Performance Measurements of ECMWF RAPS9 Benchmark on x86_64 Architectures

Dr. Dirk Merten

RAPS Workshop

Reading, November 1st 2006

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

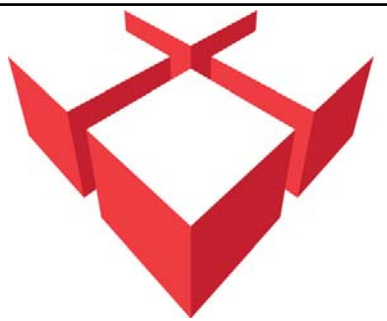

LINUX NETWORX

Architectures and Systems
Linux Networx LS-1

AMD Opteron:

- 128 Nodes, Dual Socket Dual Core 285
- 2.6 GHz
- 1 MB L2 Cache, 8 GB RAM per node
- Compiler Intel 9.1, -O3 -xW
- Infiniband Interconnect single data rate
- Scali MPI 5.2

Intel Woodcrest:

- 16 Nodes, Dual Socket Dual Core
- 3.0 GHz
- 4MB Cache, 8 GB RAM per node
- Compiler Intel 9.1, -O3 -xP
- Infiniband Interconnect single data rate
- mvapich 0.9.7

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Memory Performance

Single Process

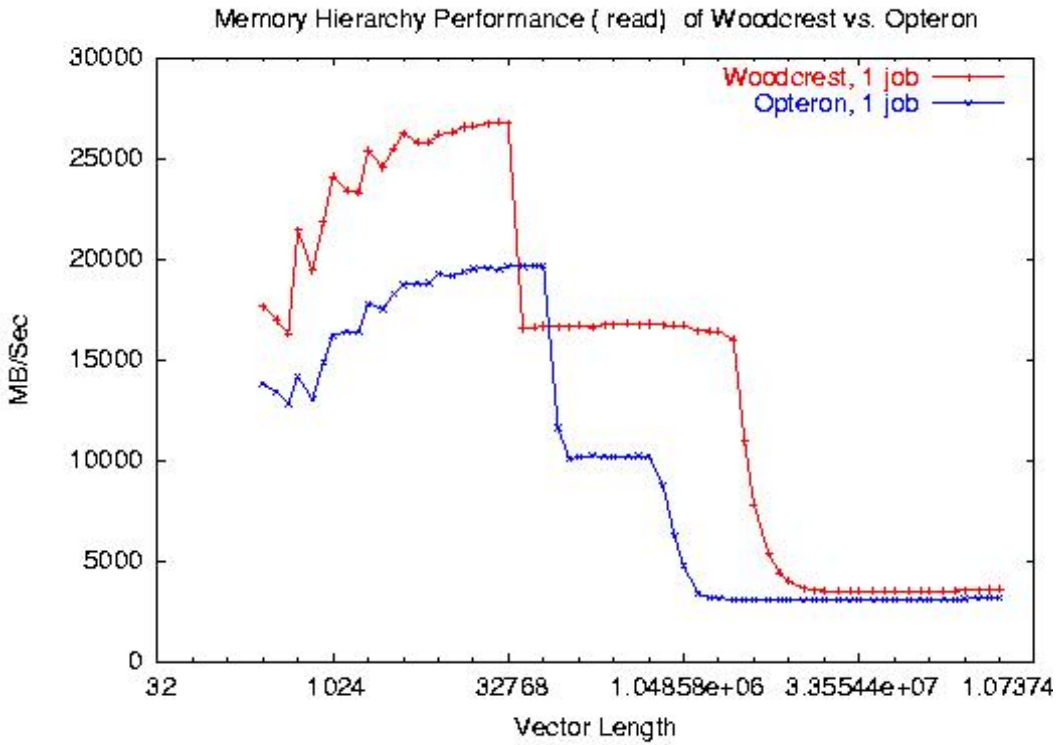

4 Processes on shared bus

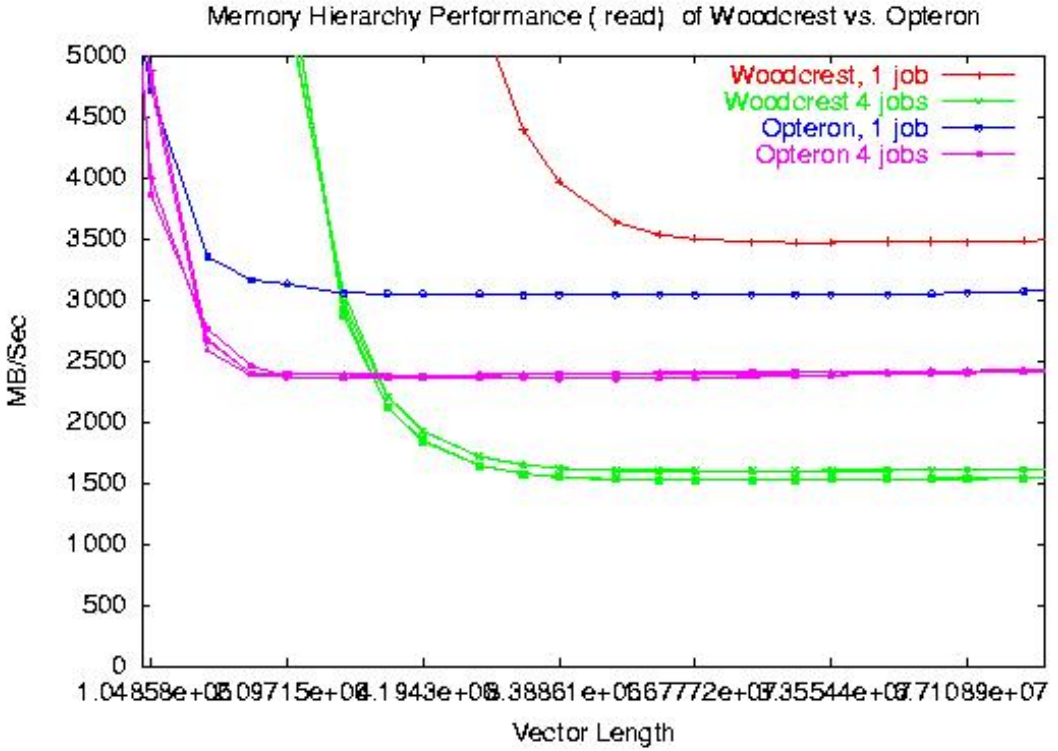

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Memory Performance

Single Process

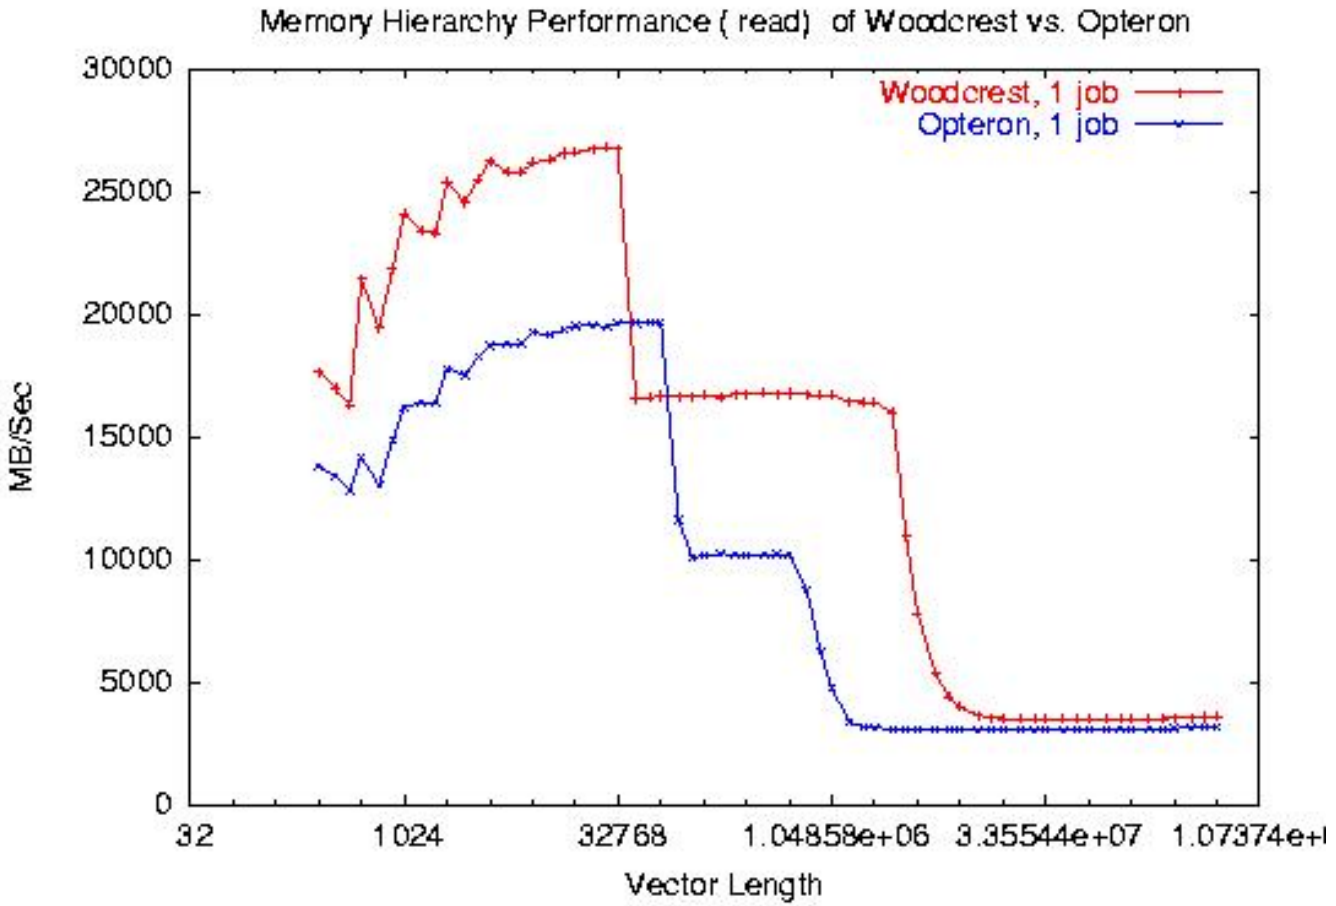

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Memory Performance

4 Processes on shared bus

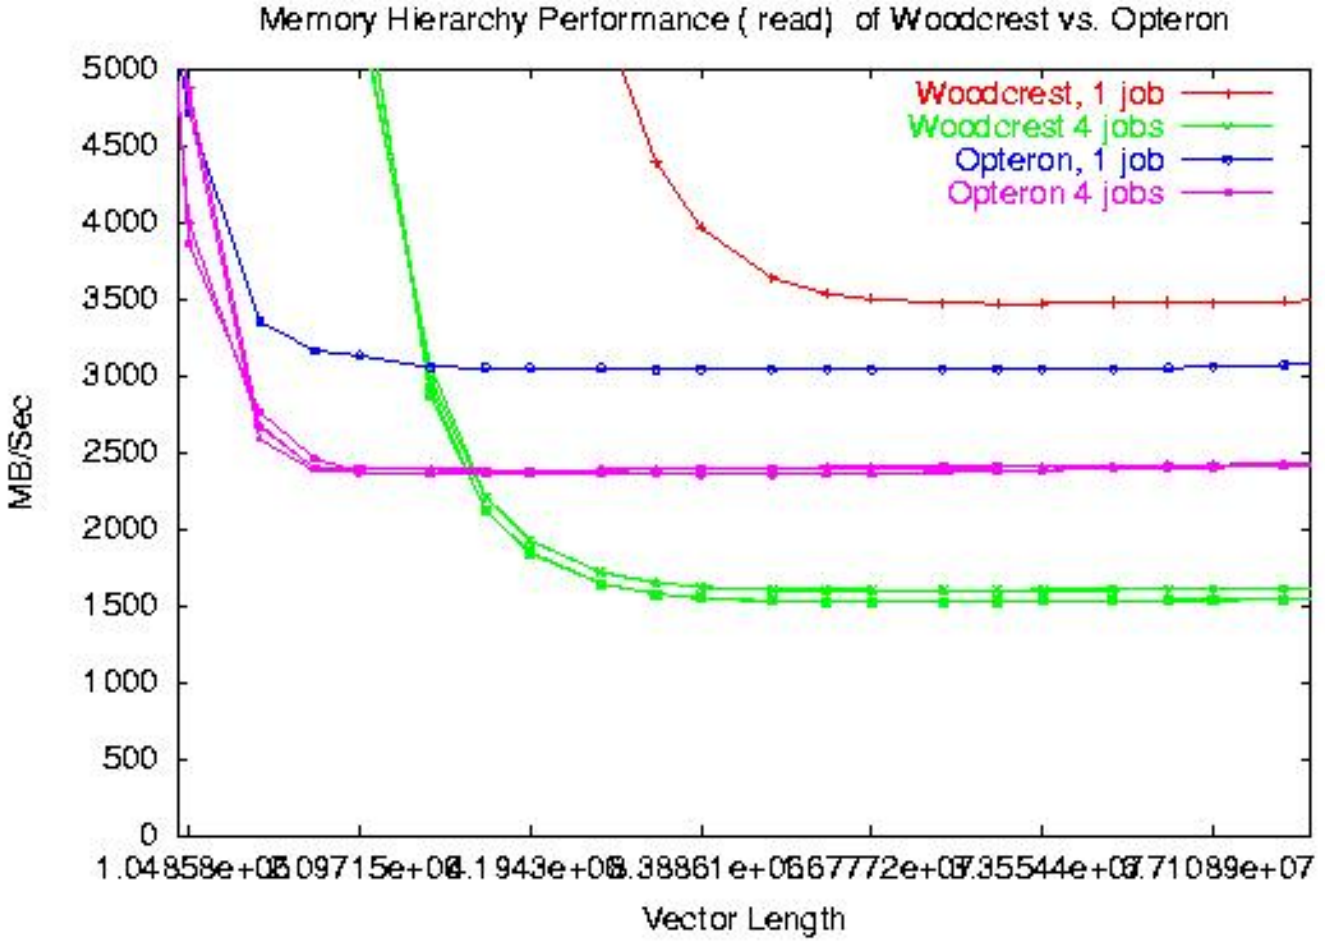

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Porting

- no changes to the code were required
  - Stack limit depends on environments
- various changes to the scripts
  - To accommodate to the environments, e.g.:
    - ‘head’ and ‘tail’ may behave differently in different environments
    - ‘Flex’ produces different code

Opteron socket E, 2.6 GHz: RAPS9 T399L62 10day

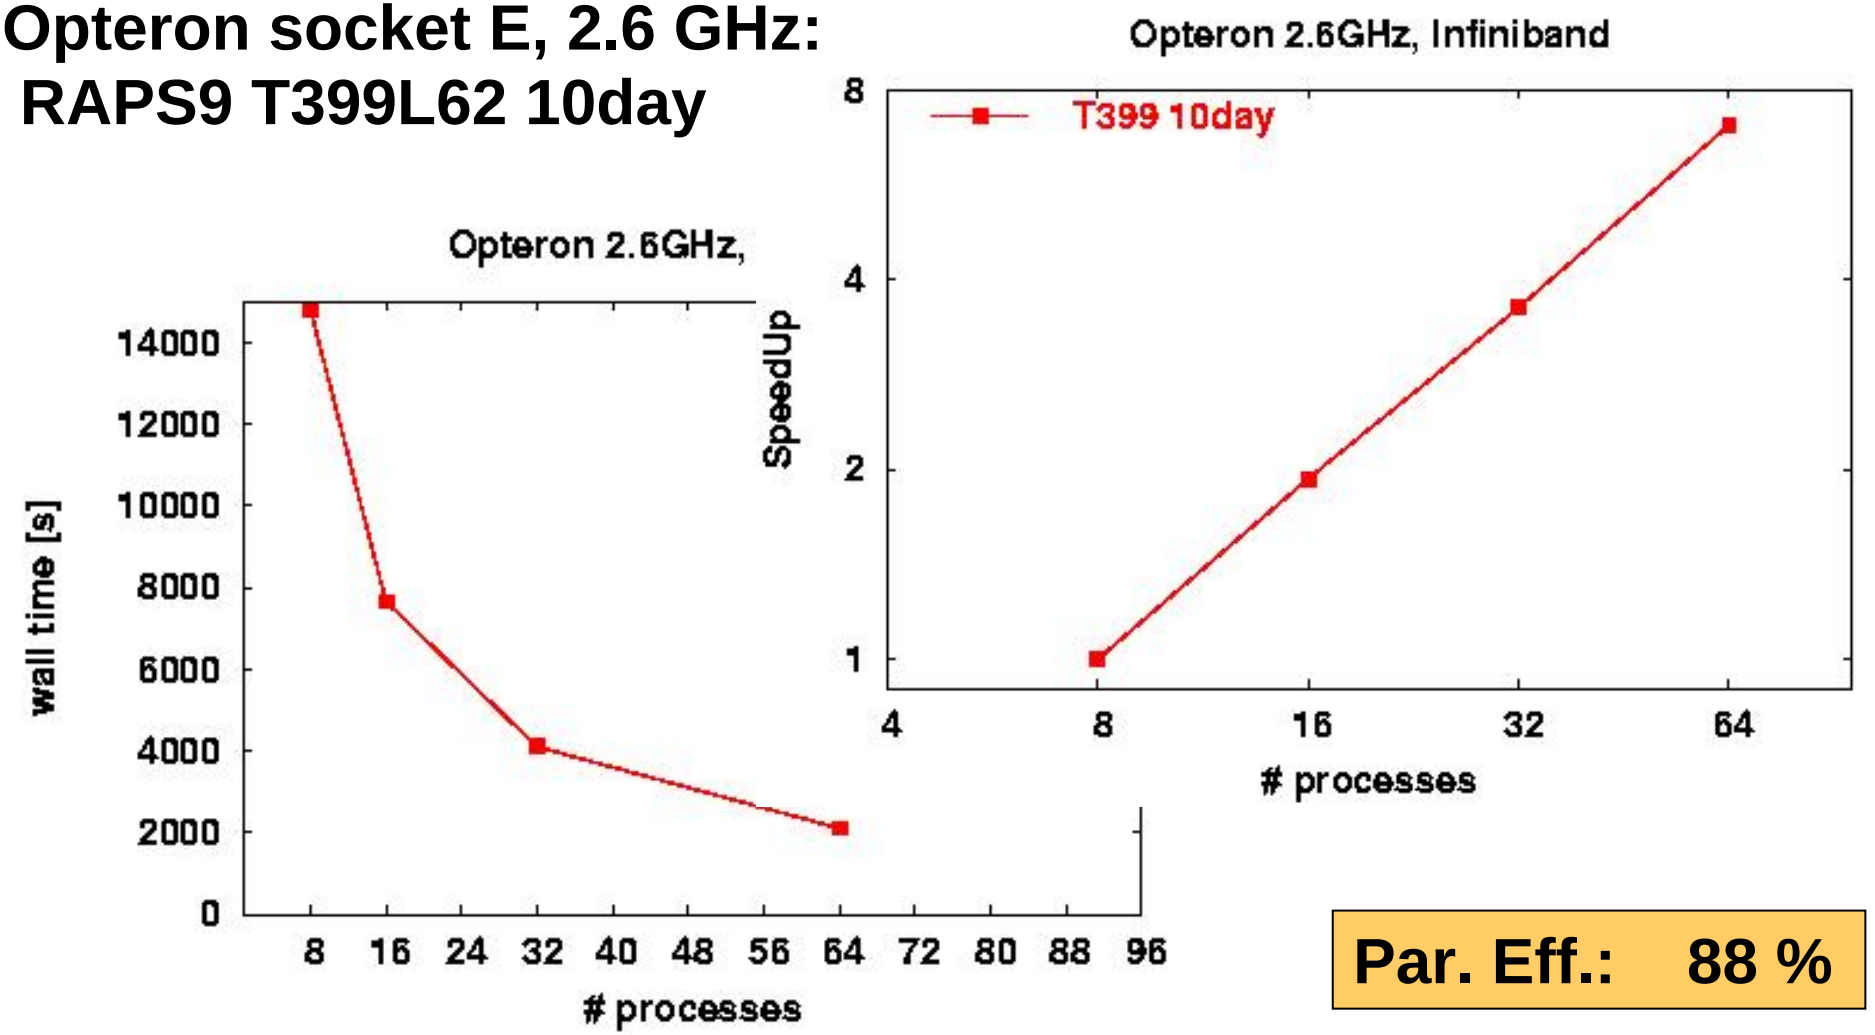

Par. Eff.: 88 %

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Woodcrest 3.0 GHz: RAPS9 T399L62 10day

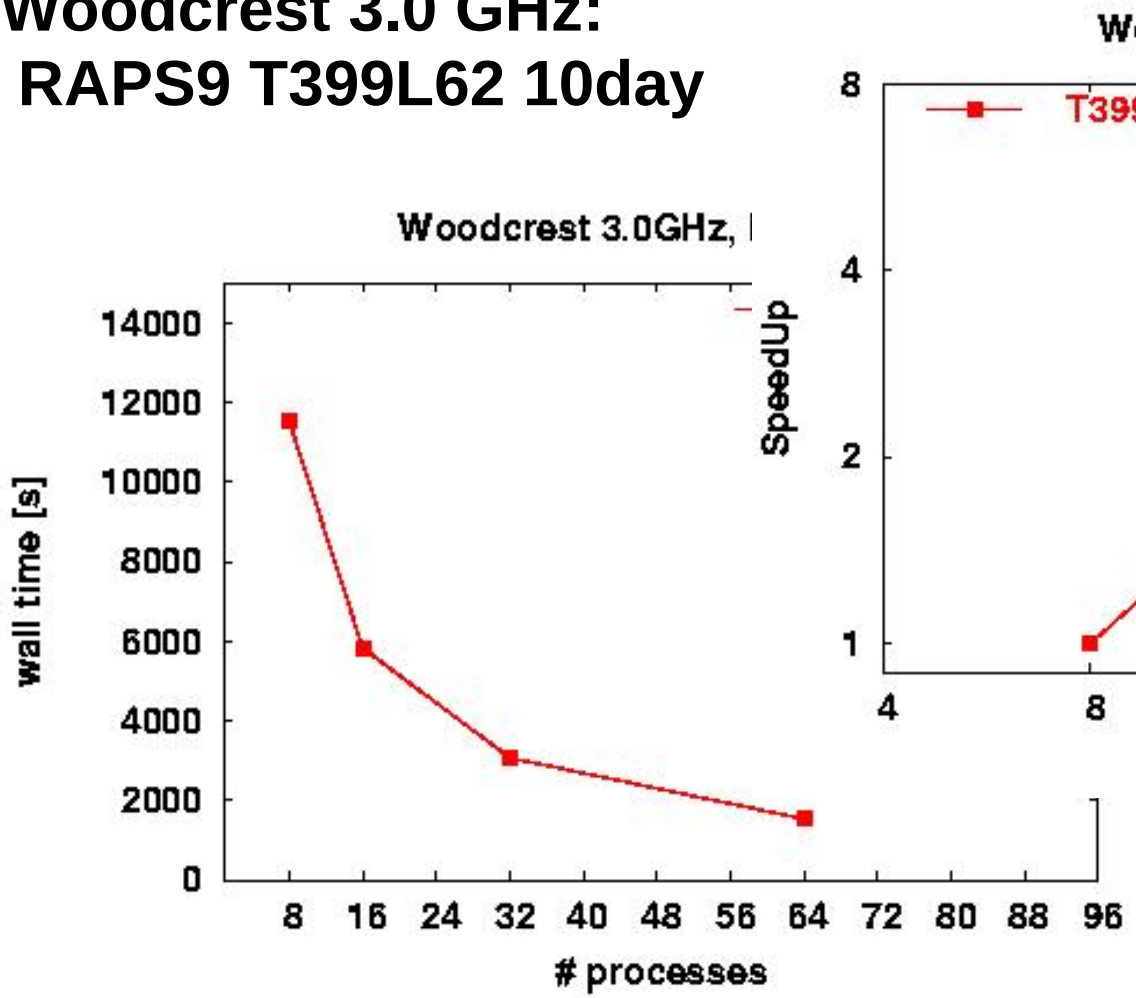

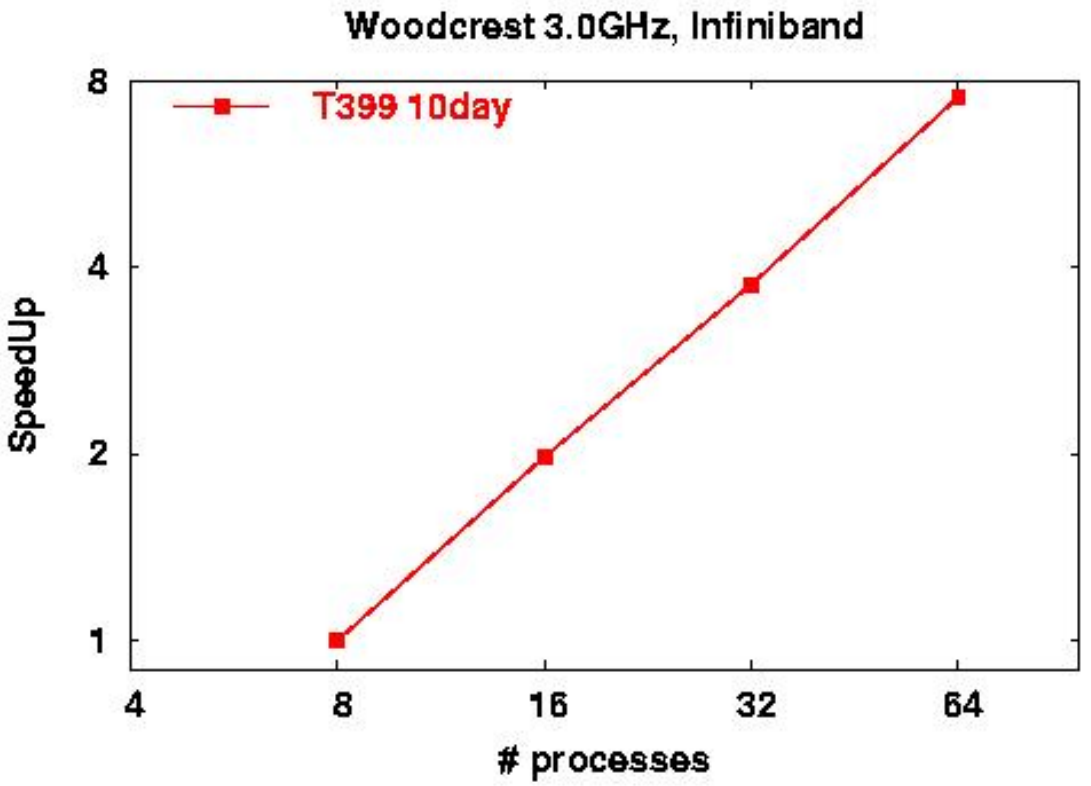

Par. Eff.: 94 %

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 T399L62

	Architecture	#procs	Walltime
T399L62 real	Opteron 2.6 GHz	32	341.5 s
T399L62 real	Woodcrest 3.0GHz	32	241.4 s
T399L62 real	IBM p690+ 1.9 GHz	8x4	282.3 s

	Architecture	#procs	Walltime
T399L62 10day	Opteron 2.6 GHz	64	2103.4 s
T399L62 10day	Woodcrest 3.0 GHz	64	1529.6 s
T399L62 10day	IBM p690+ 1.9 GHz	32x4	1032.3 s

RAPS9 T399L62

	Woodcrest 3.0 GHz		Opteron 2.6GHz	
	32 procs	64 procs	32 procs	64 procs
SUMMED TIME IN COMMUNICATIONS	36.2 s	20.5 s	34.6 s	29.9 s
SUMMED TIME IN PARALLEL REGIONS	182.3 s	93.6 s	263.0 s	130.2 s
SUMMED TIME IN I/O SECTIONS	7.9 s	7.9 s	18.0 s	12.0 s
SUMMED TIME IN SERIAL SECTIONS	8.2 s	4.7 s	20.8 s	7.6 s
SUMMED TIME IN BARRIERS	1.2 s	0.6 s	2.0 s	1.7 s
TOTAL TIME	241 s	127 s	342 s	181 s

RAPS9 T399L62 with OpenMP

	Architecture	#mpi x omp	Walltime
T399L62 real	Opteron 2.6 GHz	32x1	314.8 s
T399L62 real	Opteron 2.6 GHz	8x4	295.4 s

	Architecture	#mpi x omp	Walltime
T399L62 real	Woodcrest 3.0 GHz	32x1	241.4 s
T399L62 real	Woodcrest 3.0 GHz	8x4	229.3 s

Opteron socket E, 2.6 GHz: RAPS9 T511L60 10day

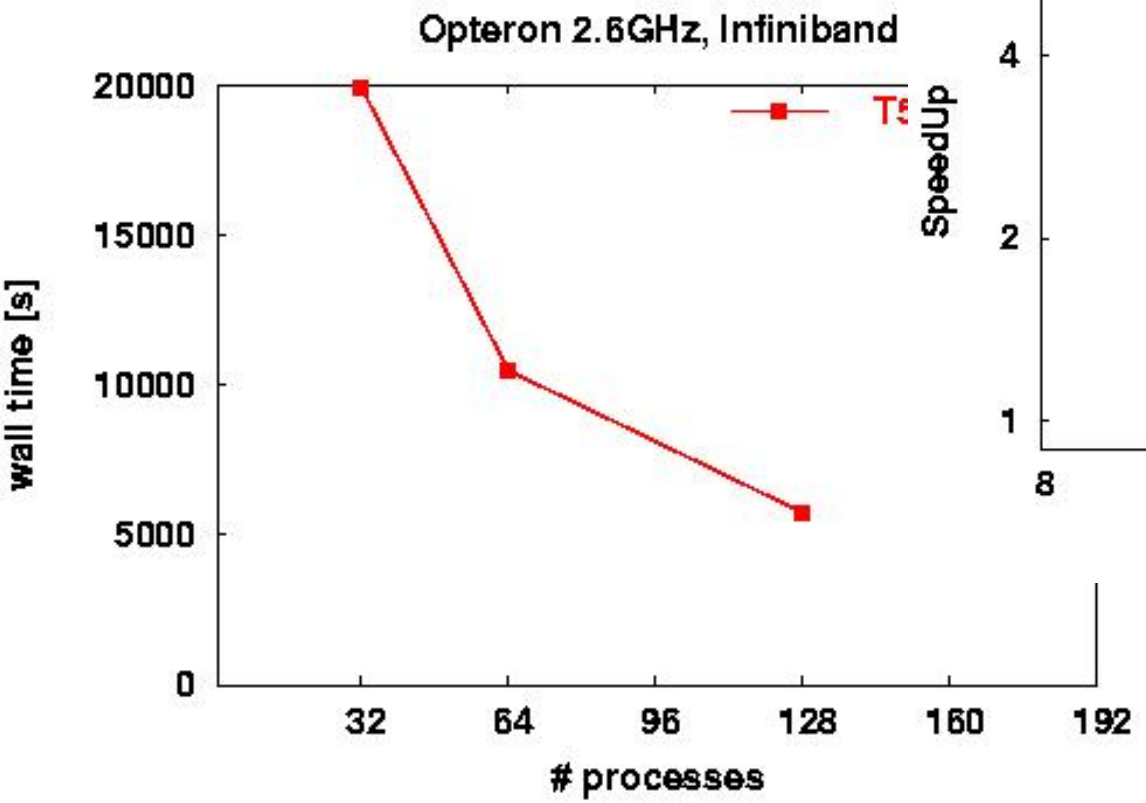

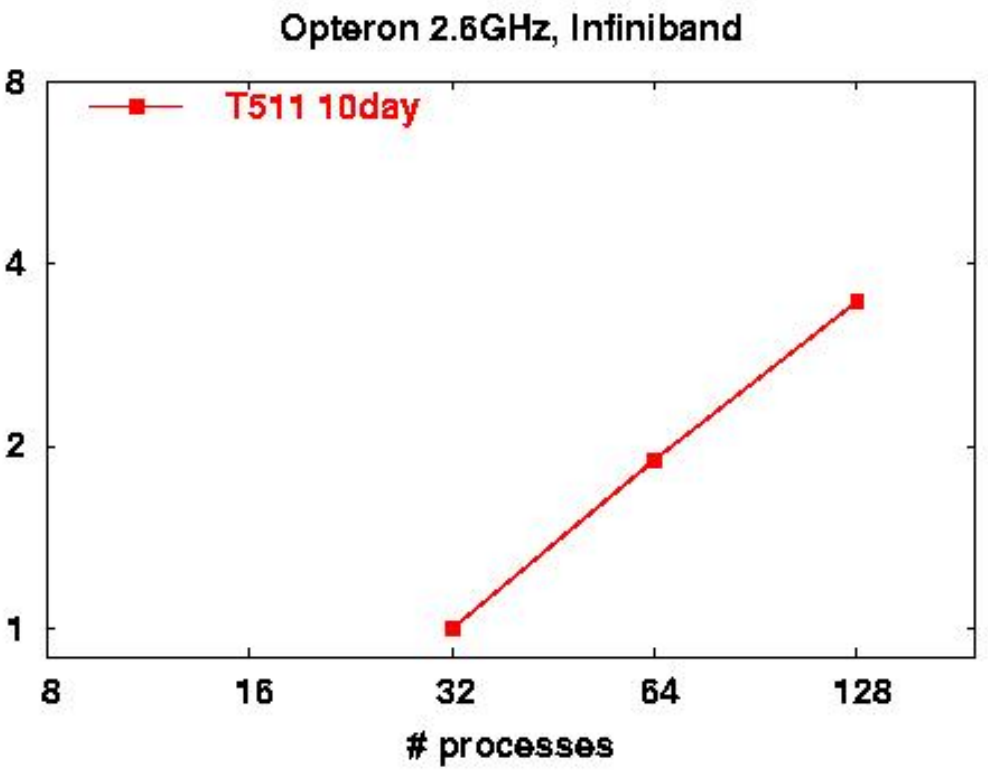

Par. Eff.: 87 %

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Woodcrest 3.0 GHz: RAPS9 T511L60 10day

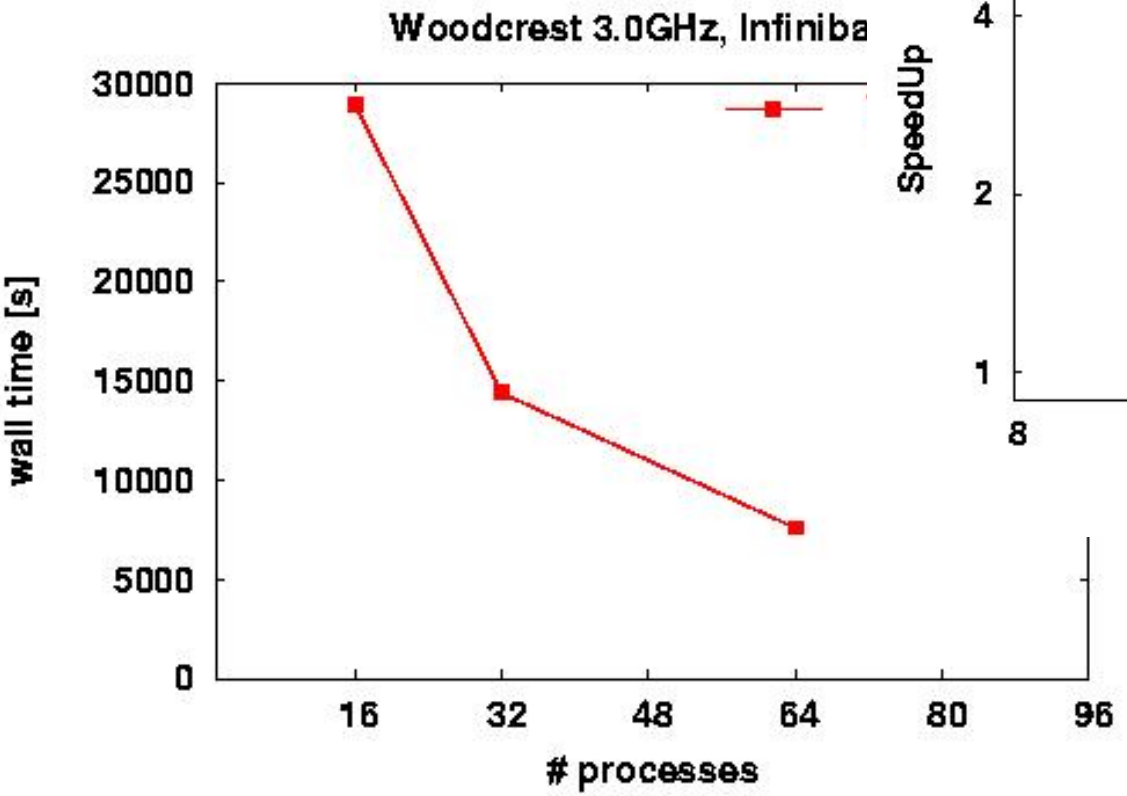

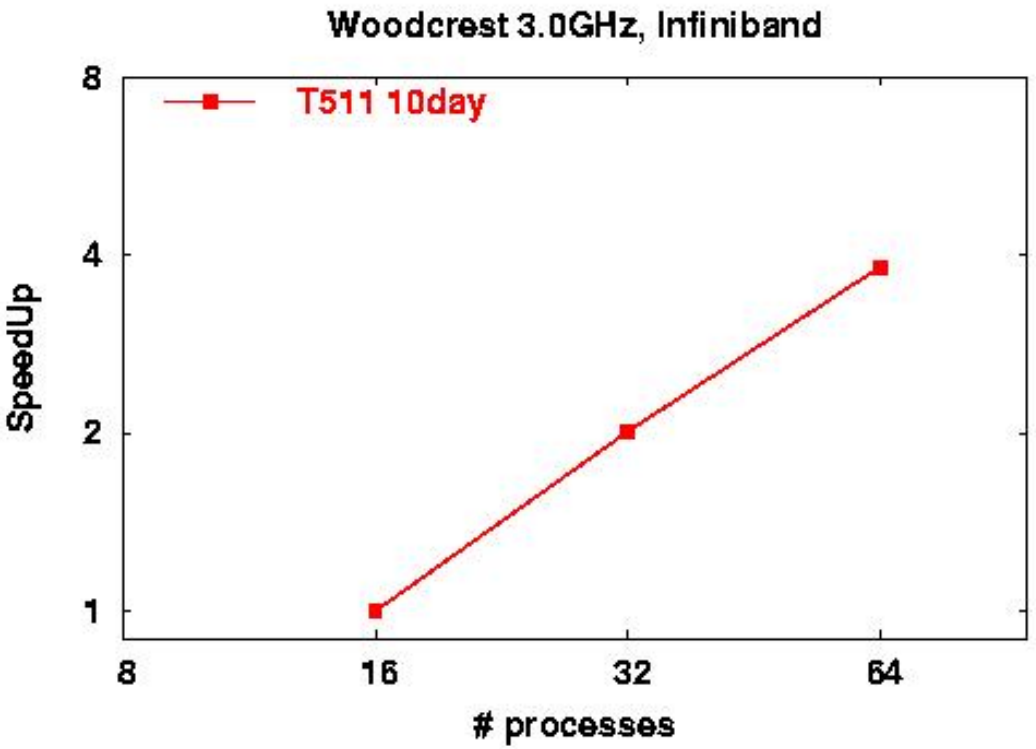

Par. Eff.: 95 %

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 T511L60

	Architecture	#procs	Walltime
T511L60 real	Opteron 2.6 GHz	64	356.7 s
T511L60 real	Woodcrest 3.0 GHz	64	236.1 s

	Architecture	#procs	Walltime
T511L60 10day	Opteron 2.6 GHz	64	10480 s
T511L60 10day	Woodcrest 3.0 GHz	64	7580 s

RAPS9 T511L60

	Woodcrest 3.0 GHz		Opteron 2.6 GHz	
	32 procs	64 procs	64 procs	128 procs
SUMMED TIME IN COMMUNICATIONS	50.1 s	33.1 s	53.4 s	36.1 s
SUMMED TIME IN PARALLEL REGIONS	320.5 s	162.3 s	234.6 s	123.1 s
SUMMED TIME IN I/O SECTIONS	16.4 s	23.3 s	36.7 s	36.5 s
SUMMED TIME IN SERIAL SECTIONS	5.7 s	10.3 s	23.3 s	46.4 s
SUMMED TIME IN BARRIERS	1.6 s	1.7 s	2.4 s	3.6 s
TOTAL TIME	396 s	236 s	357 s	247 s

RAPS9 T799L91

	Architecture	#procs	Walltime
T799L91 real	Opteron 2.6 GHz	128	820.5 s
T799L91 real	Woodcrest 3.0 GHz	64	935.6 s

	Architecture	#procs	Walltime
T799L91 10day	Opteron 2.6 GHz	128	28762 s
T799L91 10day	Woodcrest 3.0 GHz	64	36754 s

See talk of Nikita Panov,INTEL on Thursday

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 T799L91

	Woodcrest 3.0 GHz	Opteron 2.6 GHz
	64 procs	128 procs
SUMMED TIME IN COMMUNICATIONS	149.4 s	110.5 s
SUMMED TIME IN PARALLEL REGIONS	630.4 s	478.9 s
SUMMED TIME IN I/O SECTIONS	100.6 s	160.0 s
SUMMED TIME IN SERIAL SECTIONS	22.8 s	31.9 s
SUMMED TIME IN BARRIERS	4.9 s	32.1 s
TOTAL TIME	936 s	821 s

First results on 4D-Var

	Opteron 2.6 GHz	Opteron 2.6 GHz
	T159, 16 procs	T159, 32 procs
ifstraj0	708 s	421 s
ifsmin0	1075 s	602 s
ifstraj1	106 s	69 s
ifsmin1	1088 s	540 s
ifstraj2	479 s	306 s
eq8i	3456 s	1938 s

	Opteron 2.6 GHz	Woodcrest 3.0 GHz
	T511, 32 procs	T511, 64 procs
ifstraj0	1951 s	1113 s
ifsmin0	1813 s	975 s
ifstraj1	1342 s	839 s
ifsmin1	4112 s	1959 s
ifstraj2	1577 s	967 s
eqgp	10795 s	5853 s

16 → 32 : 1.8 Speed-Up

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

Summary and next steps

- IFS scales very well
- Optimiztion done ONLY via compiler options -> further optimization possible
- parallel file system necessary for improved performance
- Further performance measurements are in progress for 4D-Var, T159 and T511
- The mixed code (OpenMP + MPI) is undergoing further investigation to improve the performance
- Additional results in Nikita Panov's presentation

RAPS9 T511

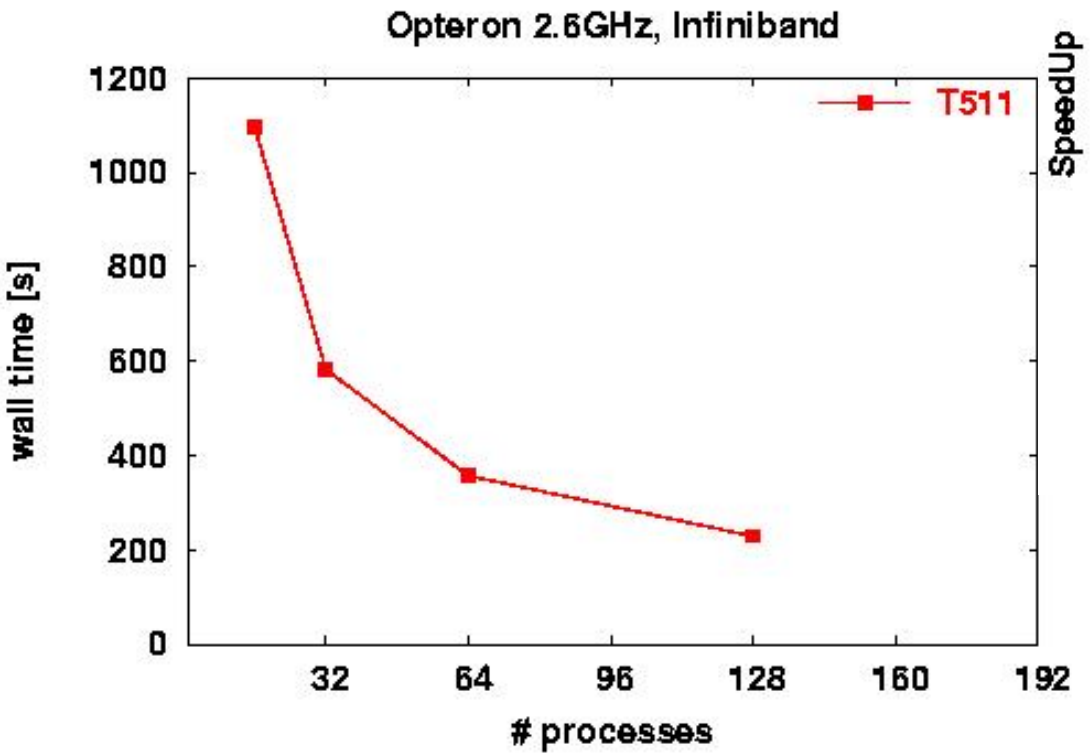

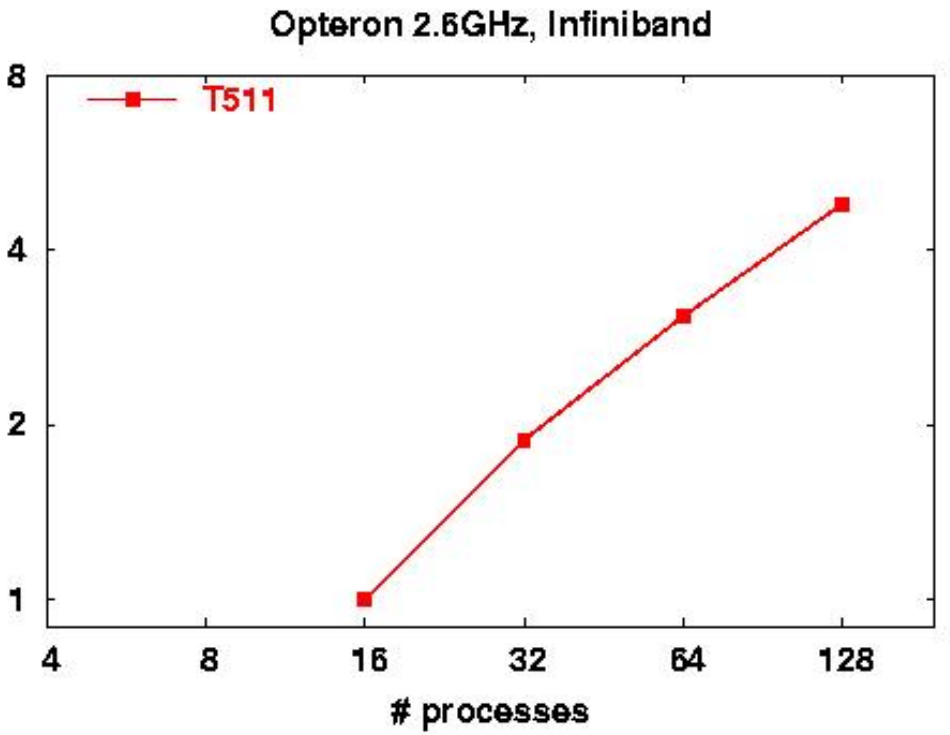

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 T511

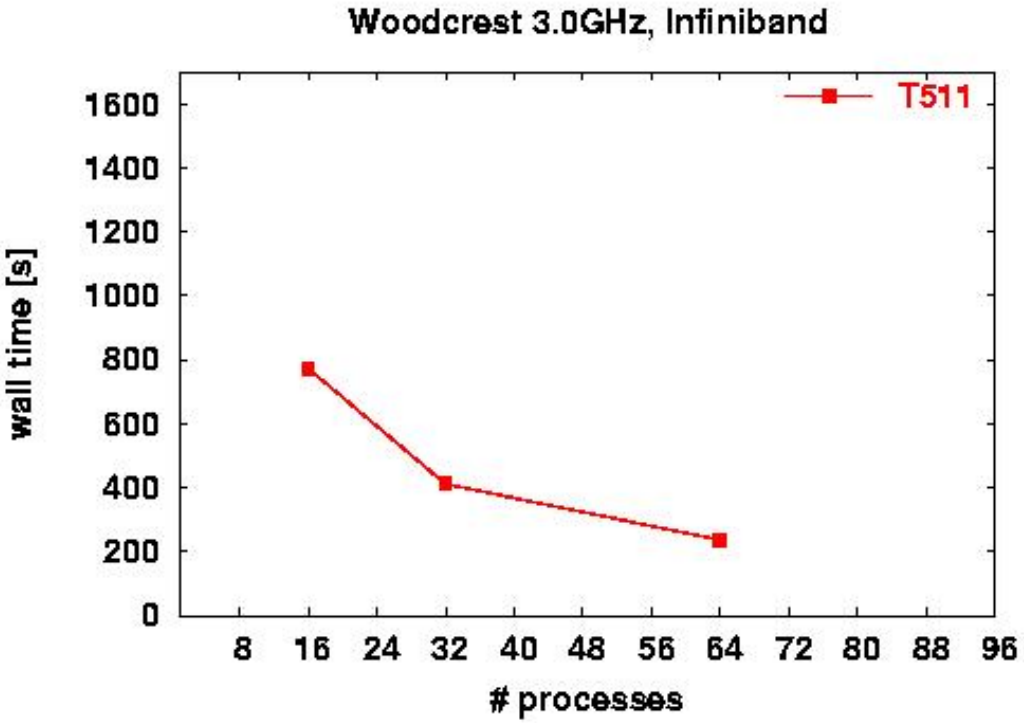

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 IFS T399

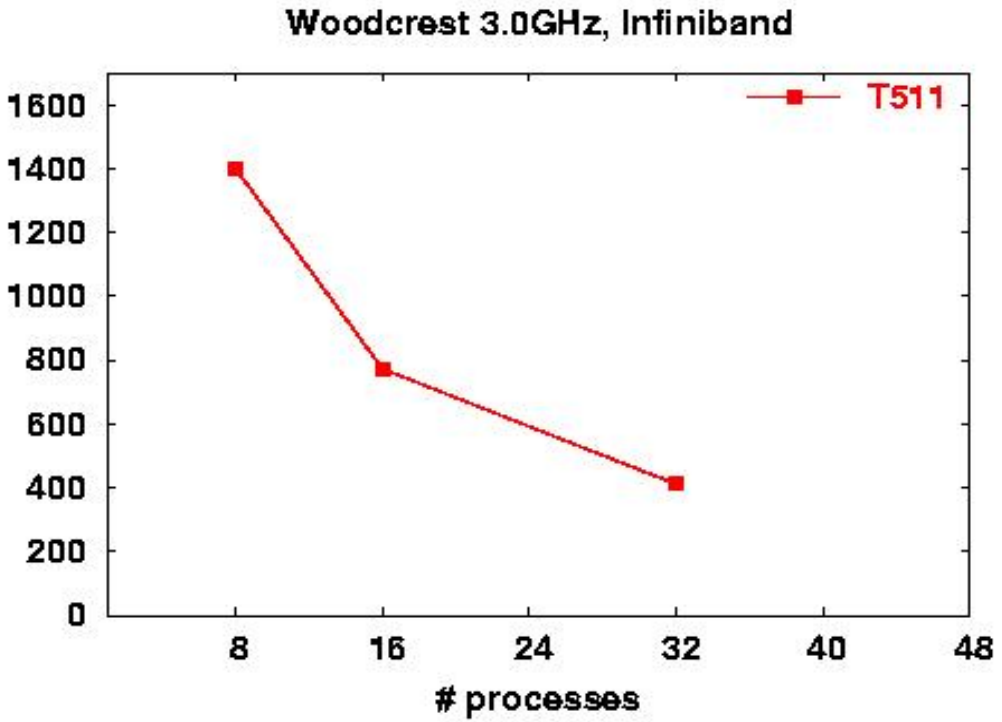

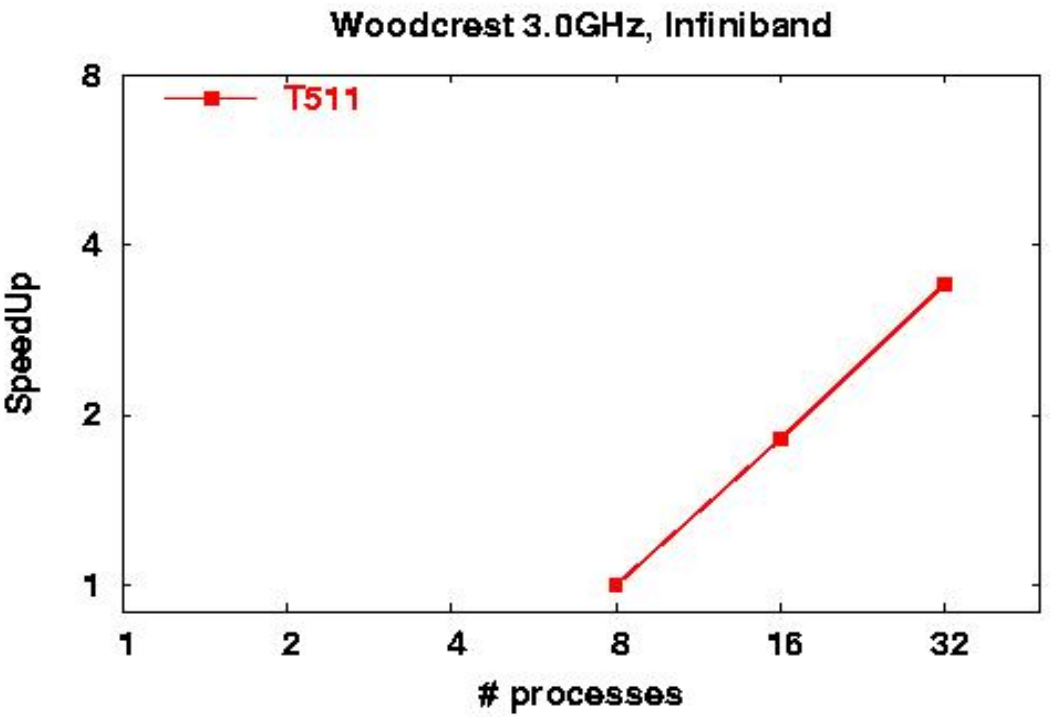

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX

RAPS9 IFS T399

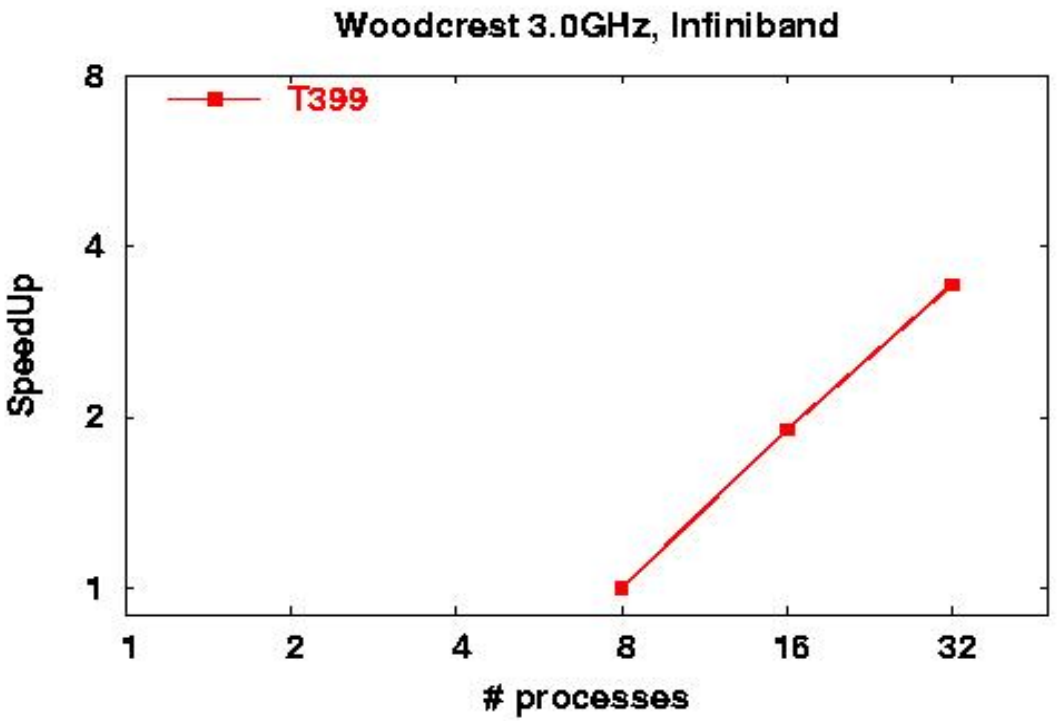

Fraunhofer ITWM Institut Techno- und Wirtschaftsmathematik

LINUX NETWORX