METEO FRANCE
Toujours un temps d'avance

MOCAGE implementation on ECMWF HPCD

Parallelisation Effort

Philippe Moinat / Vincent-Henri Peuch
and Hervé le Berre (Cerfacs)

METEO FRANCE
Toujours un temps d'avance

Summary of MOCAGE features

- **Off-line CTM. Meteo forcing by files. Up to 4 nested domains** (Global →Zoom)**. 47** (top at 5hPa) **or 60** (top at 0.02hPa) **levels.**
- **RACMOBUS chemical model** (118 species, troposhere + stratosphere, 350 reactions). Δt=15 min
- **Semi-Lagrangian Semi-Implicit advection.** Δt=30 min
- **Bechtold convection + Louis vertical diffusion**

METEO FRANCE
Toujours un temps d'avance

Initial status of MOCAGE

- **Highly vectorised, not parallelised.**
- **Operational platform: Fujitsu VPP5000 (daily forecast of the chemical weather)**

- **Memory: 2.4 Gb** (2°, global grid, 60 levels, 118 species)

METEO FRANCE
Toujours un temps d'avance

Porting of MOCAGE on the ECMWF IBM Cluster

On 1 proc.:

- **Compilation options (Makefile)**
- **Launching script (batch job)**

↳ **Very poor performance: 18 hours elapsed for a 24 hour simulation !**

(2°, global grid, 60 levels, 118 species)

➔ **Parallelisation required !!!**

METEO FRANCE
Toujours un temps d'avance

OpenMP parallelisation of MOCAGE

- **Implemented for the loops of 4 routines of the chemical model (90% of the CPU)**
- **Binding of threads to processors for optimal performance (few MPI instructions).**

METEO FRANCE
Toujours un temps d'avance

OpenMP parallelisation of MOCAGE / Results

- **Good scalability (parallelised routines)**
- **Overall speedup of 6 on 16 procs.**

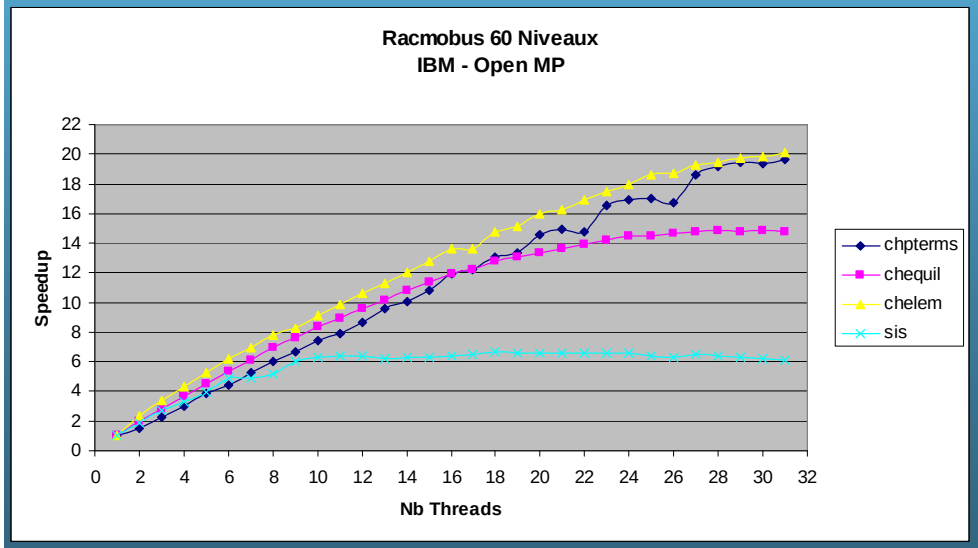

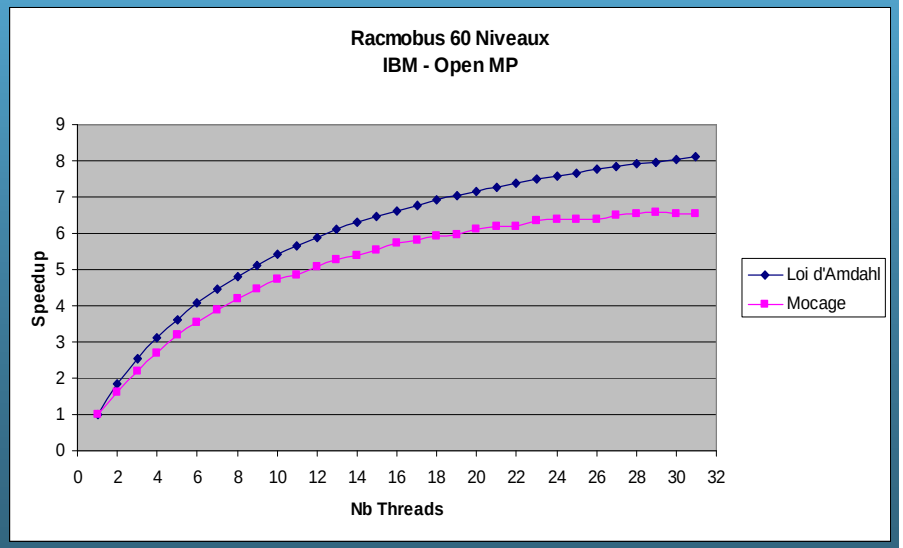

METEO FRANCE
Toujours un temps d'avance

MPI parallelisation of MOCAGE

- **Light approach (no heavy developments)**
- **Applied only to the chemical core**
- **Splitting of each vertical level**
- **Communications minimised (MPI_ALLREDUCE at the end of the parallel sequence).**
- **Drawback: whole memory duplicated on each proc.**

METEO FRANCE
Toujours un temps d'avance

MPI parallelisation of MOCAGE

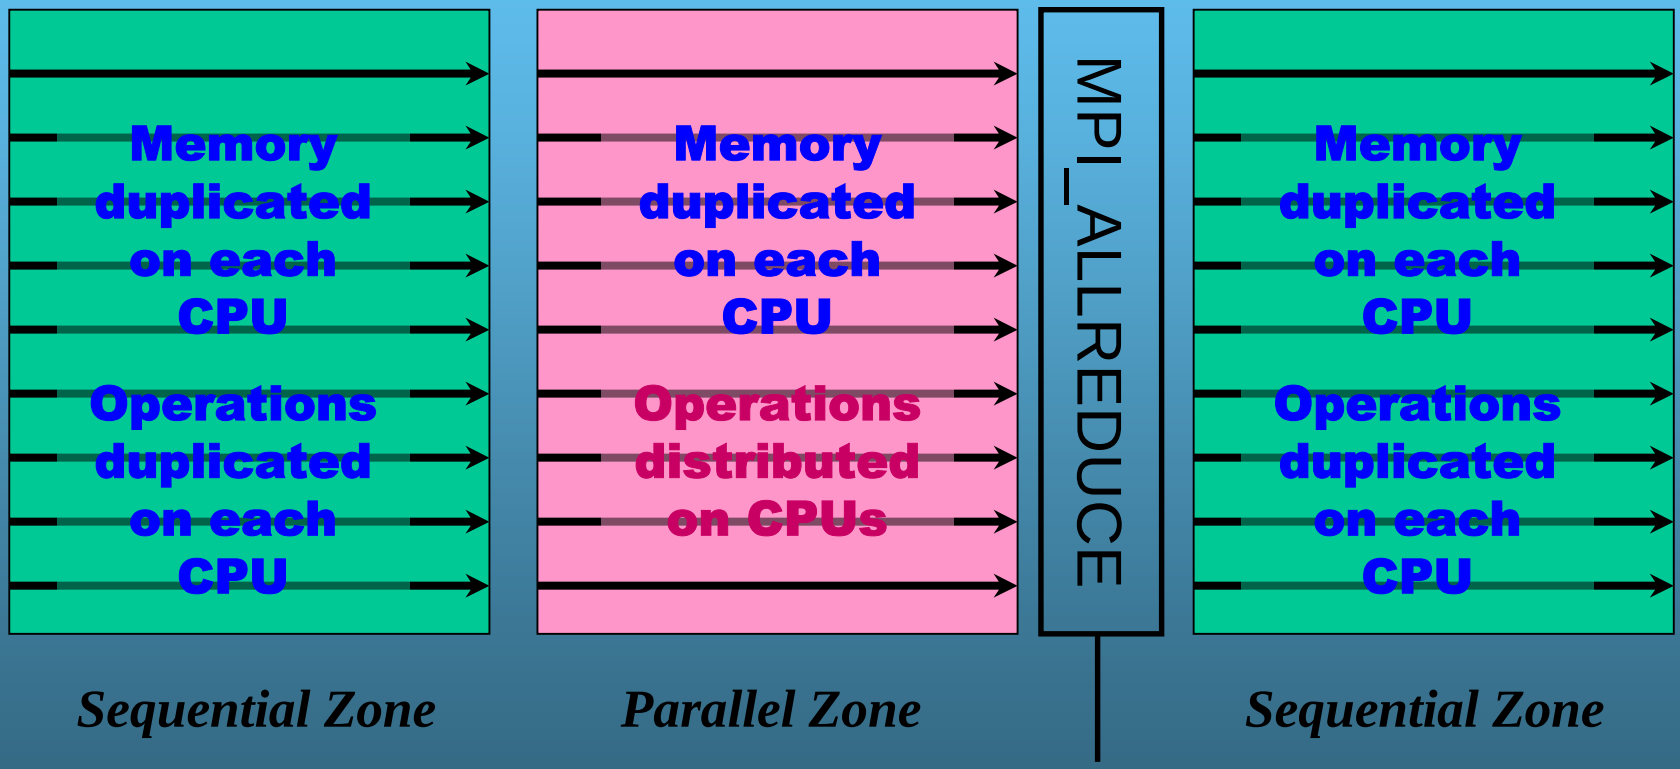

METEO FRANCE
Toujours un temps d'avance

MPI parallelisation of MOCAGE / Results

- **Overall speedup of 6 on 12 procs.**

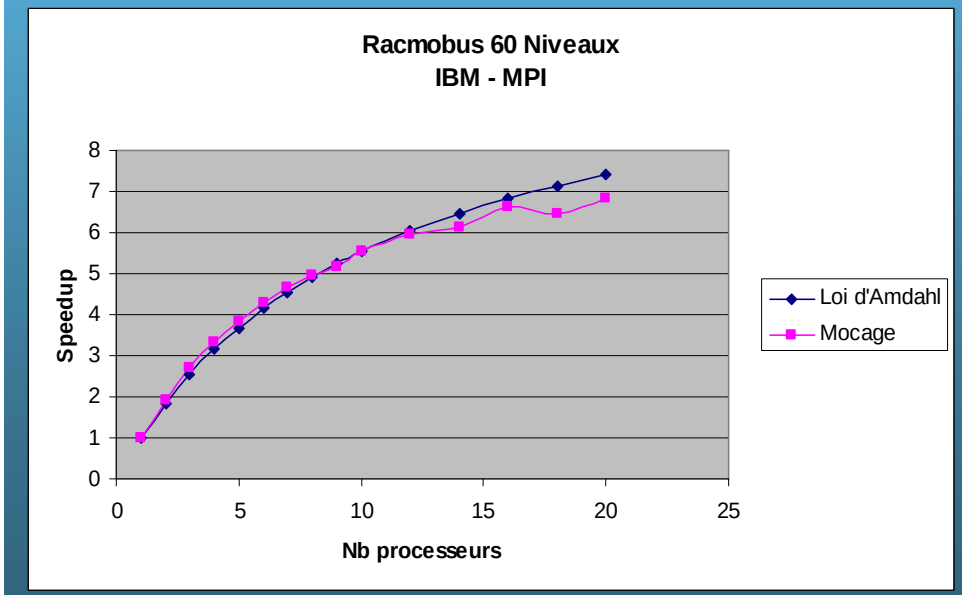

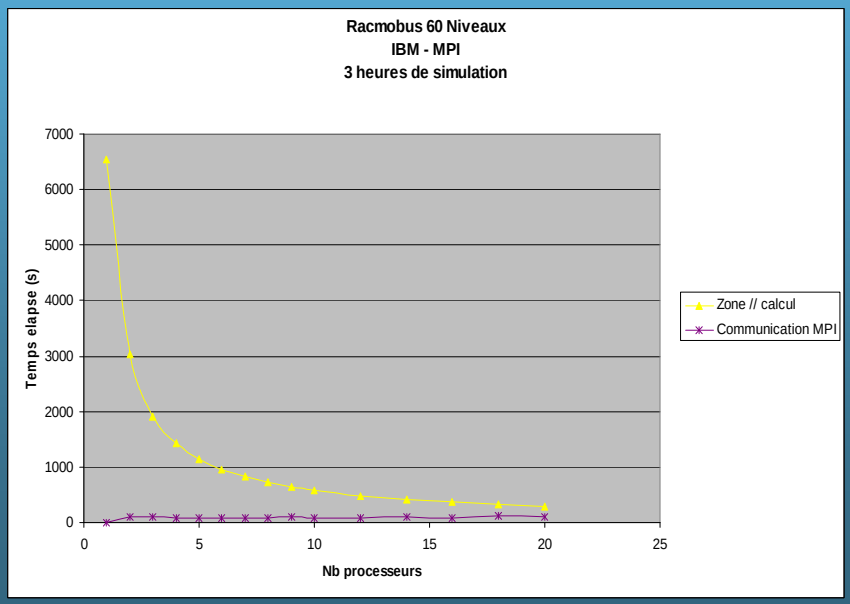

METEO FRANCE
Toujours un temps d'avance

Mixed OpenMP / MPI parallelisation of MOCAGE

- **MPI more CPU efficient but more memory consuming**
- **Optimal configuration depends on available resources (cf. coupled IFS + MOCAGE runs)**

Mocage Racmobus 60 Niveaux IBM - 3 heures de simulation		
MPI	**OMP**	**Temps Elapse (s)**
1	16	1459
2	8	1822
4	4	1265
8	2	1284
16	1	1273

(2°, global grid, 60 levels, 118 species)

METEO FRANCE
Toujours un temps d'avance

Summary

A first parallelisation effort based on OpenMP and/or MPI results in a significant improvement of MOCAGE performance on the ECMWF HPCD

- A 24 hour simulation can now be run in ~2h40min on 16 procs (speedup = 6.5)
- 1 month ~ 4 days, 1 year ~ 48 days

(2° global grid, 60 levels, 118 species, 350 reactions)

METEO FRANCE
Toujours un temps d'avance

Available options for reducing the CPU

- **RACMOBUS** (RACM + REPROBUS, 118 species, 350 reactions) **replaced by RELACS** (ReLACS + REPROBUS, 85 species, 240 reactions)
- **2°×2° replaced by T42** (nb of points / 2)
- **Factorization of the chemical computations in the PBL**

METEO FRANCE
Toujours un temps d'avance

Available options for reducing the CPU: results with the "climate" version of MOCAGE

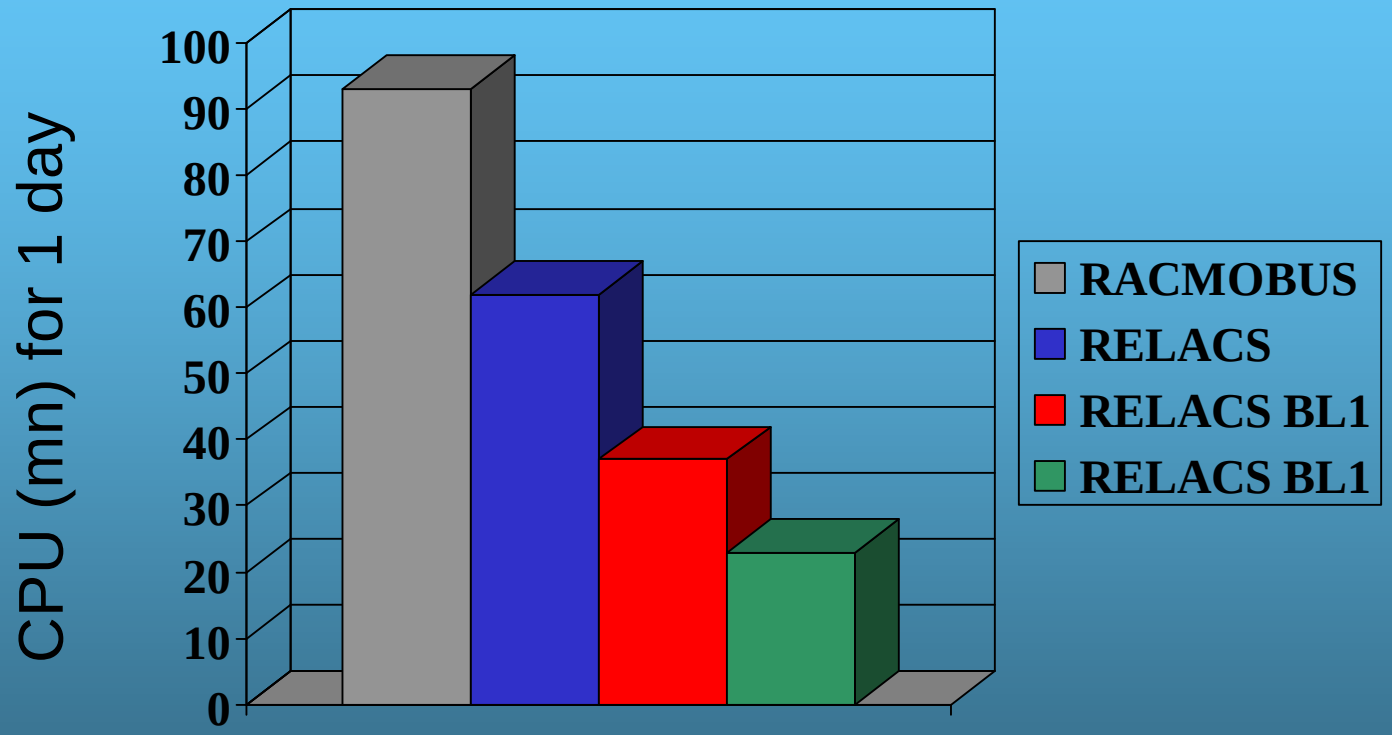

Standard version
2°x2° 47 niv.

Climate version
T42 60 niv.

METEO FRANCE
Toujours un temps d'avance

Conclusion

- **MOCAGE is now 90 % parallel**
- **Using the most demanding options, a one day run can be performed in 2h40 min**
- **Options for reducing the CPU exist and are to be further tested. A 0.25 reduction factor could be reached (1 day in 40 min).**
- **Best configuration to be decided: best compromise btw accuracy and efficiency according to the objectives.**