

Gaussian Usage on SHARCNET clusters

Jemmy Hu

SHARCNET, November, 2010

Availability (January 01, 2011 -)

| Version | Cluster | Queue |
|----------|------------------|------------------|
| G09_B.01 | orca, saw, whale | threaded, serial |
| G03_E.01 | orca, saw, whale | threaded, serial |

Recommended processors (CPUs) per G09 job at SHARCNET

| Method\Cluster | orca | whale | saw | queue |
|----------------------------|------|-------|-----|----------|
| DFT (b3lyp, etc) | 8 | 4 | 8 | threaded |
| MP (2, 3, 4) | 4 | 4 | 8 | threaded |
| CI (cis, cid, cisd, qcisd) | 1 | 1 | 1 | serial |
| CC (ccd, ccSD, ccSD(t)) | 1 | 1 | 1 | serial |

Usage on orca (torque queue, same as goblin):

sqsub -q threaded -n cpus -r runtime --mpp=memory(g) -o *.out g09 *.com
(Gaussian output is in *.log. replace 'g09' with 'g03' to run G03_E.01)

Example for serial job (1 cpu, 3 days runtime, 4g memory, inputfile: name.com)

sqsub -r 3.0d --mpp=4g -o name.out g09 name.com

Example for threaded job (4 cpus, 3 days runtime, 4g memory)

sqsub -q threaded -n 4 -r 3.0d --mpp=4g -o name.out g09 name.com

(recommend -n 4 for MPx jobs, -n 8 for DFT jobs on orca)

Usage on saw and whale (LSF queue, same as bull)

sqsub -q threaded -n cpus -r runtime -o *.log g09 *.com
(replace 'g09' with 'g03' to run G03_E.01)

Example for serial job

sqsub -r 7.0d -o name.log g09 name.com

Example for threaded job on **whale** (4 cpus)

sqsub -q threaded -n 4 -r 3.0d -o name.log g09 name.com

Example for threaded job on **saw** (8 cpus)

sqsub -q threaded -n 8 -r 3.0d -o name.log g09 name.com

General Notes:

1. cpus (-n cpus) in sqsub should match the %nproc=cpus specified in *.com file (default is 1)
2. always run serial (1 cpu) jobs for CI and CC based methods
3. use a good estimate runtime. max. is 7 days, use BLCR for longer than 7 days jobs
4. --mpp is different from %mem, see our Gaussian software page for detail descriptions
5. due to the nature of LSF scheduler, use -n 4 on whale, -n 8 on saw for DFT and MPx jobs
6. run job out of your /work directory. on whale, use /gwork may have better performance.

More detail instructions can be found at <https://www.sharcnet.ca/my/software/show/21>

Gaussian G09_B.01 Scaling Results

November, 2010

Systems:

| Name | CPUs/node | RAM/node | OS / Queue | Interconnect |
|------|--------------------------------------|----------|--------------------|--------------|
| orca | 24 (2 sockets), Opteron @ 2.2 GHz | 32.0 GB | Centos 5.x, torque | IB |
| saw | 8 (2 quad cores), Xeon @ 2.83 GHz | 16 GB | XC 4, LSF | IB |

Molecules and Methods/Models:

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

General conclusions:

1. Gaussian does not scale for **CI** (cis, cisd, qcisd) and **CC** (ccsd, etc) based methods, serial (1 cpu) job is the right choice
2. For **DFT** (b3lyp, etc), speedup for 4-cpu parallel is good, 8-cpu job is reasonable
3. For **MPx** (MP2, etc), speedup for 4-cpu job is good , 8-cpu job is poor

Recommended Maximum processors for G09 jobs at SHARCNET

| Methods/Modules | Opt | Freq | Energy |
|------------------------------|--------|--------|--------|
| DFT (b3lyp, etc) | 4 or 8 | 4 or 8 | 4 or 8 |
| MP(2, 3, 4) | 4 | 4 | 4 |
| CISD (cis, cid, cisd, qcisd) | 1 | 1 | 1 |
| CCSD (ccd, ccsd, ccsd(t)) | 1 | 1 | 1 |

Results:

B3lyp / opt + freq

| CPUs | orca | speedup | saw | speedup |
|------|------------------|---------|------------------|---------|
| 1 | 2731 (45m31s) | 1 | 1961 (32m41s) | 1 |
| 4 | 840 (14m) | 3.25 | 656 (10m56s) | 2.99 |
| 8 | 540 (9m) | 6.06 | 460 (7m40s) | 4.26 |
| 12 | 431 (7m11s) | 6.34 | | |
| 24 | 365 (6m5s) | 7.48 | | |

MP2 / opt + freq

| CPUs | orca | speedup | saw | speedup |
|------|---------------------|---------|---------------------|---------|
| 1 | 30677 (8h31m17s) | 1 | 27949 (7h45m49s) | 1 |
| 4 | 11506 (3h11m46s) | 2.67 | 9927 (2h45m27s) | 2.82 |
| 8 | 8661 (2h24m21s) | 3.54 | 7533 (2h5m33s) | 3.71 |
| 12 | 6560 (1h49m20s) | 4.68 | | |
| 24 | 8243 (2h17m23s) | 3.72 | | |

QCISD / opt+freq

| CPUs | orca | speedup | saw | speedup |
|------|--------------------|---------|--------------------|---------|
| 1 | 4290 (1h11m30s) | 1 | 4378 (1h12m58s) | 1 |
| 4 | 3384 (56m24s) | 1.27 | 5041 (1h24m1s) | 0.87 |
| 8 | 3134 (52m14s) | 1.37 | 3904 (1h5m5s) | 1.12 |

CCSD / opt+freq

| CPUs | orca | speedup | saw | speedup |
|------|---------------------|---------|---------------------|---------|
| 1 | 17807 (4h56m47s) | 1 | 16785 (4h39m45s) | 1 |
| 4 | 15237 (4h13m57s) | 1.17 | 17165 (4h46m05s) | 0.98 |
| 8 | 15953 (4h25m53s) | 1.12 | 16348 (4h32m28s) | 1.03 |