

OpenDA Course and Exercises

Nils van Velzen, Martin Verlaan, Stef Hummel

March 10, 2023

Contents

1	Double Pendulum	2
1.1	Input files	3
1.2	Simulation and postprocessing with the double pendulum model	3
1.2.1	simulation	3
1.2.2	postprocessing	3
1.3	An alternative simulation with the double-pendulum model . . .	4
1.4	An ensemble of model runs	5
2	Some basic properties of the EnKF with Lorenz 3 variable model	7
3	Steady-state Kalman Filter	8
4	Exercise 4: A black box model - Calibration	10
4.1	Wrapper configuration files	11
5	Exercise 5: A black box model - Filtering	13
5.1	Sequential simulation	13
5.2	Sequential ensemble simulation	13
5.3	parallel computing	13
5.4	ensemble kalman filter	13
6	Exercise 6: Writing your own toy model	14
7	Exercise 7: Localization	16


Installation of OpenDA

Before you can start with the exercises you must first install OpenDA. For the latest instructions, you are referred to `$OPENDA/doc/OpenDA_domumentation.pdf`, section “Installation” or the same document on our website www.openda.org.

1 Double Pendulum

Directory: exercise_double_pendulum_part1

A pendulum is a rigid body that can swing under the influence of gravity. It is attached at the top so it can rotate freely in a two-dimensional plane (x, y) . We will assume a thin rectangular shape with the mass equally distributed. A double pendulum is a pendulum connected to the end of another pendulum. Contrary to the regular movement of a pendulum, the motion of a double-pendulum is very irregular when sufficient energy is put into the system.



The dynamics of a double-pendulum can be described with the following equations. (This example was copied from https://en.wikipedia.org/wiki/Double_pendulum)

With variables $\theta_1, \theta_2, p_{\theta_1}, p_{\theta_2}$:

$$\frac{d\theta_1}{dt} = \frac{6}{ml^2} \frac{2p_{\theta_1} - 3 \cos(\theta_1 - \theta_2)p_{\theta_2}}{16 - 9 \cos^2(\theta_1 - \theta_2)} \quad (1)$$

$$\frac{d\theta_2}{dt} = \frac{6}{ml^2} \frac{8p_{\theta_2} - 3 \cos(\theta_1 - \theta_2)p_{\theta_1}}{16 - 9 \cos^2(\theta_1 - \theta_2)} \quad (2)$$

$$\frac{dp_{\theta_1}}{dt} = -\frac{1}{2}ml^2 \left(\frac{d\theta_1}{dt} \frac{d\theta_2}{dt} \sin(\theta_1 - \theta_2) + 3\frac{g}{l} \sin(\theta_1) \right) \quad (3)$$

$$\frac{dp_{\theta_2}}{dt} = -\frac{1}{2}ml^2 \left(-\frac{d\theta_1}{dt} \frac{d\theta_2}{dt} \sin(\theta_1 - \theta_2) + \frac{g}{l} \sin(\theta_2) \right) \quad (4)$$

where the x, y -position of the middle of the two segments can be computed as:

$$x_1 = \frac{l}{2} \sin(\theta_1) \quad (5)$$

$$y_1 = \frac{-l}{2} \cos(\theta_1) \quad (6)$$

$$x_2 = l(\sin(\theta_1) + \frac{1}{2} \sin(\theta_2)) \quad (7)$$

$$y_2 = -l(\cos(\theta_1) + \frac{1}{2} \cos(\theta_2)) \quad (8)$$

This model, although simple, is very nonlinear and has a chaotic nature. Its solution is very sensitive to the parameters and the initial conditions: a small difference in those values can lead to a very different solution.

The purpose of this exercise is to get you started with OpenDA. You will learn to run a model in OpenDA, make modifications to the input files and plot the results.

1.1 Input files

The input for this exercise is located in directory `exercise_pendulum_part1`.

For Linux and Mac OS X, go to this directory and start `oda_run.sh`, the main application of OpenDA. For Windows, start the main application with `oda_run_gui.bat` from the `$OPENDA/bin` directory. The main application allows you to view and edit the OpenDA configuration files, run your simulations and visualize the results.

1.2 Simulation and postprocessing with the double pendulum model

1.2.1 simulation

Try to run a simulation with the double pendulum model. You can use the configuration file `simulation_unperturbed.oda`.

1.2.2 postprocessing

For postprocessing in Matlab the results are written to the file

`simulation_unperturbed_results.m`

Next, start Matlab and load the results. We have added a routine `plot_movie` to create an intuitive representation of the data. Please type (or copy-paste):

```
1 [t,unperturbed,tobs,obs]= ...
2 load_results('simulation_unperturbed_results');
3 plot_movie(t,unperturbed)
```

Listing 1: Matlab

For postprocessing in Python the results are written to the file

`simulation_unperturbed_results.py`

These results can be loaded with:

```
1 import simulation_unperturbed_results as unperturbed
2 # use reload(unperturbed) if unperturbed was loaded before
```

Listing 2: Python initialize

We have added a routine `plot_movie` to create an intuitive representation of the data.

```
1 import pendulum as p #needed only once
2 p.plot_movie(unperturbed.model_time,unperturbed.x)
```

Listing 3: Python

To create a time-series plot in Matlab type:

```

1 subplot(2,1,1);
2 plot(t,unperturbed(1,:), 'b-');
3 ylabel('\theta_1');
4 subplot(2,1,2);
5 plot(t,unperturbed(2,:), 'b-');
6 ylabel('\theta_2');
7 xlabel('time');

```

Listing 4: Matlab

To create a time-series plot in Python type:

```

1 plt.subplot(2,1,1)
2 plt.plot(unperturbed.model_time,unperturbed.x[:,0], "b")
3 # Python counts starting at 0
4 plt.ylabel(r"$\theta_1$") # use latex for label
5 plt.subplot(2,1,2)
6 plt.plot(unperturbed.model_time,unperturbed.x[:,1], "b")
7 plt.ylabel(r"$\theta_2$")
8 plt.show()
9 # only needed if interactive plotting is off.
10 # Set with plt.ioff(), plt.ion()

```

Listing 5: Python

1.3 An alternative simulation with the double-pendulum model

Then you can start an alternative simulation with the double-pendulum model that starts with a slightly different initial condition using the configuration file `simulation_perturbed.oda`. The different initial conditions can be found in `modelDoublePendulumStochModel.xml` and `modelDoublePendulumStochModel_perturbed.xml`

Visualize the unperturbed and perturbed results in a single plot. Make a movie and a time-series plot of θ_1 and θ_2 variables. Do you see the solutions diverging like the theory predicts?

```

1 [tu,unperturbed,tobs1,obs1]=load_results('
    simulation_unperturbed_results');
2 [tp,perturbed,tobs2,obs2]=load_results('
    simulation_perturbed_results');
3 figure(1);clf;subplot(2,1,1);
4 plot(tu,unperturbed(1,:), 'b');
5 hold on;
6 plot(tp,perturbed(1,:), 'g');
7 hold off;
8 legend('unperturbed', 'perturbed')
9 subplot(2,1,2);

```

```

10 plot(tu,unperturbed(2,:), 'b');
11 hold on;
12 plot(tp,perturbed(2,:), 'g');
13 hold off;

```

Listing 6: Matlab

To create a movie with both results in python type:

```

1 import simulation_unperturbed_results as unperturbed
2 import simulation_perturbed_results as perturbed
3 p.plot_movie(unperturbed.model_time, unperturbed.x, perturbed
  .x)

```

Listing 7: Python initialize

To create a time-series plot with both results in Python type:

```

1 plt.subplot(2,1,1)
2 plt.plot(unperturbed.model_time,unperturbed.x[:,0], "b")
3 # Python counts starting at 0
4 plt.plot(perturbed.model_time,perturbed.x[:,0], "g")
5 plt.ylabel(r"$\theta_1$") # use LaTeX for label
6 plt.subplot(2,1,2)
7 plt.plot(unperturbed.model_time,unperturbed.x[:,1], "b")
8 plt.plot(perturbed.model_time,perturbed.x[:,1], "g")
9 plt.ylabel(r"$\theta_2$")
10 plt.show()

```

Listing 8: Python

1.4 An ensemble of model runs

Next, we want to create an ensemble of model runs all with slightly different initial conditions. You can do this in a number of steps:

- First create the input file `simulation_ensemble.oda` based on `simulation_unperturbed.oda`. Change the algorithm and the configuration of the algorithm.
hint: the algorithm is called `org.openda.algorithms.kalmanFilter.SequentialEnsembleSimulation`.
- Create a configuration file for the Ensemble algorithm (e.g. named `algorithm/SequentialEnsembleSimulation.xml`) with the following content:

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <sequentialAlgorithm>

```

```

3   <analysisTimes type="fromObservationTimes"></
    analysisTimes>
4   <ensembleSize>5</ensembleSize>
5   <ensembleModel stochParameter="false"
6                   stochForcing="false"
7                   stochInit="true" />
8 </sequentialAlgorithm>

```

Listing 9: XML-input for sequentialAlgorithm

Hint: do not forget to reference `algorithm/SequentialEnsembleSimulation.xml` in `simulation_ensemble.oda` and do not forget to give a different name to the output files.

- Run the new configuration with OpenDA.
- make a plot of the first and second variable of the five ensemble members in a single time-series plot

```

1 [t,ens]=load_ensemble('simulation_ensemble_results');
2 ens_th1=reshape(ens(1,:,:),size(ens,2),size(ens,3));
3 ens_th2=reshape(ens(2,:,:),size(ens,2),size(ens,3));
4 clf; subplot(2,1,1);
5 plot(t(2:end),ens_th1);
6 ylabel('\theta_1');
7 subplot(2,1,2);
8 plot(t(2:end),ens_th2);
9 ylabel('\theta_2');
10 xlabel('time');

```

Listing 10: Matlab

```

1 import ensemble
2 import simulation_ensemble_results as res
3 (t,ens)=ensemble.reshape_ensemble(res)
4 ens1=ens[:,0,:] #note we start counting at 0
5 ens2=ens[:,1,:]
6 plt.subplot(2,1,1)
7 plt.plot(t[1:],ens1,"b")
8 plt.ylabel(r"$\theta_1$")
9 plt.subplot(2,1,2)
10 plt.plot(t[1:],ens2,"b")
11 plt.ylabel(r"$\theta_2$")
12 plt.show()

```

Listing 11: Python

- Observations of θ_1 and θ_2 are available as well. Make a plot of the observations together with the simulation results.

```

1 [t,unperturbed,tobs,obs]= ...
2 load_results('simulation_unperturbed_results');
3 subplot(2,1,1);
4 plot(t,unperturbed(1,:), 'b-');
5 hold on
6 plot(tobs,obs(1,:), 'k+');
7 hold off
8 ylabel('\theta_1');
9 subplot(2,1,2);
10 plot(t,unperturbed(2,:), 'b-');
11 hold on
12 plot(tobs,obs(2,:), 'k+');
13 hold off
14 ylabel('\theta_2');
15 xlabel('time');

```

Listing 12: Matlab

```

1 import simulation_unperturbed_results as unperturbed
2 plt.subplot(2,1,1)
3 plt.plot(unperturbed.model_time,unperturbed.x[:,0], "b")
4 plt.plot(unperturbed.analysis_time,unperturbed.obs[:,0],
5          "k+")
6 plt.ylabel(r"$\theta_1$")
7 plt.subplot(2,1,2)
8 plt.plot(unperturbed.model_time,unperturbed.x[:,1], "b")
9 plt.plot(unperturbed.analysis_time,unperturbed.obs[:,1],
10          "k+")
11 plt.ylabel(r"$\theta_2$")
12 plt.show()

```

Listing 13: Python

We can see that although our simulation starts on the right track, it quickly diverges from the observations. The aim of the Ensemble Kalman filter or data-assimilation in general, is to keep the model on track.

2 Some basic properties of the EnKF with Lorenz 3 variable model

Directory: exercise_lorenz_3var_part2

In this exercise you will learn how to set up and run the EnKF method in OpenDA.

- Prepare the input files for a run with the EnKF method. Use the input files from `exercise_lorenz_3var_part1` as template. Hint: the Ensemble Kalman filter is called `org.openda.algorithms.kalmanFilter.EnKF`. The algorithm configuration file has the following content

```

1      <?xml version="1.0" encoding="UTF-8"?>
2      <EnkfConfig>
3          <ensembleSize>10</ensembleSize>
4          <ensembleModel stochParameter="false"
5                          stochForcing="false"
6                          stochInit="true" />
7      </EnkfConfig>
8

```

Listing 14: XML-input for EnKF algorithm

- Plot the ensemble mean of the first model variable and the observations. With some luck the solution should track the observations.
Tip: use the scripts `load_obs.m` and `load_ensemble.m` for reading the data into matlab (cf. Exercise1), or `load_ensemble.py` for python.
- Look at the observation input file of the StochObserver. The StochObserver does not only describe the observations but the accuracy as well. Can you make a new observation input file with similar observed values but with a 10 times larger standard deviation for the observation error. Tip: you can edit the file in OpenOffice or MS Excel or use the find and replace function of an advanced text editor.
- Repeat the run with EnKF but now for the new observations and plot the first variable of the ensemble means and the observations. What do you see and what is the reason for this behavior of the algorithm?
- The number of ensemble members controls the accuracy of the ensemble approximation. What happens if you increase the number to e.g. 100, or decrease it to 5? Use (initially) observations with a standard deviation of 5.0. Experiment as well with various standard deviations of the observations.

3 Steady-state Kalman Filter

Directory: `exercise_steady_state_filter`

In this section you will learn how to create and use a steady-state Kalman filter with OpenDA. The example model we use in this section is a 1-dimensional wave model:

$$\frac{\partial h}{\partial t} + D \frac{\partial v}{\partial x} = 0 \quad (9)$$

$$\frac{\partial v}{\partial t} + g \frac{\partial h}{\partial x} + c_f v = 0 \quad (10)$$

$$(11)$$

With $h(x, t)$ the (water) level above the reference plane, $v(x, t)$ the velocity, $D(x)$ the depth under the reference plane, g the gravitational acceleration c_f the friction coefficient and $x \in [0, L]$ the location. For our model we have selected the boundary values $v(x = L, t) = 0$ and $h(x = 0, t) = \frac{1}{5} \sin(2\pi t)$. An AR(1) model is defined on the left water level boundary.

- Look at the implementation of the model in `WaveStochModelInstance.java`, in the directory `simple_wave_model/java/src/org/openda/`. See how the state is defined and how the model is discretized. If you want you can compile the model using `ant build` as we will explain in exercise 6. However to make it easy for you, you will find the compiled version of this model, `simple_wave_model.jar` in the directory `simple_wave_model/bin`.
- The model represents a "user" model that is not part of the OpenDA distribution. Therefore you have to copy the model jar-file to the bin directory of your OpenDA installation. In this way OpenDA can find this model.
- Run the model (`waveSimulation.oda`) and visualize the model results (`plot_movie.m` or `plot_movie.py`). Do not forget to add the jar-file of the model to the `CLASSPATH` variable, or to copy the jar-file into the bin directory of your OpenDA version
- Adjust the input files in order to run the model with stochastic forcings.
- Generate water level observations from this stochastic run. We need observations at (approximately) $x = \frac{1}{4}L$, $x = \frac{1}{2}L$ and $x = \frac{3}{4}L$. You can use the script `generate_obs.m` for this task. We want to have observations at $t = 0.1, 0.2, \dots, 10.0$, (initially) select a standard deviation of 0.05.
- Run the Ensemble Kalman filter (`waveEnkf.oda`). This run will generate and write gain matrices at specified times. Find where and how this is specified in the input.
- Plot the columns of the gain matrices. (The script `plot_gains.m` or `plot_gains.py` plots the water level part of the gain matrices). What do these columns mean?
- (Re)generate the gain matrices using different numbers of ensembles. When you compare the gain matrices, what do you notice. Note: The algorithm will generate an enormous amount of output when you run the EnKF with a very large number of ensembles (e.g. 500). You can suppress the output by commenting out (or remove) the `resultWriter`-part of the oda-input file.
- Use the generated steady state gain matrices for a steady state Kalman run (`waveSteadystate.oda`). Compare the performance of:
 - a (non-stochastic) run without filtering,
 - an EnKF run with various numbers of ensembles (do not forget to reinstate the `resultWriter` if you have switched it off),
 - the various steady state gains.

you can use the scripts `plot_obs_sim.m`, `plot_obs_ens.m` and `plot_obs_steady.m` and similar routines for python.

- Generate (observations) gain matrices but now for only a single observation. Make sure that the observed values are exactly the same as in the 3 observation observer. Compare the columns of the 3-observation gain matrices to the single observation matrices. What is the main difference and why?

4 Exercise 4: A black box model - Calibration

A simple way to connect a model to OpenDA is by letting OpenDA access the input and output files of the model. OpenDA cannot directly understand the input and output files of an arbitrary model. Some code has to be written such that the black box model implementation of OpenDA can read and write these files. In this exercise you will learn how to connect an existing model to OpenDA assuming that all the input and output files of the model can indeed be accessed by OpenDA. The exercise focusses on the configuration of the black box wrapper in OpenDA.

The model describes the advection of two chemical substances. The first substance c_1 is emitted as a pollutant by a number sources. However, in this case this substance reacts with the oxygen in the air to form a more toxic substance c_2 . The model implements the following equations:

$$\frac{\partial c_1}{\partial t} + u \frac{\partial c_1}{\partial x} = -1/Tc_1 \quad (12)$$

$$\frac{\partial c_2}{\partial t} + u \frac{\partial c_2}{\partial x} = 1/Tc_1 \quad (13)$$

In the directory `exercise_4/original_model/` you will find:

1. model executable `reactive_pollution_model.py` (linux and mac) and `reactive_pollution_model.exe` (windows)
 2. `config.yaml` the model configuration file
 3. `forcings` the forcings of the model (injection of pollutant)
 4. `input` the initial model state
- Run the model, in the `original_model` directory from the command line, not using OpenDA.

The model generates the output files `c1_locA`, `c1_locB`, `c1_locC`, `c2_locA`, `c2_locB`, `c2_locC` with timeseries of substance c_1 and c_2 at three predefined locations in the model. The folders `maps` contains output files with the concentration of c_1 and c_2 on each grid point at specified times. The folder `restart` contains files that allows the model to restart; continue the computations from the point where a restart file has been written.

- Investigate the input and output files of the model

- generate a movie by running the script `plot_movie_orig.py` script from the `exercise_4 (!)` directory. This allows you to study the behaviour of the model.

4.1 Wrapper configuration files

The input and outputfiles of this model are all easy to interpret ASCII files. Therefore we do not need model specific routines to couple this model to OpenDA.

When you couple an arbitrary model to OpenDA and you want to use the black box coupler there are two approaches:

- write a pre and post processing script that translates the (relevant) model files into a more generic format that is already supported (e.g. ASCII or NetCDF).
- write your own adapter in java (dataobject) to read and write the model input and output files.

A black box wrapper configuration usually consists of three xml files. For our pollution model these files are:

1. `polluteWrapper.xml`: This file specifies how OpenDA can run the model, which input and output files are involved and dataobjects are used to interpret the model files.

This file consists of the parts:

- **aliasDefinitions**: This is a list of strings that can be aliased in the other xml files. This helps to make the wrapperxml-file more generic. E.g. the alias definition `%outputFile%` can be used to refer to the output file of the model, without having to know the actual name of that output file.
Note the special alias definition `%instanceNumber%`. This will be replaced internally at runtime with the member number of each created model instance.
- **run**: the specification of what commands need to be executed when the model is run.
- **inputOutput**: the list of 'input/output objects', usually files, that are used to access the model, i.e. to adjust the model's input, and to retrieve the model's results. For each 'ioObject' one must specify:
 - the java class that handles the reading from and/or writing to the file
 - the identifier of the ioObject, so that the model configuration file can refer to it when specifying the model variables that can be accessed by OpenDA, the so called 'exchange items' (see below)
 - optionally, the arguments that are needed to initialize the ioObject, i.e. to open the file.

2. `polluteModel.xml`: This is the main specification of the (deterministic) model. It contains the following elements:

- **wrapperConfig**: A reference to the wrapper config file mentioned above.

- **aliasValues:** The actual values to be used for the aliases defined in the wrapper config file. For instance the `%configFile%` alias is set to the value `"config.yaml"`.
 - **timeInfoExchangeItems:** The name of the model variables (the 'exchange items') that can be accessed to modify the start and end time of the period to that the model should compute to propagate itself to the next analysis time.
 - **exchangeItems:** The model variables that are allowed to be accessed by OpenDA, for instance parameters, boundary conditions, and computed values at certain locations. Each variable exchange item consists of its id, the ioObject that contains the item, and the 'element name', the name of the exchange item in the ioObject.
3. `polluteStochModel.xml`: This is the specification of the stochastic model. It contains of two parts:
- **modelConfig:** A reference to the deterministic model configuration file mentioned above `polluteModel.xml`.
 - **vectorSpecification:** The specification of the vectors that will be accessed by the OpenDA algorithm. These vectors are grouped in two parts:
 - The state that is manipulated by an OpenDA filtering algorithm, i.e. the state of the model combined with the noise model(s).
 - The so called predictions, i.e. the values on observation locations as computed by the model.

Start with a single OpenDA-run to understand where the model results appear for this configuration:

- Have a look at the files `polluteWrapper.xml`, `polluteModel.xml` and `polluteStochModel.xml`, and recognize the various items mentioned above. Start the OpenDA GUI from the `public/bin` directory and run the model by using the `Simulation.oda` configuration. Note that the actual model results are available in the directory where the black box wrapper has let the model perform its computation: `work/work0`.

Directory: `exercise_black_box_calibration_pollution`

In this exercise, we will calibrate the value of the reaction-rate constant. The algorithm used in this example is the Dud (which stands for Doesn't Use Derivative).

- Have a look at the `Dud.oda` and the configuration files it refers to. Run it from the OpenDA GUI and have a look at the results. What could you do to improve the results?
- Figure out where to change the control parameters for the calibration procedure and play around with the settings to improve your results.

Calibration runs normally take longer than a few minutes. In that case, it becomes convenient to be able to restart from a previous run.

- Adapt the configuration in such a way that you are able to restart the `Dud.oda` from the result of a previous run.

5 Exercise 5: A black box model - Filtering

5.1 Sequential simulation

We will first run our pollution model from OpenDA using the SequentialSimulation algorithm. This run is exactly the same as running the model outside openda. The difference is however that we provide a set of observations and run the model and restart the model between the observation times. Output will be available at the end in the generic OpenDa format that allows us to compare the model results with the available observations of the system.

- Run the model within OpenDA by using the `SequentialSimulation.oda` configuration. This will create the result file `sequentialSimulation.results.py`. Use the script `plot_movie.seq.py` to visualize the simulation results. The script `plot_obs.seq.py` shows the difference in time between the model results (prediction) and observed values of the system

5.2 Sequential ensemble simulation

The next step is running an ensemble of simulations. In this case we consider our main source of uncertainty the injection of pollutant c1 in the model. Similar to the sequential simulation we do not assimilate any data (yet).

- Run an ensemble forecast model by using the `SequentialEnsembleSimulation.oda` configuration. On which variable does the algorithm impose stochastic forcing? Have a look at the `work` directory, and note that the black box wrapper created the required ensemble members by repeatedly copying the template directory `stochModel/input` to `output/work<N>`.
- Compare the result between the mean of the ensemble and the results from `SequentialSimulation.oda`. Note the differences. You can use the script `plot_movie_enssim.py`.

5.3 parallel computing

Running the ensembles takes a lot of time, especially starting the model takes quite some time compared to the actual computation time. Most computers have multiple cores which we can use to propagate multiple ensemble members forward in time simultaneously.

- compare the configurations `SequentialEnsembleSimulation.oda` and `enkf.oda` which uses parallel propagation of ensemble members. Set the number of simultaneous models that corresponds to the number of cores on your computer (`maxThreads`).

5.4 ensemble kalman filter

Now let us have a look at the configuration for performing OpenDA's Ensemble Kalman Filtering on our black box model, using a twin experiment as an

example. The model has been run with the 'real' values (time dependent) for the concentrations for substance 1 as disposed by factory 1 and factory 2. This 'truth' stored in the directory `truthmodel`, and the results of that run have been used to generate observation time series at the output locations. These time series (with some noise added) have been copied to the `stochObserver` directory to serve as observations for the filtering run.

The filter run takes the original model as input, which actually is a perturbed version of the 'truth' model: the concentrations for substance 1 as disposed by factories have been flattened out to a constant value. The filter process should modify these values in such a way that the results resemble the truth as much as possible.

To do this the filter modifies the concentration at factory 2, and uses the observations downstream of factory 2 to optimize the forecast.

- Note that the same black box configuration is used for the sequential run, the sequential ensemble run, and for the EnKF run. Identify the part of the `polluteStochModel.xml` configuration that is used only by the EnKF run, and not by the others.
- Execute the Ensemble Kalman Filtering run by using the `EnKF.oda` configuration.
Check how good the run is performing, by analyzing to what extent the filter has adjusted the predictions towards the observation.
Note that the model output files in `stochModel/output/work0` only contains a few time steps. Can you explain why?
So to compare the observations with the predictions you have to use the result file produced by the EnKF algorithm which can be visualised using `plot_movie_enkf.py`.

Now let us extend the filtering process by incorporating also the concentration disposed by factory 1, and by including the observation locations downstream of factory 1.

- Make a copy of the involved config files, `EnKF.oda`, `parallel.xml`, `polluteStochModel.xml` and `timeSeriesFormatter.xml` (you could call them `EnKF2.oda`, `parallel2.xml` etc).
Adjust the files such that all references to the files are correct.
- Now adjust `polluteStochModel2.xml` and `timeSeriesFormatter2.xml` in such a way that the filtering process is extended as described above.
- Run the filtering process by using the `EnKF2.oda` configuration, and compare the results with the previous version of the filtering process.

6 Exercise 6: Writing your own toy model

Before you start:

In order to be able to compile your model you need to have a (current) version installed on your computer of:

- The Java Development Kit (JDK). You can download this from www.oracle.com¹
- Apache Ant, this is a command line tool we use for building your java code. You can download Ant from ant.apache.org.

In this exercise you will learn how to code your own model and use it in OpenDA. The directory `exercise_6` contains a template of the code for the 1-D advection model we will create in this exercise. The content of this directory is similar to the OpenDA directories you have seen in the previous exercises. The difference is that we will not use a model that is already part of the OpenDA distribution but instead our own model. The model code can be found in the directory `simple_advection_model`.

The model you will create is build as an extension of the OpenDA `simpleStochModelInstance`. This will simplify and reduce the amount of programming because a significant part of the implementation is already available. For more complex models you might need to implement all methods of the `IStochModelInstance` class.

In the directory `exercise_6/simple_advection_model/java/src/org/openda` you will find the two java source files `AdvectionStochModelFactory.java` and `AdvectionStochModelInstance.java`. The first file implements the `ModelFactory` class. The model factory is a class in OpenDA that is responsible for creating model instances (e.g. the members of an ensemble Kalman filter). The second file implements the model. This is the file you have to edit in this exercise.

- Consider the 1-dimensional advection model:

$$\frac{\partial c}{\partial t} = v \frac{\partial c}{\partial x} \quad (14)$$

where c typically describes the density of the particle being studied and u is the velocity. On the left boundary c is specified as $c_b(t) = 1 + \frac{1}{2}\sin(5\pi t)$. Discretize this model on the interval $x = [0..1]$ with velocity $v = 1$ using a 1st order upwind scheme on a grid of 51 points. The time step is chosen such that the courant number $\frac{v\Delta t}{\Delta x}$ is approximately 1.

- The deterministic model is extended into a stochastic model by adding a noise parameter ω on the left boundary. We use an AR(1) model to describe the noise.
- Code your model in `AdvectionStochModelInstance.java`. For inspiration, you will find in the same directory an implementation of the Lorenz model.
- You can compile your model by typing "ant build" in the directory `exercise_6/simple_advection_model`. This will create the file `bin/simple_advection_model.jar`.

¹Java Runtime Environment (JRE), which is installed on most computers is not sufficient since this will allow you to run java programs but it does not include the java compiler `javac` that is needed to create you own (parts of) programs

- Run the model. You can use file `advectionSimulation.oda`. In order to be able to run your model, java must be able to find the file `simple_advection_model.jar`. To accomplish this, you can copy your advection model jar-file to the bin-directory of your OpenDA installation or add the full path of your jar-file (`simple_advection_model.jar`) to the java class path variable `CLASSPATH`.
By default, the Windows scripts `oda_run_gui.bat` and `oda_run_batch.bat` use the JRE environment that is provided with OpenDA. If this JRE is incompatible with your JDK installation (Error: Unsupported major.minor version 51.0), use `oda_run_batch.bat <inputfile> -jre "location of your JDK"` to overrule the default JRE.
- Use the script `plot_movie.m` to visualize the model results. You will see that the model suffers from numerical diffusion. You can solve this by using a second order upwind method but this is not necessary for this exercise.
- Create an ensemble of model simulations and study the model uncertainty in space and time.
- The provided observation file `observations.csv` does not contain observations that correspond to the advection model. Create your own observation file for the locations $x = 0.2$, $x = 0.5$ and $x = 0.7$. Use the model to create the values. Run the model with noise on the left boundary. Optionally generate some additional noise and add it to the generated observations. You can use the script `generate_obs` to simplify the creation of the observations file.
- Using your generated observations, setup an ensemble Kalman filter run. Experiment with various numbers of ensembles, different settings of the AR(1) model.
- Experiment with different intervals between the available observations. What do you observe. Is this behavior different from the Lorenz model? Evaluate the uncertainty of the estimates.
- Experiment with only assimilating the data from one of the three locations and use the other locations as validation. What do you observe. Do all the locations have the same impact? Explain the behavior you observe.
- Use the same data as generated before and use the Kalman filter now with different values of the system noise covariance and measurement noise covariance. Explain the behaviour that you observe.

7 Exercise 7: Localization

Directory: `exercise_localization`

In this exercise you will learn about localization techniques and how to use them in OpenDA. This exercise is inspired on the example model and experiments from "Impacts of localisation in the EnKF and EnOI: experiments with a

small model”, Peter R. Oke, Pavel Sakov and Stuart P. Corney, *Ocean Dynamics* (2007) 57: 32-45.

The model we use is a simple circular advection model

$$\frac{\partial a}{\partial t} + u \frac{\partial a}{\partial x} = 0 \quad (15)$$

where $u=1$ is the speed of advection, a is a model variable, t is time and x is a space ranging from 1 to 1000 with grid spacings of 1. The computational domain is periodic in x .

In this model there are two related variables a and b where b is initialised with a balance relationship:

$$b = 0.5 + 10 \frac{da}{dx} \quad (16)$$

and propagated with an advection model similar to the one for a , i.e.:

$$\frac{\partial b}{\partial t} + u \frac{\partial b}{\partial x} = 0 \quad (17)$$

Since a and b are propagated with the same flow, the balance relationship will remain valid also for $t > 0$. The relationship between a and b is motivated by the geostrophic balance relationship between pressure (a) and velocity (b) in oceanographic and atmospheric applications.

In this experiment we will only observe and assimilate a and investigate how both a as b are updated. The ensemble is carefully constructed in order to have the right statistics. The initial ensembles are generated off line and they will be read when the model is initialised in OpenDA.

- Investigate the script `generate_ensemble.py` and figure out how the ensembles are generated.
- Run python script `generate_ensemble.py` to generate ensembles, observations and true state for a 25, 50 and 100 ensemble experiment.
- Run the experiment for 50 ensemble members (`enkf_50.oda`).
- The variables a , b can be compared to the true state using the python script `plot_results.py`.
- Run the experiment for 25 ensembles, copy the script `plot_results.py` to e.g. `plot_results_25.py` and adjust it in order to read the results from `enkf25_results.py`. (change 2nd line of the `plot_results.py` script. You will see that the 25 ensemble run is not able to improve the model.
- Create input to run a 100 ensemble experiment. Note: do not forget to change the name of the output file (section `resultWriter`) to avoid that your previous generated results are overwritten.
- Run an experiment with 25 ensembles with localization (`enkf_25_loc.oda`) and generate the plots.
- The results (for 25 ensembles) with localization should look better than the the experiment without localization.

- Investigate whether the relation between a and b is violated by the various experiments. You can use the script `check_balance.py`.
- Try changing the localization radius (initial value is 50) and see how the performance of the algorithms changes (both for results as balance between a and b). You can plot the localization weight functions for each observation location (`rho_0`, `rho_1`, `rho_2` and `rho_3`) as well.