PDAF Parallel Data Assimilation Framework

ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR-
UND MEERESFORSCHUNG

PDAF – Community Software for Ensemble Data Assimilation

Lars Nerger

Alfred Wegener Institute, Helmholtz Center for Polar and Marine Research
Bremerhaven, Germany

HELMHOLTZ
SPITZENFORSCHUNG FÜR
GROSSE HERAUSFORDERUNGEN

BOOS Annual Meeting, online, November 25, 2021

PDAF – Parallel Data Assimilation Framework

A universal tool for ensemble data assimilation …

- provide support for parallel ensemble forecasts
- provide assimilation methods (solvers) - fully-implemented & parallelized
- provide tools for observation handling and for diagnostics
- easily useable with (probably) any numerical model
- a program library (PDAF-core) plus additional functions
- run from notebooks to supercomputers (using standards: Fortran, MPI & OpenMP)
- usable for real assimilation applications and to study assimilation methods
- ~550 registered users; community contributions

Open source:
Code, documentation, and tutorial available at

http://pdaf.awi.de

L. Nerger, W. Hiller, Computers & Geosciences 55 (2013) 110-118

PDAF Application Examples

PDAF Parallel **Data Assimilation** Framework

HBM-ERGOM:
coupled physics/
biogeochemistry
coastal assimilation
(Goodliff et al., 2019)

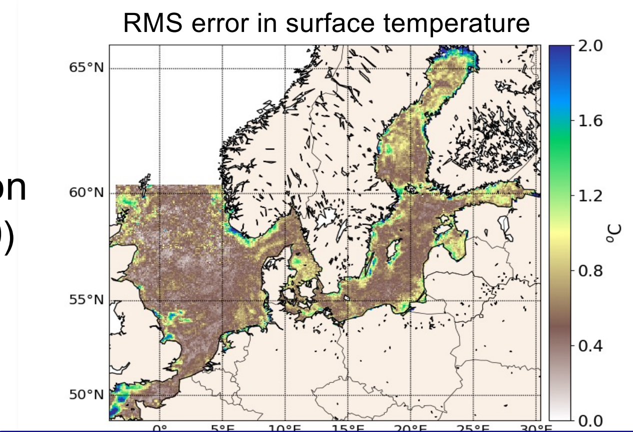

MITgcm-REcoM:
global ocean color
assimilation into
biogeochemical
model
(Pradhan et al., 2019/20)

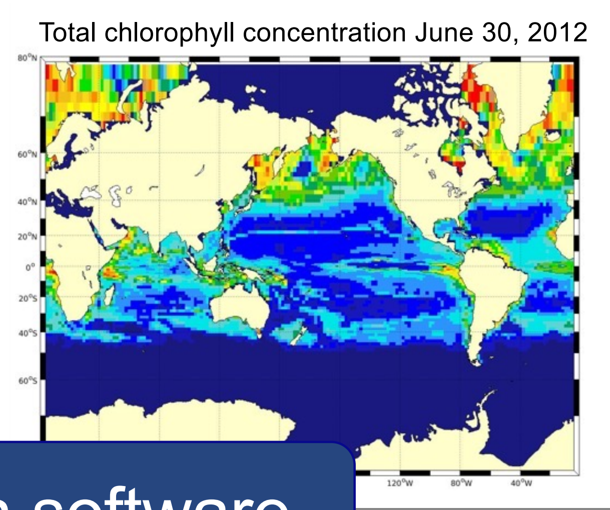

Different models – same assimilation software

AWI-CM:
coupled atmos.-
ocean assimilation
(Tang et al., 2020
Mu et al., 2020
Nerger et al., 2020)

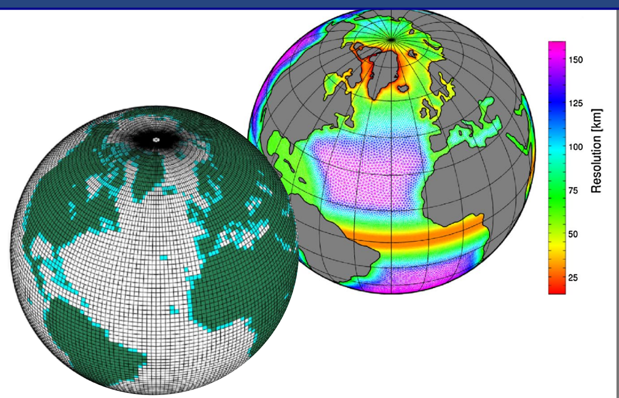

- MITgcm sea-ice assim (*operational*, NMEFC Beijing)
- CMEMS Baltic-MFC (*operational*, SMHI/DMI/BSH)
- HBM North/Baltic Seas (*operational*, BSH)
- NEMO (U Reading, P. J. van Leeuwen)
- SCHISM/ESMF (VIMS, J. Zhang)
- TerrSysMP-PDAF (hydrology, FZ Juelich, U Bonn)
- TIE-GCM (U Bonn, J. Kusche)
- VILMA (GFZ Potsdam)
- Parody geodynamo (IPGP Paris, A. Fournier)

Ensemble Data Assimilation

PDAF Parallel Data Assimilation Framework

Ensemble Kalman Filters & Particle Filters

→ Use ensembles to represent state and uncertainty

→ Propagate ensemble using numerical model

→ Use observations to update ensemble

→ EnKFs are current 'work horse'

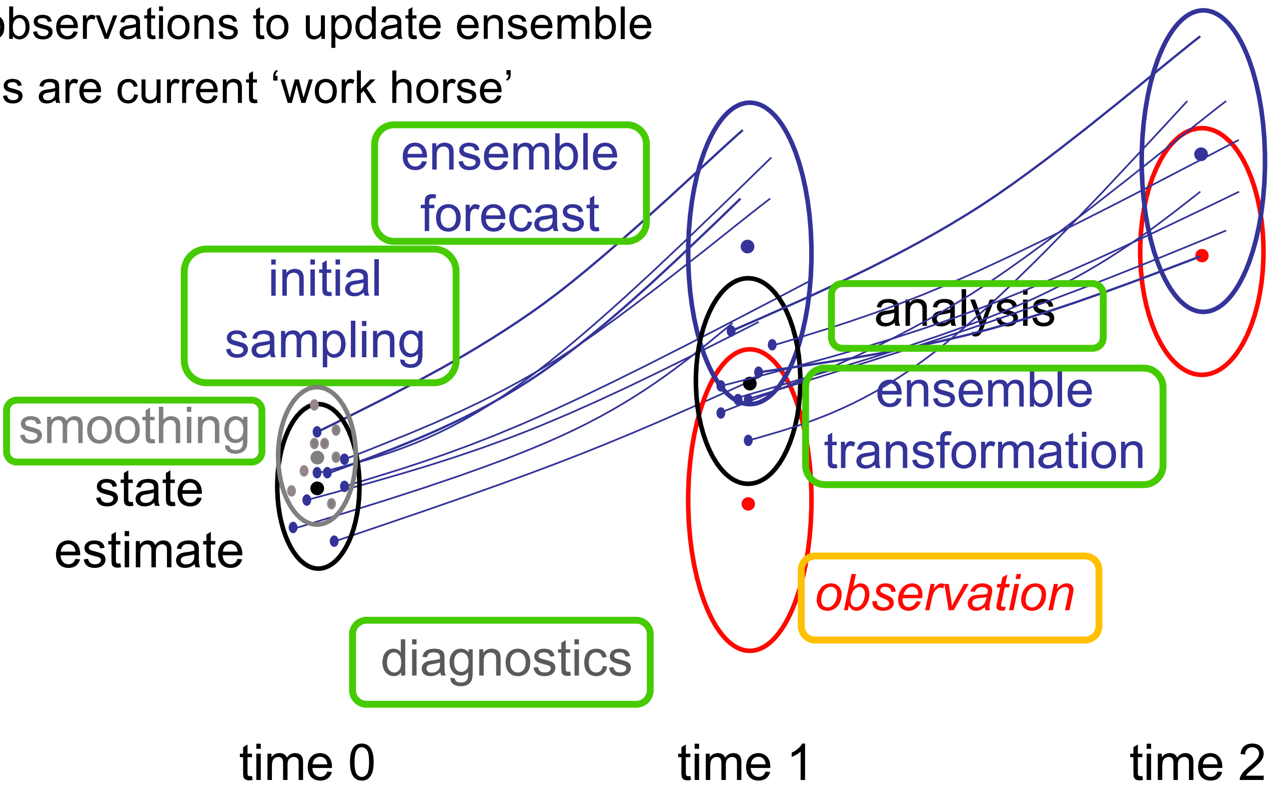

These steps can be implemented in a generic way

We can provide software including the algorithms → PDAF

Observation operators are less generic → PDAF also provides some operators and tools for obs. handling

PDAF: User-friendliness

PDAF Parallel Data Assimilation Framework

Goal: Enable easy and fast setup of a DA system,
and allow for extension to fully featured system

Assumption: Users know their model

➜ let users implement DA system in model context

For users, model is not just a time-stepping operator

➜ let users extend their model for data assimilation

Keep code simple for the user side:

- ➜ Define subroutine interfaces to DA code based on arrays
  (also simplifies interaction with languages like C/C++/Python)
- ➜ No object-oriented programming
  (most models don't use it; most model developers don't know it;
  many objects we would only have for observations – see later)
- ➜ Users directly implement case-specific routines
  (no indirect description (XML, YAML, ...) of e.g. observation layout)

Coupling Model and Assimilation Code: 2 Variants

PDAF Parallel Data Assimilation Framework

Offline

Model | **Assimilation program**

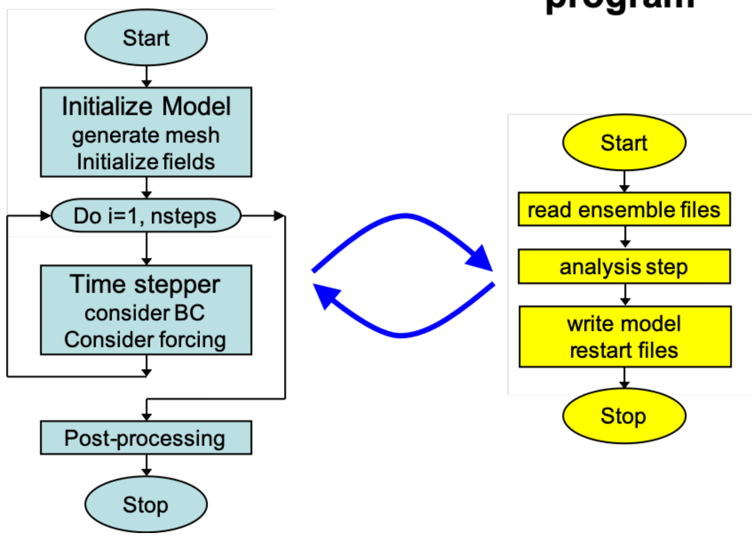

- Separate programs for model and DA
- Flexible to run
- Needs frequent model restarts and file output (less efficient than online coupling)

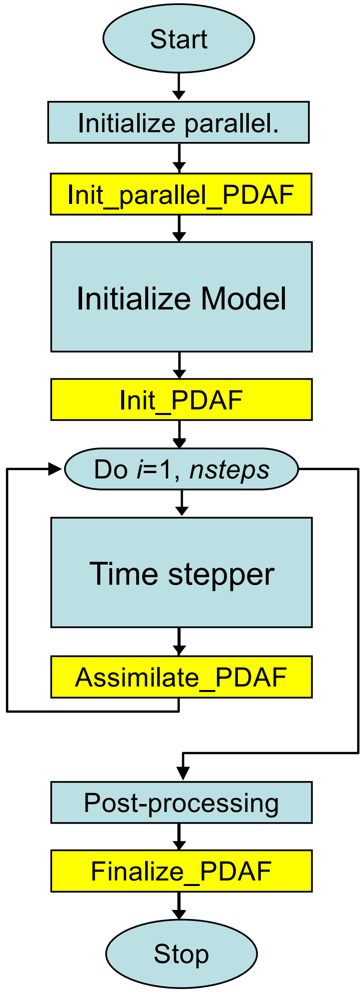

Online

- Augment model with DA functionality
- Insert 4 subroutine calls
  - Ensemble model
  - DA functionality
- Very efficient & highly scalable

Model code | DA code

Online-Coupling – Assimilation-enabled Model

Couple a model with PDAF

- Modify model to simulate ensemble of model states
- Insert analysis step/solver to be executed at prescribed interval
- Run model as usual, but with more processors and additional options

- EnOI also possible:
  - Evolve single model state
  - Prescribe ensemble perturbations

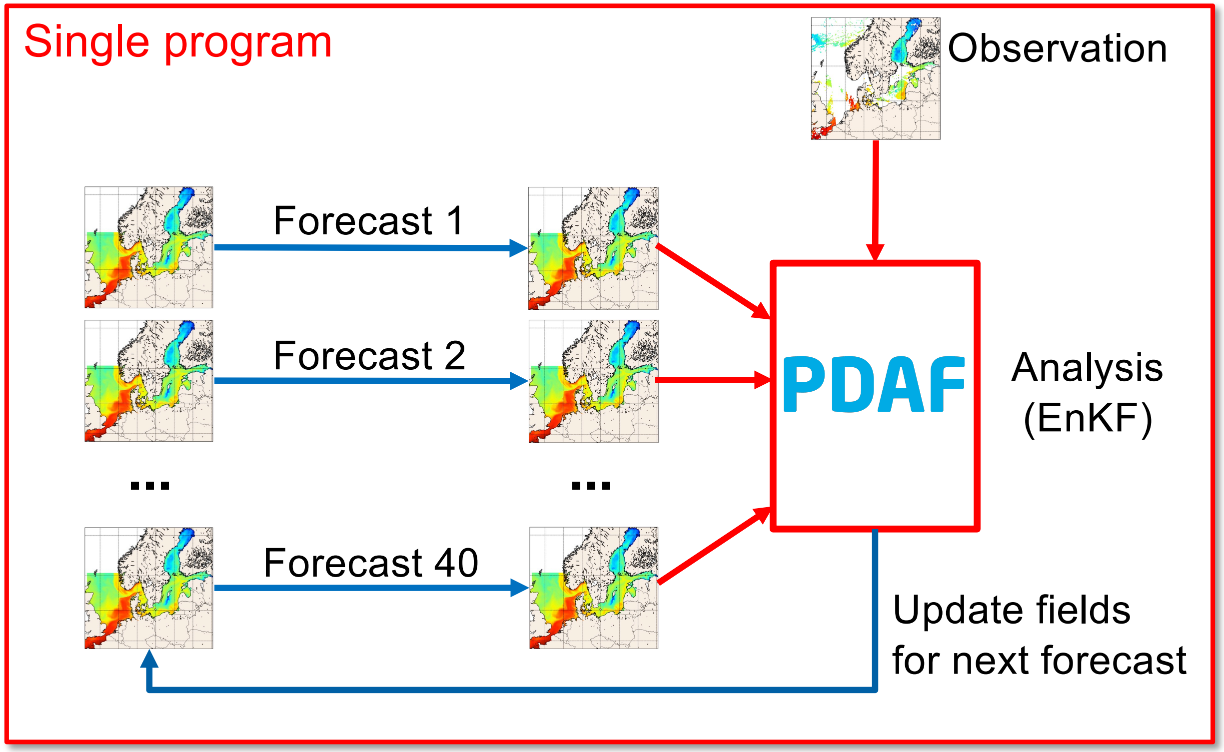

Initialize ensemble | Ensemble forecast | Analysis step in between time steps

Application types

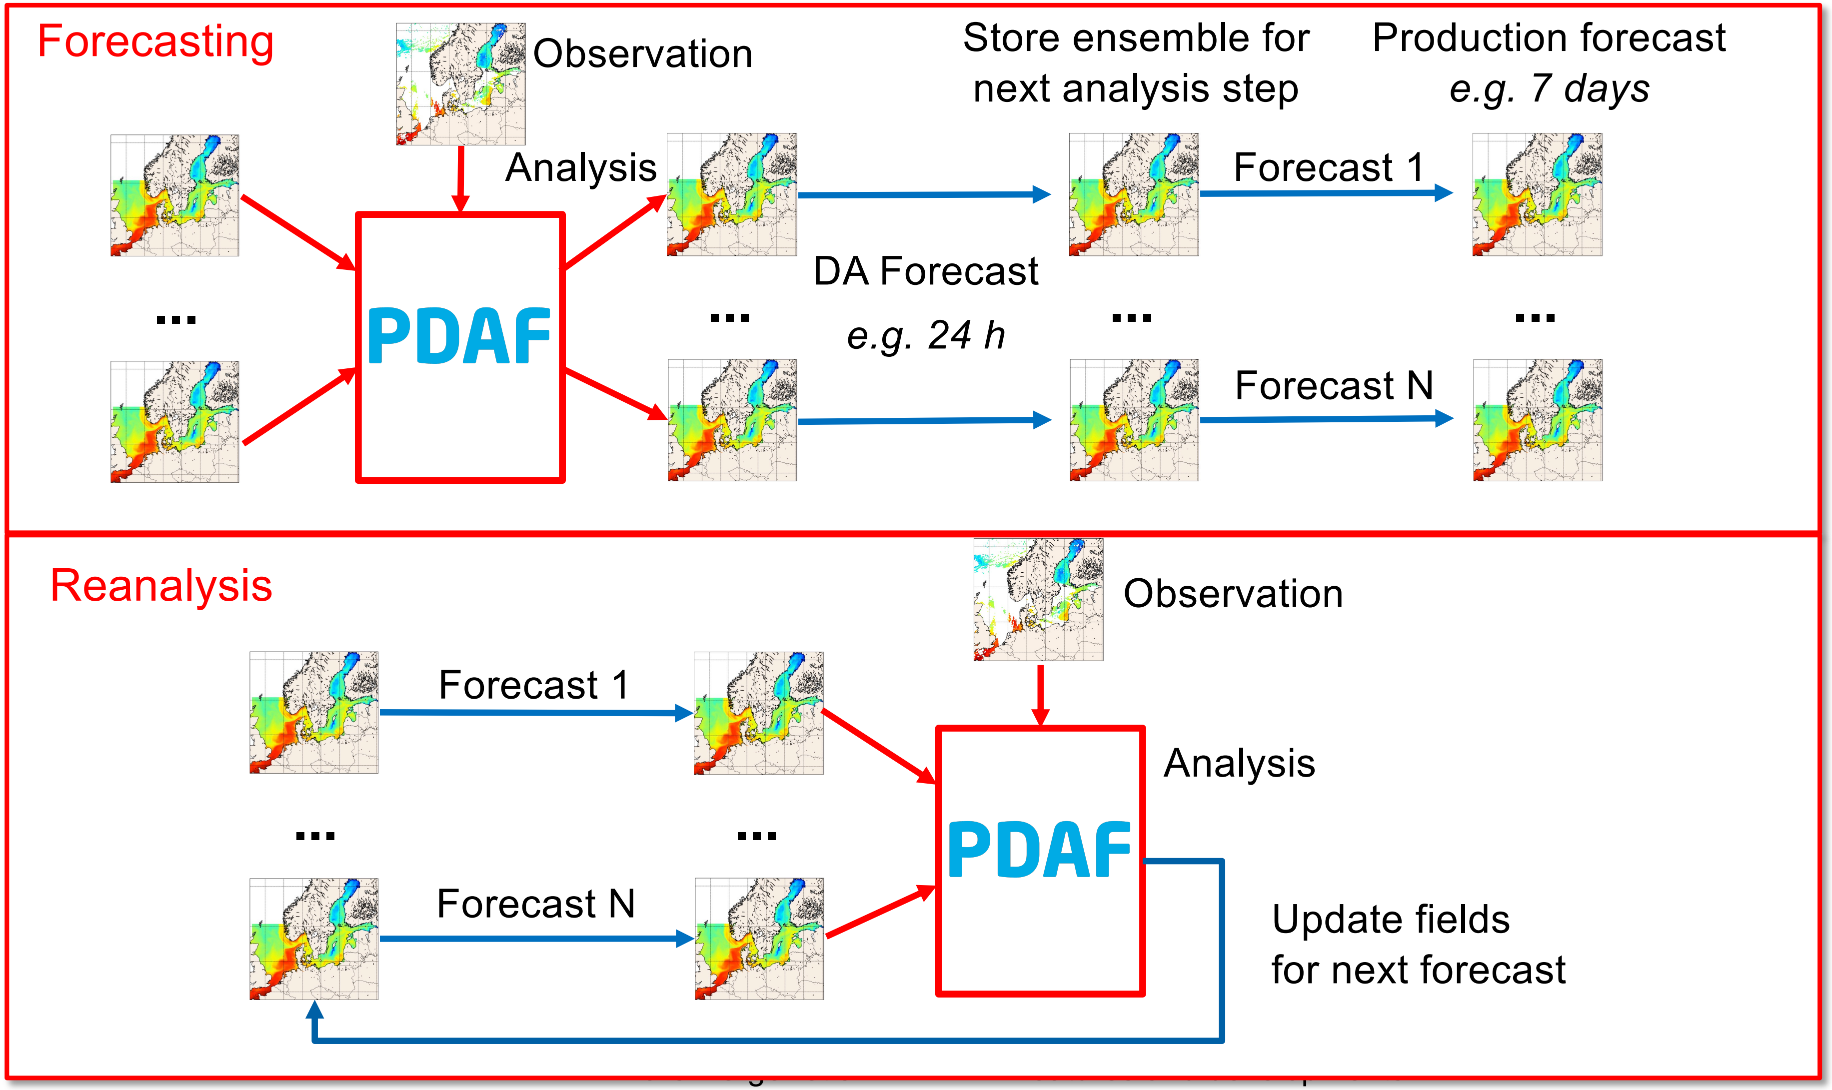

Execution times (weakly-coupled, DA only into ocean)

PDAF Parallel Data Assimilation Framework

AWI-CM (atmosphere-ocean coupled model)

Processor cores (each model instance)

- ECHAM: 72 / FESOM: 192
- Vary ensemble size
- Increasing integration time with growing ensemble size (11%; more parallel communication; worse placement)
- some variability in integration time over ensemble tasks

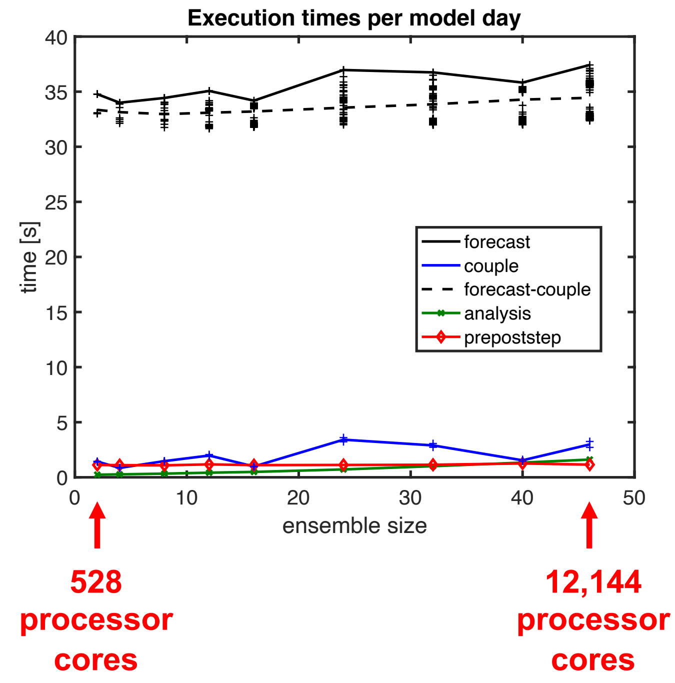

Important factors for good performance

- Need optimal distribution of programs over compute nodes/racks (here set up as ocean/atmosphere pairs)
- Avoid conflicts in IO (Best performance when each AWI-CM task runs in separate directory)

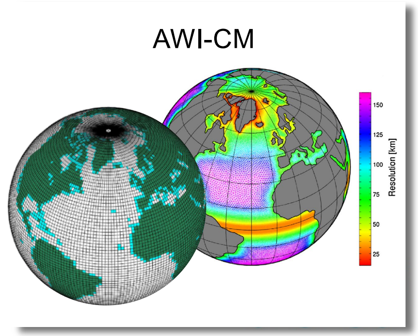

Nerger et al., GMD (2020), doi:10.5194/gmd-13-4305-2020

Online and Offline Coupling - Efficiency

Offline-coupling is simple to implement
but can be very inefficent

Example:
Timing from atmosphere-ocean
coupled model (AWI-CM)
with daily analysis step:

Model startup:	95 s	overhead
Integrate 1 day:	33 s	
Model postprocessing:	14 s	overhead
Analysis step:	1 s	

Restarting this model is ~3.5 times more
expensive than integrating 1 day

➜ avoid this for data assimilation
(in particular re-analysis applications)

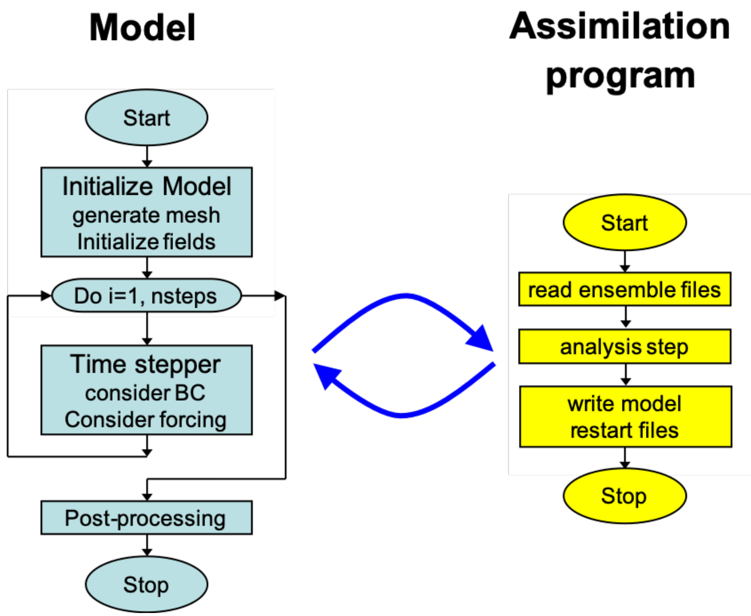

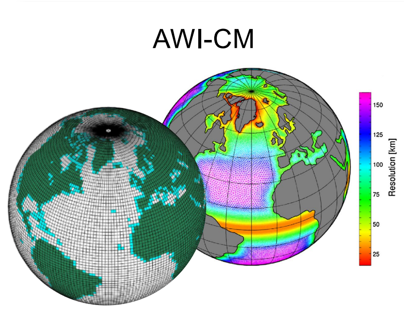

PDAF interface structure

PDAF Parallel Data Assimilation Framework

- **Model-sided Interface:** Defined calls to PDAF routines
  (called by driver program for offline coupling)

- **Case-related Interface:** User-supplied call-back routines for elementary operations:
  - transfers between model fields and ensemble of state vectors
  - observation-related operations

- **Internal Interface:** Connect to data assimilation methods

- User supplied routines can be implemented as routines of the model and can share data with it (low abstraction level)

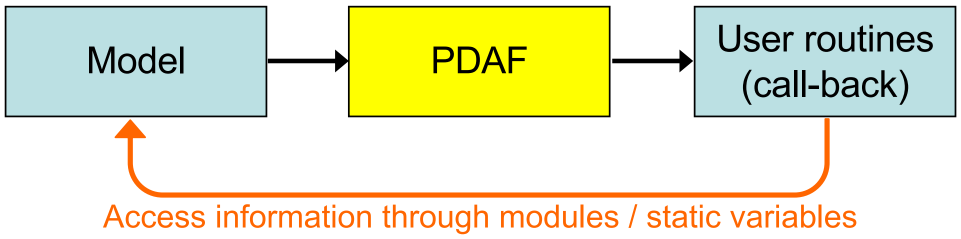

Implementing the Ensemble Analysis Step (Solver)

case-specific
call-back routines

Analysis operates
on state vectors
(all fields in one vector)

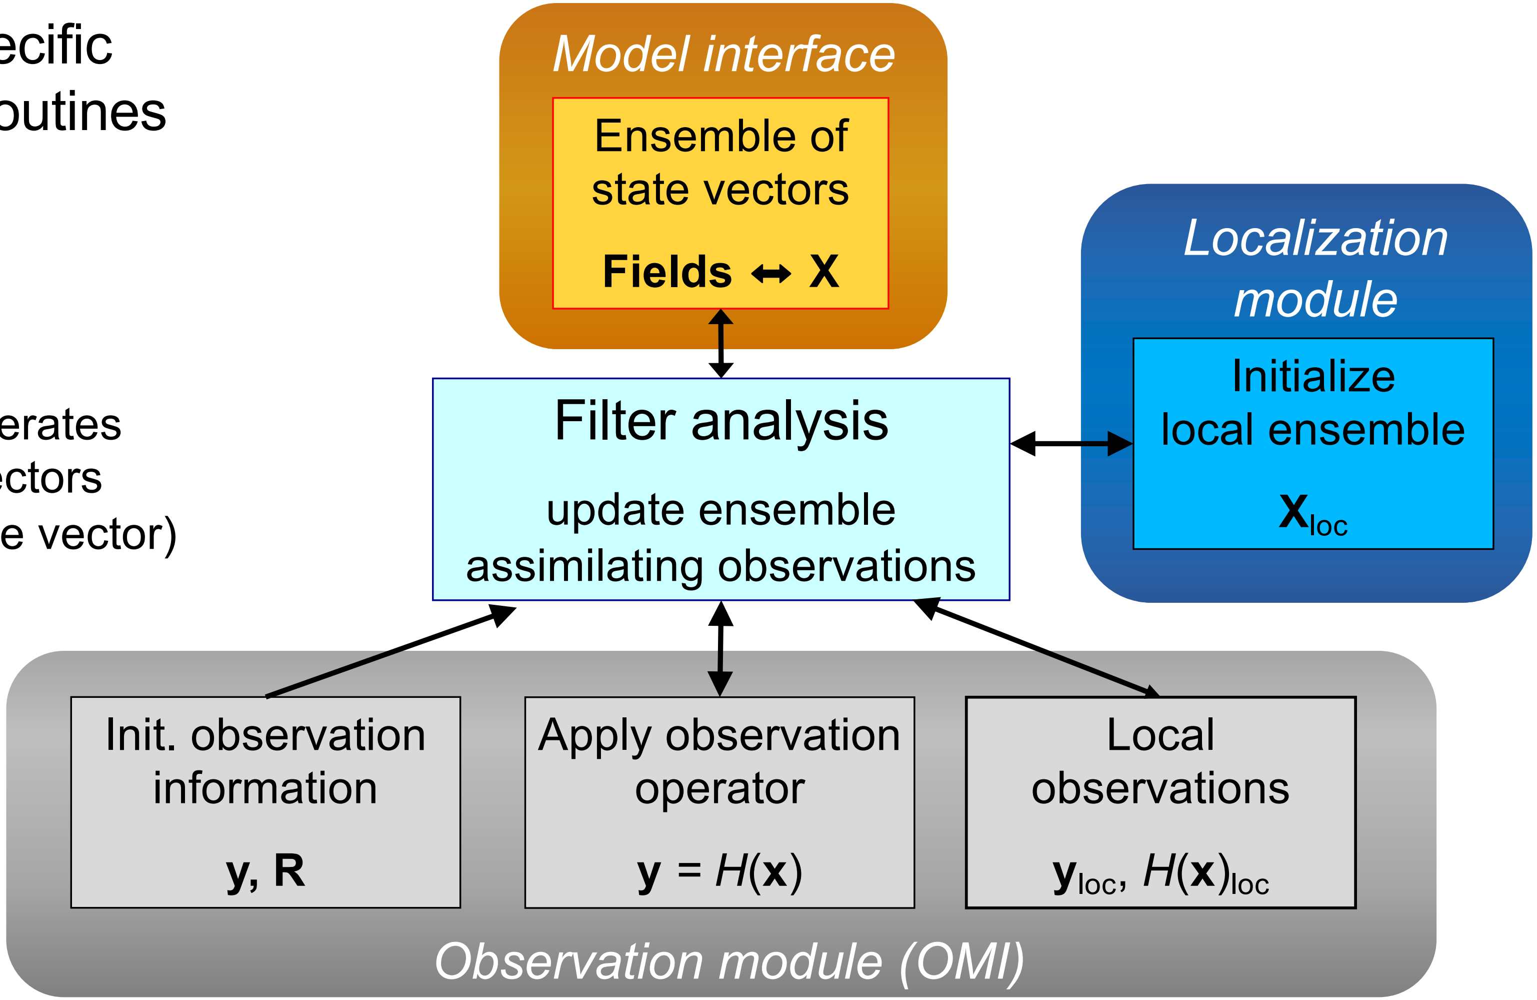

Recent and current developments

OMI: Code structure (Observation Module Infrastructure)

PDAF Parallel Data Assimilation Framework

Structure motivated by object-oriented programming. For sake of simplicity not implemented with OOP

One observation module per observation type

- separation between obs. Types
- Main implementation work in one routine
- observation operators are model independent
  - → OMI provides some generic operators
  - → Community can provide operators

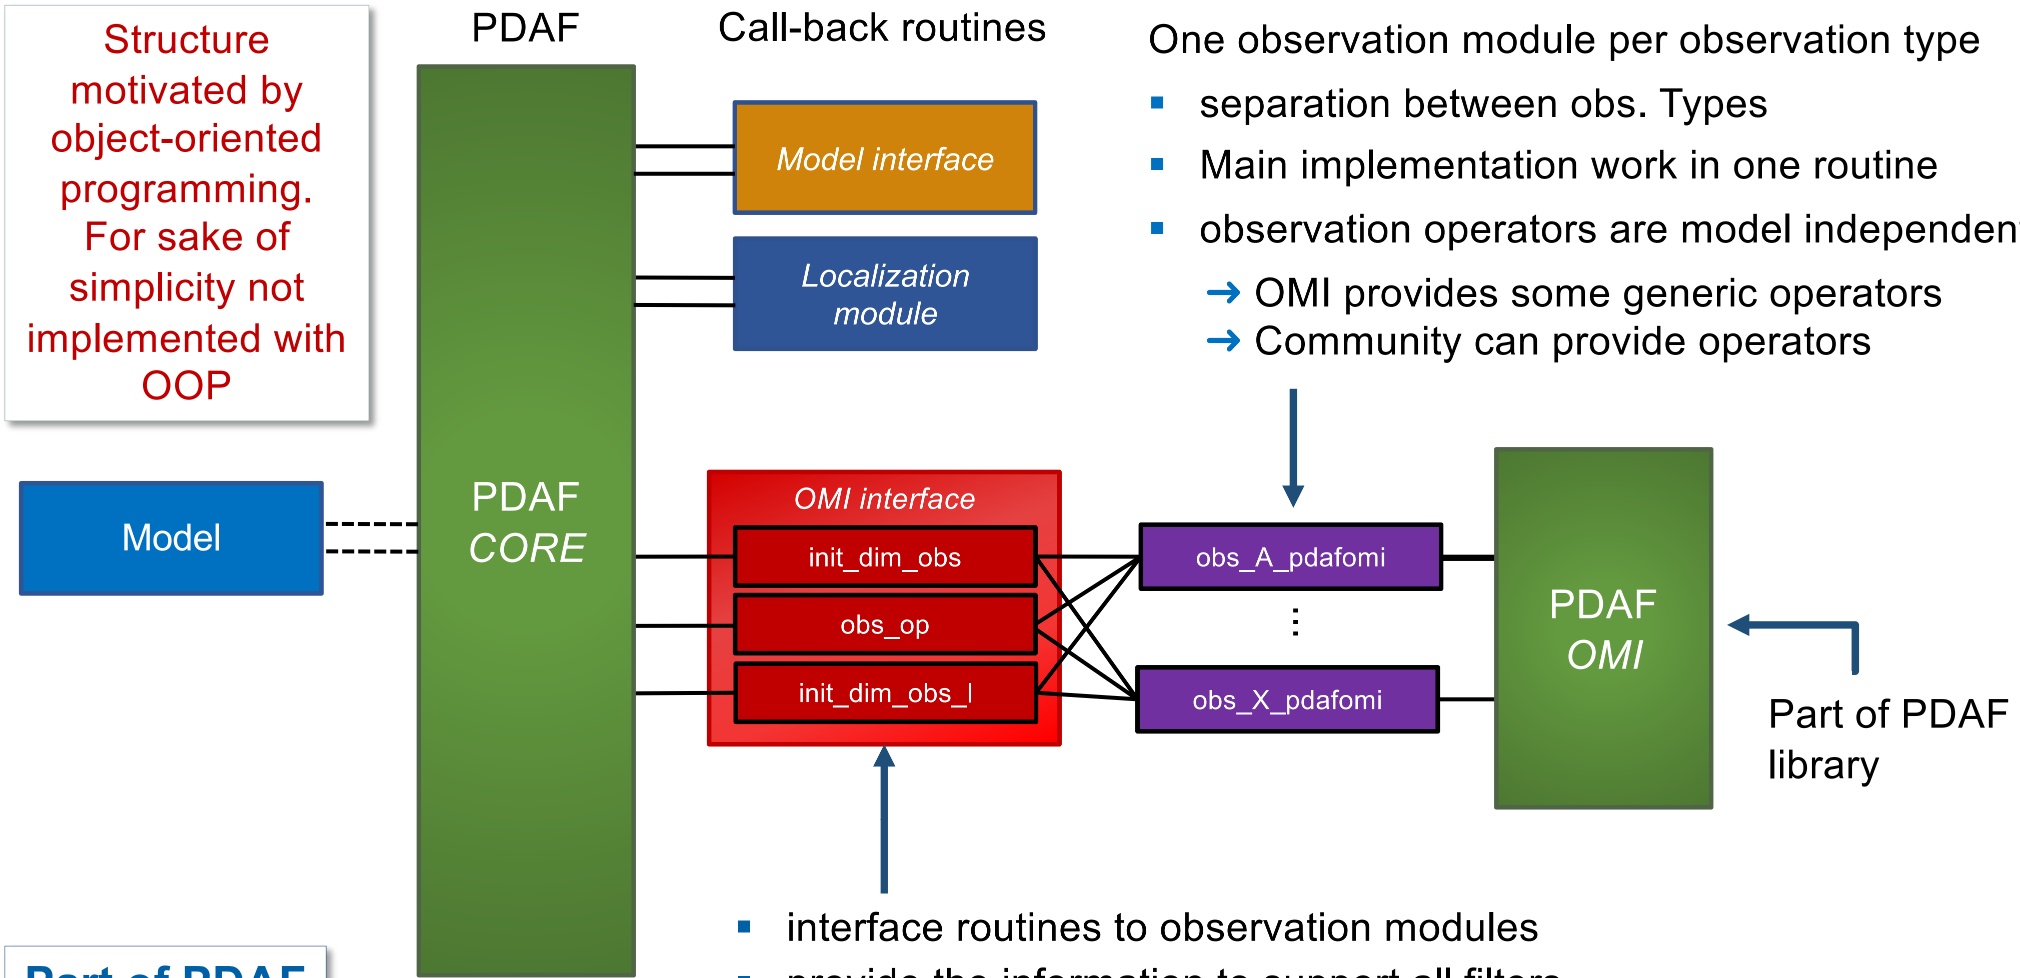

- interface routines to observation modules
- provide the information to support all filters

Part of PDAF V1.16

AWI

Strongly Coupled DA

PDAF Parallel Data Assimilation Framework

Strongly coupled DA:
Assimilate observation of component A into component B

PDAF supports strongly coupled DA:

achieved by adapting MPI communicator for the filter processes

➜ joint state vector decomposed over the processes

➜ Provide observation operator that only performs MPI communication

need innovation **d** = H(**x**) – **y**
and observed ensemble perturbations H(**X'**)

Observation operator H links different compartments

1. Compute part of **d** and H(**X'**) on process 'owning' the observation
2. Communicate **d** and H(**X'**) to processes for which observation is within localization radius

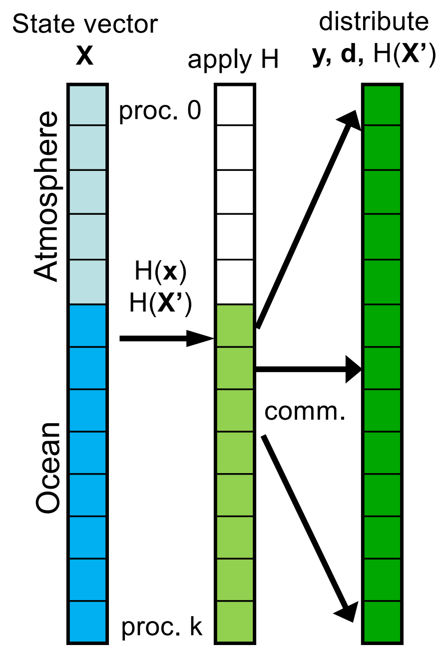

Observation handling in strongly coupled DA

Part of PDAF V1.16

Nerger, Tang, Mu (2020). GMD, 13, 4305–4321, doi:10.5194/gmd-13-4305-2020

Ensemble 3D-Var / Hybrid 3D-Var

PDAF Parallel Data Assimilation Framework

Activity in EU-project SEAMLESS

1D Prototype (in development):

- GOTM/FABM + ecosystem models
- DA functionality provided by PDAF

NEMO-ERGOM-PDAF data assimilation

Ensemble/Hybrid 3D-Var

- Some partners (PML, OGS) use 3D-Var
- Integrate in PDAF analogous to EnKFs/PFs
- Focus on infrastructure with optimizers as core

Included in upcoming release *(Dec. '21)*

- Possible extension to Ensemble 4D-Var later

SEAMLESS

www.seamlessproject.org

Services based on Ecosystem data AssiMiLation: Essential Science and Solutions

PML (UK, lead),
AWI (Germany), IGE (France),
NERSC (Norway), OGS (Italy)
Bolding&Bruggemann (Denmark)

AWI

DA Algorithms and models in PDAF

PDAF Parallel Data Assimilation Framework

PDAF originated from comparison studies of different filters

Filters and smoothers - *global and localized versions*

- EnKF (Evensen, 1994 + perturbed obs.)
- (L)ETKF (Bishop et al., 2001/Hunt et al. 2007)
- ESTKF (Nerger et al., 2012)
- NETF (Toedter & Ahrens, 2015)
- Particle filter
- *EnOI mode*

Upcoming *(Dec. '21)*:
- Ensemble 3D-Var
- Hybrid 3D-Var
- Hybrid NETF/LETKF

Model bindings

- MITgcm
- AWI-CM / FESOM

Toy models

- Lorenz-96 / Lorenz-63

Community provided:

SCHISM/ESMF

TerrSysMP-PDAF

Upcoming:
- NEMO 4 (U Reading)
- GOTM/FABM (BB ApS)

Upcoming *(Dec. '21)*:
- Lorenz-2005 II/III

PDAF: Design Considerations

PDAF Parallel Data Assimilation Framework

1. **Focus on ensemble methods**
2. **Efficiency:**
   - Direct (online/in-memory) coupling of model and data assimilation method (file-based offline coupling also supported)
   - Complete parallelism in model, DA method, and ensemble integrations
   - Provide common DA infrastructure with generic components
3. **Ease of use:**
   - require just standard compilers and libraries, no containers, etc.
   - just add subroutine calls into model code when combining with PDAF
   - model time stepper not required to be a subroutine
   - model controls the assimilation program
   - case-specific routines can be implemented like model code
   - simple switching between different filters and data sets
   - *Separation of concerns:* model, DA methods, observations

Requirements

- Fortran compiler
- MPI library
- BLAS & LAPACK
- make

- PDAF is at least tested (often used) on various computers:
  - Notebook & Workstation: MacOS, Linux (gfortran)
  - Cray XC30/40 & CS400 (Cray ftn and ifort)
  - NEC SX-8R / SX-ACE / SX-Aurora TSUBASA
  - ATOS Bull Sequana X (ifort)
  - HPE Cray Apollo (ARM)
  - Legacy:
    - SGI Altix & UltraViolet (ifort) / IBM Power (xlf) / IBM Blue Gene/Q

Summary - PDAF: A tool for data assimilation

PDAF Parallel Data Assimilation Framework

- a program library for ensemble modeling and data assimilation
- provides support for ensemble forecasts, DA diagnostics, and fully-implemented filter and smoother algorithms
- makes excellent use of supercomputers
- separation of concerns: model, DA methods, observations
- easy to couple to models and to code case-specific routines
- easy to add new DA methods
- efficient for research and operational use
- community code for DA methods and observations

PDAF adds DA functionality to models

Couple model and PDAF within days

Get DA capability in a month

Run DA in known environment

Access new DA methods by updating PDAF

Open source:
Code, documentation, and tutorial available at

http://pdaf.awi.de

L. Nerger, W. Hiller, Computers & Geosciences 55 (2013) 110-118

References

- **http://pdaf.awi.de**
- Nerger, L., Hiller, W. (2013). Software for Ensemble-based Data Assimilation Systems - Implementation Strategies and Scalability. Computers and Geosciences, 55, 110-118. doi:10.1016/j.cageo.2012.03.026
- Nerger, L., Tang, Q., Mu, L. (2020). Efficient ensemble data assimilation for coupled models with the Parallel Data Assimilation Framework: Example of AWI-CM. Geoscientific Model Development, 13, 4305–4321, doi:10.5194/gmd-13-4305-2020
- Tang, Q., Mu, L., Sidorenko, D., Goessling, H., Semmler, T., Nerger, L. (2020) Improving the ocean and atmosphere in a coupled ocean-atmosphere model by assimilating satellite sea surface temperature and subsurface profile data. Q. J. Royal Metorol. Soc., in press doi:10.1002/qj.3885
- Mu, L., Nerger, L., Tang, Q., Losa, S. N., Sidorenko, D., Wang, Q., Semmler, T., Zampieri, L., Losch, M., Goessling, H. F. (2020) Towards a data assimilation system for seamless sea ice prediction based o the AWI climate model. Journal of Advances in Modeling Earth Systems, 12, e2019MS001937 doi:10.1029/2019MS001937