SAND2008-2899C

Application Performance under Different XT Operating Systems

Courtenay T. Vaughan, John P. Van Dyke, and Suzanne M. Kelly
Sandia National Laboratories

Cray User Group
May 2008

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Background

- Cray XT3 series ran Catamount OS
  - Light Weight Kernel based on kernel developed at Sandia
- With XT4, Cray moving to Compute Node Linux (CNL)
  - tuned Linux kernel
  - added support for quad-core processors

Catamount N-Way (CNW)

- Developed as risk mitigation for ORNL with funding from DOE Office of Science
  - Jaguar being upgraded to quad-core processors
- Designed to support N cores per processor
  - Not just 4 cores per processor
  - Able to run on nodes with 1 or 2 cores per processor without recompiling
  - Able to run on a mixture of nodes

Comparison of CNL and CNW

- CNL based on Linux kernel
  - Linux supports multiple users, processes, and services
  - Undesirable features configured “off” when kernel was built
  - Tuned to minimize interrupts
- CNW designed as limited function kernel
  - Device drivers only for console output and communication with the SeaStar NIC
  - No virtual memory or unnecessary features
  - Each node supports exactly one user running one application on 1 to N cores

Tests on pre-upgrade Jaguar

- Conducted last Summer
- Jaguar was a mix of XT3 and XT4 dual-core nodes
- Specific sizes for each codes
- Results from 3 codes
  - Gyrokinetic Toroidal Code (GTC)
    - 3-d PIC code for magnetic confinement fusion
  - Parallel Ocean Program (POP)
    - ocean modeling code
  - VH1
    - a multidimensional ideal compressible hydrodynamics code

Jaguar Results

	CNL 2.0.03+	CNW 2.0.05+	Improvement
GTC			
1024 core XT3	595.6 sec	584.0 sec	2.0%
4096 core XT3	614.6 sec	593.8 sec	3.5%
20000 core XT3/XT4	786.5 sec	778.9 sec	1.0%
POP			
4800 core XT3	90.6 sec	77.6 sec	16.8%
20000 core XT3/XT4	98.8 sec	75.2 sec	31.4%
VH1			
1024 core XT3	22.7 sec	20.9 sec	8.6%
4096 core XT3	137.1 sec	117.4 sec	16.8%
20000 core XT3/XT4	1186.0 sec	981.7 sec	20.8%

Red Storm results

- Both OS based on 2.0.44
- Machine configured with 12960 nodes (25920 cores)
  - Ran with Moab scheduler for CNW
    - resulted in some bad job layout
  - Ran with interactive nodes with CNL
- Ran two codes and HPCC
  - CTH
    - shock hydrodynamics code
  - PARTISN
    - time-dependent neutron transport code

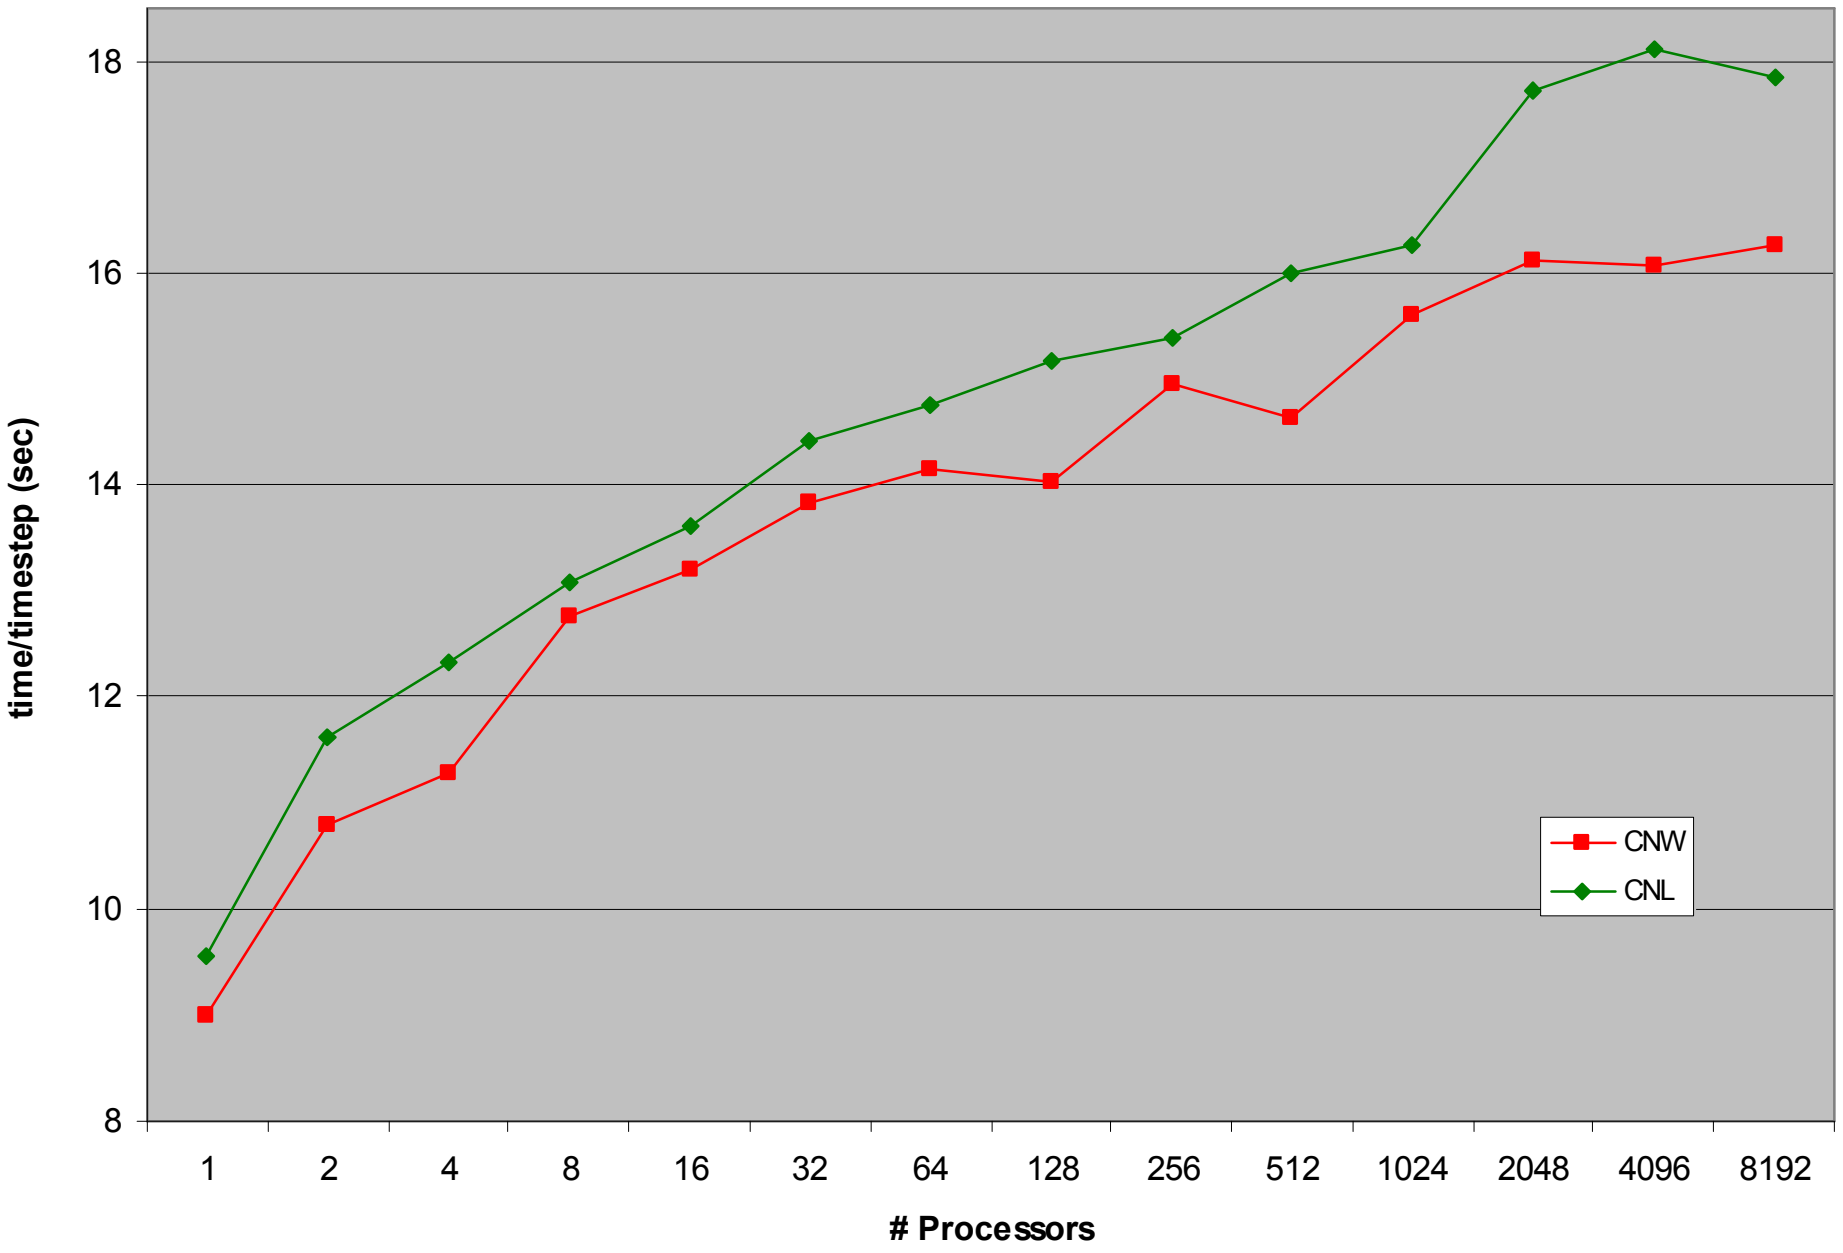
CTH 7.1 - Shaped Charge (90 x 216 x 90/proc)
time/timestep (sec)
18
16
14
12
10
8
1
2
4
8
16
32
64
128
256
512
1024
2048
4096
8192
Processors
CNW
CNL

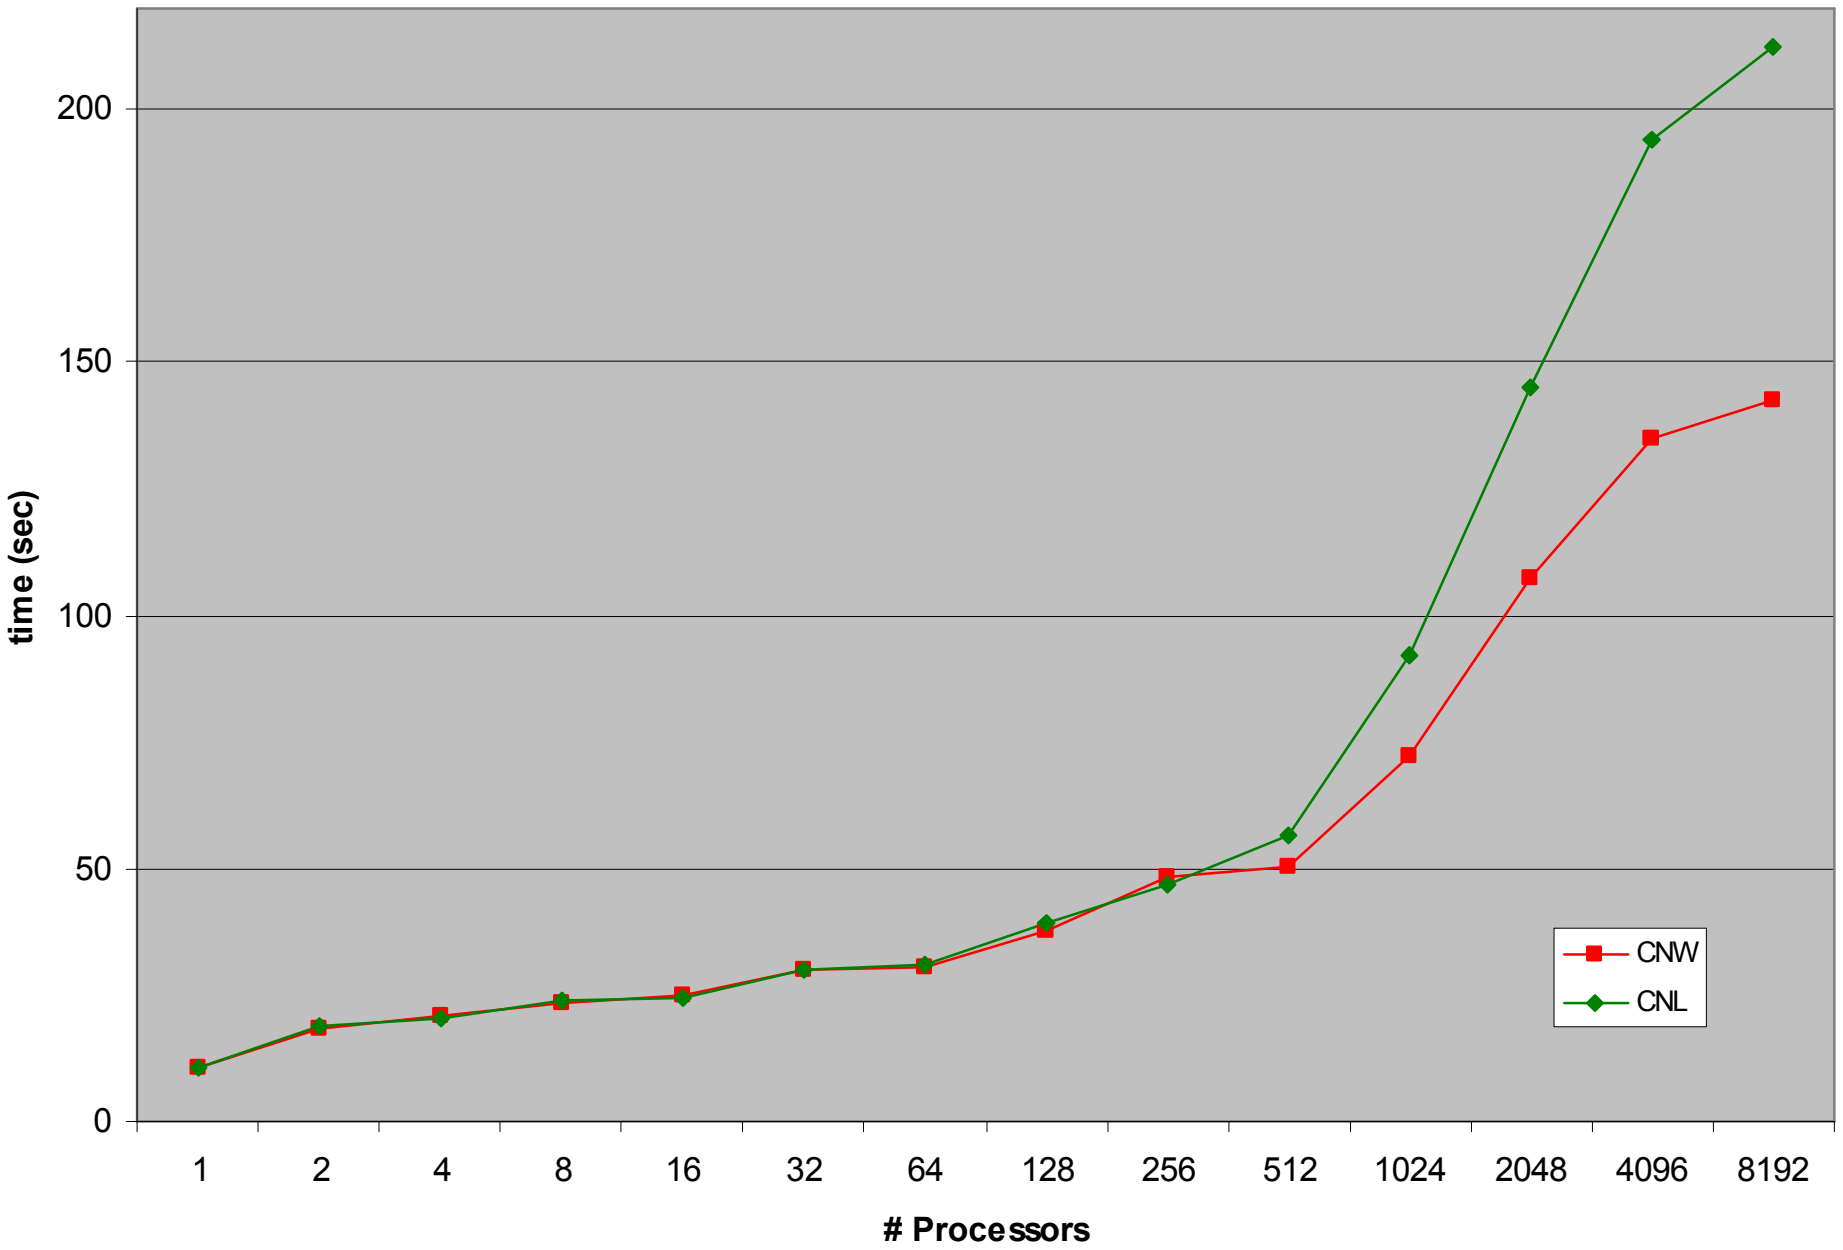
Partisn - sn timing - 24 x 24 x 24/proc
time (sec)
0
50
100
150
200
1
2
4
8
16
32
64
128
256
512
1024
2048
4096
8192
Processors
CNW
CNL

HPCC

- Series of 7 benchmarks in one package. We generally use 5 of them:
  - PTRANS - matrix transposition
  - HPL - Linpack direct dense system solve
  - STREAMS - Memory bandwidth
  - Random Access - Global random memory access
  - FFT - large 1-D FFT
- Code is C with libraries
- HPL not used for these runs
- Optimized Random Access and FFT
- Version 1.2

HPCC on 16384 cores

benchmark	units	CNL	CNW	CNW/CNL
PTRANS	GB/s	598.7	894.1	1.49
STREAMS	GB/s	24721	36499	1.48
Random Access	GUP/s	12.7	23.4	1.85
FFT	GFLOPS	1963.8	2272.2	1.16

Quad-Core System

- Machine with 4 Budapest quad-core nodes
- Running 2.0.44
- PGI 6.2.5 Compiler
- Run with Lustre filesystem

- Ran baseline HPCC version 1.0

HPCC on 16 cores (4 nodes)

Benchmark	CNL	CNW	CNW/CNL
PTRAN GB/s	1.612	2.792	1.73
HPL GFLOPS	66.55	68.02	1.02
STREAMS GB/s	31.98	35.13	1.10
Random GUPs	0.01717	0.03502	2.04
FFT GFLOPS	3.331	3.518	1.06

HPCC on 4 cores (4 nodes)

Benchmark	CNL	CNW	CNW/CNL
PTRANS GB/s	0.576	1.606	2.83
HPL GFLOPS	17.88	17.90	1.00
STREAMS GB/s	25.21	25.84	1.02
Random GUP/s	0.06445	0.11823	1.83
FFT GFLOPS	1.609	1.646	1.02

HPCC on 4 cores (2 nodes)

Benchmark	CNL	CNW	CNW/CNL
PTRANS GB/s	0.488	1.551	3.18
HPL GFLOPS	17.78	18.03	1.01
STREAMS GB/s	16.45	18.03	1.10
Random GUP/s	0.006105	0.011476	1.88
FFT GFLOPS	1.337	1.360	1.02

HPCC on 4 cores (4 nodes)

Benchmark	CNL	CNW	CNW/CNL
PTRANS GB/s	0.287	1.244	4.33
HPL GFLOPS	17.59	17.72	1.01
STREAMS GB/s	7.85	9.95	1.27
Random GUP/s	0.005984	0.011476	1.92
FFT GFLOPS	0.902	0.959	1.06

Additional Codes

- LSMS
  - electron structure
- S3D
  - combustion modeling
- PRONTO3D
  - structural analysis
- SAGE
  - hydrodynamics
- SPPM
  - 3-D gas dynamics
- UMT2K
  - unstructured mesh radiation transport

Performance on 16 cores (4 nodes)

Application	CNL seconds	CNW seconds	Improvement CNW/CNL
CTH	1513.1	1298.1	16.6%
GTC	664.9	670.6	-0.85%
LSMS	290.1	276.7	4.84%
PARTISN	499.3	491.3	1.62%
POP	153.8	151.9	1.22%
PRONTO	241.5	222.0	8.78%
S3D	1949.1	1948.9	0.01%
SAGE	267.8	234.9	14.0%
SPPM	847.8	845.0	0.33%
UMT	502.7	472.3	0.44%

Performance on 4 cores (4 nodes)

Application	CNL seconds	CNW seconds	Improvement CNW/CNL
CTH	861.4	816.7	5.47%
GTC	583.1	577.7	0.93%
LSMS	1160.6	1105.6	4.97%
PARTISN	175.1	165.5	5.75%
POP	428.0	425.5	0.61%
PRONTO	175.8	164.2	7.06%
S3D	1327.8	1282.5	3.53%
SAGE	170.0	158.9	6.94%
SPPM	294.6	293.1	0.51%
UMT	1768.8	1701.0	3.99%

Performance on 4 cores (2 nodes)

Application	CNL seconds	CNW seconds	Improvement CNW/CNL
CTH	949.7	877.8	8.19%
GTC	592.9	589.5	0.58%
LSMS	1177.3	1118.6	5.25%
PARTISN	245.5	234.4	4.77%
POP	440.1	435.7	1.01%
PRONTO	186.8	175.0	6.74%
S3D	1482.2	1439.7	2.95%
SAGE	179.9	165.3	8.85%
SPPM	297.3	295.2	0.71%
UMT	1816.2	1760.4	3.17%

Performance on 4 cores (1 node)

Application	CNL seconds	CNW seconds	Improvement CNW/CNL
CTH	1219.5	1037.8	17.51%
GTC	622.8	622.4	0.06%
LSMS	1208.1	1144.6	5.55%
PARTISN	447.1	441.9	1.16%
POP	467.3	464.3	0.66%
PRONTO	209.1	195.1	7.18%
S3D	1937.3	1940.4	-0.16%
SAGE	233.4	190.2	17.47%
SPPM	301.1	297.8	1.11%
UMT	1944.6	1827.6	6.40%

Summary

- We developed a version of Catamount for quad-core and beyond
- Most applications at scale on dual-core systems run better with CNW than with CNL
  - Difference gets bigger with larger numbers of cores
- On our 4 quad-core system, most applications perform somewhat better with CNW
  - Different applications react differently
- Need to do a large scale test with quad-core processors to see if the effects are cumulative