Gaussian does not scale for CI and CC based methods.

3. G09 is about 2 times faster than G03 for DFT, CI and CC based methods.

Maximum processors for G09 jobs

(In practice, in order to run more jobs on a system, smaller cpus/size jobs are recommended)

[1] Silky (SMP machine)

Methods/Modules	Opt	Freq	Energy
HF	16	16	16
DFT (B3LYP, etc)	16	16	16
MP(2, 3, 4)	8	8	8
CISD (cis, cid, cisd, qcisd)	1	1	1
CCSD (ccd, ccsd, ccsd(t))	1	1	1

[2]Saw (2 quad-core nodes)

Methods/Modules	Opt	Freq	Energy
HF	8	8	8
DFT (B3LYP, etc)	8	8	8
MP(2, 3, 4)*	8	8	8
CISD (cis, cid, cisd, qcisd)	1	1	1
CCSD (ccd, ccsd, ccsd(t))	1	1	1

*due to the 1 node per job LSF nature, run 8-way MP2 on saw is fine. If a node can be shared by multiple jobs (torque on hound), 4-way MP2 jobs are recommended.

[3] bull, goblin (and other 4-core node XC clusters)

Methods/Modules	Opt	Freq	Energy
HF	4	4	4
DFT (B3LYP, etc)	4	4	4
MP(2, 3, 4)	4	4	4
CISD (cis, cid, cisd, qcisd)	1	1	1
CCSD (ccd, ccsd, ccsd(t))	1	1	1

Results on saw, molecule-I

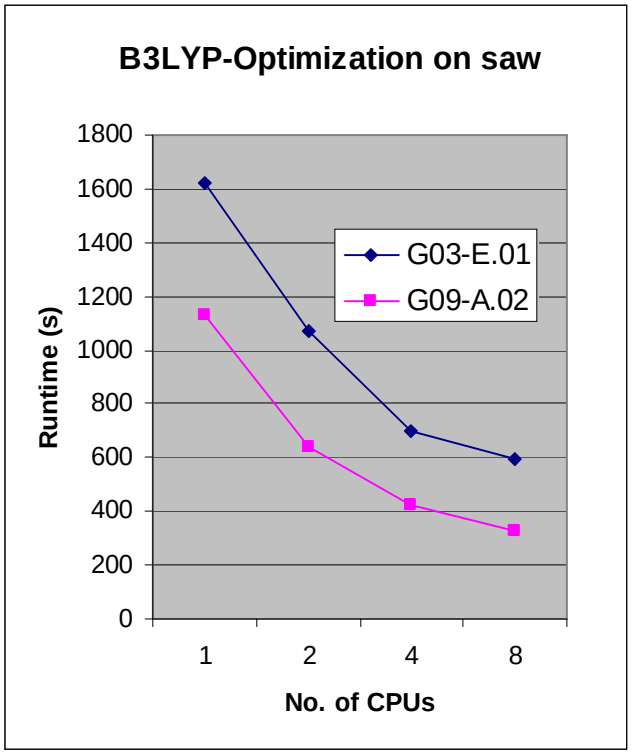

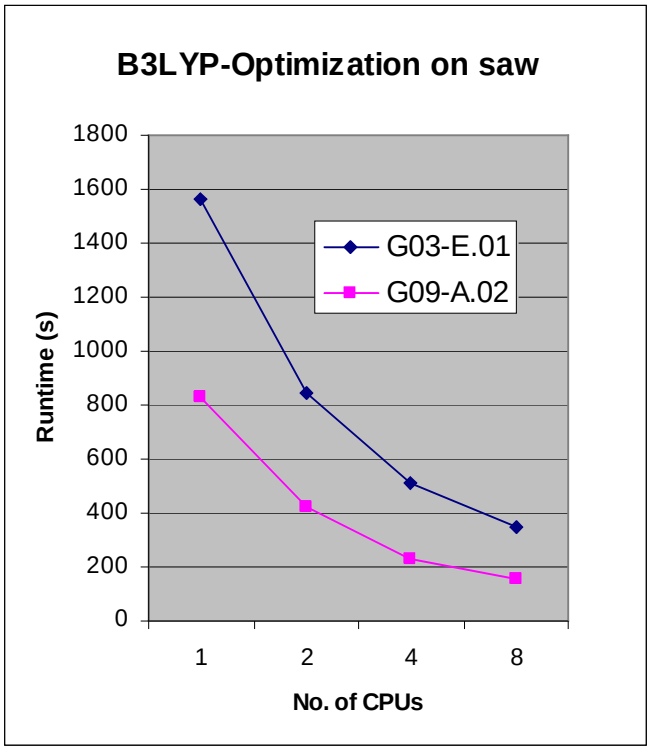

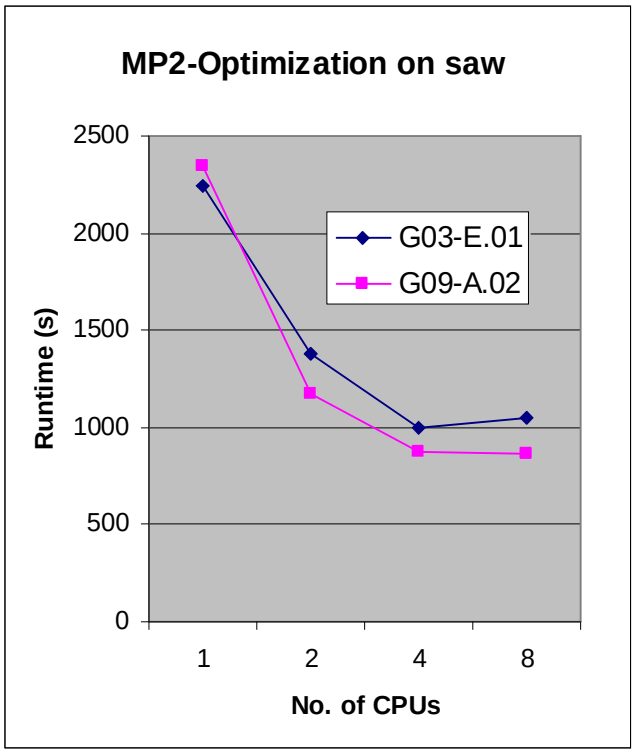

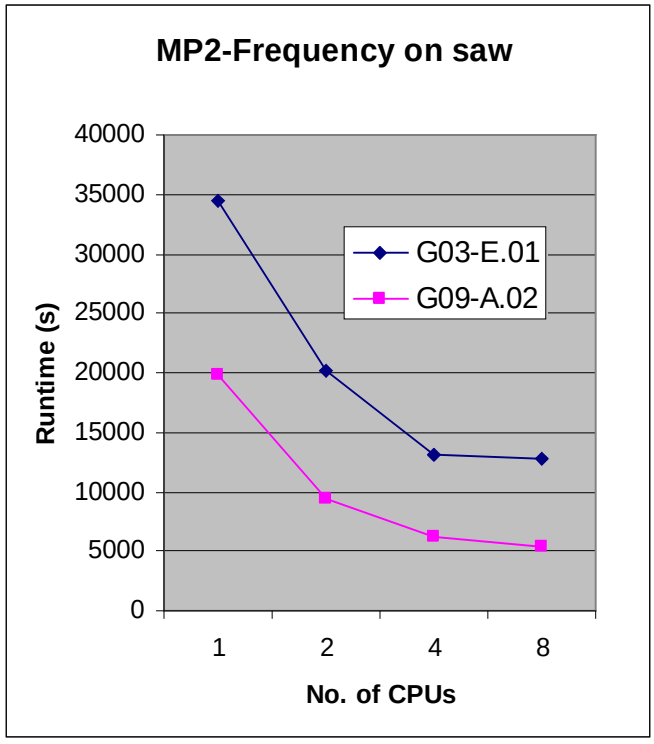

Molecule-III

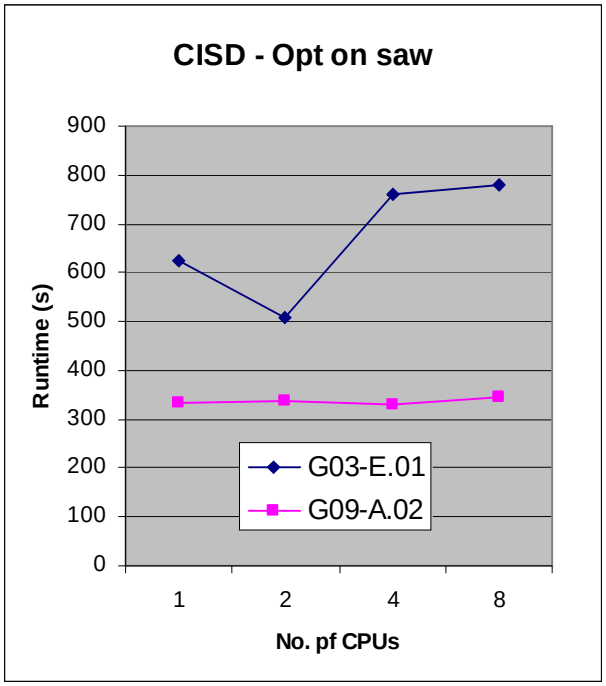

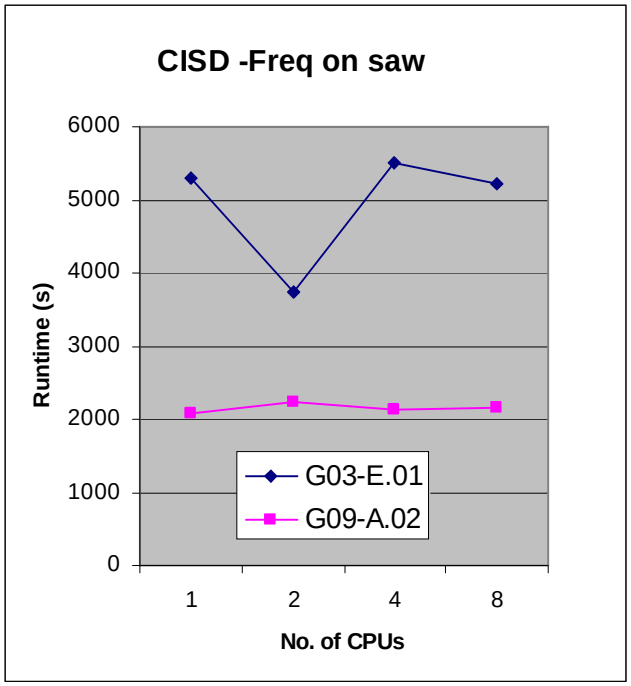

Molecule-IV

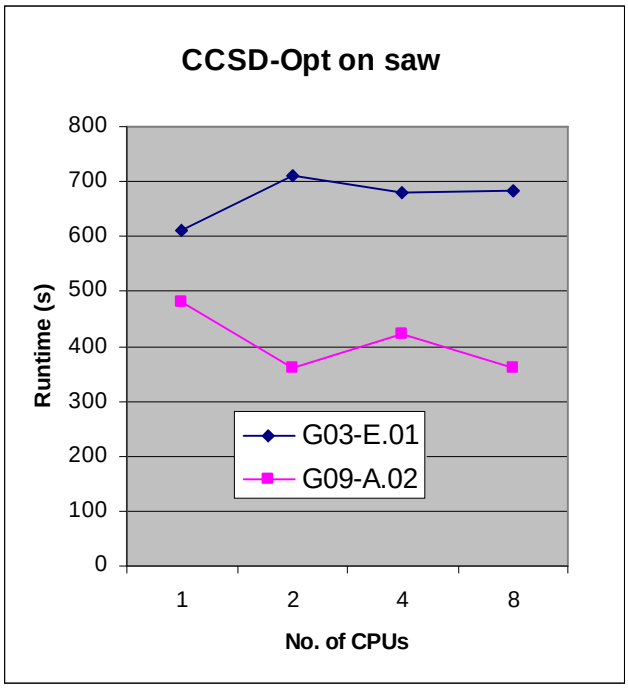

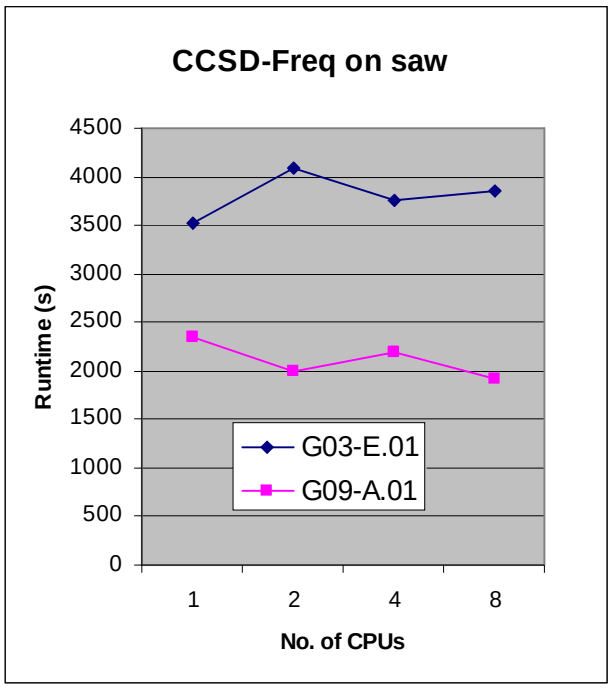

Results on narwhal, molecule-I

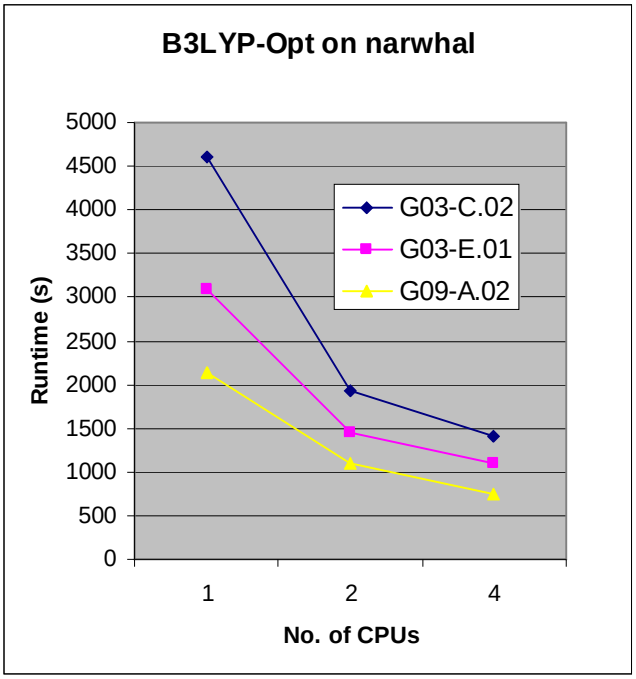

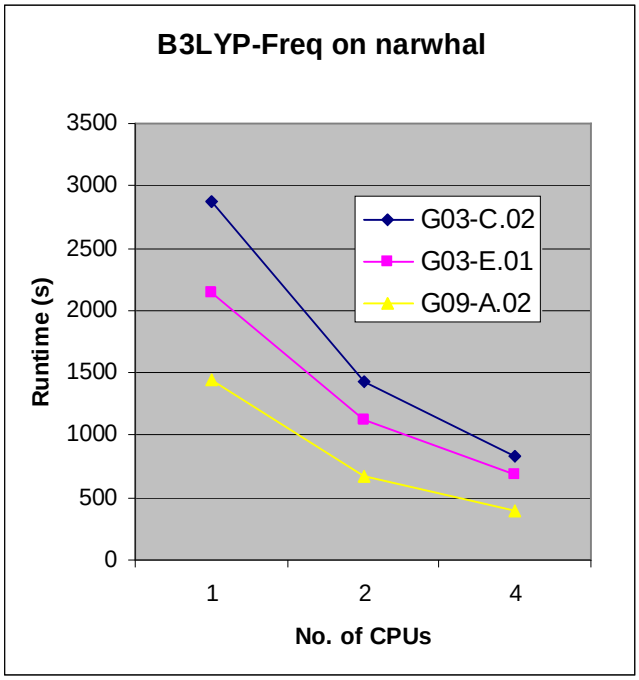

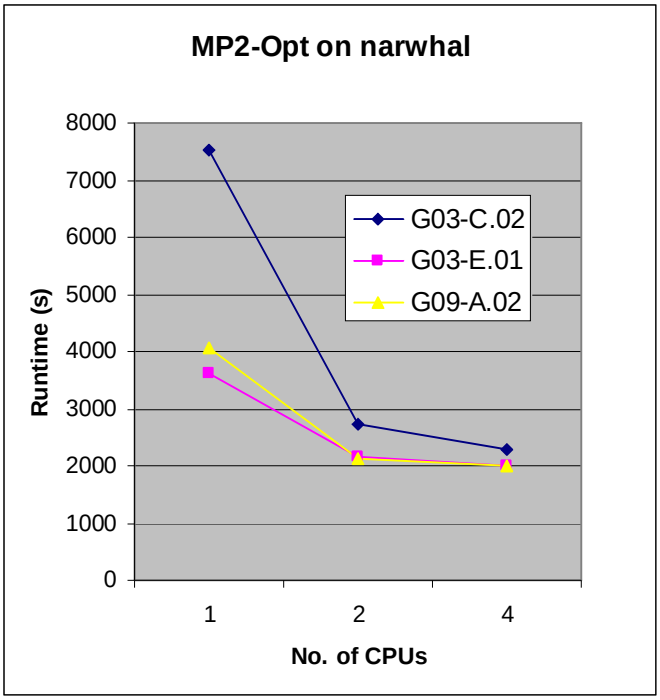

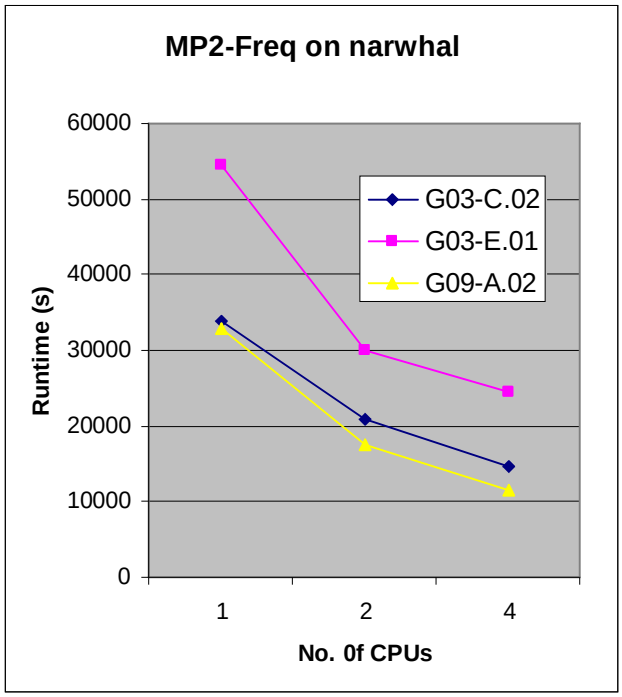

Results on silky, molecule-I

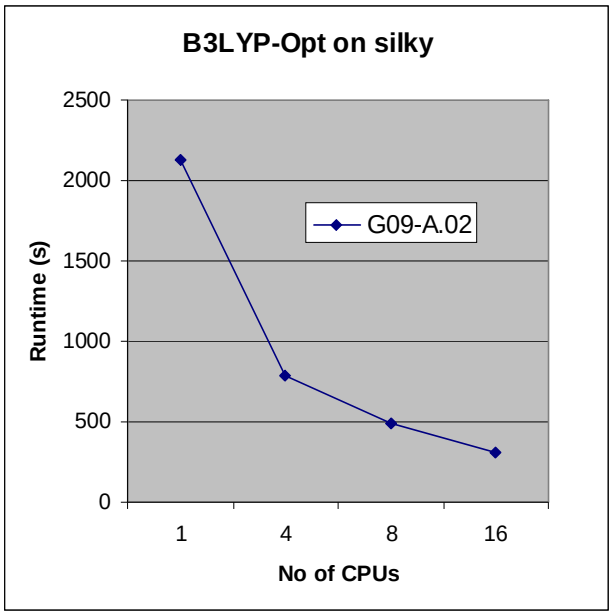

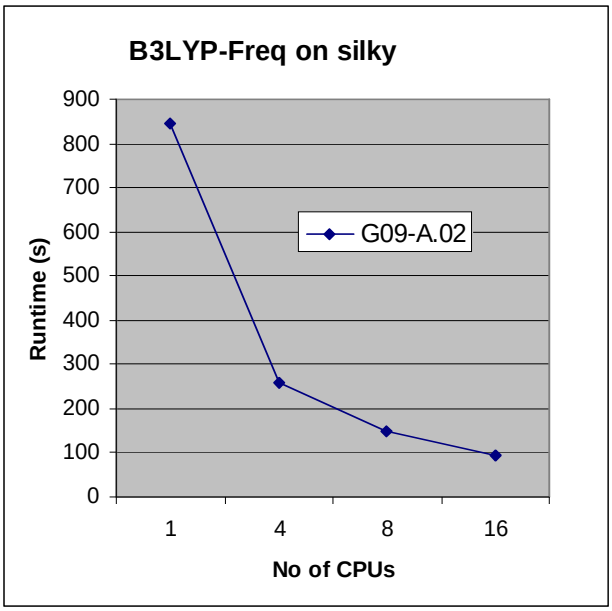

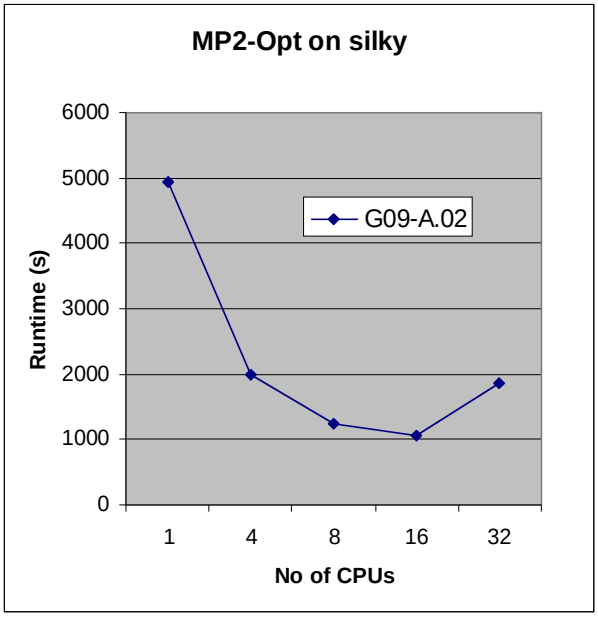

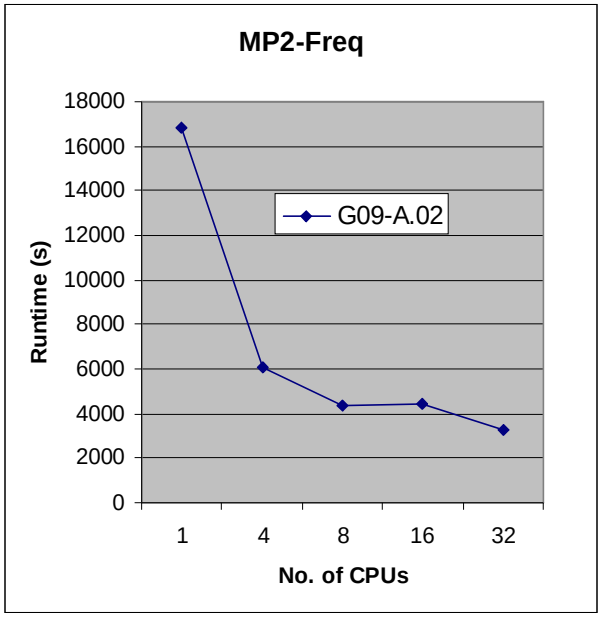

Molecule: WH(CO)(NO)(PMe3)3

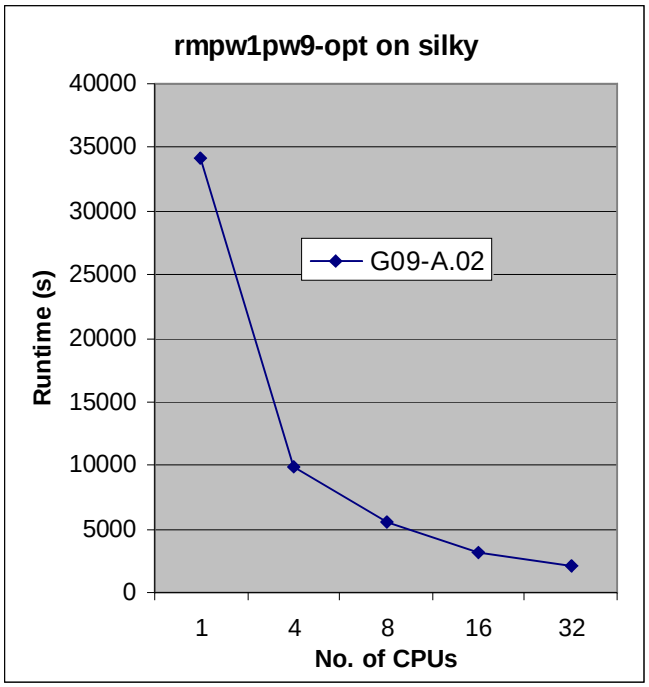

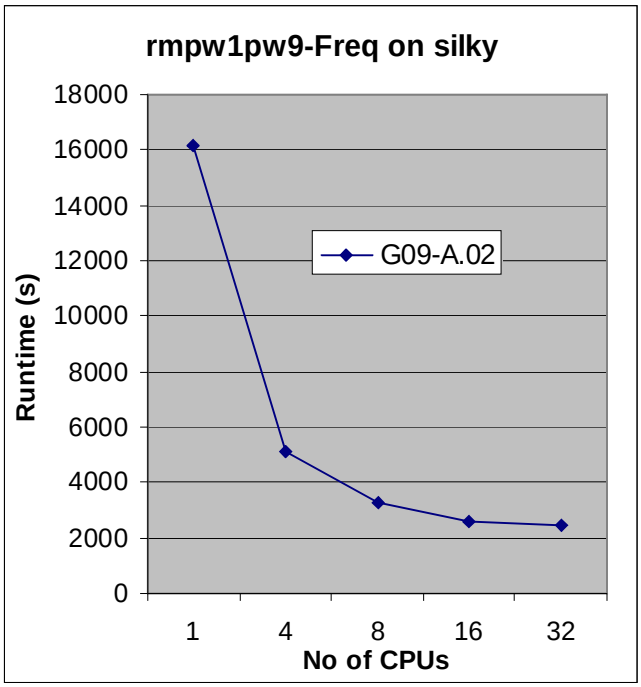

Results: saw, Molecule - I

B3lyp / opt

CPUs	G03-E.01		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64
1	1621 (27m1s)	1	(18m3s)		1129 (18m49s)	1	(26m4s)
2	1071 (17m51s)	1.51	(10m39s)		639 (10m39s)	1.77	(12m32s)
4	700 (11m40s)	2.32	(7m)		423 (7m3s)	2.67	(8m1s)
8	597 (9m57s)	2.72	(5m37s)		326 (5m26s)	3.46	(6m19s)

B3LYP / Freq

	G03-E.01		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64
1	1561 (26m1s)	1	(13m23s)		832 (13m52s)	1	(20m)
2	847 (14m7s)	1.84	(7m)		420 (7m)	1.98	(9m26s)
4	513 (8m33s)	3.04	(3m54s)		233 (3m53s)	3.57	(5m)
8	345 (5m45s)	4.52	(2m38s)		157 (2m37s)	5.30	(3m6s)

MP2 / opt

	G03-E.01		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64
1	2246 (37m26s)	1	(32m41s)		2341 (39m1s)	1	(50m45s)
2	1377 (22m57s)	1.63	(19m43s)		1174 (19m34s)	1.99	(23m50s)
4	997 (16m37s)	2.25	(14m32s)		870 (14m30s)	2.69	(15m58s)
8*	1051 (17m31s)	2.14	(14m27s)		866 (14m26s)	2.70	(15m53s)

MP2 / Freq

	G03-E.01		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64
1	34517 (9h35m17s)	1	(5h15m1s)		19836 (5h30m36s)	1	(9h6s)?
2	20192 (5h36m32s)	1.71	(2h33m25s)		9355 (2h35m52s)	2.12	(3h58m13s)
4	13054 (3h37m34s)	2.64	(1h43m14s)		6300 (1h45m)	3.15	(2h20m20s)
8*	12849 (3h34m9s)	2.69	(1h37m23s)		5455 (1h30m55s)	3.64	(1h45m39s)

Results: CISD, Molecule-III

Cluster: saw, G03-E.01/G09-A.01, 6-311g(2df,p)

	G03-E.01				G09-A.01			
	Opt		Freq		Opt		Freq	
	Run time	Speedup						
1	626 (10m26s)	1	5303 (1h28m23s)	1	332 (5m32s)	1	2073 (34m33s)	1
2	507 (8m27s)	1.23	3732 (1h2m12s)	1.42	337 (5m37s)	0.99	2244 (37m24s)	0.92
4	762 (12m42s)	0.82	5504 (1h31m44s)	0.96	328 (5m28s)	1.01	2130 (35m30s)	0.97
8	780 (13m)	0.80	5220 (1h27m)	1.02	344 (5m44s)	0.97	2160 (36m)	0.96

Results: CCSD, Molecule - IV

Cluster: saw, G09-A.01

	Opt		Freq		Opt		Freq	
	6-311g*		6-311g*		6-311g(2df,p)		6-311g(2df,p)	
1	480 (8m)	1	2355 (39m15s)	1	3318 (55m18s)	1	16785 (4h39m45s)	1
2	360 (6m)	1.33	2003 (33m23s)	1.18	2501 (41m41s)	1.33	12734 (3h32m14s)	1.32
4	422 (7m2s)	1.14	2185 (36m25s)	1.08	3114 (51m54s)	1.07	15953 (4h25m53s)	1.05
8	360 (6m)	1.33	1920 (32m)	1.22	2834 (47m14s)	1.17		

Cluster: saw, G03-E.01

	Opt+Freq		Opt		Freq	
	6-31g*		6-311g*		6-311g*	
	Run time	Speed up				
1	34m45s	1	610 (10m10s)	1	3518 (58m38s)	1
2	34m16s	1	711 (11m51s)	0.86	4096 (1h8m16s)	0.86
4	34m11s	1	681 (11m21s)	0.90	3759 (1h2m39s)	0.94
8			682 (11m22s)	0.89	3864 (1h4m24s)	0.91

Gaussian does not scale for CI or CC based methods, but G09-A.01 is about 2 times faster than G03-E.01 for the CISD and CCSD jobs (6-311g* results)

Cluster: narwhal, Molecule - I

MP2 / Opt

CPUs	**G03-C.02**		**G03-E.01**		**G09-A.02, amd64**	
	Runtimes(s)	**speedup**				
1	7522 (2h5m22s)	1	3629 (1h29s)	1	4066 (1h7m46s)	1
2	2744 (45m44s)	2.74	2157 (35m57s)	1.68	2140 (35m40s)	1.9
4	2285 (38m5s)	3.29	2006 (33m26s)	1.81	2001 (33m21s)	2.03

MP2 / Freq

	G03-C.02		**G03-E.01**		**G09-A.02, amd64**	
1	33799 (9h23m19s)	1	54521 (15h8m41s)	1	32799 (9h6m39s)	1
2	20962 (5h49m22s)	1.61	29931 (8h18m51s)	1.82	17444 (4h50m44s)	1.88
4	14751 (4h5m51s)	2.29	24402 (6h46m42s)	2.23	11461 (3h11m1s)	2.86

B3LYP / OPT

	G03-C.02		**G03-E.01**		**G09-A.02, amd64**		**G09-A.02, em64t**
1	4614 (1h16m54s)	1	3097 (51m37s)	1	2146 (35m46s)	1	(35m35s)
2	1925 (32m05s)	2.40	1451 (24m11s)	2.13	1091 (18m11s)	1.97	(18m32s)
4	1403 (23m23s)	3.29	1092 (18m12)	2.84	744 (12m24s)	2.88	(12m42s)

B3LYP / Freq

	G03-C.02		**G03-E.01**		**G09-A.02, amd64**		**G09-A.02, em64t**
1	2880 (48m)	1	2144 (35m44s)	1	1441 (24m1s)	1	(26m29s)
2	1435 (23m55s)	2.0	1125 (18m45s)	1.91	676 (11m16s)	2.13	(12m4s)
4	829 (13m49s)	3.47	683 (11m23s)	3.14	401 (6m41s)	3.59	(7m6s)

Cluster: silky

Molecule – I, benchmark-1

Molecule – II, Dmitri's sample (#rmpw1pw91/genecp nosymm opt freq)

DFT / opt

CPUs	G09-A.02, ia64-b, M2 (rmpw1pw9)			G09-A.02, ia64-b, M1	(B3LYP)	G09-A.02, ia64-s, M1	
1	34126 (9h28m46s)	1		2130 (35m30s)	1	33m8s	
2				1315 (21m55s)	1.62	22m54s	
4	9843 (2h44m3s)	3.47		785 (13m5s)	2.71	12m31s	12m52s, 12m40s
8	5566 (1h32m46s)	6.13		485 (8m5s)	4.39		
16	3144 (52m24s)	10.85		305 (5m5s)	6.98		
32	2053 (34m13s)	16.62					

DFT / Freq

	G09-A.02, ia64-b, M2 (rmpw1pw9)			G09-A.02, ia64-b, M1	(B3LYP)	G09-A.02, ia64-s, M1	
1	16131 (4h28m51s)	1		847 (14m7s)	1	14m14s	
2				480 (8m)	1.76	8m1s	
4	5086 (1h24m46s)	3.18		260 (4m20s)	3.26	4m9s	4m21s, 4m32s
8	3286 (54m46s)	4.91		150 (2m30s)	5.65		
16	2557 (42m37s) (1.36)	6.31		93 (1m33s)	9.11		
32	2424 (40m24s)	6.65					

MP2 –M1

		Opt			Freq	
1		4946 (1h22m26s)	1		16821 (4h40m21s)	1
4		1990 (33m10s)	2.49		6106 (1h41m46s)	2.75
8*		1232 (20m32s)	4.01		4346 (1h12m26s)	3.87
16		1051 (17m31s)	4.71		4452 (1h14m12s)	3.78
32		1861 (31m1s)	2.66		3285 (54m45s)	5.12

Cluster: hound, Molecule – I (NFS storage file system, results are meaningless)

B3lyp / opt

CPUs	G09-A.01, amd4		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64	
	Runtime(s)	Speedup						
1	43m34s		24m15s		25m1s		34m2s	
4	17m14s		9m14s		9m13s		16m13s	
8	12m25s		6m42s		6m40s		9m	
16	11m29s		6m8s		6m7s		7m36s	
32	7m4s							

B3LYP / Freq

	G09-A.01, amd64		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64	
1	13m42		17m50s		36m4s		24m30s	
4	8m48s		5m1s		4m57s		6m49s	
8	5m		2m49s		2m49s		3m40s	
16	15m42s		2m11s		(2m10s)		2m51s	
32	2m44s							

MP2 / opt

	G09-A.01, amd64		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64	
1	1h57m47s		1h14m22s		1h19m39s		1h5m44s	
4	42m40s		19m7s		26m53s		32m12s	
8*	1h15m53s		15m		22m22s		41m53s	
16	1h47m34s		20m9s		20m3s			
32	2h14m26s							

MP2 / Freq

	G09-A.01, amd64		G09-A.01, em64t		G09-A.02, em64t		G09-A.02, amd64	
1	10h57m49s		11h4m34s		17h31m48s		10h45m8s	
4	6h30m40s		6h23m45s		13h14m48s		5h2m	
8*	7h30m30s		7h23m16s		12h52m7s		3h47m38s	
16	3h5m42s		5h19m28s		5h33m29s			
32	4h46m42s							

Input files

%mem = 2GB for B3LYP
%mem = 4GB for MP2 computations
%mem = 2GB for CISD
%mem = 4GB for CCSD computations

%nproc varies from 1, 2, 4, 8, 16 to 32 threads/cpus depending on the node structures

Molecule – I, (H2PCH2CH2PH2)PdCl2(CH3)2 for B3LYP and MP2

It is from Gaussian test job 445, the geom. and basis sets can be found in test445.com in the directory
/opt/sharcnet/gaussian/g09/tests/com
or
/opt/sharcnet/gaussian/g03/tests/com

The following leading lines have been added above the geom. inputs
(%nproc varies for scaling tests)

%nosave
%mem=2GB
%chk=benchmark-b3lyp-1
%nproc=1
#p b3lyp/gen 6d opt freq (for B3LYP computations)
[#p mp2/gen 6d opt freq (for MP2 computations)]

Gaussian Test Job 445:
(H2PCH2CH2PH2)PdCl2(CH3)2 benchmark optimization

0 1

……

Molecule: WH(CO)(NO)(PMe3)3 for rmpw1pw91

%chk=test4CPUsSilky.chk
%mem=256MW
%nproc=4
#opt rmpw1pw91/genecp nosymm

WH(CO)(NO)(PMe3)3 test calculation using 4 CPUs

0 1

W	0	4.015400	11.910400	1.804600
P	0	5.427200	10.103200	0.878300
P	0	2.127200	11.316300	0.275900
P	0	3.060800	14.148200	2.150600
N	0	3.183400	11.120800	3.216900
O	0	2.582000	10.654200	4.167000
C	0	5.621600	12.556800	2.763300
O	0	6.571400	12.942500	3.288400

C	0	5.473600	9.866900	-0.925400
H	0	4.562000	9.742700	-1.262100
H	0	6.010700	9.073800	-1.137300
H	0	5.874200	10.654200	-1.349200
C	0	5.135100	8.417900	1.498600
H	0	5.097800	8.432400	2.477400
H	0	5.864500	7.832800	1.207500
H	0	4.285900	8.082800	1.143800
C	0	7.205000	10.287300	1.206200
H	0	7.348400	10.426000	2.165400
H	0	7.550800	11.058700	0.709700
H	0	7.675900	9.478300	0.918900
C	0	0.466300	11.807000	0.846100
H	0	0.331800	11.497900	1.765100
H	0	-0.211200	11.405400	0.262600
H	0	0.386500	12.783600	0.815000
C	0	2.187900	11.976100	-1.412800
H	0	2.371400	12.938200	-1.380300
H	0	1.326700	11.822900	-1.854800
H	0	2.897500	11.525300	-1.914600
C	0	1.777300	9.553400	-0.016900
H	0	1.722600	9.083900	0.840900
H	0	2.495500	9.163400	-0.557600
H	0	0.924700	9.465300	-0.492600
C	0	2.379500	15.069000	0.752600
H	0	3.030400	15.079200	0.019500
H	0	2.186300	15.989300	1.030700
H	0	1.553200	14.638500	0.449700
C	0	1.736700	14.236900	3.393700
H	0	2.017000	13.763100	4.204700
H	0	0.924700	13.820900	3.037500
H	0	1.555100	15.174500	3.612000
C	0	4.205600	15.398400	2.788000
H	0	4.672900	15.043000	3.573100
H	0	3.704100	16.201600	3.042700
H	0	4.858600	15.628100	2.095200
H	0	4.593600	12.828400	0.376900

H C N O P 0
6-31g(d,p)
****
W 0
sdd
****

W 0
sdd

--Link1--
%chk=test4CPUsSilky.chk
%mem=512MW
%nproc=4
#freq geom=check guess=read rmpw1pw91/genecp nosymm

WH(CO)(NO)(PMe3)3 test calculation using 4 CPUs

0 1

H C N O P 0

6-31g(d,p)
****
W 0
sdd
****

W 0
Sdd

Molecule – III for CISD Opt and Freq

%NoSave
%chk=ch3oh_cisd-4
%mem=2GB
%nproc=4
#p cisd/6-311g(2df,p) opt freq

Gaussian Test Job 58:
MEOH opt, freq STD MOD cisd

0 1
C
O 1 CO
H 1 CH 2 T
H 1 CH 2 T 3 T 1
H 1 CH 2 T 3 T -1
H 2 OH 1 T 3 180.

CO 1.43
CH 1.09
OH 0.96
T 109.471221

Molecule –IV, for CCSD Opt and Freq

%NoSave
%chk=ch3ch2_ccsd-8
%mem=4GB
%nproc=8
#p ccsd/6-311g* opt freq

Gaussian Test Job 684:
Ethyl radical CCSD opt+freq

0 2
C1
C2 C1 CC
H1 C1 CH C2 T
H2 C1 CH C2 T H1 T 1
H3 C2 CH C1 T H1 180.
H4 C2 CH C1 T H3 120.
H5 C2 CH C1 T H3 240.

CC 1.54
CH 1.09
T 109.471221