

Bringing the OpenMI to Life

A layman's report



Background and aims of the OpenMI-LIFE project

This is the story of a project which set out to reduce the problem of 'unintended consequences'. It started in the water domain but came up with a highly imaginative solution that has wide application. The solution is called the OpenMI and it has the potential to open up enormous opportunities in any field, where people need to understand how one action might impact on another.

The work arose out of the Water Framework Directive (WFD) and, more particularly, its demand that managers think about the wider consequences of their plans before they implement them. While this is self-evidently sensible, history has taught us that it is actually very difficult to put into practice. The switch to bio-fuels and its unforeseen impact on basic food production is but the latest example of how easy it is for well intended actions to go horribly wrong. So, can we do better?

For many years managers and engineers have used models to help them find the best decisions. Thus, if they were considering deepening the approaches to a harbour, they would use a model of the river channel to try out different channel widening and deepening schemes and hence find the best solution. Today, it is no longer acceptable to solve the navigation problem in isolation. Dredging the channel may well have disastrous effects for the local fishermen or upset the balance of the area's ecosystems.

We have fish models and we have ecosystem models. However, for a long time the problem has been how to put them together in order to see if dredging the harbour entrance might damage the local fishing industry or wildlife. The difficulty is that the models come from different organizations, run on different computers and are often based on very different principles. At first sight the problem looks impossibly complex.

Take the question, 'What are the implications of climate change for the likely cost of flood damage?' How can you join something that models the climate to a river model and that to an economic model? In principle, however, the answer is simple. The economic model just needs to know how deep, how often and where properties will be flooded. This the hydraulic model can provide. To compute the pattern of water levels, the hydraulic model needs to know when and how much water will appear in the river channels and this can be derived from the rainfall patterns output by the climate model.

The problem of joining them together therefore comes down to one of how to pass data between different models. Often this is quite straightforward but it becomes more complicated when the models run at different time steps or represent the world in different ways. Rivers are often represented by simple lines but the ocean and the atmosphere are usually modelled in three dimensions.

Feedback loops can also cause difficulties. What happens in one model can affect what happens in another; for example, the chemistry of pond water affects plant growth and the plant growth affects the chemistry of the water, so the models have to be run side by side, exchanging data at each timestep.

A more complex problem that arises in many different contexts can be illustrated by river flow calculations. To calculate the water level in one river section, you need to know the flow into the section from the upstream river; the catch is that to calculate that flow you need to know the level in

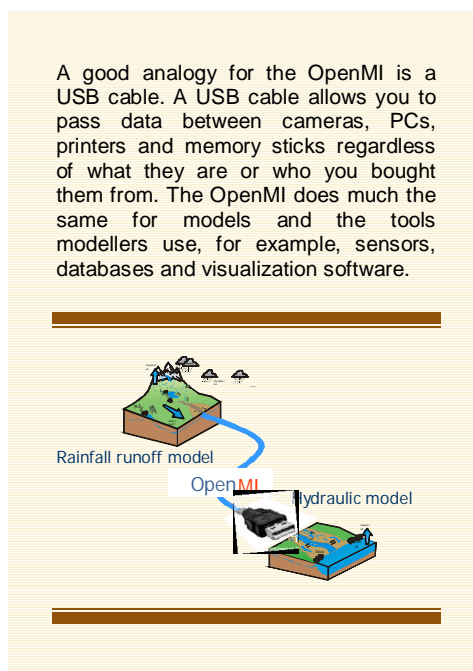


the original section. So, the data exchange mechanism must be able to handle a trial and error process that will enable the calculation process to converge on the answer. A further and mildly amusing question is what happens when your model fish gets eaten or a new one is born! There are plenty more such challenges.

To cut a long story short, a group of fourteen leading European companies, government research labs and universities, using their own and European funding, came up with a first solution to this difficult problem – the Open Modelling Interface standard also known as the OpenMI. They did so by looking at a large array of different models through half closed eyes and discovering that they all looked very much alike. Having discovered the shared properties of all models, it was then possible to find a way that any model could communicate with any other relevant model. The tricky bit was to do so in a way that could ultimately be used by anyone on anything from a home PC to the most powerful supercomputer.

At this point in the story the OpenMI was a research output. OpenMI-LIFE has been about testing, demonstrating and evaluating the OpenMI's use under operational conditions on real-world problems. It has also been about creating a long-term support organization, the OpenMI Association, to take the OpenMI forwards into the future. The LIFE project's final objective was to run a major dissemination campaign to raise awareness of the OpenMI around the world. In summary, the project's objectives were to:

- Prove it works
- Create a long-term support organization
- Tell the modelling world about it



OpenMI-Life

Proving the OpenMI

Currently joined up or 'integrated modelling' is an activity that is largely undertaken in research establishments. The aim of this part of the LIFE project was to find out if the OpenMI had made integrated modelling into a technique that government agencies and consultants could now use. Did they have the skills, did they have the equipment, was it useable, was it well enough supported and did it produce better results?

Two groups of water management authorities were found; one responsible for managing the international river basin of the Scheldt on the Franco-Belgian-Dutch border and the other responsible for advising the Greek authorities managing the Pinios basin in Greece. Between them they undertook seven case studies. Each was a real operational problem and required a solution that took account of how several processes might interact with each other. Two are briefly described below, one from each basin. A description of each will be given first and then the tests and results will be described.



The Scheldt example

The river Scheldt has its source in France and flows through Belgium, finally entering the Netherlands before flowing into the North Sea – see Figure 1 and Figure 2. It covers an area of 36,416 km², which makes it one of the smaller European river basin districts. The Flemish and Dutch parts of the Scheldt are mainly flat. Along the coast there are polders – areas of reclaimed land. It is one of Europe's most densely populated and industrialised areas with a population of 13 million. The strongest industrial sectors are the food and the metallurgical industries. Other important sectors are the chemical and the textile industries. Agriculture accounts for only 60% of the land use and comprises both arable and cattle farming. The main woodlands are to be found in the north-east, where the soil is poor and unsuited for agriculture.



Figure 1 The Scheldt estuary



Figure 2 The international Scheldt river basin

This small international basin is therefore an ideal place for integrated river management trials; within it can be found examples of nearly every management challenge. Most require many organizations to collaborate in order to find a solution. Each has its own models to aid decision making, though, so far, it has proved difficult both for technical and institutional reasons to bring them together. The participating partners were interested to see if the OpenMI would lessen these problems.

studies evaluated integrated modelling and the OpenMI as a way of improving the effectiveness of capital works and the operational procedures proposed by sewerage and river managers. Flows in the sewers and rivers are often interdependent; changes to the sewer network or to the operating policy can impact river flows, and hence the flood risk. Sewer network changes can also alter river water quality. Similarly, changes to the river control structures or their operating policies, usually for better flood management, change the water level in the rivers, which, in turn, can affect the levels in the sewers and possibly change the frequency of overflow spills and hence pollution events.

The case study linked models of the Leuven sewer system to a model of the Dijle River under a range of operating scenarios – see Figure 3. It was a demanding test because of the large number of points at which the sewers connect to the river system.

The example taken from the Scheldt case

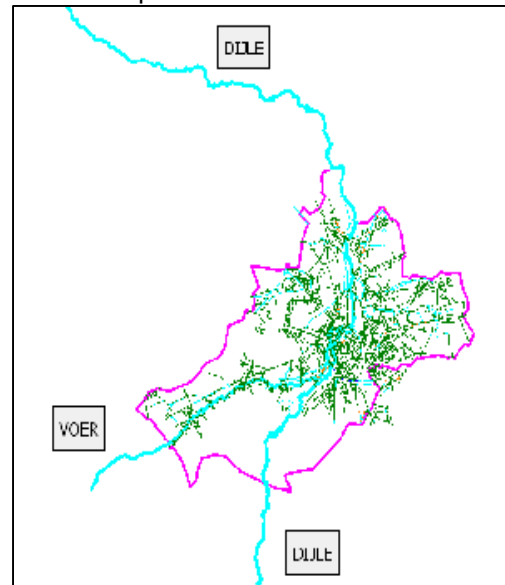


Figure 3 The River Dijle and the sewer network in the area of Leuven



The Pinios example

The Pinios basin is situated in the region of Thessaly. It has an area of 14,000 km² and a population of 760,000. Agriculture is the major source of income and employment. The cultivated area is ~4,000 km² and the basin represents the main agricultural production area of Greece. Crops produced in the plain area are cotton, wheat and maize, and apples, apricots, cherries and olives; grapes are cultivated at the foothills of the eastern mountains. Thessaly is not heavily industrialized, food processing being the most prominent activity. Other important sources of income include breeding, fisheries, forestry and tourism.

The sustainability of Thessaly's natural and built environment depends entirely on the quantity and quality of water in the Pinios basin. Available surface water and groundwater resources have to satisfy the needs of farmers, industry and local municipalities and, in addition, support the local environmental ecosystems. Changing water availability may have a wide range of impacts from wildlife disturbance to limiting economic growth and from inhibiting land use development to changing human mobility patterns. Current simulation models developed by a variety of interested parties provide "limited snapshots" of the water resources situation but do not show the "bigger picture". It is hoped that integrated modelling will enable those "snapshots" to be linked together. Managers can then assess the wider implications of proposed policies before implementing them.

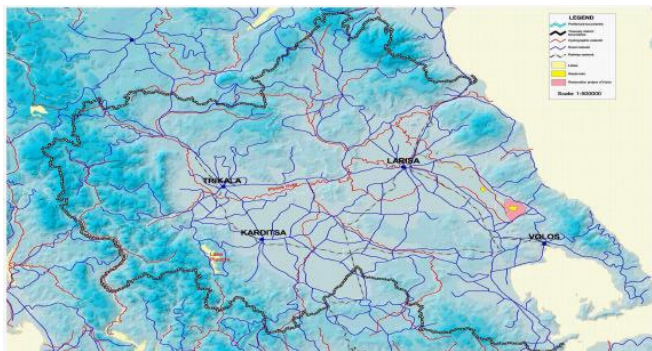


Figure 4 The Pinios river basin



Figure 5 Smokovo reservoir

The Pinios case study assessed the impact of climate change on the reliability of Smokovo reservoir (Figure 5) by linking a model of how rainfall turns into runoff to a reservoir management model. Located approximately 20km to the east of the Smokovo reservoir is Lake Xiniada (~36.6 km²). Today it serves mainly local irrigation needs. During the second phase of the study, the benefits of a possible restoration of the drained Lake Xiniada were considered by linking two reservoir management models to the rainfall runoff model and to an additional "Rule" component.

More comprehensive descriptions of the case studies can be found at www.openmi-life.org.

Specific points tested

Together the seven use cases tested:

- The benefits of integrated modelling – does, for example, treating the sewer/river system as a single entity produce greater insights into how it works than treating it as two separate entities; is it now feasible for organizations outside academia to use integrated modelling; will it one day be possible to search out ‘unintended consequences’ with greater success than now?
- The OpenMI’s ability to link models from different domains, e.g. river and sewer engineering
- The OpenMI’s ability to link models from the same domain but from different suppliers, e.g. two hydraulic models supplied by different companies
- The support organization’s ability to help solve integration problems when the models came from different suppliers; was it difficult to find out where the problem lay – model A, model B or in the interface?
- The OpenMI’s ability to link models at a large number of nodes
- The impact of linking on performance – do linked models run more slowly?
- The scientific validity of linking different models – does linking a hydraulic model to an economic model produce meaningful results?
- Sensitivity testing – did having a standard interface make it easier to replace one model by another equivalent model and hence make it easier to discover whether the choice of model influenced the answer?
- The OpenMI’s ability to link models whose geographic extent and spatial representation of the real world were different
- The OpenMI’s ability to link models running with different time steps, e.g. one runs on a daily timestep and the other on a 5-minute time step

Results

An early discovery in addressing problems requiring an integrated approach to their management was that there is an overwhelming human temptation for modellers to jump immediately to the technical problem of how to join models together before the issues have been defined and hence the need for integrated modelling understood. An important project output is therefore a method for approaching a problem requiring integrated modelling. Prior to the project, the expectation was that it would be possible to re-use existing proven models. Experience from OpenMI-LIFE shows that once models are linked, it becomes easy to create situations where models are pushed beyond their design limits. Much rewriting will be required. It is also clear that for integrated modelling to take off, a large pool of easily linkable models must be created. A great deal of research and development into areas such as how to calibrate and validate linked models will be required as well. The resolution of these and many other topics will require significant resources. The conclusion of the OpenMI-Life project is that the way most likely to mobilize the creativity and initiative needed is to follow the ‘open source’ route and create a collaborative community of practice.

Somewhat contrary to expectations, the OpenMI proved to be considerably ahead of the requirement in terms of its capabilities. Its use was within the capacity of the agencies taking part but it was clear that much could and needed to be done to make its use much simpler and ultimately invisible to the user. As a result of the study, it was possible to see that one day it would only be necessary to state the problem and it should be possible to find the models and data, link and run them. However, there is a huge amount of research between here and there. At present someone has to have thought of an outcome to be able to assess its impact. It is further possible to conceive of systems that could seek out the wider implications of a particular course of action. Coming back to the present, the next stage in the development of the OpenMI is to ensure that it can connect models running on PCs to models running in other environments, especially the high performance environment. Version 2.0 is



addressing this issue. However, the important conclusion from the project's point of view is that the OpenMI works.

Creating the OpenMI Association

It was apparent from the outset that organizations outside the research world would not take up the OpenMI if it was not supported. One of the first acts of OpenMI-LIFE was therefore to create the OpenMI Association as a legal entity to take ownership of and to promote the OpenMI. Its main task is the support and future development of the standard. It is an open, international, independent and democratic organization.

The Association's work is done by three groups. An Executive Committee defines the aims and strategy. Technical and Dissemination committees implement the strategy and provide vital information and feedback. For further information please go to www.openmi.org.

Telling the world

Although the OpenMI is generic and can link almost any appropriate models together, its origins lay in water and therefore that was the natural initial focus for the OpenMI Association's promotion work. However, that area was soon expanded to cover the land and water environments. While the internet has global reach, the vast majority of models come from Europe and the US and our promotional efforts have been focused there. Our aim is to convince the key players that the OpenMI works and that integrated modelling has become a usable technology. OpenMI-LIFE and now many other projects have provided us with credible demonstrations of the former and OpenMI-LIFE of the latter. The main means of communication is the internet but papers and conferences are still important. Video conferencing has yet to replace face-to-face meetings when it comes to thrashing out ideas.

Key achievements of OpenMI-LIFE

The key achievements of OpenMI-LIFE are:

- The demonstration and evaluation of the OpenMI under seven different sets of operational conditions in the Scheldt and Pinios basins
- The establishment of the OpenMI Association to take over the long-term ownership, support and development of the OpenMI
- The development of version 2.0 of the OpenMI which is out for international review at the time of writing – it simplifies the OpenMI and paves the way for linking to models running in the high-performance computing environment
- The dissemination of information about and the promotion of the OpenMI on a worldwide basis
- The gaining of worldwide recognition for the OpenMI by major organizations such as the US Environmental Protection Agency (US EPA), the US Department of Agriculture (US DA), the US Army Corps of Engineers, Consortium of Universities for the Advancement of Hydrologic Science Inc. (CUAHSI) and Community Surface Dynamics Modelling System (CSDMS)
- The commercial take-up of the OpenMI by leading European model developers
- A vastly increased understanding of the potential of integrated modelling
- An understanding of the problems that must be solved to fully unlock and make that potential accessible.



- A strategy for the forward development, support and funding of the OpenMI and the Association.

Technical and commercial application

Perhaps, the most convincing indicator of the success of the OpenMI-LIFE project is that all the commercial partners are now firmly committed to the use of the OpenMI in their commercial products. There has also been significant take up of the OpenMI in other EU research projects, for example, NitroEurope, Sensor, Seamless, Evoltree and EFORWOOD. A list of known OpenMI-compliant models can be found at www.OpenMI.org.

Where the OpenMI is probably achieving its greatest recognition is in the USA. The US Department of Agriculture has made its Open Modelling System (OMS) OpenMI-compliant and the US EPA recently announced that its new integrated modelling framework would also be compliant. Two widely-used hydraulic and groundwater models, HEC-RAS and MODFLOW, are being made linkable by the US Army Corps of Engineers.

Two National Science Foundation projects are contributing in a major way to the development of the OpenMI, one in the high-performance computing area and the other in the development of web services.

The emergence of modelling platforms where modellers can explore the possibilities of integrated modelling is an interesting development which has been greatly facilitated by the OpenMI. Two new ones have emerged during the LIFE project, one from HR Wallingford and one from Deltares.

Environmental benefits

Although the project provides no direct environmental benefits, the indirect benefits are potentially huge. Understanding how processes interact and being able to predict the result of interactions is central to our understanding of how the earth system works as a whole. Without such understanding, we will not be able to respond to environmental change, whether it is the result of man's activity or natural events. For the foreseeable future, models will have a key role to play in increasing our understanding and ability to predict. Since, self evidently, it is not feasible to build a model of everything, it will be necessary to link new and existing models of different aspects of global processes. It is here that the benefit of a standard such as the OpenMI will be found. The return on the investment in developing individual models can be hugely increased when they become linkable to other models.

Transferability

The OpenMI was designed to be generic. There is no aspect of it which is domain-specific. Care has also been taken to ensure that it is not based on any particular modelling concept. There are no geographical constraints on the use of the OpenMI. Version 2.0 and the move towards making models into web services should reduce or eliminate the problems of linking models running on different computing systems.

These actions should ensure the transferability of the OpenMI from the water domain where it started to other domains. There are strong reasons for believing that a high degree of success has been achieved here. The OpenMI has already been applied to economic, agricultural, forestry and chemical quality models. And projects are in progress that will apply it to the study of ecosystem services. It has also been used for linking to databases. There is no theoretical reason why it should not move from these domains into others such as transport, energy or health. Each of these would represent a small logical progression away from water. To this end the strategy will be to promote demonstration projects in as many domains as practicable.



Conclusions

New technologies generally emerge in the relatively sheltered environment of the research laboratory. OpenMI-LIFE was designed to find out if the OpenMI and integrated modelling could be applied in the outside world. Its aim was to seek out the problems and find out if they could be solved now or in the future. Many problems were found – not only technical and scientific but also institutional ones. At times it was hard going but that was the price of having the excitement of being an early adopter. The reward was finding that the problems could be overcome; there were no show stoppers. The OpenMI does stand up outside academia. That is not to say that the present version is the ultimate answer; it isn't. In ten years' time, it or its successors will be light years ahead compared to the OpenMI today.

The important conclusions relate more to integrated modelling than to the OpenMI. The OpenMI at this moment is probably functionally ahead of the current operational world's requirement but the future requirements are foreseeable. This is why Version 2.0 is under development now. In particular, it addresses the needs of modellers working in the high-performance computing environment and the need to be able to join their models to models in other environments, for example, .NET or Java. It also looks forward to models and data being provided as linkable web services.

By making it relatively easy to join models together, the OpenMI and the LIFE project have shown that integrated modelling is now a tool that an agency with good technical and scientific back-up can apply. What the model developers have learned is that integrated modelling will have to be made very much easier to use, if it is to become accessible first to any agency and ultimately the public. The range of linkable models and other components will have to reach critical mass. Finding and linking models must become a process of which most users need never be aware. On the modelling side, it is clear there are many issues to be addressed. When models are linked it is very easy to create situations where one or other is taken outside its design limits. While it is easy to join sewer models to river models, the question "How do sewer flows impact river flows" cannot be answered until the sewer and river engineers reconcile the way they assess the carrying capacities of rivers and sewers; presently, one uses design storms while the other takes a risk-based approach. This problem frequently occurs when models are linked across disciplinary boundaries.

So, the outcome of OpenMI-LIFE is that the OpenMI has been shown to be a viable model-linking standard and has been shown to work under arduous conditions. The project has created and left in place a support organization with a committed membership. A vigorous dissemination campaign has ensured the world knows about it and as a result there is strong support for the OpenMI outside Europe and especially in the US. The reward at the end of OpenMI-LIFE was to see its work being taken up by two US government agencies, the US DA and the US EPA, and by leading European commercial developers.

The project partners

The project partners were: The Natural Environment Research Council, UK (NERC); DHI, DK (DHI); Deltares, NL; MWH Soft (formerly Wallingford Software (WSL)), UK (MWH); the Flemish Environment Agency, BE (VMM and VMM-AWA); Flanders Hydraulics, BE (FH); Aquafin BE; University of Liege, BE (Ulg); the National Technical University of Athens, GR (NTUA); and the University of Thessaly, GR (UTH).

Acknowledgement

The OpenMI Association's members would like to acknowledge the contribution of the European Commission in co-funding the HarmonIT and OpenMI-LIFE projects. In particular, we would like to thank the Commission's staff for their sustained encouragement and support over many years.

Further information

Further information on OpenMI-LIFE can be found on the project website, www.OpenMI-Life.org.

Information on the OpenMI Association and the Open Modelling Interface (OpenMI) can be found on www.openmi.org.

