

Contents

- Project Status
- HarmonIT & OpenMI
- OpenMI Terms explained
- Models, quantities & links
- Running models & exchanging data
- The OpenMI Environment
- Interview with Roger Moore
- How do I make my model OpenMI compliant
- Need a helping hand?
- Frequently Asked Questions
- Model migration report
- The HarmonIT Consortium

OpenMI: A new era in integrated water management

OpenMI stands for Open Modelling Interface and Environment. It aims to deliver a standardized way to link water related computational models that exchange data on a time basis. It thus enables process interaction being represented more accurately, compared to sequential linkages.

One of the major challenges of the moment is putting integrated water management into practice. It is challenging because it requires water managers, before selecting a policy, to try and anticipate all its consequences for the environment and society. This implies not only understanding how individual catchment processes work but also how they will interact one with another. For example, in deriving a programme of measures to bring a river up to 'good ecological status', the options might include changing farm subsidies, raising the price of fertiliser and requiring farmers to leave buffer strips between cropped fields and the river. All three policies could have important implications not only for farming but for other sections of society as well. To identify which policy will achieve good ecological status and be of most benefit to society as a whole, the manager will not only need to be able to predict the response of land and the river processes but also their interaction with each other and the economy.

To make these predictions, the water management community uses a huge range of models. However, few of them are designed to be coupled and exchange information. The OpenMI addresses this issue by providing a mechanism, in the form of a standard interface, which allows any model complying with the standard to request and receive data from any other compliant model. For example, a river model could obtain surface runoff data from a rainfall runoff model.

Project Status

The OpenMI is just entering its fourth year of development. A standard defining an interface that will allow models to request and receive data has been developed. To help model developers and users migrate new and existing models to the standard a software environment has been created. This provides a range of tools to aid the migration, linking and running of linked models. During 2005, the OpenMI will under go extensive testing. In the process, a wide range of existing models will be migrated and linked. In parallel with this activity, the HarmonIT team will be working to build an organisation to support the OpenMI into the future. Three training courses on model migration will be given in the first quarter of 2005.

HarmonIT & OpenMI

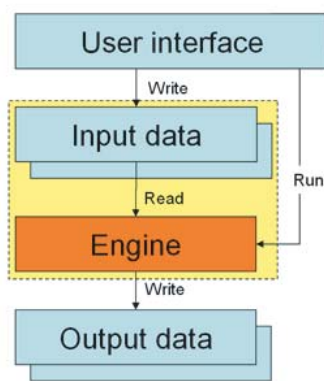
HarmonIT is the project that is developing the OpenMI. It is a highly ambitious four-year research project, whose aim is to develop an Open Modelling Interface and Environment, (the OpenMI) that will allow the linking of water related models. It involves 14 European organisations including three major commercial modelling software vendors. The project is led by the Centre of Ecology and Hydrology (UK) and is supported by the European Commission under the Fifth Framework Programme. It contributes to the implementation of the Key Action "Sustainable Management and Quality of Water" within the Energy, Environment and Sustainable Development. Contract no: EVK1-CT-2001-00090.

Want to know more?
Become a registered user now!

<http://www.openmi.org>

Just visit the project website and delve into a wealth of information on the OpenMI architecture, model linking and state-of-the art software technology. Keep in touch with the latest developments in integrated water management!

Model application



Some OpenMI terms explained...

The term **model application** encompasses all the parts of a modelling system, example applications being Mike11, RIBASIM or InfoWorks RS. Typically, such systems consist of a **user interface** and an **engine**. Usually, the engine is a generic representation of a process and comprises the calculations involved in modelling that process. The user supplies information through the user interface, enabling the user interface to find or generate input data for the engine. These data will describe a specific scenario in which the process is to be simulated. An example scenario might be the Rhine during a time of extreme rainfall. The user can run the engine by, for example, pressing a button in the user interface. This starts the engine which reads the input, performs calculations and finally outputs the results to files or displays.

When an **engine** has read its input it becomes a **model**. Hence, an engine representing the generic process of water flowing in an open channel becomes a model of the Rhine, once it has read in the data describing the channel network and the other boundary conditions, for example, rainfall data defining the scenario to be simulated. In other words, a model is an engine populated with data.

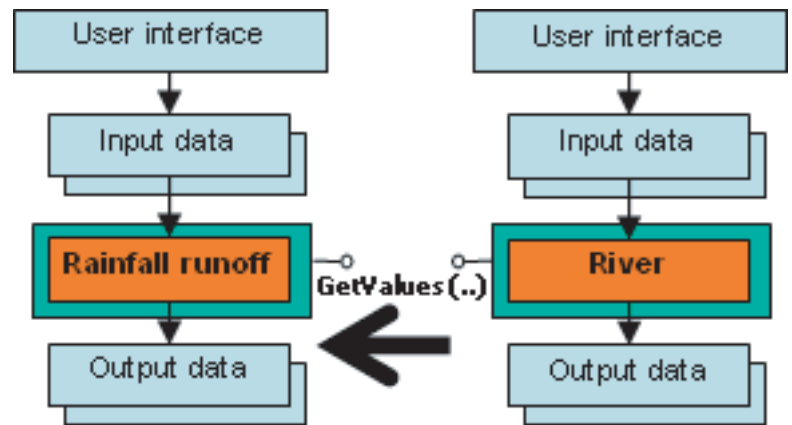
If the code of an engine can be instantiated separately and has a well-defined **interface** by which it can **accept** and **provide** data, then it becomes an **engine component**. The key to enabling models to exchange data lies in standardising the design of this interface. When an engine component implements such a standard interface, it becomes a **linkable component**.

OpenMI goes Open Source!


The OpenMI Standard, the code of the OpenMI Environment and the supporting documentation will be soon available as Open Source under the LGPL license. Find more about it by visiting the OpenMI website <http://www.openmi.org>

Models, quantities and links

A model is a simplification of the real world and different models simplify the world in different ways, for example, as vectors for rivers, grids for groundwater and the atmosphere and agents for modelling people and fish. What they have in common is that they all simulate how a set of **quantities**, e.g. the concentration of nitrate, will respond to a specified set of circumstances at a number of locations in the model which are referred to as **elements**. When two models are linked, a quantity in one model is matched to the equivalent quantity in the other model. Hence, 'runoff' in a rainfall runoff model might be matched to 'lateral inflow' in a river model. Similarly, the calculation points in the two models are also linked. Here though, it may well be necessary to match one element in one model to several elements in the other or vice versa. Such matchings are called **links**. To handle the aggregation and interpolation implicit in such links, the OpenMI environment provides a number of **data operations**.

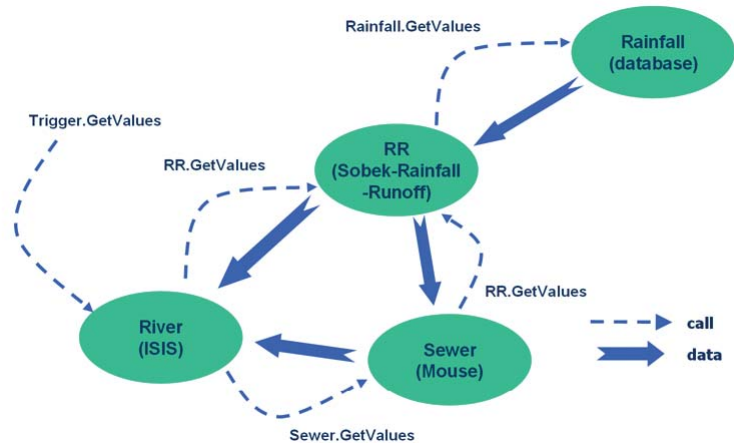


The figure above shows two engines that have been migrated. The important point to note is that the overall structure of the original applications containing the applications remained unchanged. They can still be run as before. The difference is that the option now exists, in this case, for the river model to obtain runoff information from the rainfall runoff model.



Running models and exchanging data

With the OpenMI, there is no framework into which models are plugged. Instead, one model is made the trigger. Starting this model causes the modelling process to begin. When the trigger model reaches a point where it needs data from another model, it uses the `GetValues` method to request information about the value of a particular quantity, at a particular point in time, for a particular link. The figure on the right shows a typical calling sequence between models.

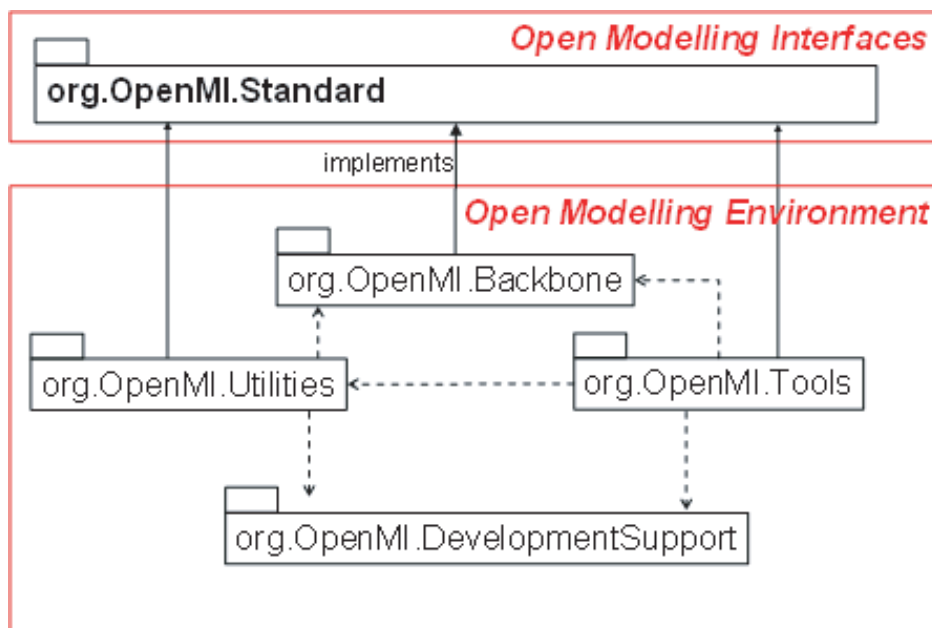


The OpenMI Environment

The interfaces of the OpenMI architecture, i.e. the Open Modelling Interfaces, are specified in the namespace `org.OpenMI.Standard`. Software components that implement and use these interfaces properly are called OpenMI compliant.

A default implementation of these interfaces is provided in the `org.OpenMI.Backbone` package. This Backbone is part of the OpenMI environment being developed by the HarmonIT project. In addition, the environment provides utilities to support wrapping of legacy code and to configure and deploy components (all part of the `org.OpenMI.Utilities` namespace) together with a general support package (`org.OpenMI.DevelopmentSupport`) which is not part of the OpenMI standard but is used by the environment and front-end tools to enable interaction with its users (the `org.OpenMI.Tools` namespace). Utilization of the OpenMI environment is not obligatory to develop OpenMI components.

Where wrapping is chosen as means of implementing of the system, the architecture imposes only few rules on how the wrapper should be implemented.




HarmonIT in HydroInformatics 2004



Pictured left to right: Peter Gijbers (WL|Delft Hydraulics), David Fortune (Wallingford Software), Roger Moore (CEH - Project Coordinator), Johan Van Assel (Aquafin), Flemming Damgaard Christensen, Jan Borge Gregersen (DHI Water & Environment), Isabella Tindall (CEH - Technical Manager) at the HydroInformatics 2004 Conference, Singapore

HydroInformatics is one of the main world forums for reporting advances in modelling and the HarmonIT project took full advantage of it to present its work on the OpenMI. In a special session devoted to the OpenMI, five papers covered all aspects of the architecture, the supporting software, model migration and the application of the OpenMI in decision support. The discussion was lively and extremely supportive of the approach adopted. Most gratifying was the fact that so many speakers, including the keynote speaker, referenced the OpenMI. Outside the meeting a large stand provided detailed information on the OpenMI standard and six leading members of the project were on hand to answer questions, of which there were many. HarmonIT flyers and OpenMI posters were distributed to all conference attendees. Dissemination is seen to be an important part of the project and, from this point of view, the conference was regarded as extremely successful, in that it enabled the project to reach many of the key players in the modelling field.



HarmonIT in iEMSs 2004

The original plan for an OpenMI workshop evolved into an OpenMI discussion forum attended by 30+ people. It was a lively affair from which a number of points emerged. It was clear that people were beginning to understand the purpose of the OpenMI and to see it as something that was potentially useful. There was much relief that the OpenMI would not be platform dependent and that there could be both Java and .NET versions. There were a number of questions about the US High Level Architecture (HLA) and whether or not this had been adopted in the US. From discussions outside the meeting, it would seem that as yet there has been little take up. It was explained that models running within a framework could use the OpenMI - they are not mutually exclusive. Representatives from a number of potential FP6 projects were present and the use of the OpenMI by these projects was discussed. It was clear that the OpenMI and some of the other current CatchMod projects would have much to offer FP6. Since the iEMSs meeting three of these projects have been successful, AquaStress, Newwater and Seamless and all three intend to use the OpenMI in their integrated modelling work.

Interview with Roger Moore

— How did you get involved?

Well, one wet afternoon in October 2000, the Director put his head round the door and asked if I was doing anything. As I was having a cup of tea at the time, I replied "No". Whereupon he said, "Great, could you put a bid into Brussels". Mistakenly, I thought he meant what he said. Having never worked in Europe, I had no idea that were the proposal to succeed, I had just committed myself to running a major international project for four years.

— What does being the Project Coordinator mean?

HarmonIT is slightly unusual in that three of the key partners are normally fierce commercial competitors. However, in this instance, they had identified something that they all needed but which none of them, acting on their own, had the resources to develop. Therefore, they had to collaborate and, hence, my first role has been to act as a neutral chairman.

With respect to the project as a whole, being Co-ordinator is about leadership and management. Fortunately, it doesn't mean having to have the ideas. In fact, not being involved in the ideas has been a great advantage, since I can focus on achieving the end without being distracted by the means. I have an excellent science team. My job is to make sure that they have the right resources. I also represent the project to the Commission and make sure that the world is aware of the projects results.

— How have you found the project so far?

I have been lucky in my career, in that most of the projects upon which I have worked have been great fun and very exciting. HarmonIT has been no exception. What makes it so enjoyable is that the people involved are enthusiastic, committed and mutually supportive when there is a problem. I believe that what has made it exciting has been the combination of commercial and academic partners. The original ideas came from one of the academic partners while the drive and motivation to translate them into reality has come from the commercial partners.

— What are the problems of directing the project consortium?

Far fewer than expected. During the bid process, I was regaled with many nightmare stories about the problems of persuading partners from all over Europe to work together, or as one person put it, even to work at all. However, I suspect problems are an attitude of mind, if you don't see them, you don't have them. Alternatively, you can always adopt the American approach and perceive them as opportunities.

— How do you envisage the future of the OpenMI?

It looks as though, by the end of HarmonIT, we will have a viable first version of the OpenMI. It won't be perfect and there will be lots of scope for doing things better. However, if it does work, and even if it doesn't, the OpenMI will be a remarkable European achievement. The great challenge ahead is to take the project from the research phase to a financially self sustaining Standard. To do this we propose to make the Standard Open Source, build a user community, raise the funds to bring the present work up to commercial standards and encourage organisations to build a range of goods and services around the OpenMI.

— Would you do it again?

Without a moment's hesitation. What other project enables you to meet and work with such a wide range of people, not just from Europe but across the world. It has been incredibly demanding but worth every minute.

Model migration report

HarmonIT commercial partners are currently migrating their state-of-the art models to the OpenMI standard and new OpenMI-compliant versions will be soon available in the market.

DHI - Water & Environment will make its OpenMI-compliant models available with its 2005 B release expected in December 2005. This will include:

- Mike Basin - conceptual water management model
- Mike 11 - Hydrodynamic river model
- Mike SHE - Ground water model
- Mike Urban - Sewer system model

WL|Delft Hydraulics' is at the following stage of model migration

- SOBEK-RR/CF/RTC is available in beta version (i.e. technically fully working but no helpdesk support yet) and will be included in the next release, expected in the second half of 2005.
- DelftFEWS (flood forecasting system): The OpenMI adaptor is expected to be operational this year .
- Work is ongoing on HYMOS (hydrological database), RIBASIM (water allocation modelling) and other WL|Delft Hydraulics software packages including Wanda and Delft3D

Wallingford Software has migrated the following modelling systems and they are currently being tested:

- InfoWorks RS - river basin modelling system (flow and water quality)
- Isis - river model (flow)
- InfoWorks CS - sewer modelling system (flow and water quality)
- SULIS - 3D lake and estuary model (flow and water quality)

The Dutch Institute for Inland Water Management and Waste Water Treatment (RIZA) has also migrated the following models:

- DM (water distribution model)
- Mozart (soil physical processes, unsaturated zone)
- Agricom (cost model)

The Center for Ecology & Hydrology in the UK is migrating CLASSIC, a hydrological continuous flow simulation model which runs at either a daily or hourly timestep. It is expected to be ready in the autumn of 2005.

How do I make my model OpenMI compliant?

OpenMI defines an interface to create direct access to a model. In order to make this possible, the original engine needs to be turned into an engine component and the engine component needs to implement an interface through which the data inside the component is accessible. OpenMI defines a standard interface to be implemented by engine components. When an engine component implements this interface it becomes an OpenMI compliant linkable component. A similar pattern can be applied for databases or other kinds of data sources. By turning them into components and implementing the OpenMI interface they become linkable components that provide direct access to its data at run time.

The easiest way to make a model compliant with OpenMI is to contain it in a suitable wrapper. The wrapper controls the runtime activity of pulling data across links. The HarmonIT project provides a number of "smart wrappers" that already contain most of the tedious (and difficult) tasks to be performed in the wrapper.

To become an OpenMI linkable component, a model has to:

- be able to expose information (what, where) to the outside world on the modelled variables which it can provide, or which it is able to accept;
submit to run-time control by an outside entity;
- be structured in that initialization is separate from computation. Boundary conditions must be collected in the computation phase and not during initialization;
- be able to provide the values of the modelled variables for the requested points in time and space;
- be able to respond to a request, even when the component itself is time independent; if such response requires data from another component, the component should be able to pass on the time as well in its own request.
- In the case some values in the returned value set are missing, this should be flagged in the value set.
- In the case that an entire value set is unavailable, an exception needs to be thrown. Be aware that such exceptions will stop the entire computation process and thus should be prevented whenever possible.

Need a helping hand?

At first hand it may seem a major challenge to turn a model engine into an OpenMI compliant LinkableComponent.

However it is not so difficult. The OpenMI provides guidelines for migration of models and a great number of software tools and utilities that will make migration easier. These tools and utilities can be used by anyone who is migrating a model but are not required in order to comply with the OpenMI standard.

For existing model engines, wrapping is the recommended technology for migration. If a model engine is, for instance, a numerical model programmed in Fortran this engine can be compiled into a dynamic link library (dll). If this dll is organized with a specific Win32API, an OpenMI utility wrapper class (SmartWrapper) can be used. This class will take care of all the bookkeeping associated with handling links, and all the interpolation in time and space.

It is estimated that migration of models will take between a few weeks to a few months of work, depending on how well the program code has been structured.

If you need further help in migrating your models, you can also attend the OpenMI Training Courses. Find more about them by visiting the website <http://www.harmonit.org/Workshops/training/>.

Do you need more information?

If you have any queries regarding the OpenMI Standard and Architecture, if you would like more help migrating your model or if you are interested in joining the OpenMI User Group, don't hesitate to contact the HarmonIT Consortium by e-mailing

openmi@openmi.org

Frequently Asked Questions

What is the OpenMI environment ?

The OpenMI environment is the implementation offered by the HarmonIT project of the OpenMI Standard interfaces. The environment includes utilities to wrap (model) engine components (including facilities for buffering, spatial mapping, temporal mapping), facilities to create, store and deploy your composition of linked models and front end tools to interact with the system and see what happens.

Is it necessary to buy a complicated integrated modelling system?

No. It is not necessary to buy a new integrated modelling system. You can just use your own models with their own database & visualisation systems, as long as they are wrapped to become OpenMI compliant.

The wrapping facilities offered by the OpenMI environment may help you if you want. But, if you want to buy a complicated integrated modelling system, do buy an OpenMI-compliant one !

Is there a choice of technology?

Yes. The interfaces are formulated in platform independent primitives. HarmonIT offers an extensive implementation of the OpenMI environment in C# running on .NET (including spatial utilities and front end) and a less extensive implementation in Java. Any software language can be used for computational cores, but may require wrapping in a .NET language or Java.

What is the ambition of OpenMI ?

The ambition of OpenMI is to become a standard for on-line memory based linking water related model components. However, its design accommodates a much wider domain than water only, as long as spatial and temporal variability are relevant. Although stochastic data can be exchanged, OpenMI has been designed with deterministic simulation models in mind.

Is it essential to use the OpenMI environment including its utilities and tools?

No. They are just there to help. You can use your own environment if you like, but your system must implement the org.OpenMI.Standard interface to be OpenMI compliant.

What is the difference between OpenMI and HarmonIT ?

HarmonIT (IT Frameworks) is the name of the EU 5FP co-funded R&D project that developed the first version of the OpenMI interfaces and environment. OpenMI is thus the deliverable of the project. (i.e. the interface specification and the implementation in the OpenMI environment).

When will the standard be ready for use?

Although the HarmonIT project finishes by the end of 2005, the most important deliverables will be completed around the summer of 2005.

The software release, including documentation, is expected in June 2005 and will be released under Lesser GPL license conditions (for libraries) and MIT license conditions (for executables). The release versions will be made available on www.openmi.org.

An open project workshop is scheduled for September 27, 2005 in Munich (Germany) to discuss the experiences on the migration to and application of OpenMI so far.

If I use the OpenMI environment offered by HarmonIT, how much time will it take to wrap a model?

This depends on your code. Our experience varies from as little as a week for a well structured engine (i.e. not yet an engine component), to several weeks for the more complicated ones. This guideline, based on using the OpenMI environment, should provide a quick start.

What will happen to OpenMI after the HarmonIT project is over?

We hope it will become an established standard that will be widely applied, developed and maintained in the long term.

The project team is still working to make this happen and all help suggestions are welcome.

OpenMI Training Courses

The HarmonIT partners are running a series of workshops/training courses in 2005 for individuals and organizations who would like to learn more about the OpenMI.

Three training courses have been held so far at different European venues and more are being planned in Europe and the US. To find out more about the Training Courses and to register online go to:

<http://www.harmonit.org/Workshops/training>

Who are we?

The Centre for Ecology and Hydrology	UK
Institute for Inland Water Management and Waste Water Treatment (RIZA)	Netherlands
DHI - Water & Environment	Denmark
WL Delft Hydraulics	Netherlands
HR Wallingford Group	UK
University of Dortmund	Germany
Water Research Institute of the Italian National Research Council	Italy
National Technical of University of Athens	Greece
Water Research Centre	UK
DHI Hydroinform a.s.	Czech Republic
Povodi Labe s.p.	Czech Republic
Hydroprojekt a.s.	Czech Republic
Alterra	Netherlands
Cemagref	France