
OpenMI 'in a nutshell'

OpenMI Association, Technical Committee

January 2010

1. Introduction	1
2. Standard	2
2.1. Definitions	4
2.1.1. Model	4
2.1.2. Value	4
2.1.3. Spatial	5
2.2. Composition building	5
2.3. Runtime	6
2.3.1. Pull driven	6
2.3.2. Loop driven	7
2.4. Compliance	7
3. Tools and Utilities	8
4. Chronology	9
5. References	10

This article, updates an earlier paper [1], and contains a brief overview of OpenMI, a software standard which facilitates the linking of software based applications to realise integrated modelling.

OpenMI was originally developed for integrated water resource management, but its application has been extended to other domains of environmental management, such as agriculture and economics.

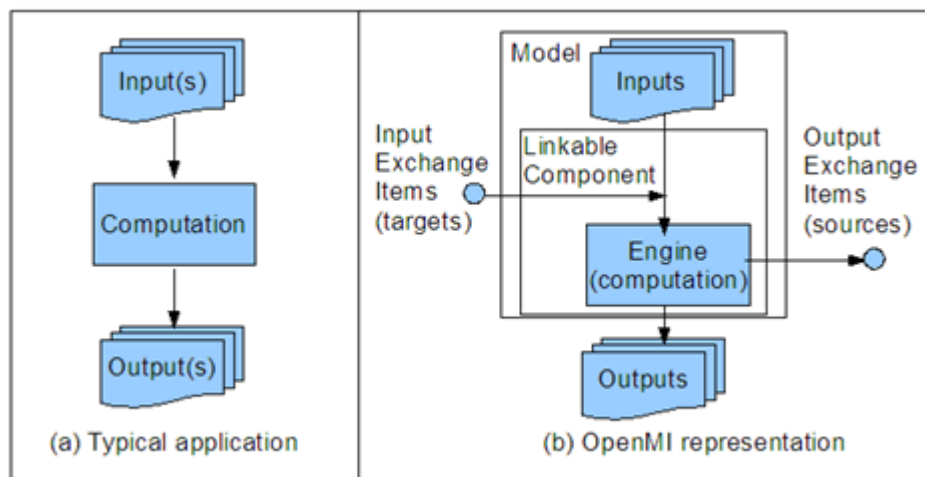
OpenMI is maintained by the OpenMI Association and this document, along with other more detailed documentation, can be obtained from <http://www.openmi.org>

1. Introduction

OpenMI is primarily concerned with integrated software systems. These systems can only be developed and maintained if they are based on a collection of dynamically interlinked applications. So the Association has developed a unified approach to linking applications which it promotes as the OpenMI Standard. The standard is intended for use with both legacy and new applications and has been widely adopted.

Large software packages typically consist of a number of discrete phases (figure 1.a), the most typical being

- Generation of input file(s) (often by a User Interface)
- Computation
- Generation of output file(s) (results)

Figure 1. Terminology

OpenMI focuses on the computation phase of the application, and the code base that performs that computation is referred to as the *engine*. OpenMI also defines a *linkable component* as an engine that has been adapted to provide dynamic data exchange. Dynamic data is transferred between components by the linking of *output exchange items and input exchange items* (sometimes also referred to as *targets* and *sources* for brevity). A *model* is a specific instance of a component that also includes input data (figure 1.b). For instance, a ground water model would be a ground water component with input data for a specific geographic region of interest. This input data can often be considered as default data, since the act of linking the model could result in the overriding of this data with dynamic values from another source. When two or more models are linked together they are considered to be in a *composition*.

This approach breaks the pattern of importing all input data before computation and not writing any outputs until computation has completed. To reiterate, an OpenMI model will

- Only request inputs immediately prior to their use in the engine
- Provide outputs immediately after they have been computed in the engine

In order to achieve this, the application needs to be turned into an engine component that implements this dynamic data transfer in a way that other engines can also recognize and utilize.

OpenMI addresses these issues by defining a set of software interfaces (the *OpenMI Standard*), which if implemented around the engine, will allow other similarly modified engines to transfer data dynamically between each other. If an engine implements this standard in an approved way (i.e. as recognised by the Association) the application is deemed to be *OpenMI Compliant*. Users of a compliant model will know that they can perform dynamic data exchange with that engine. However, they will not necessarily know whether the data exchange is of a type or form that would be useful to their particular interests. Hence, an additional task of OpenMI is to define a uniform way of describing the available targets and sources.

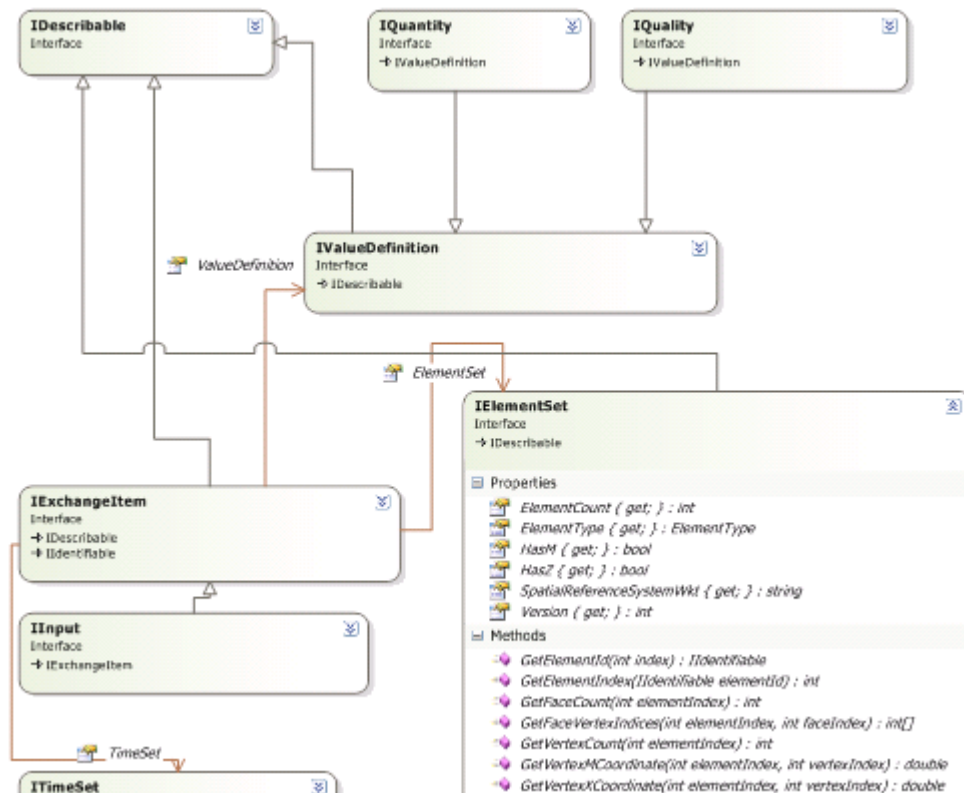
The standard is intended to enable on-line (memory based) data exchange between models on a time (e.g. a time stamp or a time span) basis. Given that different models will be running on different temporal and spatial domains, the mechanism for data conversions between domains also needs to be prescribed. Consequently, the standard must not only specify definitions, but also provide guidelines as to how those definitions are implemented.

The standard is not limited to only connecting time-based engines. Version 2 of the standard has specifically concentrated on simplifying the mixing of temporal and non temporal (e.g database) models within a single composition.

2. Standard

The standard is essentially defined by a set of interfaces that a component must implement. These are described using Unified Modeling Language (UML) diagrams (figure 2).

Figure 2. Example of OpenMI UML diagrams



Additional implementation requirements are provided as embedded comments with the interface definitions (figure 3). Currently the Association provides implementations for two software *frameworks*, Microsoft .NET [2] and SUN Java [3]. To be OpenMI compliant, the component must implement at least one of these frameworks..

It should also be stressed that this does not preclude the use of engines written in procedural languages such as C or FORTRAN (or indeed other languages such as C++, Matlab etc.). However, these engines must use intermediate code (*wrappers*) that provide the bridge between themselves and their chosen framework. Many examples exist that do this, and the wrappers are often made openly available for general use, frequently in the form of code libraries (often referred to as *Software Development Kits, SDK's*).

To reiterate, the OpenMI Standard is defined to be framework neutral; it is only specific implementations that might focus on particular frameworks and platforms. The Association actively promotes the uptake and use of OpenMI with different programming languages, frameworks and platforms.

Figure 3. Implementation Guidelines

The screenshot displays a Java IDE interface for the `ILinkableComponent` interface. On the left, a navigation pane lists various classes and interfaces. The main window shows the source code for the `ILinkableComponent` interface, which extends `IIdentifiable` and `IDescribable`. A summary of the interface is provided, stating that it is the key interface in the OpenMI standard and that any compliant component must implement it. The summary also includes a note about the interface's location in the `OpenMI.Standard2` package.

2.1. Definitions

The standard interfaces prescribe specific requirements as to how exchange items should be defined. This allows users, whether programmers using the interfaces directly or browsing indirectly via a Graphical User Interface (GUI) to get a consistent view on what is available from different components.

This is important as users need to be able to apply engineering judgement as to what is compatible, e.g. Can a source exchange item be linked with a target exchange item, from a different component, in a way that is meaningful from both physical and modeling viewpoints.

2.1.1. Model

To allow an OpenMI component to be usable by other components (i.e. be linkable), it is necessary that it can be found and its OpenMI specific information extracted. The standard requires that a compliant component must provide meta-data about itself. This consists of

- How to locate a code library that implements the OpenMI Standard `ILinkableComponent` interface.
- The type name of a class within that code library that can be used to create an instance of the `ILinkableComponent` interface.
- Any additional values required for the execution of the model via that `ILinkableComponent` instance.

The meta-data must be provided as XML, This XML must validate against a specific OpenMI prescribed XSD schema. A model is NOT compliant if it cannot be successfully validated against this schema. Typically this XML is provided in a file (known as an *OMI file*), this however is not a requirement for compliance.

2.1.2. Value

Value definitions are subdivided into Quantities and Qualities

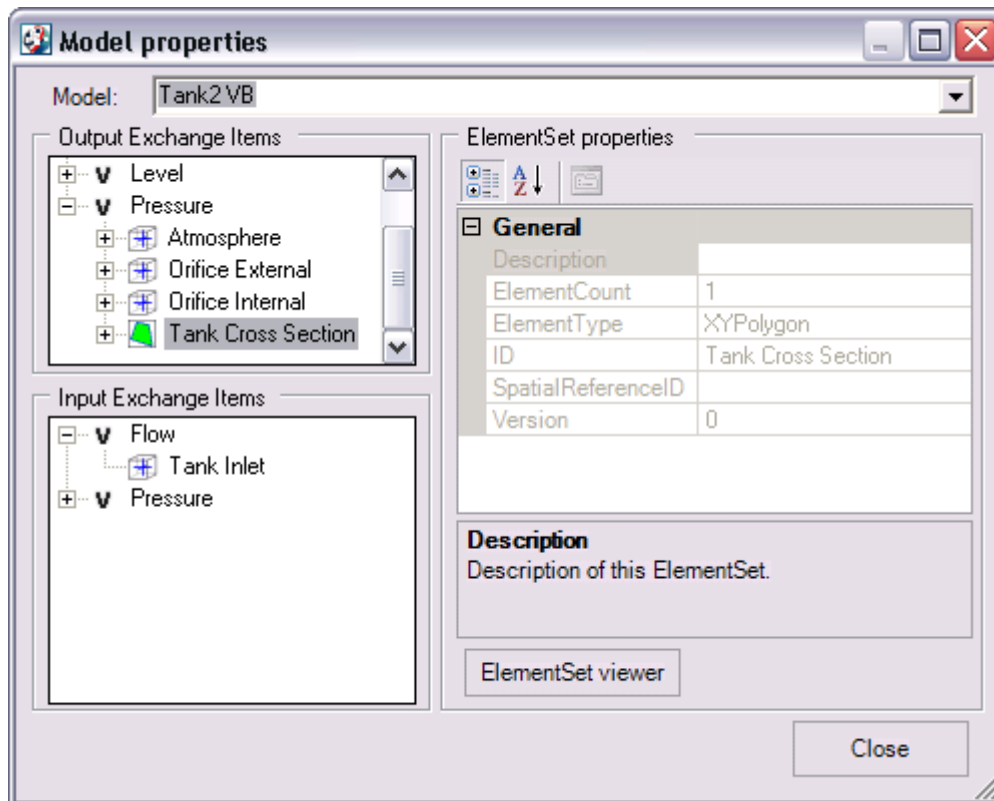
- Quantities

For values that can be expressed using SI units, e.g. flow in m³/s

- Qualities

For values that are qualitative, e.g. Bird population referred too by categories "Safe", "Endangered" or "Critical"

Figure 4. GUI representation of Value Definitions



2.1.3. Spatial

Typically values will be associated with a spatial domain. When this is not the case the transferred values are considered to be *ID based*. Spatial domains might be based on points (rain gauge locations), polylines (river centre lines), polygons (catchment) or polyhedrons (porous rock strata). OpenMI defines an *Element set* to encapsulate the spatial representation. An element set can contain one or many elements, but must not mix element types, e.g. it cannot contain both points and polylines together.

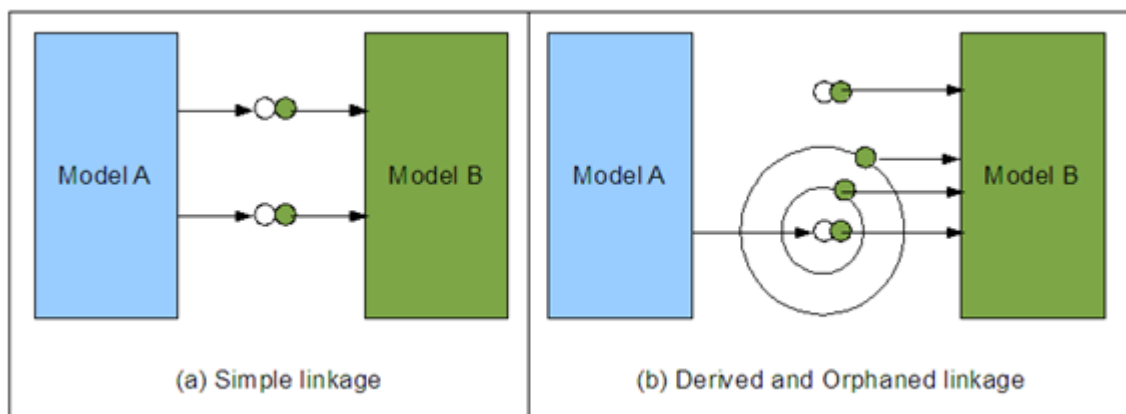
An element set can be interrogated to obtain the defining coordinates which may be plotted via a UI to assist the user in making linkage decisions when building a composition. Coordinates can be geo-referenced against a specific GIS "well known type" definition.

An element set could be a vast structure, say all the cells in a computation grid. For reasons of efficiency, the standard therefore does not explicitly define an element. Specific element information is extracted indirectly via the element set interface.

2.2. Composition building

When a user links two or more models together they are creating a *composition*. Linkages are created by connecting output exchange items (sources) from one model to input exchange items (targets) on another (figure 5a).

An exchange item is defined by the pairing of a value definition with a spatial definition.

Figure 5. A Compositions with linkage

The onus is on the user who links models together to ensure that the linkages are physically and numerically meaningful. Clearly in many instances the sources made visible on one component will not be compatible with the targets on another e.g.

- A river model with a leakage source on a polyline element set representing the river centre line.
- A ground model with an inflow target on a polygon element set representing the 2D surface of the computational grid.

For this reason, OpenMI allows for the generation of *derived* sources. A derived source is also an output exchange item; it takes the existing exchange item and modifies it so it can be connected to a different target (figure 5b). Derived exchange items can be provided either by the source component or by a 3rd party source library. For example the *OpenMI Association Technical Committee (OATC)* provides an example implementation of a 3rd party library of simple conversions between different element set types.

The main advantages of derived sources are

- Derived sources that involve considerable computational expense can be repeatedly used on multiple targets without duplication of that expense.
- Exchange only occurs between something that is a source to something that is a target. So no additional connection mechanism has to be defined, i.e. no need for OpenMI to define a link interface.

The standard also allows orphaned sources. An orphaned source has no providing model attached to it. This is specifically useful for linking with simple data sources that involve no computation, e.g. Data files, Databases, Web services etc. In these cases the provider of the orphaned source has a much simpler implementation task as they do not have to implement a component.

2.3. Runtime

During runtime a composition must be driven to ensure the correctly timed exchange of data between models. The standard allows for two approaches, *pull* driven and *loop* driven.

2.3.1. Pull driven


This is based on the 'request & reply' mechanism according to Buschmann et al. (1996) [4]. A component is defined as the primary driver of the composition, update is repeatedly called on this primary component until it reaches completion. On each update call, before the component performs its computation, it updates all its active targets until they can provide the required input data, if necessary, these targets will in turn call update on their connected source components and so on. This update and compute mechanism then propagates data around all the components in the composition in a sequential manner.

To avoid grid-lock, a scenario where two components require inputs from each other simultaneously, there is an additional requirement that if a component is requested for output which it cannot provide by computation; then it must still supply a value, even if it has to guess. Typically this results in an extrapolation from previous computed values. Apart from the primary component, other components will only progress their computation if values have been requested. So it is quite possible that at the end of a composition run, not all models have proceeded to the same final temporal moment.

This approach has the advantage that, whatever the composition, it will always progress to completion without any additional knowledge about the composition being solved. This makes it ideal for implementation by generic GUI's. Version 1 of the standard only focused on this mode of operation.

The primary disadvantage is that is a purely synchronous process, i.e. Single threaded. One component will not progress until another has finished. Note, this does not preclude a component utilising multi-threading techniques itself during its computation phases, merely that the inter component communication is sequential. Figure 6 shows a simple three component composition; model A cannot ask models B and C to progress simultaneously from time t1 to t2 even though they have no data interdependences. Model A must ask B to progress, wait until it returns and then request C to do likewise (or visa versa). Actually, for this simple example, a clever implementation could resolve this issue generically, but in real life, compositions rapidly become too complicated to be so optimised without additional user runtime input.

Figure 6. Sequential composition execution



For this reason version 2 of the standard allows for an additional driving mode. The loop approach.

2.3.2. Loop driven

This approach assumes the presence of an external controller with additional knowledge about the composition being solved. Therefore, currently, it is largely of benefit for programmers interacting directly with the interfaces rather than users running indirectly using a generic UI.

The loop driven approach allows the controller to update components in any order. It is up to the controller to protect against and resolve any grid-lock situations.

This approach is only available from version 2 of the standard onwards.

2.4. Compliance

As mentioned earlier, it is a condition of compliance that the implementer must provide meta-data about the model. Apart from this model specific information, additional general information is also required. The purpose of this additional meta-data is to provide an overview of the components purpose, usage and implementation. For compliance, a supplier should communicate this meta-data as XML to the Association. The Association will then review, and if accepted, publish the data on the Association's web site to promote its availability. The meta-data provides information such as

- What the component models and its utility

- Who owns the component and its availability for usage by others
- Which exchange items the component exposes
- Which computational frameworks and platforms it implements

Figure 7. Compliance meta-data presentation on Association website

The screenshot shows a webpage titled 'Compliant Software' with a navigation bar containing 'Home / Users' and a 'Print' button. Below the title, there is a section 'About the listed components' with a paragraph explaining that the OpenMI Association maintains a list of compliant components but does not test them, placing responsibility on the providers. To the right of this text is a callout box titled 'To submit a new OpenMI compliant model' which asks if the user wants their model added to the list and provides instructions to submit an XML file to the OpenMI Association Technical Committee. Below this is a table with three columns: Provider, Component, and Description. The table lists three components: ZOOMQ3D from the British geological survey & University of Birmingham, and two versions of InfoWorks (CS 10.0 and RS 10.0) from Wallingford Software Ltd. Each component entry includes a note about compliance with the OpenMI Standard 1.4 .Net and links to XML and HTML info.

Provider	Component	Description
British geological survey & University of Birmingham	ZOOMQ3D Compliant to: OpenMI Standard 1.4 .Net XML Info , HTML info	Finite-difference groundwater flow model
Wallingford Software Ltd	InfoWorks CS 10.0 Compliant to: OpenMI Standard 1.4 .Net XML Info , HTML info	Hydrological modeling for the urban water cycle
	InfoWorks RS 10.0 Compliant to: OpenMI Standard 1.4 .Net XML Info , HTML info	Flow simulation for rivers, channels and floodplains

3. Tools and Utilities

As part of the development of the standard, the *OpenMI Association Technical Committee (OATC)* developed a number of tools, libraries, utilities and unit tests that it used to implement and verify the standard. There is no requirement to use these, they are not part of the standard, and indeed the Association does not formally support them. However, they are made available as open source utilities to assist users in developing their implementations. The two primary utilities that are available are the SDK and Configuration Editor.

- OATC Software Development Kit (SDK)

A library of C# classes that provide implementations, examples and unit tests for the standard.

- OATC Configuration Editor


A simple editor that allows users to add, link and run OpenMI compositions.

These utilities can be obtained from the OpenMI Source Forge project site, using one of these two methods:

- <http://sourceforge.net/projects/openmi>
- svn co <https://openmi.svn.sourceforge.net/svnroot/openmi> openmi



Figure 8. <http://sourceforge.net/projects/openmi/files/>

SourceForge.net > Find Software > OpenMI > Browse Files














 **OpenMI** by adharper, hartnack, hrajagers, jan_gregersen, jgr-dhigroup, ...

Summary | **Files** | Support | Develop

The Open Modeling Interface and Environment (OpenMI) defines a standardized way to exchange data between computational models that run simultaneously. OpenMI aims to enhance process interaction representation in integrated environmental modelling.

Download Now!
Oatc.ConfigurationEditor_... (1.2 MB)  OR [View all files](#) 

Browse Files for OpenMI

File/Folder Name	Platform	Size	Date ↓	Downloads	Notes/Subscribe
Newest Files					
 Oatc.SDK_1.4.0.0.zip		4.8 MB	2007-12-20	1,124	
 Oatc.ConfigurationEditor_1.4.0.0.msi		1.2 MB	2007-12-20	1,922	
All Files					
▶  SDK		4.8 MB	2007-12-20	1,124	 
▶  Configuration Editor		1.2 MB	2007-12-20	1,922	 
▶  Superceeded (Deprecated)		16.4 MB	2006-08-30	1,967	 

4. Chronology

- 2001 -2005

OpenMI initially developed as part of the HarmonIT[4] project. An EEC co-funded research project,

Key milestones

- Release of version 1 of the OpenMI standard

- 2005 -2010

OpenMI further developed as part of the OpenMI-Life[5] project, an EEC co-funded research project,

Key milestones

- Formation of the OpenMI Association to protect the long-term future of the Standard.
- Version 1 of the standard was tested on several large scale real world projects.
- Version 2 of the standard was developed and submitted to peer review prior to formal ratification by the Association.

5. References

1. Based on an original paper "The OpenMI Standard in a nutshell", Sept 2007 by Peter J.A. Gijbers (WL | Delft Hydraulics) and Jan B. Gregersen (DHI Water & Environment)
2. .NET Framework, msdn.microsoft.com/netframework
3. Java Software Platform, www.java.com
4. "A System of patterns", F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, M. Stal. (ISBN 0 471 95869 7) 1996.
5. HarmonIT research project co-funded by the European Commission (Contract EVK1-CT-2001-00090)
6. OpenMI-Life project as part of the EU-LIFE Environment program (Contract no : LIFE06 ENV/UK/000409).