

COUPLED PROBLEMS 2011
Coupled Solution Strategies II

Kos (Greece), June 20th-22nd 2011

**CERFACS-ONERA Open-PALM:
an open source dynamic parallel coupler**

Andrea Piacentini

and the PALM team: Th. Morel, F. Duchaine, A. Thévenin

CERFACS - Toulouse (France)

www.cerfacs.fr/globc/PALM_WEB

- × Rationale and Genesis
- × Design
- × Implementation
- × Technical challenges
- × Some applications
- × Discussions

Dynamic coupling = a coupling where the components execution scheduling and the data exchange patterns cannot be entirely defined before execution

Dynamic coupling = a coupling where the component execution scheduling and the data exchange patterns cannot be entirely defined before execution

The PALM rationale:

Historical

Implementation of data assimilation suites: dynamic aspects of data assimilation algorithms themselves

Current

Multi-physics coupling in flexible configurations

Future

MPP, high performances, self-tuning, resilient applications

Data assimilation is a technique aiming to improve numerical models skills by the use of observational data.

It is based on computationally expensive algebraic algorithms involving the model, the observation treatments, the statistical characterisation of the errors on both sides.

In 1996, the MERCATOR operational oceanography project faced the problem to set-up a new operational suite with Data Assimilation for Research and Operations in an evolving configuration.

Instead of hard-coding data assimilation routines in the model, or vice-versa, they decided to **couple** **model** + **observations handling** + **error statistics** + **algebra** in a flexible and computationally effective way

Some data assimilation algorithms are based on an iterative minimization.

This implies the repeated execution of the tasks. The total number of iterations is not known beforehand.

In some configurations some tasks are activated only if some observations are available at run-time.

This implies the conditional execution of some tasks.

⇒ DYNAMIC COUPLING

- Process management

- Buffered communications

- Object versioning

DYNAMIC COUPLING

Process management

- Starts and synchronizes the tasks
- Handles algorithms (DO and WHILE loops, IF and CASE switches)

DYNAMIC COUPLING

Process management

- Starts and synchronizes the tasks
- Handles algorithms (DO and WHILE loops, IF and CASE switches)

Buffered communications

- Non blocking on the production side
- Order of production and reception is not relevant
- Cumulated objects (linear combinations on the fly)
- Permanently stored objects for repeated receptions

DYNAMIC COUPLING

Process management

- Starts and synchronizes the tasks
- Handles algorithms (DO and WHILE loops, IF and CASE switches)

Buffered communications

- Non blocking on the production side
- Order of production and reception is not relevant
- Cumulated objects (linear combinations on the fly)
- Permanently stored objects for repeated receptions

Object versioning

- Last In Only Out paradigm
- Coherence of the components produced by a parallel task with loose synchronization

Model integration and some data assimilation tasks can run simultaneously.

This requires the handling of concurrent tasks parallelism. Like in other couplers it is just the first level of parallelism.

Models and/or assimilation codes are themselves parallel.

This requires the handling of the execution of parallel codes and the management of their data exchanges, including the remapping between codes with different distributions. It is the second level of inner parallelism.

Data assimilation algorithms require linear algebra (operations on very large vectors and matrices, usually sparse, and effective minimisations).

This suggested to include in PALM an algebra toolbox interfacing the most effective linear algebra libraries in the form of pre-defined generic "entities" (*units*) that can be coupled with other codes.

One of the aims of coupling is the reuse of legacy codes.

To make it simple we had to grant (reasonably) minimal intrusiveness.

Three main assumptions:

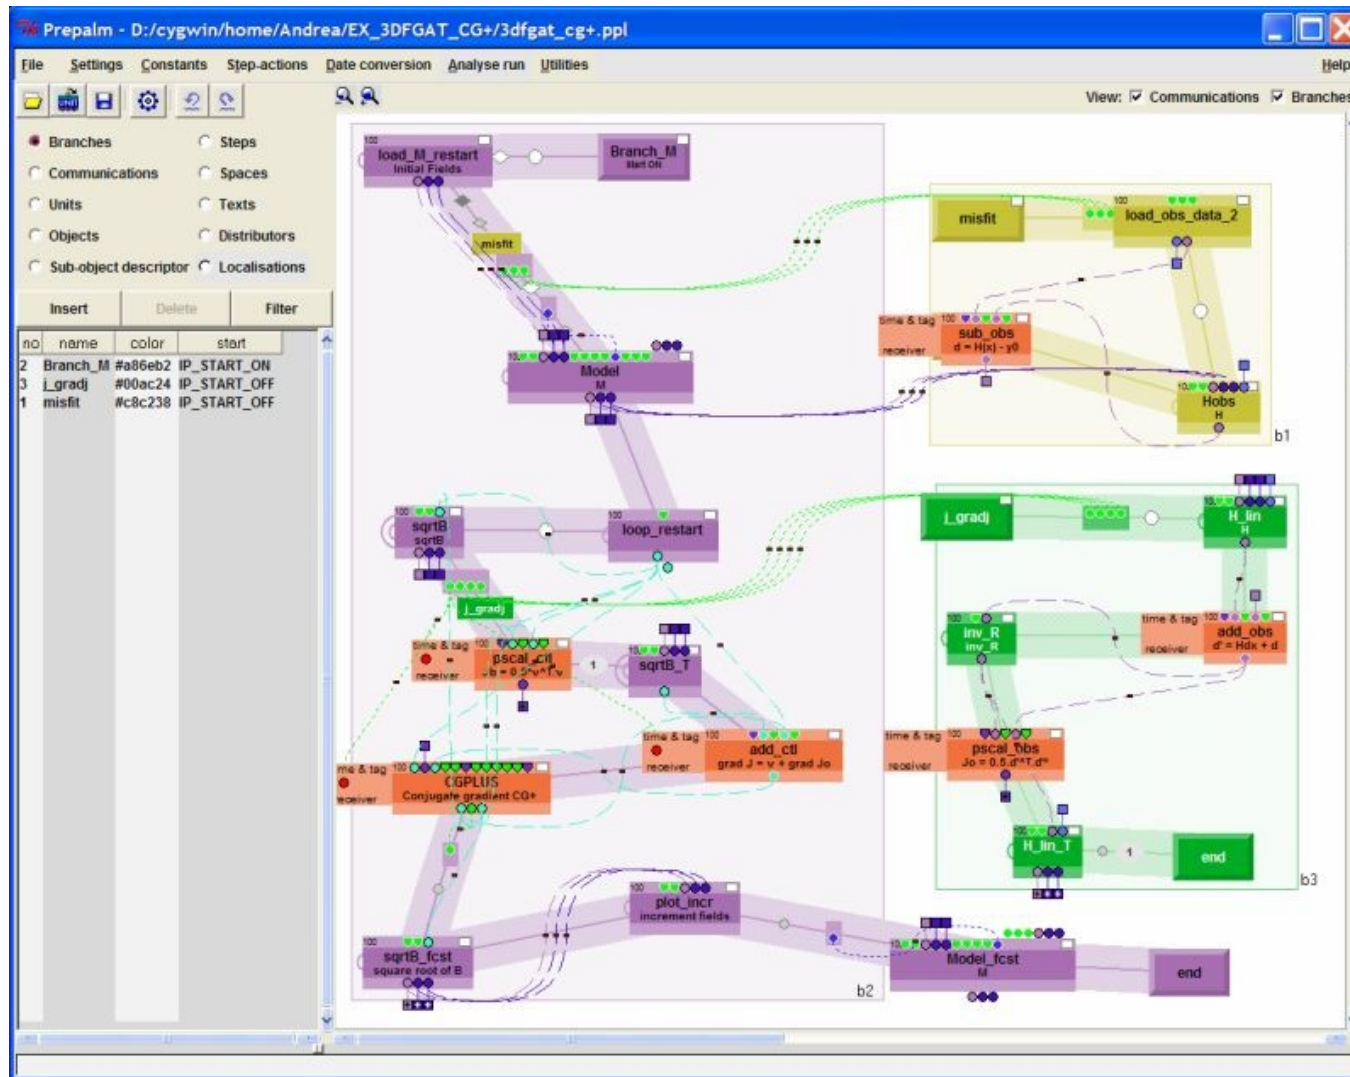
1) The "end point" communications paradigm: the producer of an object does not know anything about the recipients (if any) and the other way round. The coupler makes the matching.

One of the aims of coupling is the reuse of legacy codes.

To make it simple we had to grant (reasonably) minimal intrusiveness.

Three main assumptions:

- 1) The "end point" communications paradigm: the producer of an object does not know anything about the recipients (if any) and the other way round. The coupler makes the matching.
- 2) A reduced set of APIs complemented by a very detailed Graphic User Interface.



One of the aims of coupling is the reuse of legacy codes.

To make it simple we had to grant (reasonably) minimal intrusiveness.

Three main assumptions:

1) The "end point" communications paradigm: the producer of an object does not know anything about the recipients (if any) and the other way round. The coupler makes the matching.

2) A reduced set of APIs complemented by a very detailed Graphic User Interface.

3) Multi language APIs for the most common compiled and interpreted languages used for geophysics modelling

- F77, F90, C, C++ bindings
- SWIG interfaces for Python, Perl, Java, Tcl/Tk, Octave (Matlab), ...
- API's for precompiled codes via dynamic libraries

The same tool for

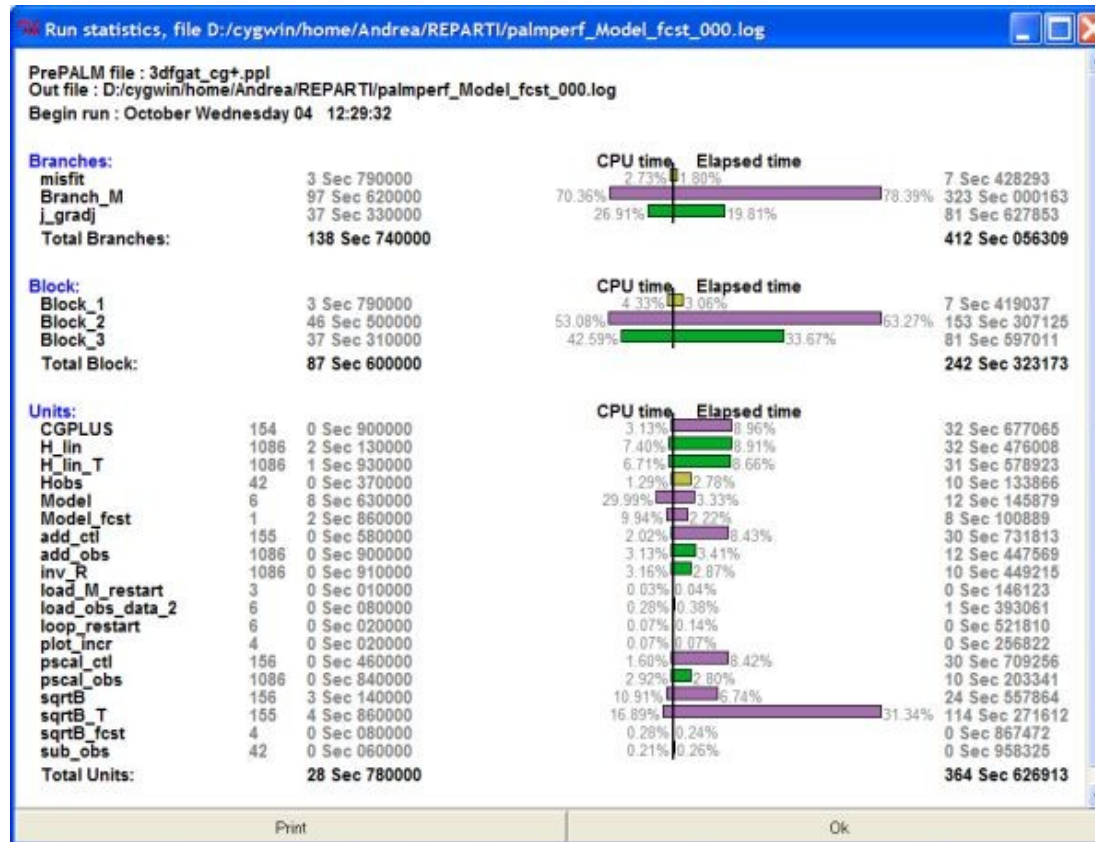
Operational Use

- Performances
- Run-time monitoring

The same tool for

Operational Use

- Performances
- Run-time monitoring
- Post mortem performances analyses



The same tool for

Operational Use

- Performances
- Run-time monitoring
- Post mortem performances analyses

Research Use

- User friendliness
- Portability (\Rightarrow standard solutions)

The same tool for

Operational Use

- Performances **MPI based**
- Run-time monitoring
- Post mortem performances analyses

Research Use

- User friendliness
- Portability (\Rightarrow standard solutions) **MPI based**

PALM has been adopted for different applications dealing with flexible coupling configurations.

A typical example are the shape optimization applications in computational fluid dynamics, coupling heavy [proprietor] simulation codes with algebraic iterative optimization algorithms.

For data assimilation there is no need of generic grid-to-grid interpolation tools (at most it is user provided code).

Multi-physics applications (including climate modelling), using independently developed components, have arouse the need of an efficient grid-to-grid parallel interpolation tool in PALM.

Last generation models with automatic load balancing, data redistribution and self-tuning capabilities are defining the next challenges for the use of couplers on MPP machines.

We focus on the [successive] implementation choices for the PALM driver and libraries.

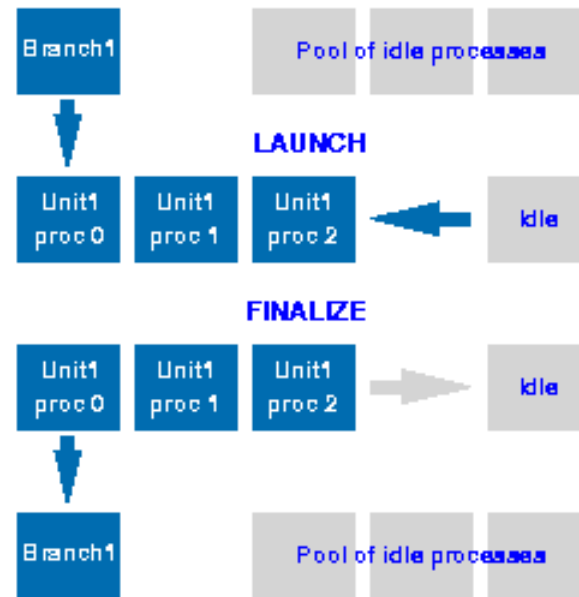
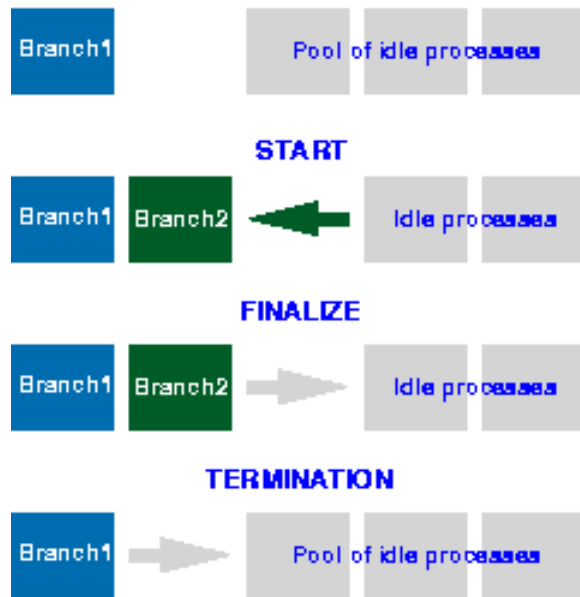
The GUI is coded in Tcl/Tk and would deserve a separate speech.

In 1997 lack of complete and robust MPI2 implementations:

⇒ MPI1 emulation based on a pool of idle processes.

PALM_SP (a.k.a. PALM_RESEARCH).

Implementation



In 1997 lack of complete and robust MPI2 implementations:

⇒ MPI1 emulation based on a pool of idle processes.

PALM_SP (a.k.a. PALM_RESEARCH).

Intended for functional tests, it is still used for some applications (e.g. N.R.T. air-quality forecasts in the EU funded MACC project)

Some interesting features that could be recovered in the next Open-PALM versions.

In 2003 release of the first version of a full MPMD (Multiple Programs Multiple Data) MPI2 based coupler:

PALM MP

Dynamic process management via `MPI_Comm_Spawn` + a scheduler.

Option to merge into a single executable (a block) the coupled components that are started in a sequence.

End-point high bandwidth communication scheme, with the driver acting as a broker (useful for dynamic coupling and for monitoring).

Since then some achieved and some under development enhancements:

The possibility to interface commercial black-box codes (such as Fluent, Abaqus MSC/MARC) by the use of external dynamic libraries and/or a socket based layer

Since then some achieved and some under development enhancements:

The possibility to interface commercial black-box codes (such as Fluent, Abaqus MSC/MARC) by the use of external dynamic libraries and/or a socket based layer

The implementation of a simplified working mode entirely compliant with the MPI-1 library (MPMD mpirun extension)

Since then some achieved and some under development enhancements:

The possibility to interface commercial black-box codes (such as Fluent, Abaqus MSC/MARC) by the use of external dynamic libraries and/or a socket based layer

The implementation of a simplified working mode entirely compliant with the MPI-1 library (MPMD mpirun extension)

The optimisation of repeated well synchronised communications that don't require the intervention of the driver

Since then some achieved and some under development enhancements:

The possibility to interface commercial black-box codes (such as Fluent, Abaqus MSC/MARC) by the use of external dynamic libraries and/or a socket based layer

The implementation of a simplified working mode entirely compliant with the MPI-1 library (MPMD mpirun extension)

The optimisation of repeated well synchronised communications that don't require the intervention of the driver

The enhancement of the parallel algebra toolbox that now includes the CWIPI interpolation library from ONERA for the grid to grid remapping

Since then some achieved and some under development enhancements:

Starting from January 2011 PALM has become open source with the name Open-PALM.

It is the most suitable environment to accept collaborations and contributions on the coupler development.

Difficult trade-off between a centralised and a fully distributed approach.

Process management is a key issue for dynamic coupling, but it implies some extra constraints. We've already implemented the MPMD MPI1 extension (mpirun with >1 executables). How to deal with automatic load balancing? Separate phases of reorganization (higher overhead) and simulation (lower overhead) and/or resurrect the link of >1 codes in a single executable (*cf.* PALM_SP).

Difficult trade-off between a centralised and a fully distributed approach.

Process management is a key issue for dynamic coupling, but it implies some extra constraints. We've already implemented the MPMD MPI1 extension (mpirun with >1 executables). How to deal with automatic load balancing? Separate phases of reorganization (higher overhead) and simulation (lower overhead) and/or resurrect the link of >1 codes in a single executable (*cf.* PALM_SP).

Very effective communications on MPP configurations. Trade-off between flexibility and monitoring on one side and performances on the other.

Difficult trade-off between a centralised and a fully distributed approach.

Process management is a key issue for dynamic coupling, but it implies some extra constraints. We've already implemented the MPMD MPI1 extension (mpirun with >1 executables). How to deal with automatic load balancing? Separate phases of reorganization (higher overhead) and simulation (lower overhead) and/or resurrect the link of >1 codes in a single executable (*cf.* PALM_SP).

Very effective communications on MPP configurations. Trade-off between flexibility and monitoring on one side and performances on the other.

Role of the coupler code itself: make it parallel or let most of the tasks to the units and/or the system?

Integration of the communication and of the interpolation layers: CWIPI developed by ONERA. Exchange of fields defined on any kind of non-structured mesh. Surface and volume non conservative interpolations (clouds) + callback of user defined interpolations.

Based on robust industrial MPI wrappers (by EDF).

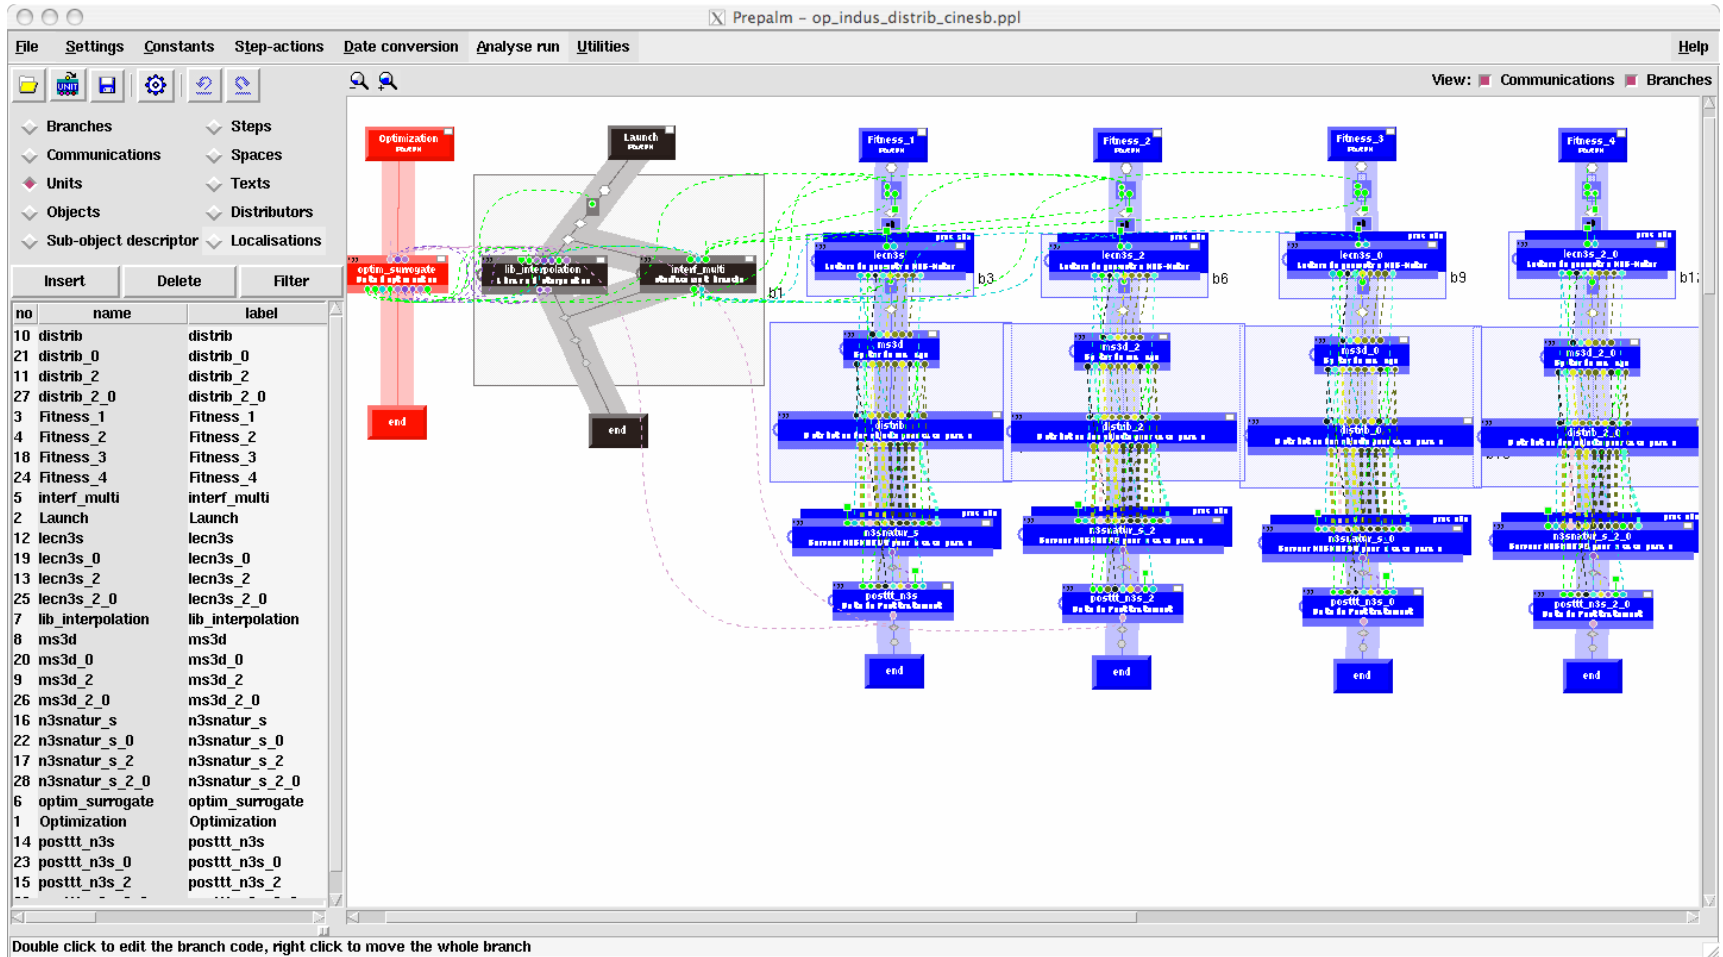
Very good scaling tested up to 256 processors. Still under investigation for MPP figures.

Basic bricks to implement conservative interpolations, but still to do and test.

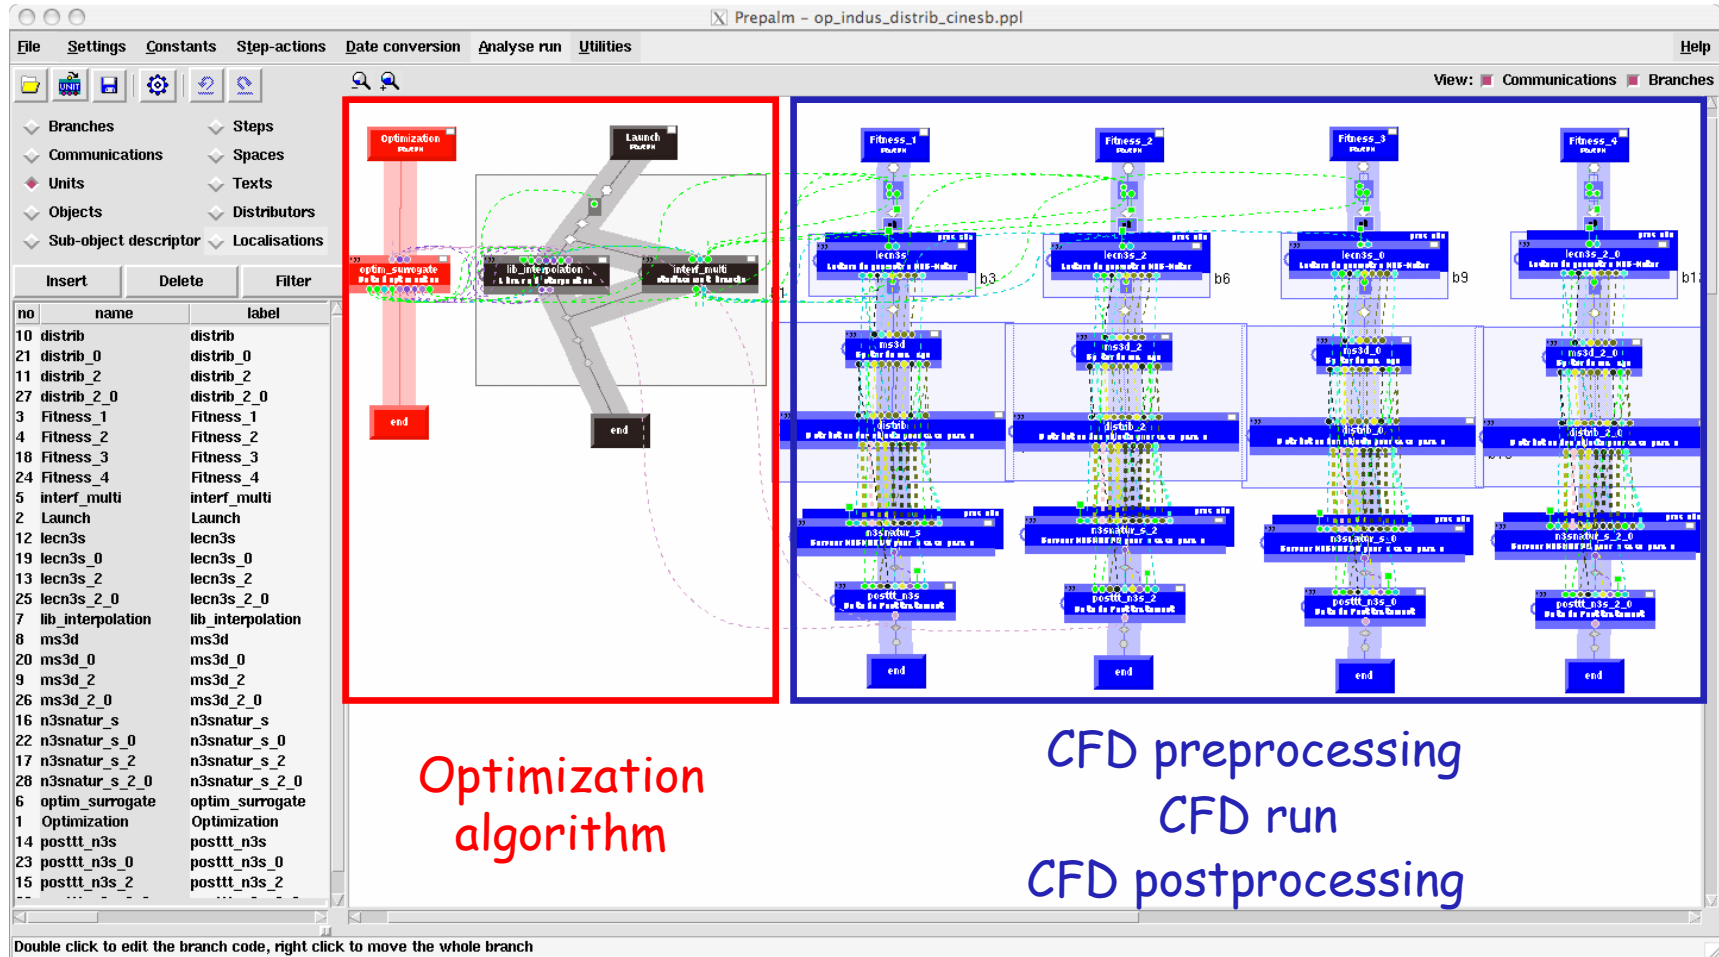
Projects of assessment of the CWIPI (and its evolutions) layer for climate couplings.

Some application canvases to stress the importance of the
aforementioned features

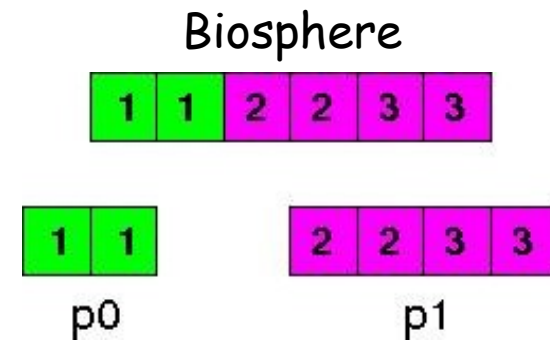
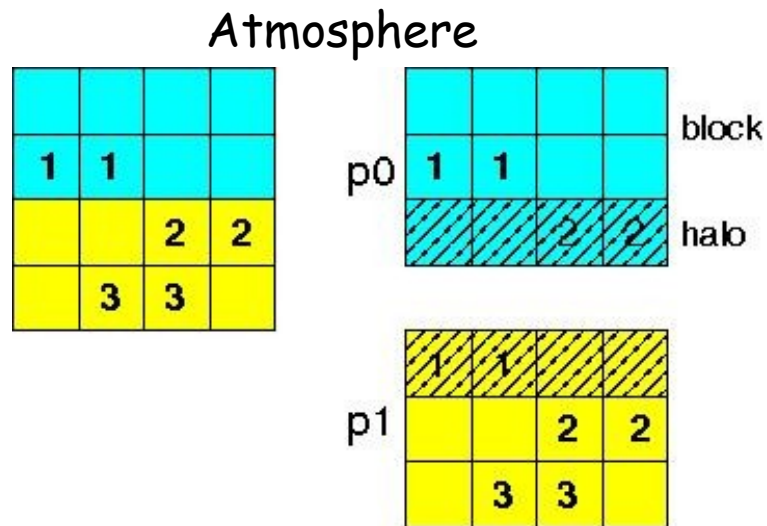
Dynamic coupling for a combustor cooling system optimization



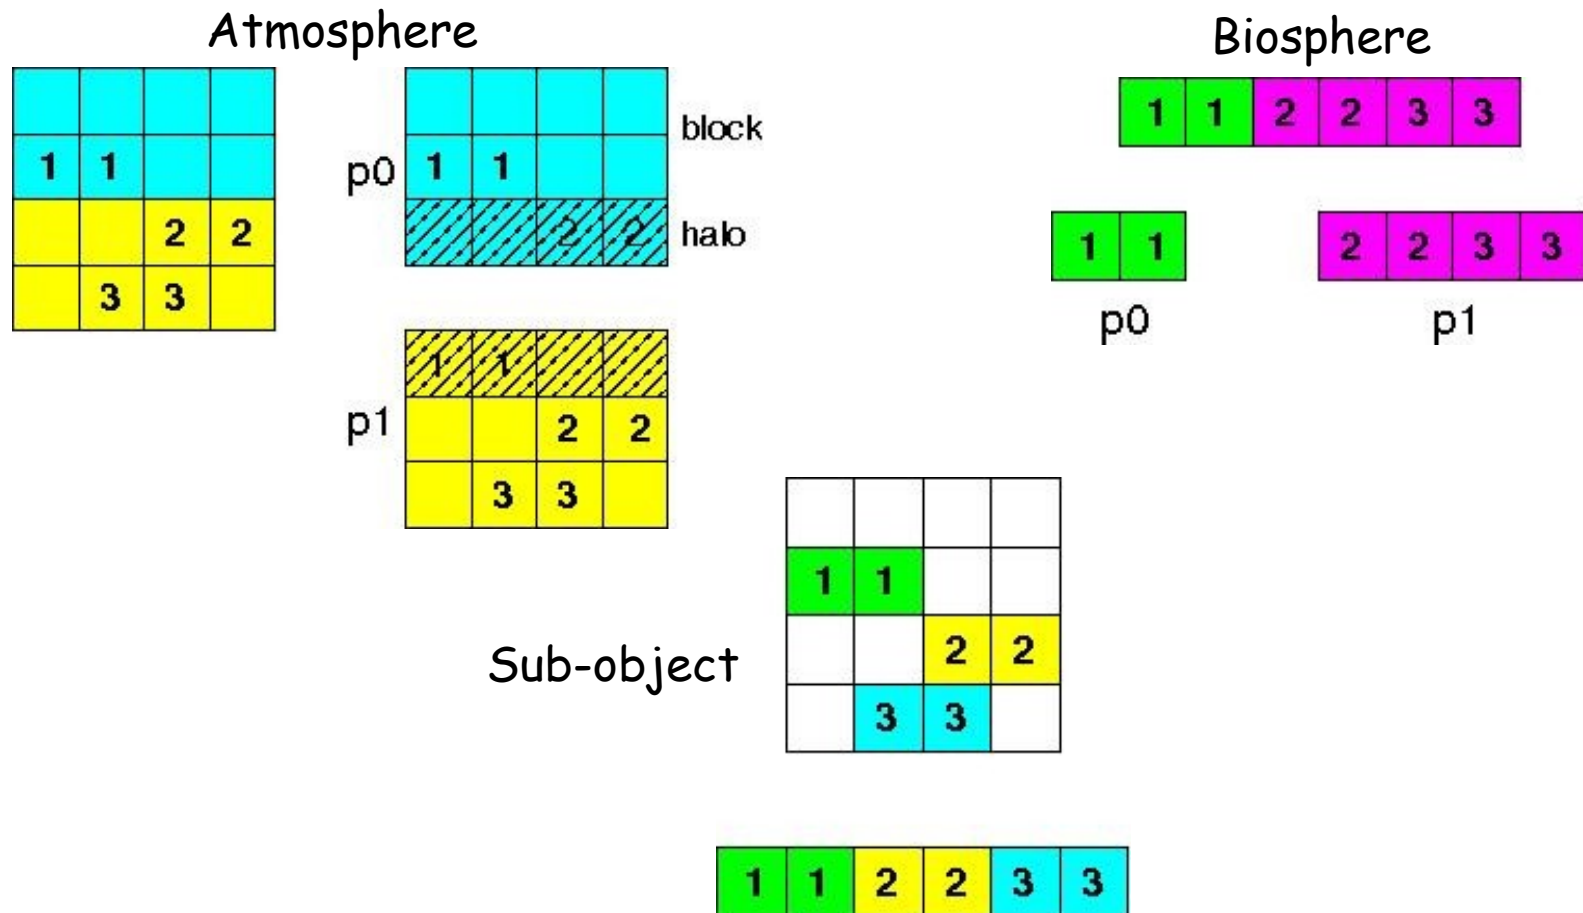
Dynamic coupling for a combustor cooling system optimization



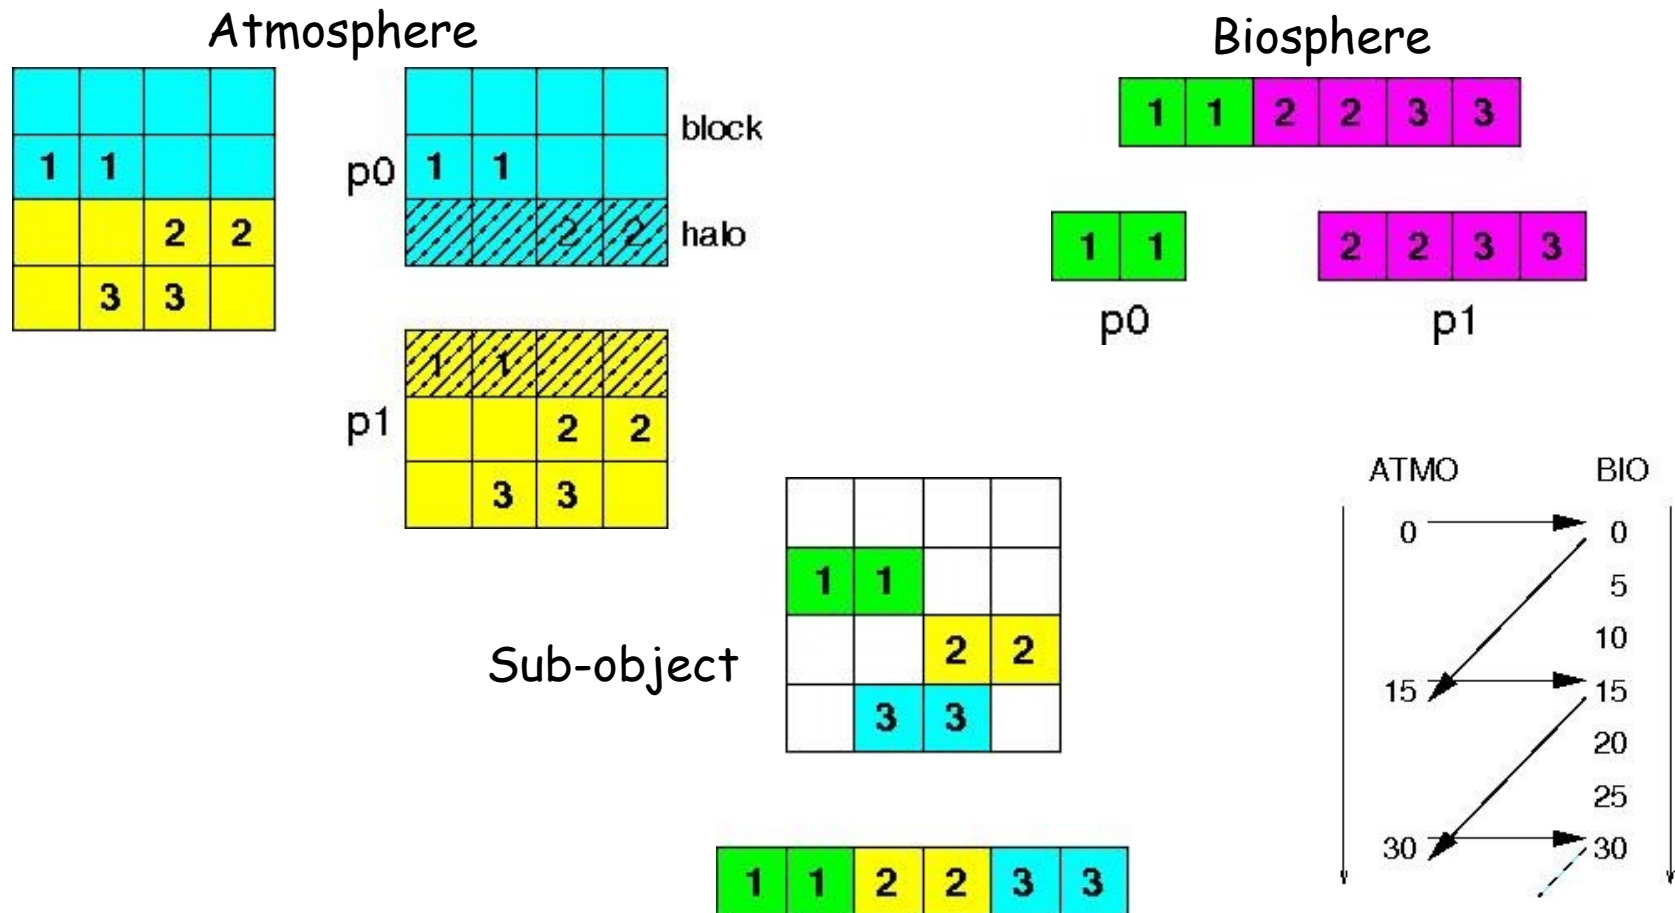
Flexible communication scheme (remapping, sub-objects, time shifts) for atmosphere biosphere coupling

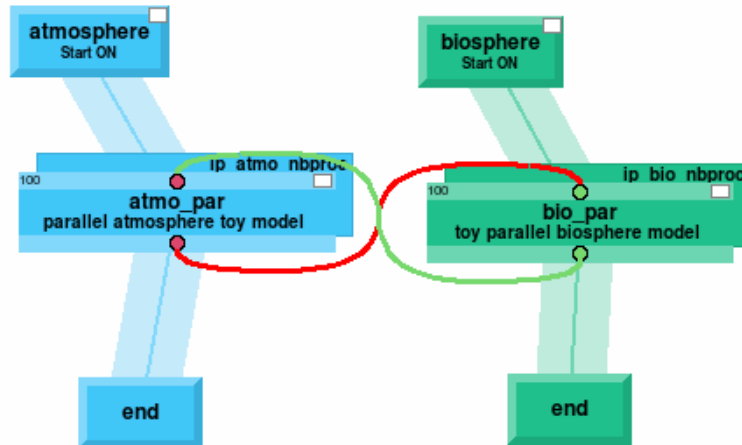


Flexible communication scheme (remapping, sub-objects, time shifts) for atmosphere biosphere coupling



Flexible communication scheme (remapping, sub-objects, time shifts) for atmosphere biosphere coupling





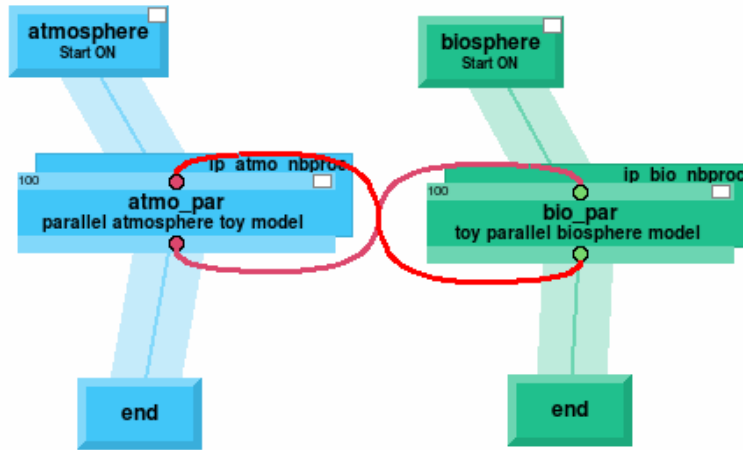
Change communication properties

Unit source name : atmo_par
Source Object : array_par_out.atmo_par
Source Distributor : reg_out_atmo_dist.atmo_par
Sub-object descriptor : vector_subobject

Unit target name : bio_par
Target object : vector_par.bio_par
Target Distributor : uneven_2_4_dist.bio_par
Sub-object descriptor : IDENTITY

Time list : 0:ip_nhours*3600-900:900
Local. assoc. : AUTOMATIC
Palm debug status : PL_NO_DEBUG
Palm track : PL_NO_TRACK
Data managment : MEMORY
Optimisation : PL_NO_BLOCKING

Cancel Update



Change communication properties

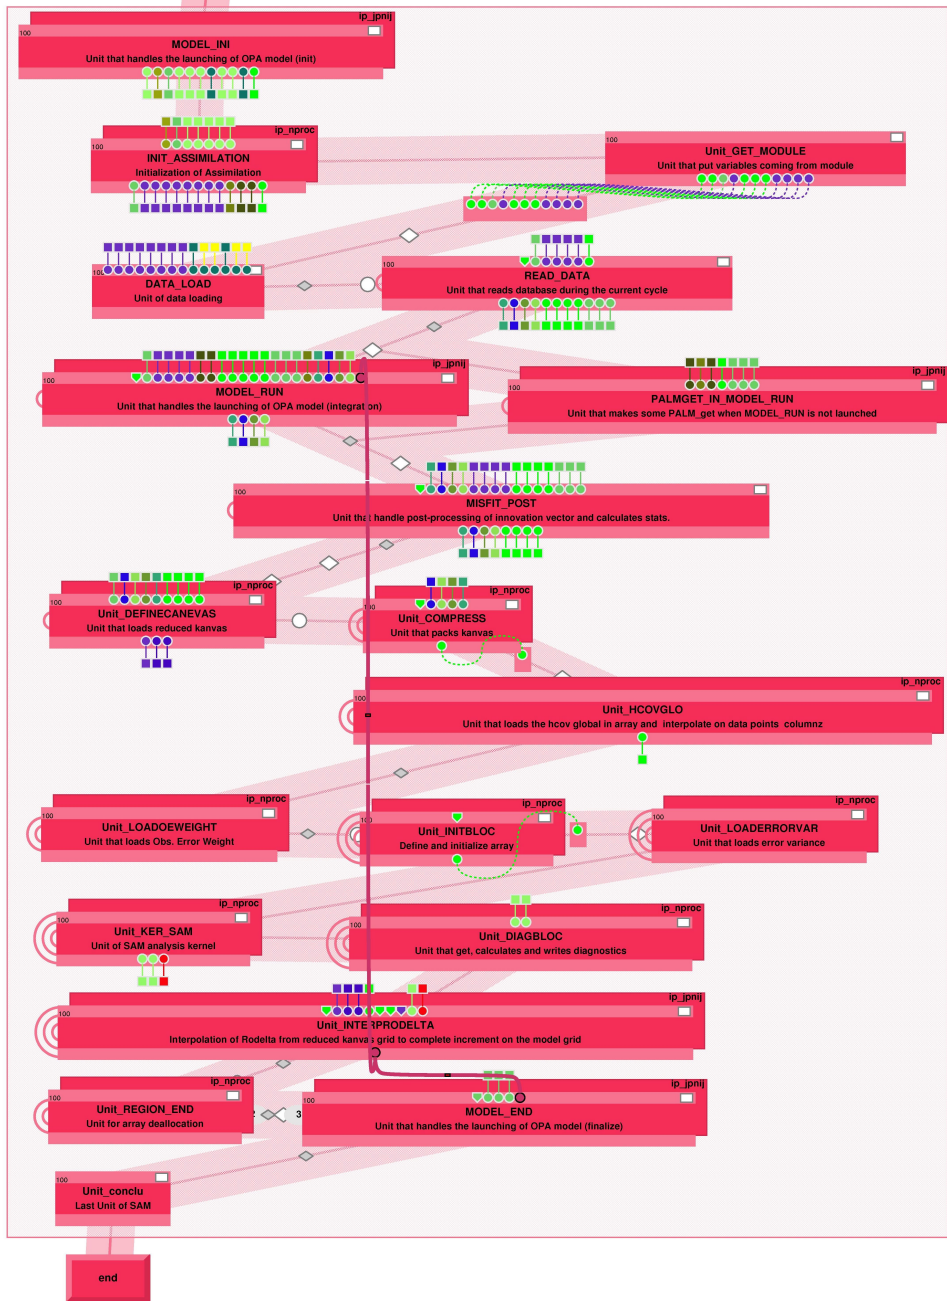
Unit source name : bio_par
Source Object : vector_par.bio_par
Source Distributor : uneven_2_4_dist.bio_par
Sub-object descriptor : IDENTITY

Unit target name : atmo_par
Target object : array_par_in.atmo_par
Target Distributor : regw_atmo_dist.atmo_par
Sub-object descriptor : vector_subobject

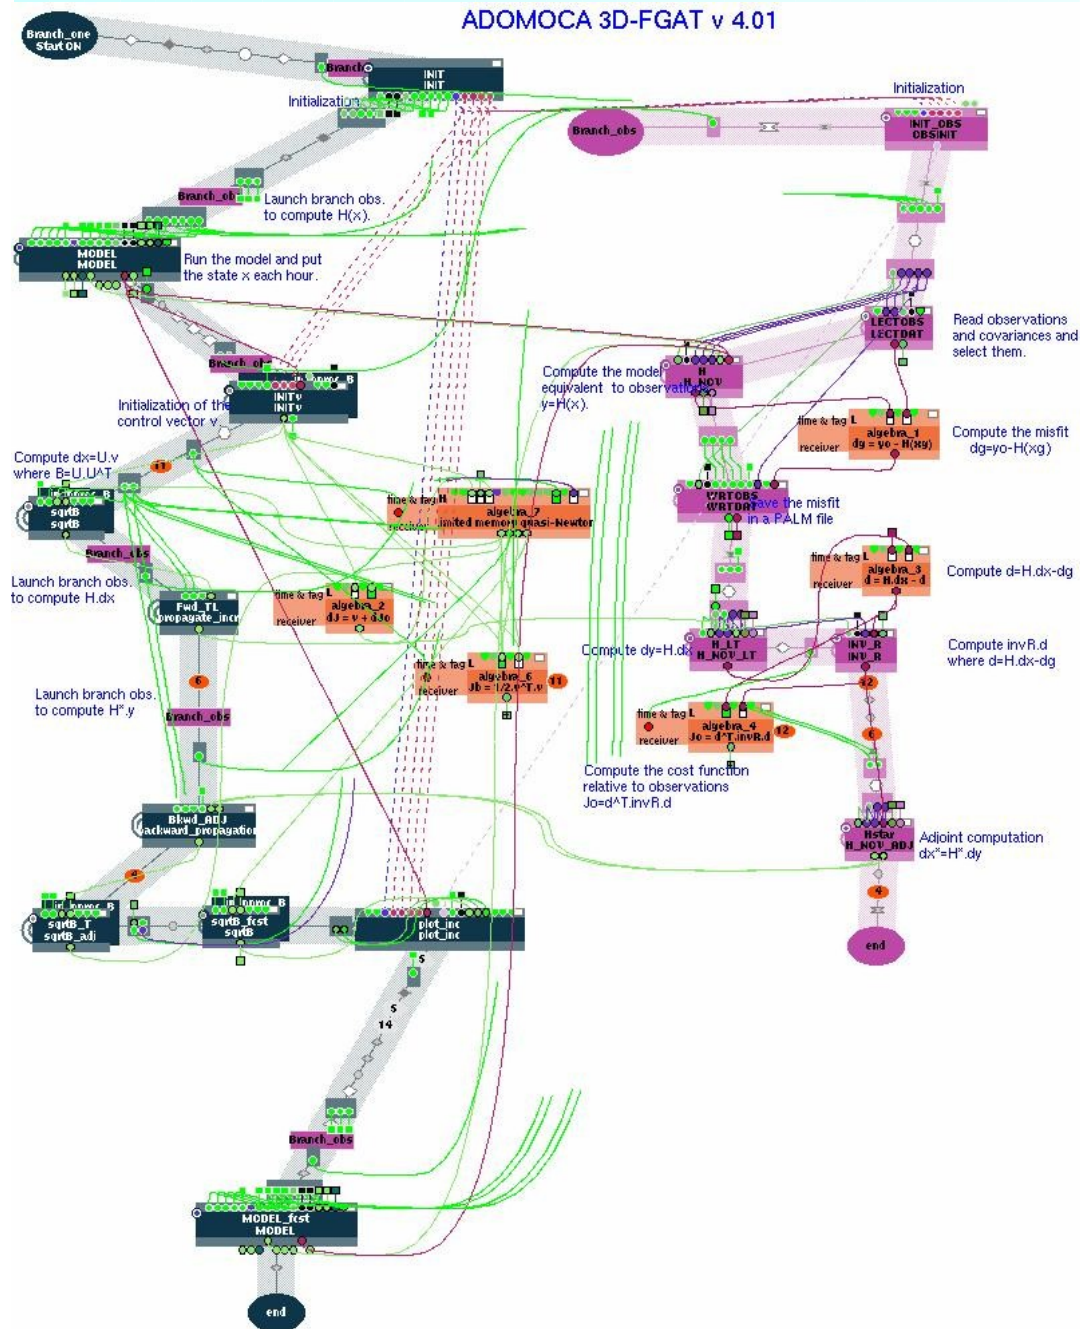
Time list : 0:ip_nhours*3600-900:900|i+900

Local. assoc. : AUTOMATIC
Palm debug status : PL_NO_DEBUG
Palm track : PL_NO_TRACK
Data managment : MEMORY
Optimisation : PL_NO_BLOCKING

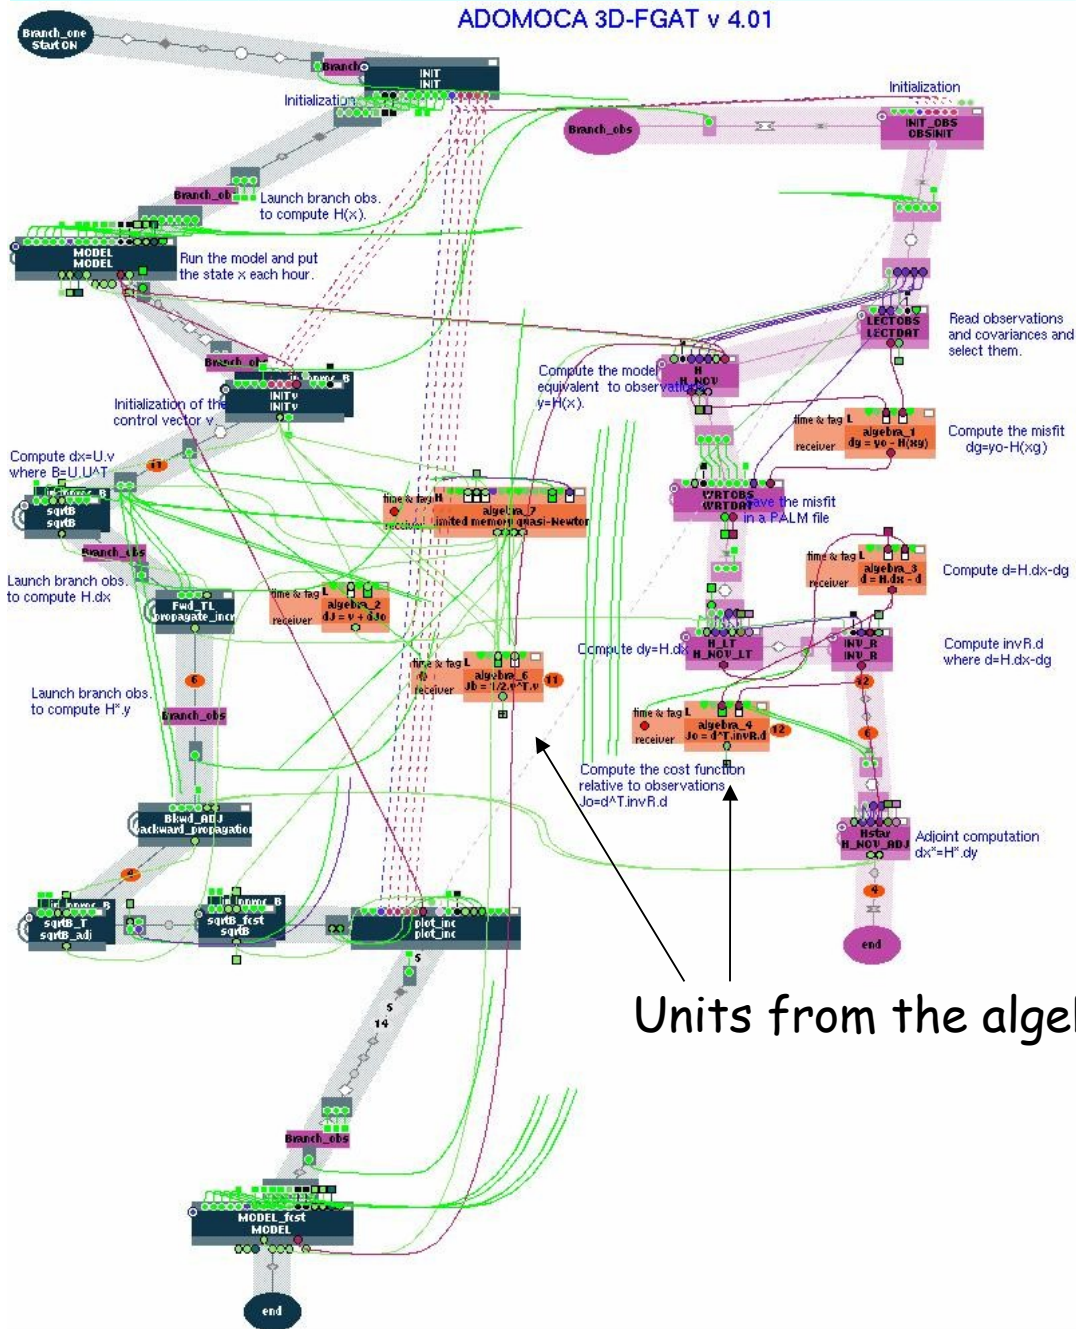
Cancel Update



Complex algorithm
High performances
... as required!

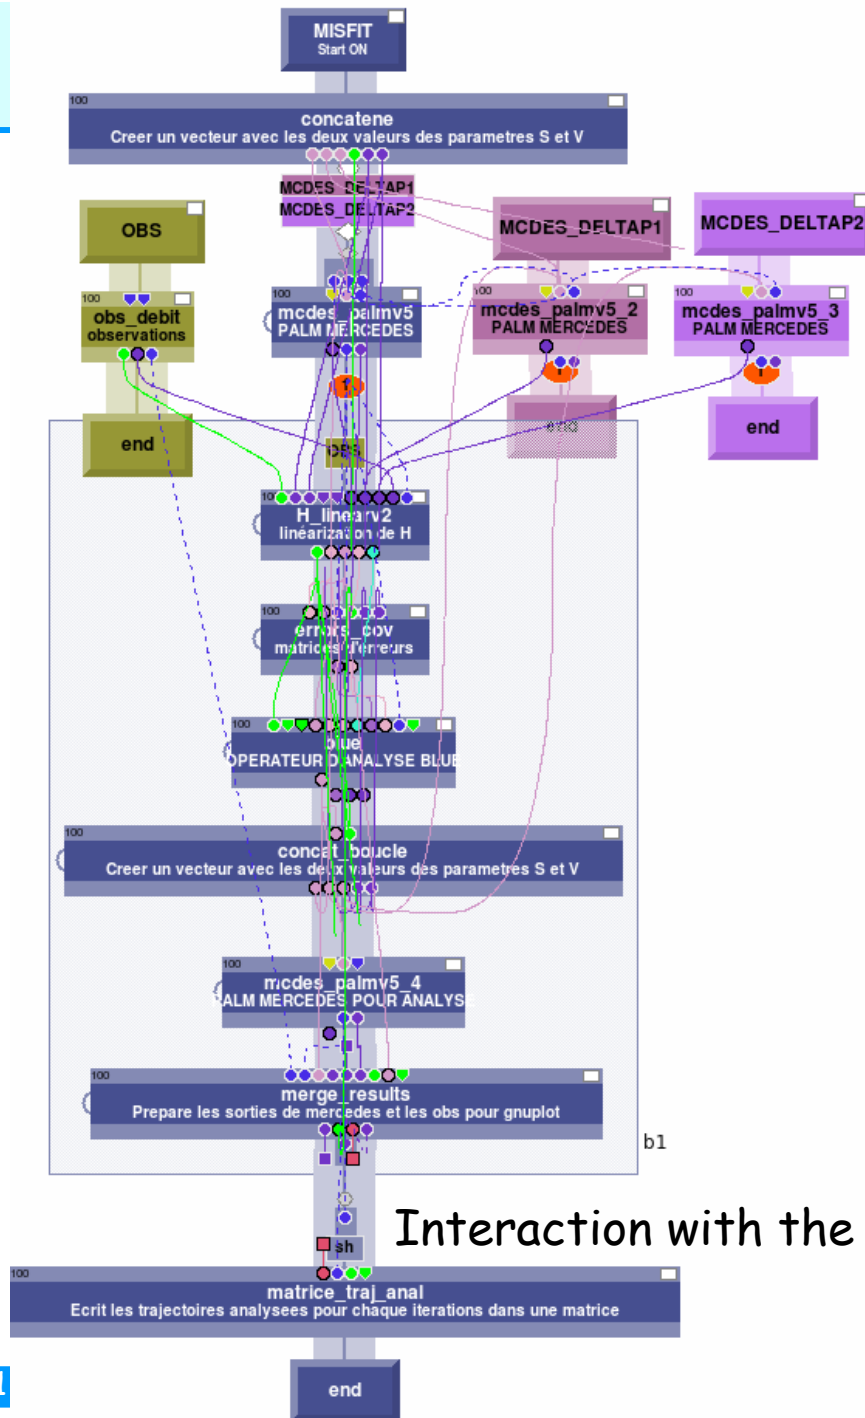


Still using PALM_SP
And V6 is more complex!



Still using PALM_SP And V6 is more complex!

Units from the algebra toolbox



Code reuse for parallel perturbed runs

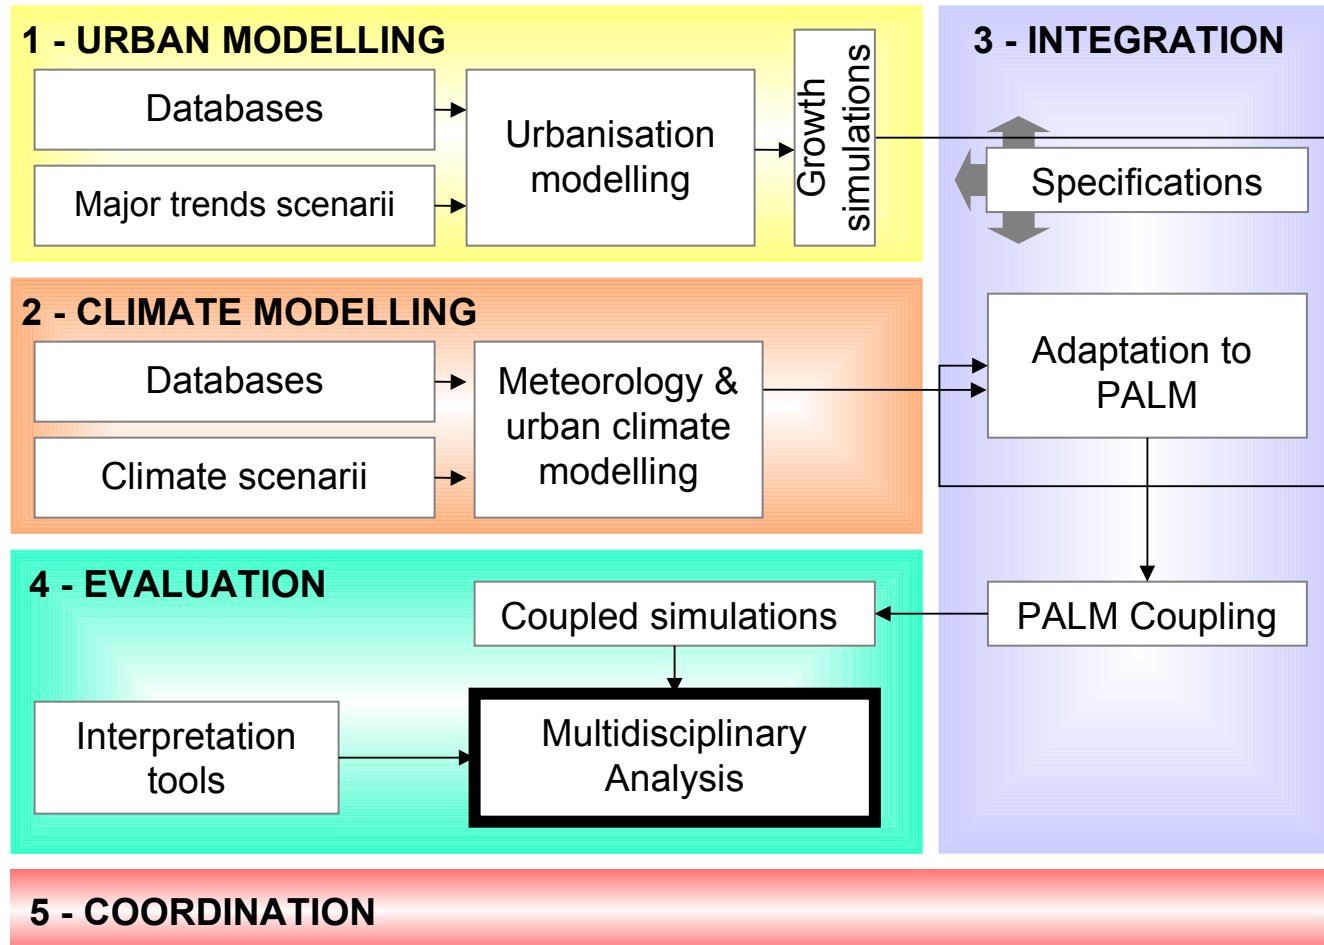
Interaction with the system for file handling

Urban scale climate evolution



ACCLIMAT project

The ACCLIMAT project has been funded with the support of the STAE-Toulouse Cooperation Foundation



Urban scale climate evolution

