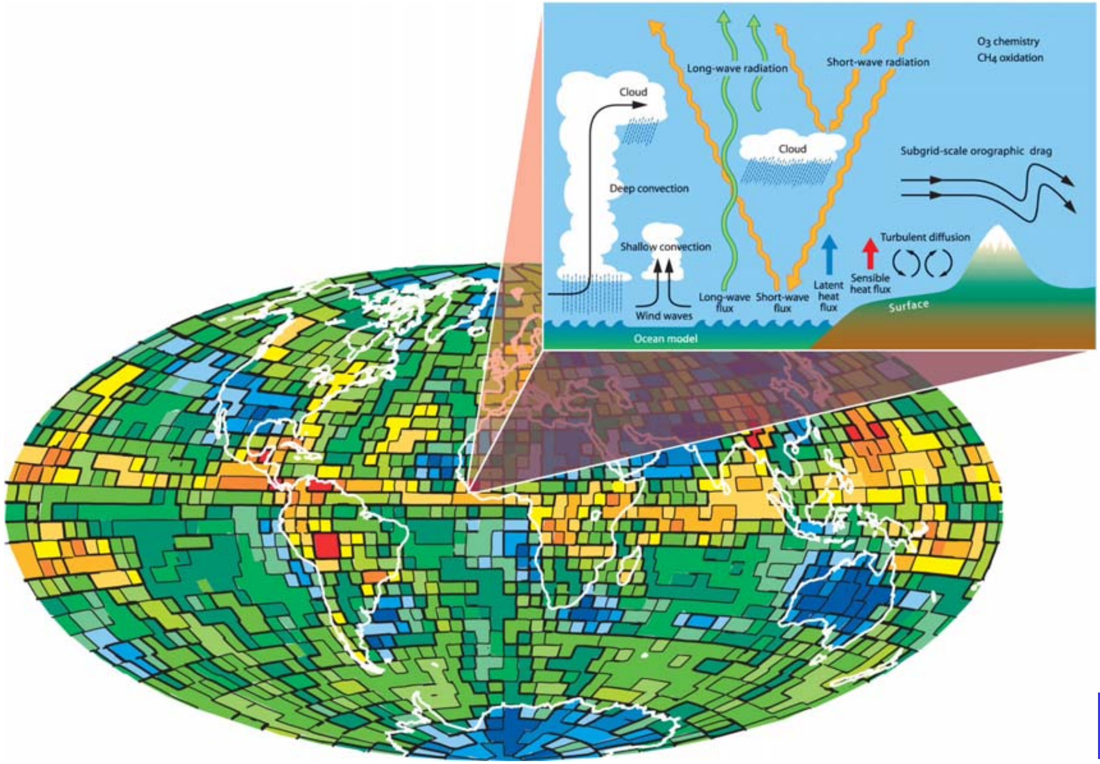

IFS Scaling

George Mozdzynski and Mats Hamrud
ECMWF

Outline

- Motivation
  - Running future IFS applications on systems with 100K cores and tight operational deadlines
- Initial focus on IFS model (not 4D-Var)
- Run model for different numbers of tasks
- Compare run with n/2 tasks with n tasks
- Analyze instrumentation counters
- What are the main reasons for non-perfect scaling
  - MPI communications ?
  - Load imbalance ?
  - O/S jitter ?
  - Other ?
- Make recommendations to improve scalability

T399 model (EPS resolution)

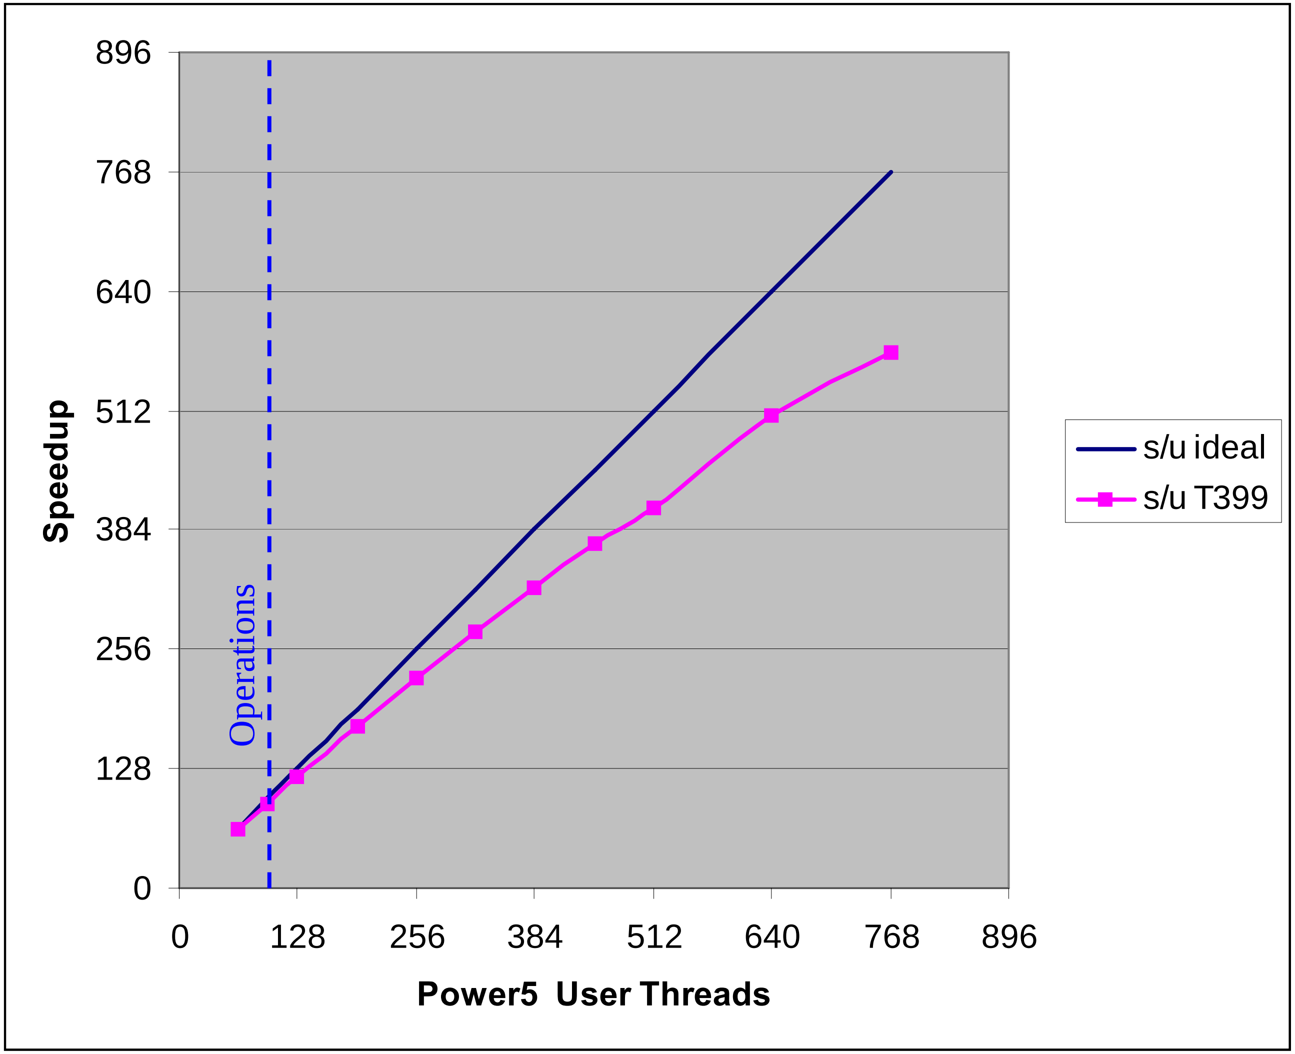

ECMWF

T799 model

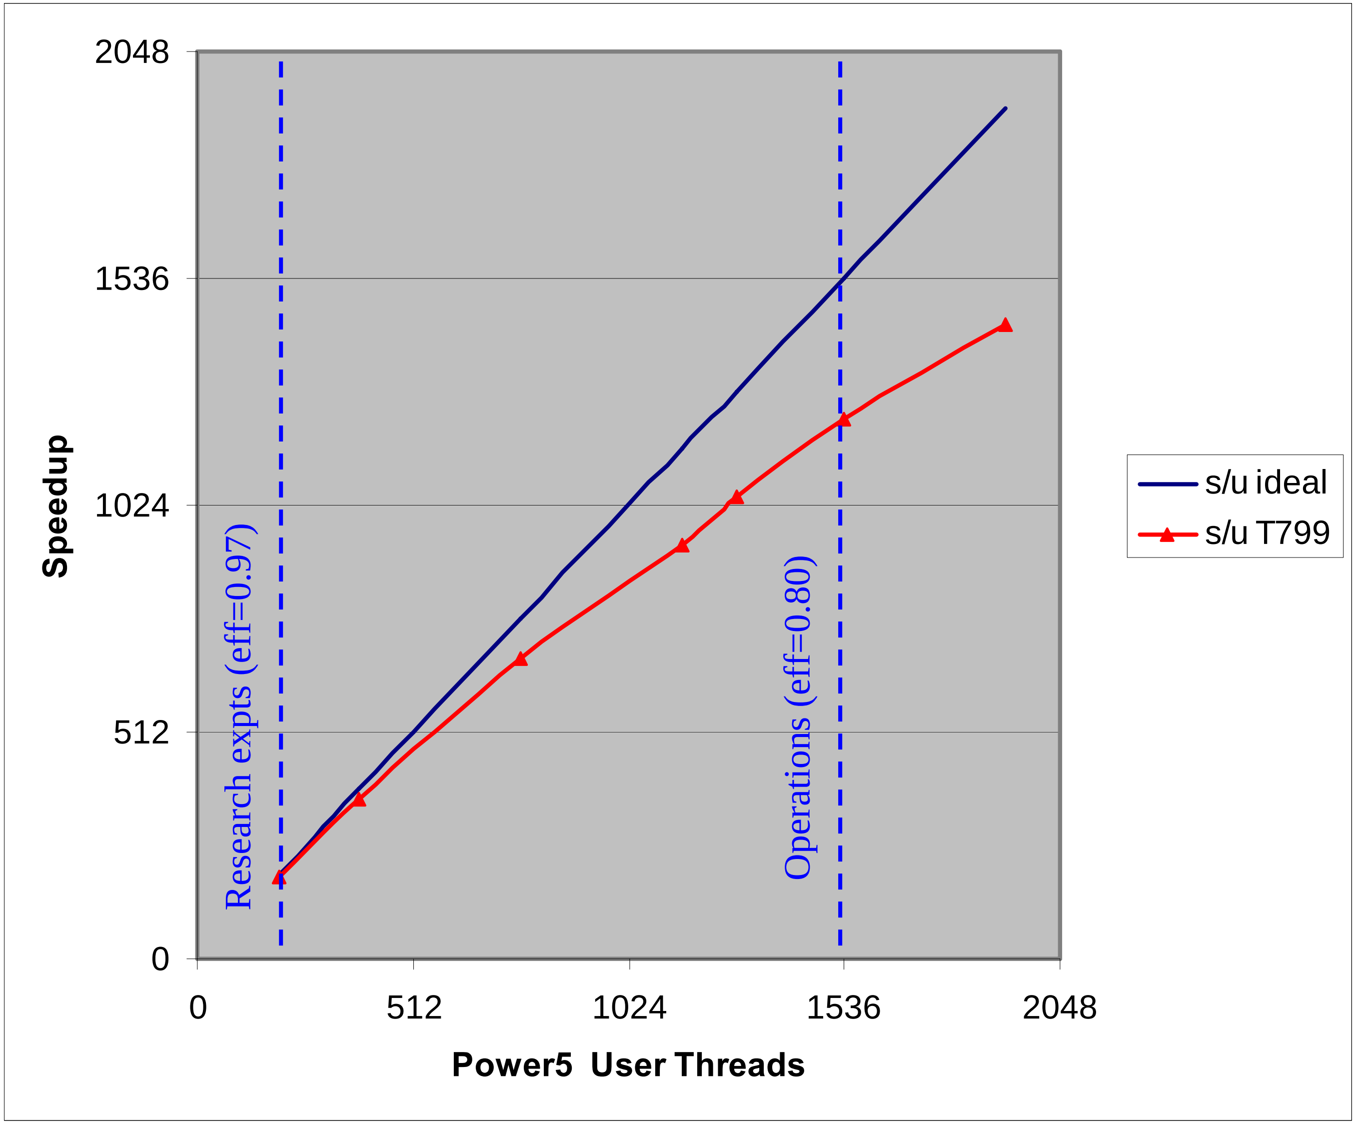

ECMWF

Calculating SpeedUp and Efficiency
T799L91 model, 2 day forecast

parallel	serial
649849.3	115.4

CPUS (User threads)	Actual Wall Time	Calculated Wall Time	Calculated SpeedUp	Calculated efficiency
192	3505.3	3500.0	185.4	96.6
384	1794.6	1807.7	362.2	94.3
768	958.5	961.6	678.1	88.3
1152	695.7	679.5	934.3	81.1
1280	623.2	623.1	1042.9	81.5
1536	533.2	538.5	1219.0	79.4
1920	453.7	453.9	1432.6	74.6
1		649964.7		

T399 v T799 Scaling

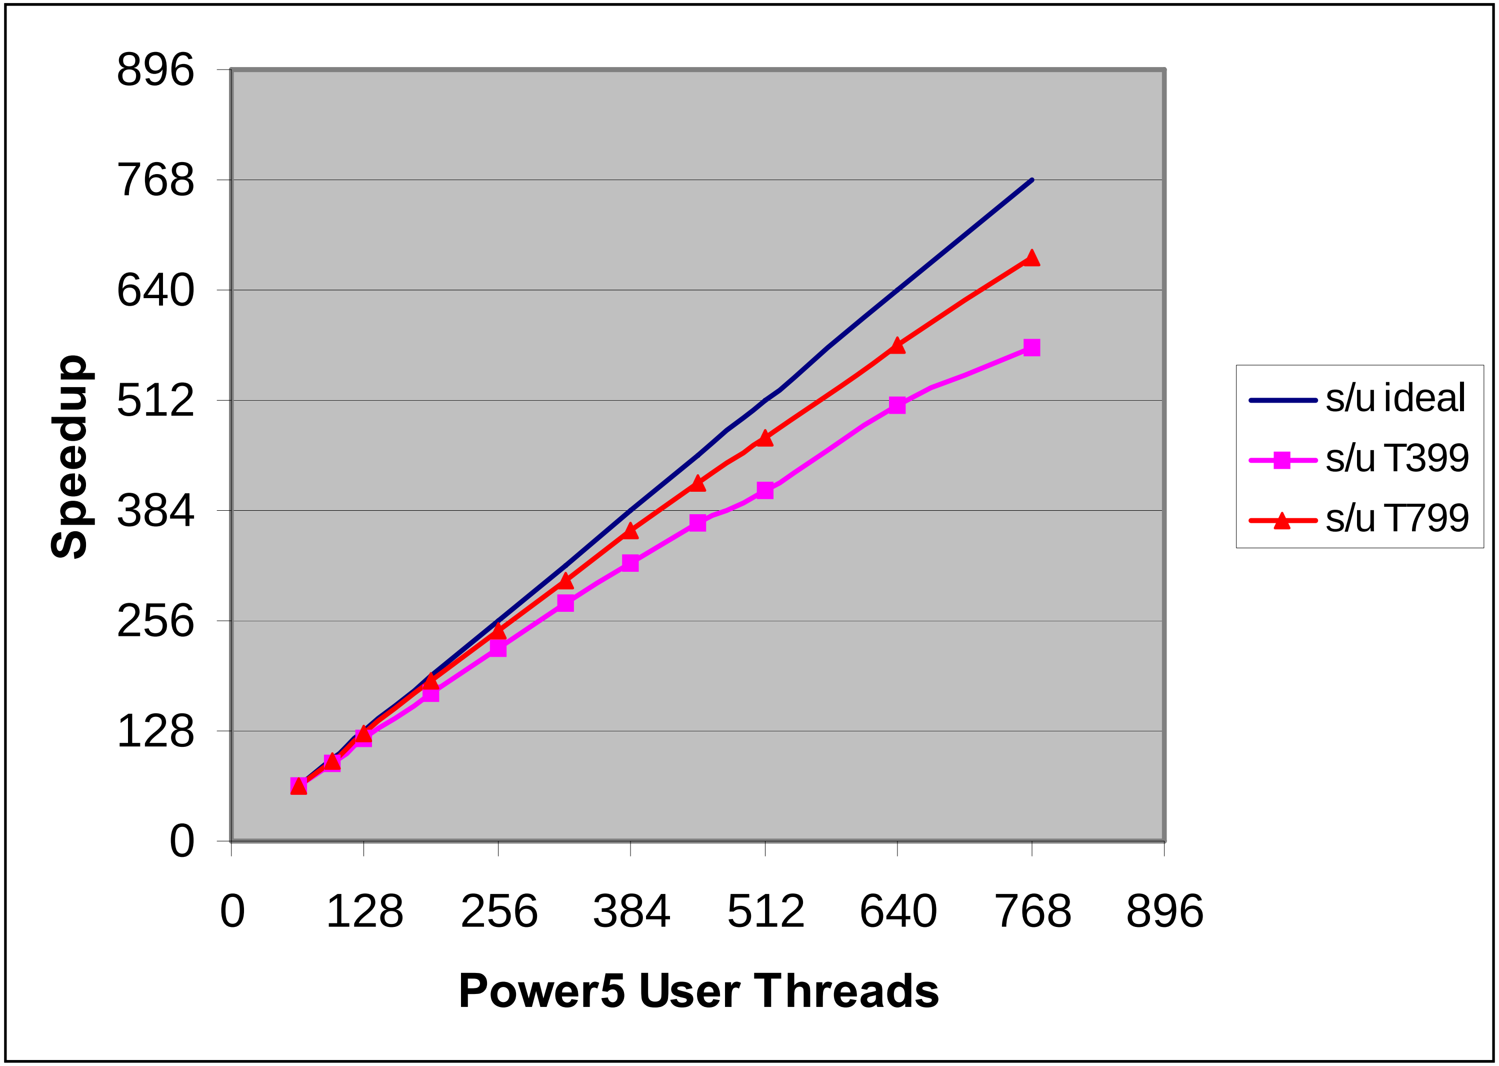

ECMWF

IFS instrumentation (GSTATS package)

- About 2000 counters
- High level summary
- Counter groups (MPL, OMP, SER, IO, BAR)
- LDETAILED_STATS=T
  - Detailed printout per task
  - Summary per task
  - Summary per counter group
- LBARRIER_STATS=T
  - Extra barriers to accurately time areas of load imbalance

GSTATS (high level summary)

STATS FOR ALL PROCESSORS

NUM	ROUTINE		CALLS	MEAN(ms)	MAX(ms)	FRAC(%)	UNBAL(%)
0	CNT0	- COMPLETE EXECUTION	1	2631728.6	2631881.8	100.01	0.01
1	CNT4	- FORWARD INTEGRATION	1	2609593.6	2610044.0	99.18	0.02
8	SCAN2MDM	- GRID-POINT DYNAMICS	1285	123.8	126.3	6.17	2.07
9	SPCM	- SPECTRAL COMP.	1200	34.2	46.8	2.13	36.72
10	SCAN2MDM	- PHYSICS	1201	728.5	827.9	37.78	13.64
11	IOPACK	- OUTPUT P.P. RESULTS	63	82.9	90.2	0.22	8.90
12	SPNORM	- SPECTRAL NORM COMP.	1203	2.5	3.4	0.16	34.65
13	SCAN2MDM	- RADIATION CALC.	241	1796.2	1845.9	16.90	2.77
14	SUINIF		1	11506.4	11715.0	0.45	1.81
17	GRIDFPOS	IN CNT4	21	119.9	120.4	0.10	0.41
18	SUSPECG		1	3759.6	3965.7	0.15	5.48
19	SUSPEC		1	3780.1	3985.8	0.15	5.44
24	SUGRIDU		1	6756.9	6766.9	0.26	0.15
25	SPECRTDM		1	86.1	90.6	0.00	5.26
26	SUGRIDF		1	882.5	900.1	0.03	1.99
37	CPGLAG	- SL COMPUTATIONS	1201	227.2	242.3	11.06	6.63
38	WAM	- TOTAL COST OF WAVE MODEL	1200	90.9	91.9	4.19	1.09
51	SCAN2MDM	- SL COMM. PART 1	1201	52.6	88.5	4.04	68.39
54	SPCM	- M TO S/S TO M TRANSP.	1200	31.1	40.7	1.86	31.16
56	SPNORM	- SPECTRAL NORM COMM.	1203	3.3	5.3	0.24	60.13
65	RADINTG	- SL COMM. OUTPUT	241	147.1	228.5	2.09	55.37
66	RADINTG	- SL COMM. INPUT	241	35.7	82.1	0.75	129.69
89	SCAN2MDM	- SL COMM. PART 2	1201	43.7	50.0	2.28	14.37
102	LTINV_CTL	- INVERSE LEGENDRE TRANSFORM	1225	61.4	64.4	3.00	4.78
103	LTDIR_CTL	- DIRECT LEGENDRE TRANSFORM	1243	36.1	37.8	1.78	4.48
106	FTDIR_CTL	- DIRECT FOURIER TRANSFORM	1243	6.5	7.8	0.37	20.49
107	FTINV_CTL	- INVERSE FOURIER TRANSFORM	1225	16.6	20.1	0.94	21.25
140	SULEG	- COMP. OF LEGENDRE POL.	2	207.2	226.4	0.02	9.28
152	LTINV_CTL	- M TO L TRANSPOSITION	1225	58.4	64.3	2.99	10.01
153	LTDIR_CTL	- L TO M TRANSPOSITION	1243	34.2	59.0	2.79	72.50
157	FTINV_CTL	- L TO G TRANSPOSITION	1225	63.7	86.2	4.01	35.36
158	FTDIR_CTL	- G TO L TRANSPOSITION	1243	99.8	184.8	8.73	85.14
190	SUTRLE	- COMMUNICATE LEG.POL.	1	247.4	374.0	0.01	51.13

GSTATS, detailed stats per task
(only showing counters > 1 percent, for task 87)

```
TIMING STATISTICS:PROCESSOR=  87
STARTUP COST    0.1 SECONDS
 NUM     ROUTINE                                  CALLS    SUM(s)   AVE(ms) STDDEV(ms)  MAX(ms)   FRAC(%)
 509 MPL SLCOMM1_COMMS                             1203      34.4      28.6       1.8      42.3      1.31
 604 MPL GATHERV IN GPNORM1                        6019      43.2       7.2       6.4     474.4      1.64
 759 BAR BARRIER IN SLCOMM2                         480      35.1      73.2      64.5     215.5      1.33
 761 BAR BARRIER IN TRGTOL                         1243      64.5      51.9      28.3     225.6      2.45
 805 MPL TRLTOG_COMMS                              1225      75.2      61.4       3.8     121.8      2.86
 806 MPL TRLTOM_COMMS                              1243      29.5      23.7       4.5      61.8      1.12
 807 MPL TRMTOL_COMMS                              1225      64.2      52.4       2.9      65.2      2.44
1001 OMP PHYSICS                                   1201     885.2     737.0      33.1    1082.9     33.63
1004 OMP CALL_SL      1                            1201     137.5     114.4       1.1     119.1      5.22
1005 OMP CALL_SL      2                            1201     146.5     122.0      11.1     138.0      5.57
1025 OMP CPG          1                            1201     153.5     127.8       7.8     388.3      5.83
1029 OMP SPCM                                      1200      49.6      41.3       0.0      42.3      1.88
1210 OMP RADINTG-RADLSW                             241     366.3    1519.9      27.5    1611.5     13.92
1431 OMP WAMODEL 2                                 1200      60.5      50.5       0.8      52.6      2.30
1645 OMP LTDIR_CTL    - DIRECT LEGENDRE TRANSFORM  1243      46.3      37.3       4.3      68.5      1.76
1647 OMP LTINV_CTL    - INVERSE LEGENDRE TRANSFORM 1223      77.0      63.0       1.5      67.0      2.93

SUMMED TIME IN COMMUNICATIONS    =      329.2 SECONDS  12.51 PERCENT OF TOTAL
SUMMED TIME IN PARALLEL REGIONS =     2062.0 SECONDS  78.35 PERCENT OF TOTAL
SUMMED TIME IN I/O SECTIONS      =        8.3 SECONDS   0.31 PERCENT OF TOTAL
SUMMED TIME IN SERIAL SECTIONS   =       15.7 SECONDS   0.60 PERCENT OF TOTAL
SUMMED TIME IN BARRIERS          =      180.0 SECONDS   6.84 PERCENT OF TOTAL
FRACTION OF TOTAL TIME ACCOUNTED FOR    98.61
```

GSTATS analysis, 192 to 384 tasks (counters with lost time > 2 secs)

```
Output file 1 - f0y7.192x4.lbarrier_stats=t
Output file 2 - f0y7.384x4.lbarrier_stats=t
Hoped for speedup factor -  2
Summed time of job 1 4760.08 s
Summed time of job 2 2633.92 s
Achieved speedup factor 1.807
```

Id	Descriptor	Calls	Time1(s)	Time2(s)	Speedup	Lost time(s)
509	MPL SLCOMM1_COMMS	1203	40.3	33.8	1.19	13.7
512	MPL SLCOMM2A_COMMS	2402	23.1	18.2	1.27	6.7
604	MPL GATHERV IN GPNORM1	6019	34.5	44.8	0.77	27.6
616	MPL ALLGATHERV IN MPBCASTSCFLD	1200	25.8	17.5	1.47	4.6
658	MPL SIGCHECK	1201	4.1	5.3	0.77	3.2
694	MPL MPDECOMP BCAST	10	7.2	6.7	1.07	3.1
757	BAR BARRIER IN SLCOMM1	1201	32.1	25.8	1.24	9.7
759	BAR BARRIER IN SLCOMM2	480	60.4	41.5	1.46	11.3
761	BAR BARRIER IN TRGTOL	1243	157.9	93.5	1.69	14.5
763	BAR BARRIER IN TRLTOM	1243	14.8	12.8	1.16	5.4
764	BAR BARRIER IN TRMTOL	1225	7.2	7.5	0.96	3.9
766	BAR BARRIER IN TRSTOM	1200	23.2	17.0	1.36	5.4
803	MPL TRGTOL_COMMS	1243	37.7	26.3	1.43	7.4
805	MPL TRLTOG_COMMS	1225	85.7	62.3	1.38	19.4
806	MPL TRLTOM_COMMS	1243	57.3	30.8	1.86	2.1
807	MPL TRMTOL_COMMS	1225	112.6	64.1	1.76	7.8
810	MPL GATH_SPEC_CONTROL_COMMS	42	16.5	18.9	0.87	10.7
1001	OMP PHYSICS	1201	1680.9	863.4	1.95	22.9
1029	OMP SPCM	1200	59.3	38.0	1.56	8.4
1121	OMP SLCOMM2a 2	1201	10.1	7.9	1.28	2.9
1210	OMP RADINTG-RADLSW	241	709.8	363.4	1.95	8.5
1431	OMP WAMODEL 2	1200	106.4	57.9	1.84	4.8
1645	OMP LTDIR_CTL - DIRECT LEGENDRE	1243	85.8	44.9	1.91	2.0
1710	IO- CNT4 IFLUSHFDB	21	4.0	11.6	0.34	9.6
1771	IO- MPDECOMP I/O	7	5.9	5.3	1.11	2.3
1902	UND SUOYOMB	1	4.4	6.1	0.72	3.9
1905	SER UPDTIM update	11	3.7	4.3	0.86	2.5

```
Total time lost=253.93 s
```

GSTATS analysis
(lost time, T799 from 192 to 384 tasks)

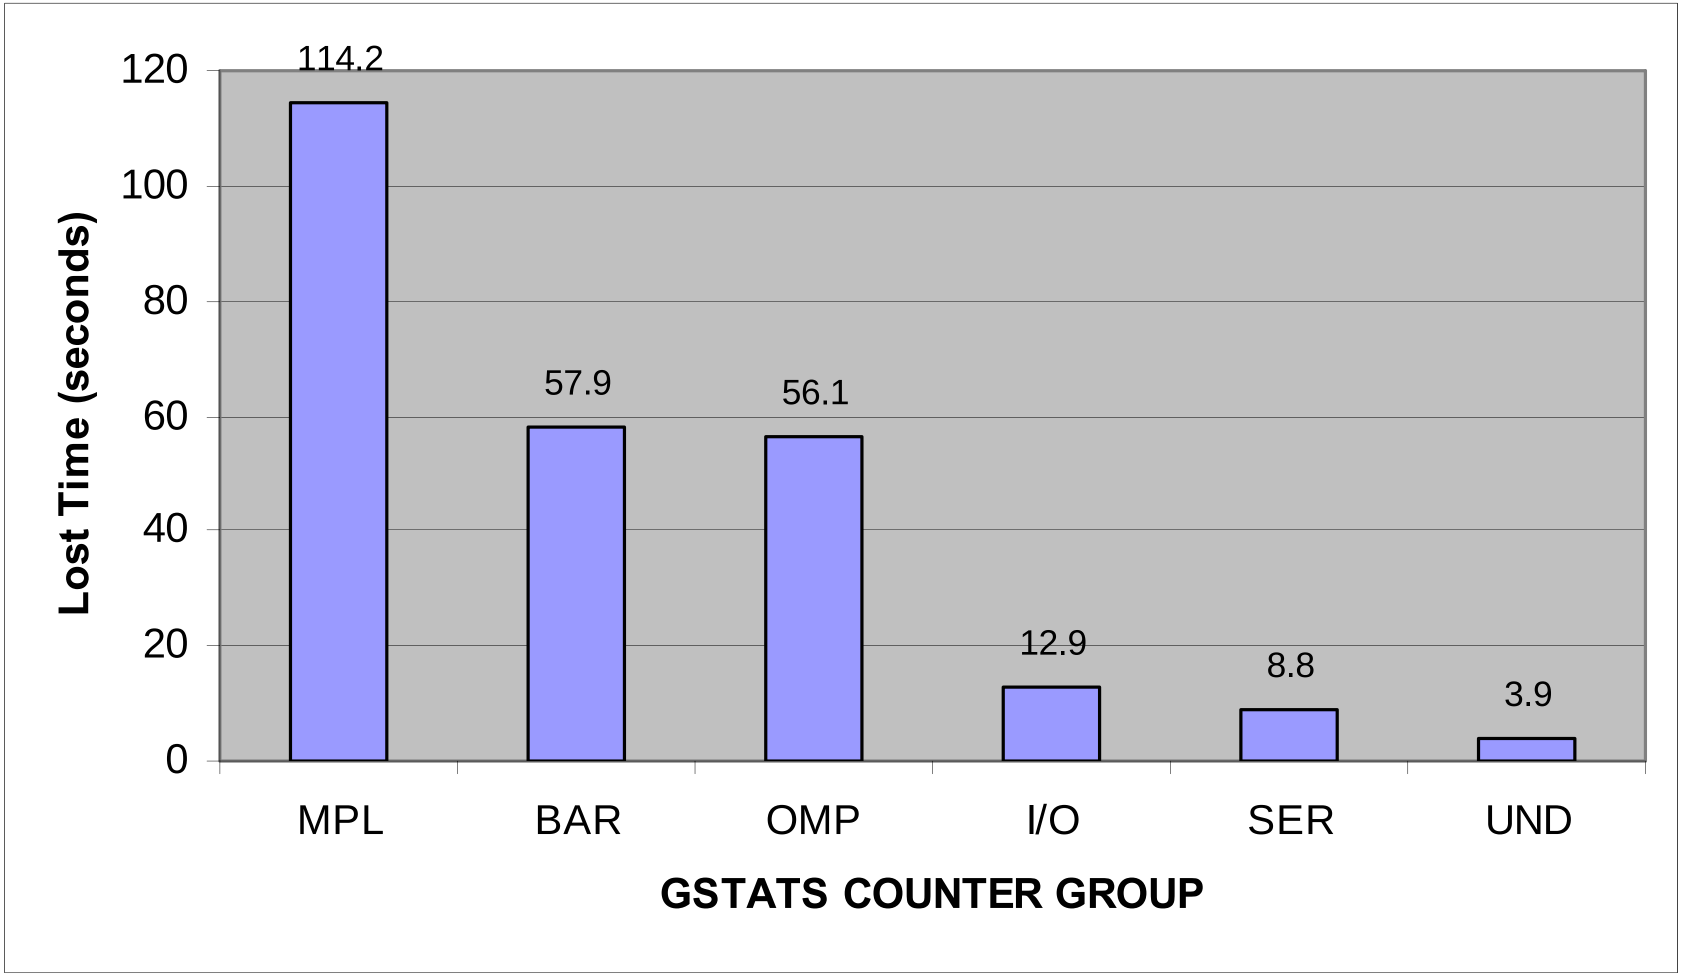

GSTATS analysis, 192 to 384 tasks (counters with lost time > 2 secs)

Output file 1 - f0y7.192x4.lbarrier_stats=t
Output file 2 - f0y7.384x4.lbarrier_stats=t
Hoped for speedup factor - 2
Summed time of job 1 4760.08 s
Summed time of job 2 2633.92 s
Achieved speedup factor 1.807

Id	Descriptor	Calls	Time1(s)	Time2(s)	Speedup	Lost time(s)
509	MPL SLCOMM1_COMMS	1203	40.3	33.8	1.19	13.7
512	MPL SLCOMM2A_COMMS	2402	23.1	18.2	1.27	6.7
604	MPL GATHERV IN GPNORM1	6019	34.5	44.8	0.77	27.6
616	MPL ALLGATHERV IN MPBCASTSCFLD	1200	25.8	17.5	1.47	4.6
658	MPL SIGCHECK	1201	4.1	5.3	0.77	3.2
694	MPL MPDECOMP BCAST	10	7.2	6.7	1.07	3.1
757	BAR BARRIER IN SLCOMM1	1201	32.1	25.8	1.24	9.7
759	BAR BARRIER IN SLCOMM2	480	60.4	41.5	1.46	11.3
761	BAR BARRIER IN TRGTOL	1243	157.9	93.5	1.69	14.5
763	BAR BARRIER IN TRLTOM	1243	14.8	12.8	1.16	5.4
764	BAR BARRIER IN TRMTOL	1225	7.2	7.5	0.96	3.9
766	BAR BARRIER IN TRSTOM	1200	23.2	17.0	1.36	5.4
803	MPL TRGTOL_COMMS	1243	37.7	26.3	1.43	7.4
805	MPL TRLTOG_COMMS	1225	85.7	62.3	1.38	19.4
806	MPL TRLTOM_COMMS	1243	57.3	30.8	1.86	2.1
807	MPL TRMTOL_COMMS	1225	112.6	64.1	1.76	7.8
810	MPL GATH_SPEC_CONTROL_COMMS	42	16.5	18.9	0.87	10.7
1001	OMP PHYSICS	1201	1680.9	863.4	1.95	22.9
1029	OMP SPCM	1200	59.3	38.0	1.56	8.4
1121	OMP SLCOMM2a 2	1201	10.1	7.9	1.28	2.9
1210	OMP RADINTG-RADLSW	241	709.8	363.4	1.95	8.5
1431	OMP WAMODEL 2	1200	106.4	57.9	1.84	4.8
1645	OMP LTDIR_CTL - DIRECT LEGENDRE	1243	85.8	44.9	1.91	2.0
1710	IO- CNT4 IFLUSHFDB	21	4.0	11.6	0.34	9.6
1771	IO- MPDECOMP I/O	7	5.9	5.3	1.11	2.3
1902	UND SUOYOMB	1	4.4	6.1	0.72	3.9
1905	SER UPDTIM update	11	3.7	4.3	0.86	2.5

Total time lost=253.93 s

GATHERV IN GPNORM1

```
GPNORM  HUMIDITY             AVERAGE                   MINIMUM                   MAXIMUM
        AVE    0.178473217723406E-02 0.999999999999757E-08 0.230986213507641E-01
GPNORM  LIQUID WATER         AVERAGE                   MINIMUM                   MAXIMUM
        AVE    0.328000416002415E-05 -.575982404132924E-19 0.206569592945375E-02
GPNORM  ICE WATER            AVERAGE                   MINIMUM                   MAXIMUM
        AVE    0.155138686301991E-05 -.296461531539005E-19 0.188392347144291E-02
GPNORM  CLOUD FRACTION       AVERAGE                   MINIMUM                   MAXIMUM
        AVE    0.503635639758255E-01 -.114491749414469E-15 0.100000000000000E+01
GPNORM  OZONE                AVERAGE                   MINIMUM                   MAXIMUM
        AVE    0.258063427899260E-05 0.999653412623365E-09 0.166043853080967E-04
```

- **Grid-point norms performed every time step (shown above)**
- **General case: multiple fields (i.e. levels) are gathered using MPI_GATHERV, all tasks sending to 91 tasks**
- **Dominant case (above): only 1 field is gathered using MPI_GATHERV, all tasks sending to 1 task, repeated 5 times per time step**

Grid-Point Norms (new approach)

- **Grid-point norms now done using a 2-D parallel approach**
- **Grid point variable is redistributed from (subset of grid-points, all levels) to (subset of full latitudes, subset of levels)**
- **This is exactly the redistribution that takes place going from grid space to Fourier space (so we can reuse TRGTOL)**
- **Partial sums computed in this new distribution (all tasks are now used in the general and dominant cases)**
- **E-W (NPRTRV) sums are message passed to E-W master**
- **N-S (NPRTRW) sums are message passed from the E-W masters to the global master**
- **Where total tasks = NPRTRW * NPRTRV**

Grid point norms

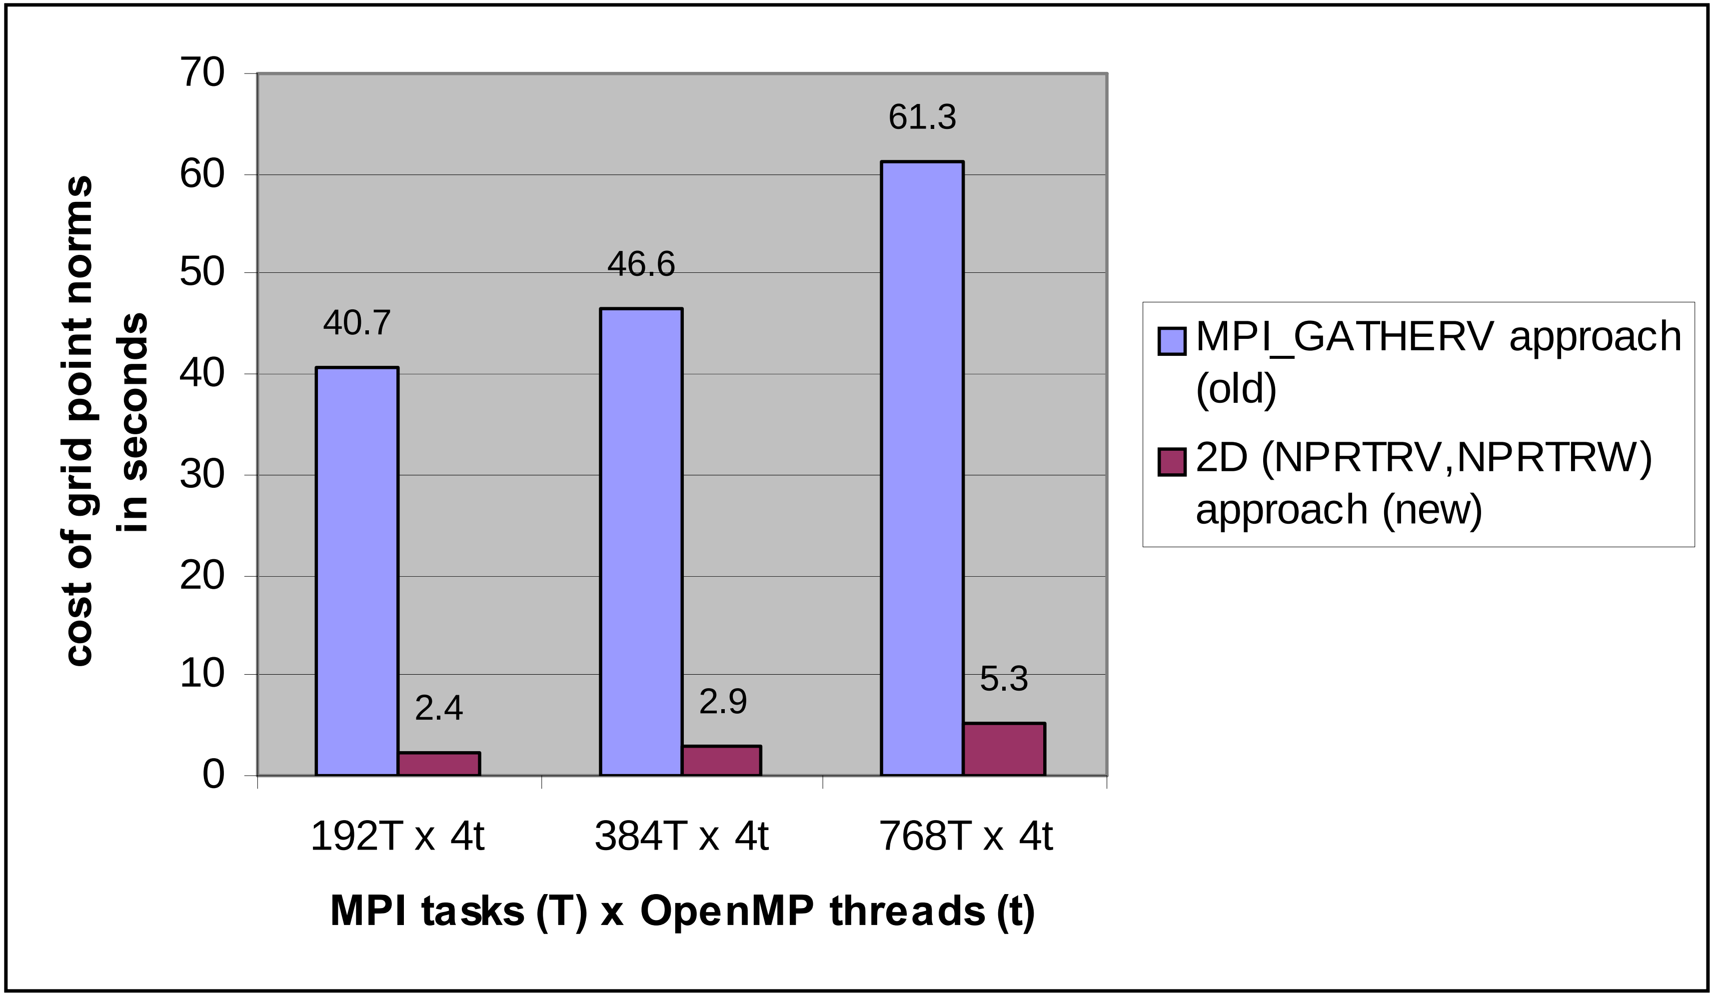

ECMWF

GSTATS analysis, 192 to 384 tasks (counters with lost time > 2 secs)

```
Output file 1 - f0y7.192x4.lbarrier_stats=t
Output file 2 - f0y7.384x4.lbarrier_stats=t
Hoped for speedup factor -  2
Summed time of job 1 4760.08 s
Summed time of job 2 2633.92 s
Achieved speedup factor 1.807
```

Id	Descriptor	Calls	Time1(s)	Time2(s)	Speedup	Lost time(s)
509	MPL SLCOMM1_COMMS	1203	40.3	33.8	1.19	13.7
512	MPL SLCOMM2A_COMMS	2402	23.1	18.2	1.27	6.7
604	MPL GATHERV IN GPNORM1	6019	34.5	44.8	0.77	27.6
616	MPL ALLGATHERV IN MPBCASTSCFLD	1200	25.8	17.5	1.47	4.6
658	MPL SIGCHECK	1201	4.1	5.3	0.77	3.2
694	MPL MPDECOMP BCAST	10	7.2	6.7	1.07	3.1
757	BAR BARRIER IN SLCOMM1	1201	32.1	25.8	1.24	9.7
759	BAR BARRIER IN SLCOMM2	480	60.4	41.5	1.46	11.3
761	BAR BARRIER IN TRGTOL	1243	157.9	93.5	1.69	14.5
763	BAR BARRIER IN TRLTOM	1243	14.8	12.8	1.16	5.4
764	BAR BARRIER IN TRMTOL	1225	7.2	7.5	0.96	3.9
766	BAR BARRIER IN TRSTOM	1200	23.2	17.0	1.36	5.4
803	MPL TRGTOL_COMMS	1243	37.7	26.3	1.43	7.4
805	MPL TRLTOG_COMMS	1225	85.7	62.3	1.38	19.4
806	MPL TRLTOM_COMMS	1243	57.3	30.8	1.86	2.1
807	MPL TRMTOL_COMMS	1225	112.6	64.1	1.76	7.8
810	MPL GATH_SPEC_CONTROL_COMMS	42	16.5	18.9	0.87	10.7
1001	OMP PHYSICS	1201	1680.9	863.4	1.95	22.9
1029	OMP SPCM	1200	59.3	38.0	1.56	8.4
1121	OMP SLCOMM2a 2	1201	10.1	7.9	1.28	2.9
1210	OMP RADINTG-RADLSW	241	709.8	363.4	1.95	8.5
1431	OMP WAMODEL 2	1200	106.4	57.9	1.84	4.8
1645	OMP LTDIR_CTL - DIRECT LEGENDRE	1243	85.8	44.9	1.91	2.0
1710	IO- CNT4 IFLUSHFDB	21	4.0	11.6	0.34	9.6
1771	IO- MPDECOMP I/O	7	5.9	5.3	1.11	2.3
1902	UND SUOYOMB	1	4.4	6.1	0.72	3.9
1905	SER UPDTIM update	11	3.7	4.3	0.86	2.5

Total time lost=253.93 s

Semi-Lagrangian Communications (current)

- **Scaling only 1.2 when doubling to operational task count (192 to 384 tasks)**
- **Each task requires a HALO of neighbouring grid points**
- **HALO width is constant**
  - **max wind speed (400 m/s) * time-step (720 secs @ T799)**
  - **Halo volume > partition volume @ 384 tasks!**
- **Full HALO only required for u, v, w wind vectors**
  - **On-demand scheme for other interpolated variables**
- **Approach has worked well for o(10) to o(100) tasks**
- **Relatively expensive for o(1000) or greater tasks?**

Scaling Semi-Lagrangian communication

- Investigating a ‘HALO-lite’ approach
- Interpolations to determine departure point and mid-point to be computed by the MPI task(s) that ‘own’ these points
- Each MPI task will still have a stencil-width halo for interpolations at the boundary of a task’s partition
- Extra cost to perform reproducible sum (when required) for 4D-VAR adjoint interpolations in ifsmin
  - Running 4D-VAR in reproducible mode costs an extra 10% today mainly due to use of double width halo (see paper in Nov 2000 workshop proceedings)
  - Use of reproducible sum should cost less than double halo approach, it will definitely be more scalable!
- An extra cost due to load imbalance?

Operational v Research T799 forecasts

- Operational forecasts ...
- Use 384 tasks x 4 threads = 1536 ‘user threads’
- Write many more fields to the Fields Data Base (FDB)
- Generate more diagnostics (DDH)
- Write restart files @ 2, 5, and 8 days
- Takes over 10 minutes longer than a typical research experiment running on same number of nodes
- An opportunity for a little optimisation?

Why does writing restart files take so long!

- Takes 300 to 400 secs total (quite variable)
- Writing/reading restart files is now instrumented!
- Each task write 5 files per restart timestep
  - At (end of) forecast day 2, 5, and 8
  - All files written to a single directory
  - Total files = 384 tasks x 5 files x 3 = 5760 files
- At day 5, day 2's files are deleted by each task
  - Using close(file,status='delete')
- At day 8, day 5's files are deleted by each task
- Can we do better?

Why does writing restart files take so long!

- **Starting time 300 – 400 secs**
- **OK, don't delete restart files**
  - **do this later outside of critical time window**
- **Now takes 100 – 200 secs**
- **Combine 5 files per task to 1 file per task**
- **Now takes 58 secs**
- **Now touch all the restart files in a simple perl script before we start execution (perl script takes < 1 sec)**
- **Now takes 26 secs**
- **Lesson: avoid writing 'large' numbers of files in parallel into a single GPFS directory**
- **Fix: touch files before going parallel, don't delete files in parallel**

GSTATS analysis, 192 to 384 tasks (counters with lost time > 2 secs)

```
Output file 1 - f0y7.192x4.lbarrier_stats=t
Output file 2 - f0y7.384x4.lbarrier_stats=t
Hoped for speedup factor -  2
Summed time of job 1 4760.08 s
Summed time of job 2 2633.92 s
Achieved speedup factor 1.807
```

Id	Descriptor	Calls	Time1(s)	Time2(s)	Speedup	Lost time(s)
509	MPL SLCOMM1_COMMS	1203	40.3	33.8	1.19	13.7
512	MPL SLCOMM2A_COMMS	2402	23.1	18.2	1.27	6.7
604	MPL GATHERV IN GPNORM1	6019	34.5	44.8	0.77	27.6
616	MPL ALLGATHERV IN MPBCASTSCFLD	1200	25.8	17.5	1.47	4.6
658	MPL SIGCHECK	1201	4.1	5.3	0.77	3.2
694	MPL MPDECOMP BCAST	10	7.2	6.7	1.07	3.1
757	BAR BARRIER IN SLCOMM1	1201	32.1	25.8	1.24	9.7
759	BAR BARRIER IN SLCOMM2	480	60.4	41.5	1.46	11.3
761	BAR BARRIER IN TRGTOL	1243	157.9	93.5	1.69	14.5
763	BAR BARRIER IN TRLTOM	1243	14.8	12.8	1.16	5.4
764	BAR BARRIER IN TRMTOL	1225	7.2	7.5	0.96	3.9
766	BAR BARRIER IN TRSTOM	1200	23.2	17.0	1.36	5.4
803	MPL TRGTOL_COMMS	1243	37.7	26.3	1.43	7.4
805	MPL TRLTOG_COMMS	1225	85.7	62.3	1.38	19.4
806	MPL TRLTOM_COMMS	1243	57.3	30.8	1.86	2.1
807	MPL TRMTOL_COMMS	1225	112.6	64.1	1.76	7.8
810	MPL GATH_SPEC_CONTROL_COMMS	42	16.5	18.9	0.87	10.7
1001	OMP PHYSICS	1201	1680.9	863.4	1.95	22.9
1029	OMP SPCM	1200	59.3	38.0	1.56	8.4
1121	OMP SLCOMM2a 2	1201	10.1	7.9	1.28	2.9
1210	OMP RADINTG-RADLSW	241	709.8	363.4	1.95	8.5
1431	OMP WAMODEL 2	1200	106.4	57.9	1.84	4.8
1645	OMP LTDIR_CTL - DIRECT LEGENDRE	1243	85.8	44.9	1.91	2.0
1710	IO- CNT4 IFLUSHFDB	21	4.0	11.6	0.34	9.6
1771	IO- MPDECOMP I/O	7	5.9	5.3	1.11	2.3
1902	UND SUOYOMB	1	4.4	6.1	0.72	3.9
1905	SER UPDTIM update	11	3.7	4.3	0.86	2.5

Total time lost=253.93 s

Wave model (WAM) / IFS communications

- **Wave model and IFS model have different grids**
  - Each time step, IFS sends 5 fields to WAM and WAM returns 1 field to IFS
- **Initial implementation**
  - neither requires knowledge of the other's grid
  - used MPI_ALLGATHERV
  - every task received a full field and took the bits for their task
- **IFS model input to Wave model was 'optimised'**
  - MPI_ALLGATHERV only on the first call to the wave model
  - thereafter only the exact data needed is sent using MPI_ISEND, MPI_RECV, MPI_WAITALL
- **Wave model output to atmos model**
  - continued to use MPI_ALLGATHERV on all calls

IFS / WAM communications

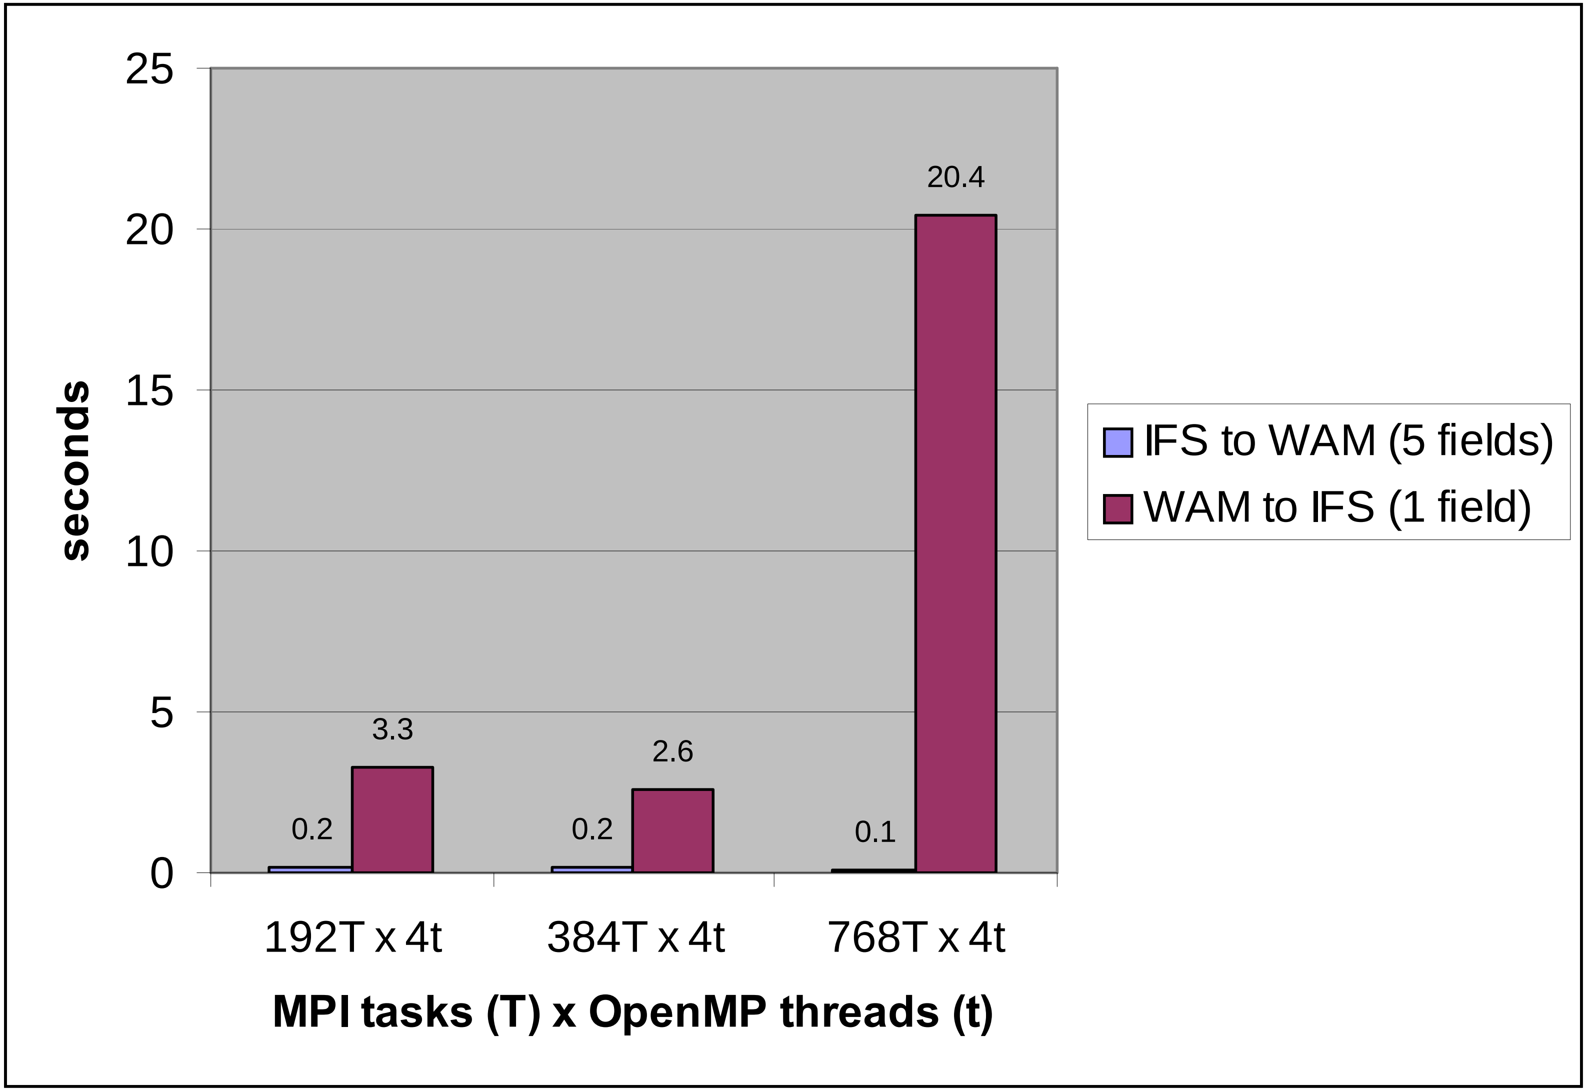

Improved WAM to IFS communications (J. Bidlot)

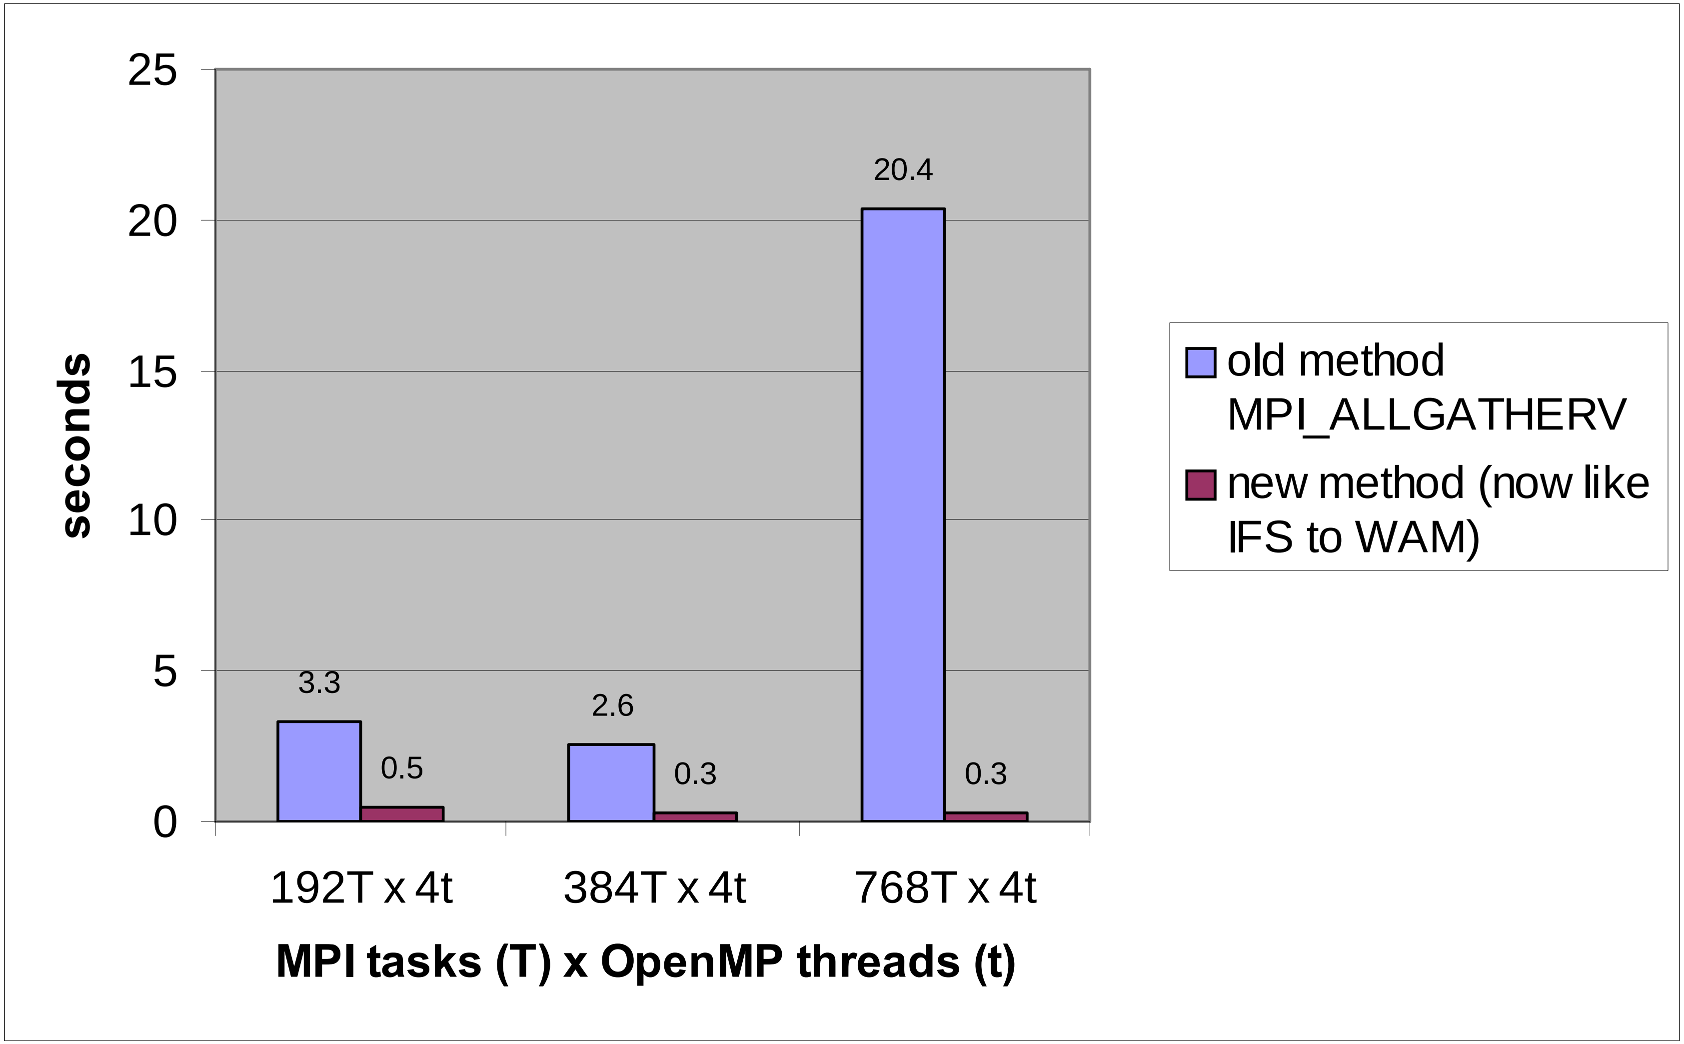

ECMWF

GSTATS analysis, 192 to 384 tasks (counters with lost time > 2 secs)

Output file 1 - f0y7.192x4.lbarrier_stats=t
Output file 2 - f0y7.384x4.lbarrier_stats=t
Hoped for speedup factor - 2
Summed time of job 1 4760.08 s
Summed time of job 2 2633.92 s
Achieved speedup factor 1.807

Id	Descriptor	Calls	Time1(s)	Time2(s)	Speedup	Lost time(s)
509	MPL SLCOMM1_COMMS	1203	40.3	33.8	1.19	13.7
512	MPL SLCOMM2A_COMMS	2402	23.1	18.2	1.27	6.7
604	MPL GATHERV IN GPNORM1	6019	34.5	44.8	0.77	27.6
616	MPL ALLGATHERV IN MPBCASTSCFLD	1200	25.8	17.5	1.47	4.6
658	MPL SIGCHECK	1201	4.1	5.3	0.77	3.2
694	MPL MPDECOMP BCAST	10	7.2	6.7	1.07	3.1
757	BAR BARRIER IN SLCOMM1	1201	32.1	25.8	1.24	9.7
759	BAR BARRIER IN SLCOMM2	480	60.4	41.5	1.46	11.3
761	BAR BARRIER IN TRGTOL	1243	157.9	93.5	1.69	14.5
763	BAR BARRIER IN TRLTOM	1243	14.8	12.8	1.16	5.4
764	BAR BARRIER IN TRMTOL	1225	7.2	7.5	0.96	3.9
766	BAR BARRIER IN TRSTOM	1200	23.2	17.0	1.36	5.4
803	MPL TRGTOL_COMMS	1243	37.7	26.3	1.43	7.4
805	MPL TRLTOG_COMMS	1225	85.7	62.3	1.38	19.4
806	MPL TRLTOM_COMMS	1243	57.3	30.8	1.86	2.1
807	MPL TRMTOL_COMMS	1225	112.6	64.1	1.76	7.8
810	MPL GATH_SPEC_CONTROL_COMMS	42	16.5	18.9	0.87	10.7
1001	OMP PHYSICS	1201	1680.9	863.4	1.95	22.9
1029	OMP SPCM	1200	59.3	38.0	1.56	8.4
1121	OMP SLCOMM2a 2	1201	10.1	7.9	1.28	2.9
1210	OMP RADINTG-RADLSW	241	709.8	363.4	1.95	8.5
1431	OMP WAMODEL 2	1200	106.4	57.9	1.84	4.8
1645	OMP LTDIR_CTL - DIRECT LEGENDRE	1243	85.8	44.9	1.91	2.0
1710	IO- CNT4 IFLUSHFDB	21	4.0	11.6	0.34	9.6
1771	IO- MPDECOMP I/O	7	5.9	5.3	1.11	2.3
1902	UND SUOYOMB	1	4.4	6.1	0.72	3.9
1905	SER UPDTIM update	11	3.7	4.3	0.86	2.5

Total time lost=253.93 s

Load imbalance

- **Static imbalance**
  - Distribution of grid points, spectral waves, atmospheric levels are never perfectly distributed at large numbers of tasks/threads
  - Will we have more cores than grid points?
  - The first minimisation step of the current operational 4D-Var uses a T_L95 spectral resolution with 13280 grid points
  - OpenMP dynamic scheduling has no advantage when there is only one unit of work per thread, we ideally need 100’s of such units
- **Dynamic imbalance**
  - IFS physics
  - Wave Model

IFS physics computational imbalance (T799L91, 384 tasks)

~11% imbalance in physics, ~5% imbalance (total)

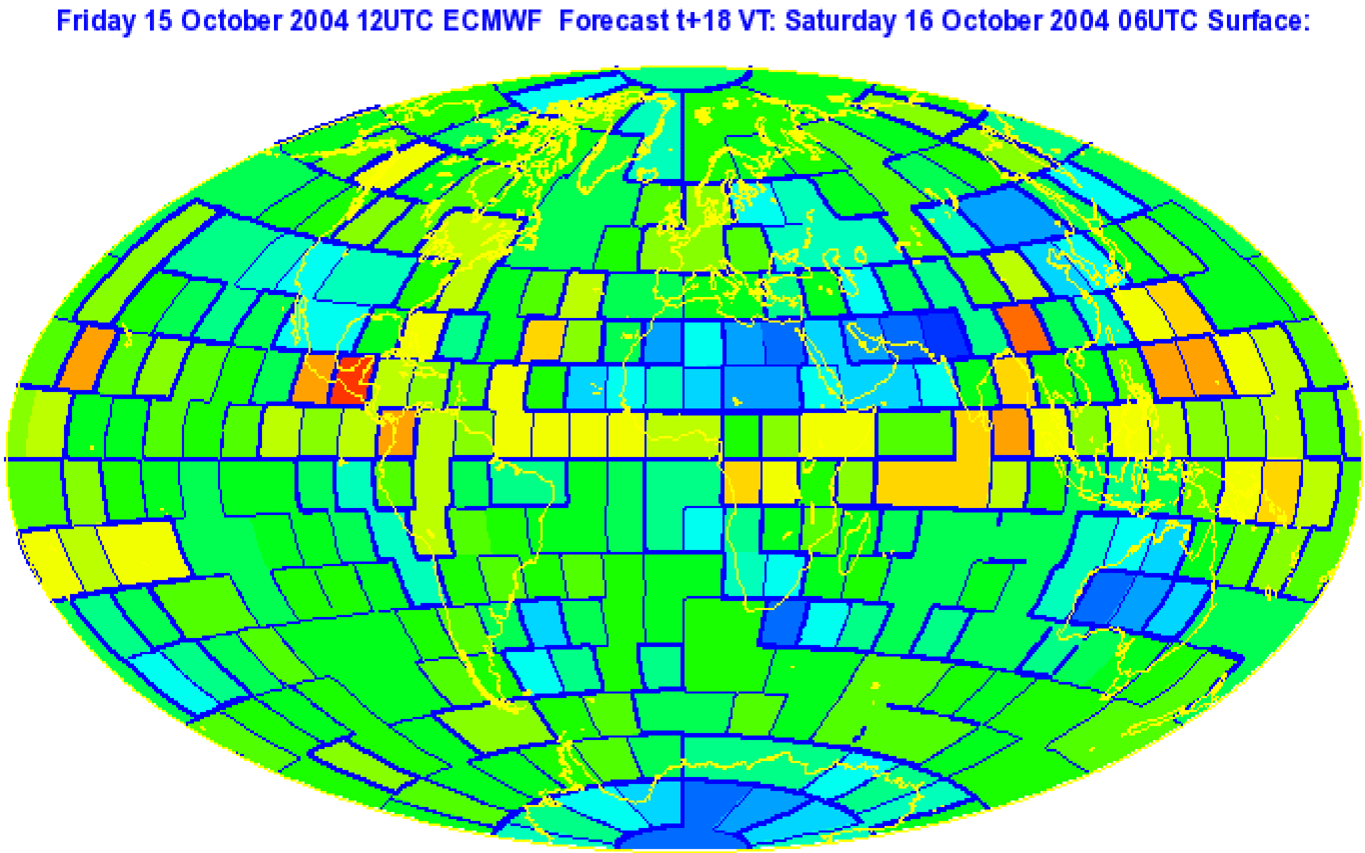

ECMWF

384 tasks (physics cost, 6 hour sampling, 48 hrs)

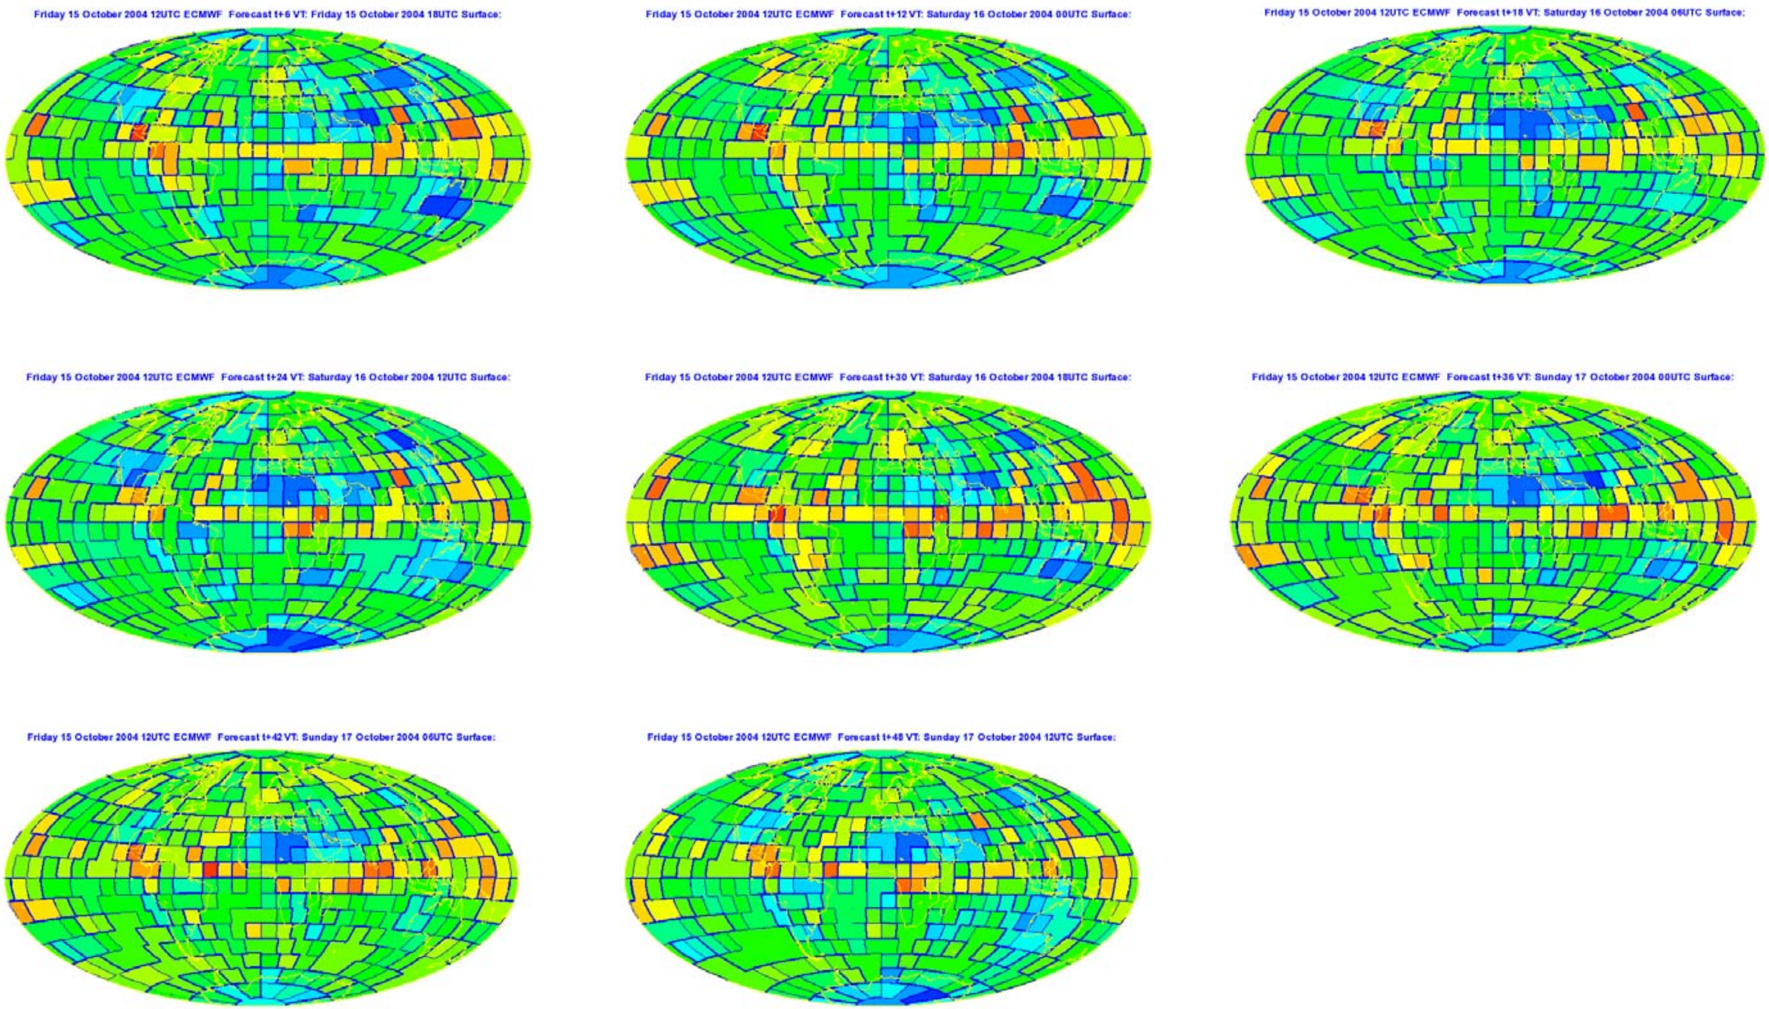

IFS physics load balancing (work in progress)

- Redistribute grid columns between tasks
- Redistribute before physics, restore after physics
- Redistribution in task groups (n tasks per group)
  - Separate MPI communicators for performance
  - Grouping possibilities
    - East West (intra node?)
    - North South
    - Regional mixing
    - Random mixing

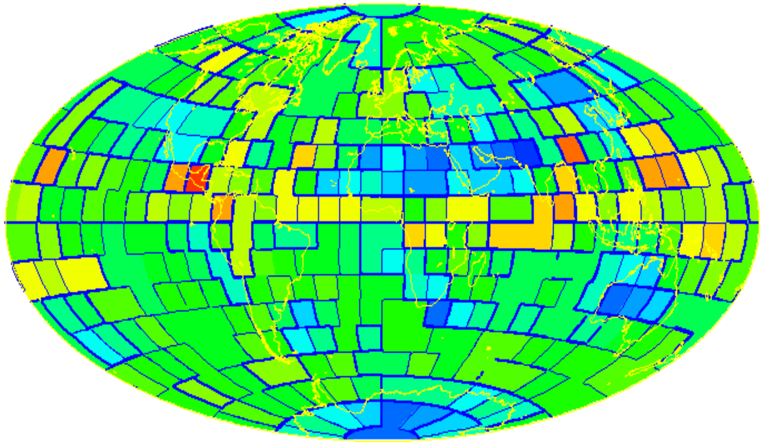

- Combine with existing coarse physics routine (w/o interp.)
- Overheads are high!
  - In (P→Z) / out (Z→P), comms buffer pack/unpack, comms, memory

IFS physics load balancing (next?)

- **Instead of redistributing all grid columns we load balance grid columns based on cost of earlier time step**
  - **Most costly tasks send some columns to least costly tasks**
- **Determine cost every n time steps**
  - **optimal n to be found by experimentation**
- **Predicted improvement in physics ~ 6%**
- **Predicted overall improvement ~ 2.5%**
- **Additional code before/after calling physics is hard to hide!**
- **IFS physics load balancing worth the effort?**

When the work is done ...

- **How much more scalable will the IFS model be when all presented today is implemented?**
- **What metric?**
  - **Propose calculated efficiency at operational task count (slide 5)**
  - **10 day forecast in less than 1 hour**
  - **Today T799L91 model uses 48 Power5 nodes**
  - **384 tasks x 4 threads**
- **IFS model was 79.4 % efficient (at the start of this work)**
- **The new efficiency will be ...**
  - **Reported at the next RAPS meeting**
  - **By then a T1279L91 model on Power6**