

A Distributed Data Component for the Open Modeling Interface

T. Bulatewicz^{a,*}, D. Andresen^a, S. Auvenshine^b, J. Peterson^c, D. R. Steward^b

Kansas State University, Manhattan, KS, 66502, USA

^a*Dept. of Computing and Information Sciences*

^b*Dept. of Civil Engineering*

^c*Dept. of Agricultural Economics*

Abstract

Data management is a fundamental part of environmental modeling and simulation. This is particularly true for the types of interdisciplinary, interconnected models required to address water resources challenges faced by society, such as our case study of a depleting aquifer in an agriculturally important area. Model input data are often obtained from online data services and output data uploaded to them for purposes such as storage or distribution. Enabling linked models to directly communicate with such services can simplify this process. We have developed an Open Modeling Interface (OpenMI) data component that retrieves input data for model components from standards-based web services and delivers output data to them. The adoption of standards for both model component input-output interfaces

*Corresponding author, 234 Nichols Hall, Kansas State University, Manhattan, KS 66502, USA, 505-490-9681

Email addresses: tombz@ksu.edu (T. Bulatewicz), dan@ksu.edu (D. Andresen), auvenshi@ksu.edu (S. Auvenshine), jpeters@ksu.edu (J. Peterson), steward@ksu.edu (D. R. Steward)

and web service application programming interfaces make it possible for the component to be reconfigured for use with different linked models and various web services. The data component employs three techniques tailored to the unique design of the OpenMI that enable efficient operation: caching, prefetching, and buffering, making it capable of scaling to large numbers of simultaneous simulations executing on a computational grid. In this work we present the design of the component, an evaluation of its performance, and a case study demonstrating how it can be incorporated into modeling studies. The results of the performance study indicate that it is capable of scaling to large numbers of simulations (tested up to 1000) incurring no delay when delivering data and an average delay of 0.3 to 3.78 s per time step when retrieving data. The techniques of caching and prefetching were effective in reducing or eliminating this delay in cases in which simulations used identical input data or when data could be retrieved from the web services before it was requested by a model component.

Keywords: OpenMI, Data management, Web services, Integrated modeling

1. Introduction

Data management is a fundamental part of environmental modeling and simulation. Within the context of integrated environmental modeling, the input data required by a set of computer models is typically collected from a variety of sources and assembled into a set of input files that are deployed with the model programs to a desktop computer or to a compute cluster composed of high-performance computers connected via a fast network. These sources often include several different online (i.e. Internet-connected) data repositories.

9 ries provided by various government, academic, and private agencies. The
10 data typically varies spatially and temporally and may be consumed from
11 input files by a model once during initialization or throughout the execution
12 of a simulation. The output data from computer models follow the reverse
13 path as the output files are collected and aggregated before being uploaded
14 to online services. These services may provide the means to archive data,
15 publish data publicly, share data within and across institutions, or analyze
16 and visualize data.

17 Linked (or coupled) models are composed of independent models that
18 cooperate to collectively perform simulations where each model consumes a
19 set of input files and produces a set of output files. The preparation of these
20 sets of input files (including the retrieval of data from online sources) and
21 the processing of the output files sets (including the delivery of data to on-
22 line services) is typically performed manually or through ad-hoc automation
23 techniques such as scripting. This may require a substantial effort in devel-
24 oping and configuring the necessary software and scripts for both processing
25 the model-specific input and output files and for communicating with online
26 services, both of which may require changes or additional software develop-
27 ment each time the integrated model is changed (e.g. adding or removing
28 models) or the online services it relies on.

29 Enabling linked models to directly communicate with online services can
30 simplify the management of model data by avoiding the intermediary use of
31 data files and obviating the need for manual data processing tasks and ad-
32 hoc scripting. Through the adoption of standards, in both model component
33 input-output interfaces and web service application programming interfaces,

34 general-purpose data components can facilitate the exchange of data between
35 model components and web services. It is advantageous to place such web
36 service functionality into data components rather than directly into model
37 components because it allows for more efficient operation (e.g. avoiding du-
38 plicate data retrieval by different components) and minimizes the software
39 complexity of the model components.

40 We have developed a distributed data component that conforms to the
41 Open Modeling Interface (OpenMI) (Gregersen et al., 2007) that both pro-
42 vides input data to model components retrieved from standards-based web
43 services and delivers model output data to such services on each time step.
44 By operating on a time step basis, the data component enables model com-
45 ponents to consume input data, such as measurement data from sensor net-
46 works, and distribute output data in real-time. This also supports compu-
47 tational steering scenarios in which model output is monitored and inputs
48 are manipulated as necessary as a simulation is being performed. The data
49 component employs three techniques tailored to the unique design of the
50 OpenMI that enable efficient operation: caching, buffering, and prefetching.
51 This work unifies our previous efforts (Bulatewicz and Andresen, 2011, 2012)
52 and includes improvements to the software design that achieve a significant
53 increase in scalability. It also provides an integral part of an interdis-
54 plinary modeling study in which we are integrating models of groundwater,
55 economic decision making, and crop production to investigate the impact
56 of policy on irrigated agricultural systems. The following sections position
57 this work within the context of existing research and introduce the aspects
58 of the OpenMI relevant to understanding the design and implementation of

the data component. We then present the design of the data component in Section 2, an evaluation of its performance in Section 3, and a demonstration of how it may be incorporated into an integrated modeling study in Section 4.

1.1. Related work

This work lies at the intersection of component-based modeling, web services, and grid computing. The synergy between web services and modeling and simulation was recognized quickly as web standards emerged (Chandrasekaran et al., 2002). Web services can provide a means for both remotely controlling the execution of computer models running on servers or computational grids (Castronova et al., 2013a; Goodall et al., 2011; Horak et al., 2008; Pullen et al., 2005) and enabling desktop or grid-based models to exchange input and output data with online services. In the latter case an online service may be composed of a suite of Internet applications and/or a collection of databases.

One class of online services that is well-suited for exchanging data with computer models is workflow management systems which are frameworks to setup, execute, and monitor scientific workflows composed of web services, such as Taverna (Hull et al., 2006) and VisTrails (Bavoil et al., 2005). Such systems could provide workflows that pre-process or post-process model data or conduct simulations whose input or output data is utilized by models. Another class of online services are data-centric and provide data storage (e.g. archiving) and retrieval (e.g. public access or sharing within or across institutions). Examples include the Integrated Rule-Oriented Data System (iRODS) (Rajasekar et al., 2006) which is a file-based distributed

84 data storage system, the Consortium of Universities for the Advancement of
85 Hydrologic Science, Inc. (CUAHSI) Hydrologic Information System (HIS)
86 (Maidment, 2008; Tarboton et al., 2009) which facilitates the management
87 of hydrologic data, Globus Online (Globus Online, 2013) which provides on-
88 line managed data storage based on GridFTP (Globus Toolkit, 2013), and
89 HDF5WS (Shasharina et al., 2006) which provides access to HDF5 data files.

90 Web services provide a means for these online application and data ser-
91 vices to achieve interoperability with one another and with client applications
92 running on desktop computers and compute clusters. Standards for web ser-
93 vices and the data encodings they use make it possible for independent ap-
94 plications to interpret exchanged data in a meaningful way. In the context
95 of environmental modeling in which data is spatial-temporal in nature, the
96 standards published by the Open Geospatial Consortium (OGC) for location-
97 based information and services are of particular relevance. For example, the
98 Web Feature Service (WFS) Standard (Vretanos, 2010) defines how geospa-
99 tial data may be accessed from a web service and utilizes the Geographic
100 Markup Language (GML) (Portele, 2007) Standard. Within the domain of
101 hydrology, the CUAHSI HIS WaterOneFlow web service Application Pro-
102 gramming Interface (J. S. Horsburgh and Whitenack, 2009) defines how time
103 series hydrological observations data may be accessed and utilizes the Water
104 Markup Language (WaterML) encoding standard (Zaslavsky et al., 2007).

105 The fundamental data model upon which these services and encodings are
106 based (consisting of quantities, times, and locations) is generally compatible
107 with the data model employed by the OpenMI for the exchange of data be-
108 tween components making interoperability between services and components

possible (Castronova et al., 2013b). Several OpenMI components have been developed that retrieve time series data from WFS web services (OpenMI Association, 2010). In a related work, Castronova et al. (2013b) enabled a desktop application to retrieve input data from WaterOneFlow web services and store them in a local database which could then be accessed by model components via a general-purpose data-access component.

Our work complements these efforts in two ways. First, our data component is not only capable of retrieving data from web services but delivering data to them as well. Second, the data component is not limited to use on desktop computers but may also be used on high-performance compute clusters. The prototype implementation is compatible with WaterOneFlow web services and is being extended to support additional standards. In our previous work (Bulatewicz and Andresen, 2011, 2012) we developed independent components for retrieving data from web services and delivering data to them. This work unifies our earlier efforts into a single component and includes fundamental changes to the software design to scale to significantly higher numbers of simultaneously executing simulations.

1.2. The Open Modeling Interface

The Open Modeling Interface (OpenMI) Standard (Gregersen et al., 2007) defines how software components may exchange spatial-temporal data with one another and coordinate their execution. Components that possess the capabilities defined by the interface can be linked together and exchange data, typically on each time step, as they carry out simulations. These capabilities are implemented as functions (specifically, object methods and properties) within the source code of a component that either provide descriptive infor-

134 mation about the component (such as its inputs and outputs) or support its
135 execution (such as performing initialization or exchanging data).

136 Each input and output is formalized as an *exchange item* that describes
137 the properties of a domain quantity such as its name, units, and spatial dis-
138 tribution. The way in which a quantity is spatially distributed is formalized
139 as an *element set* that is composed of a list of *elements* each of which has a
140 textual identifier, spatial shape (point, line, or polygon) and geographic co-
141 ordinates. When configuring a linked model, called a *composition*, a scientist
142 uses a visual software tool (the OpenMI Configuration Editor application -
143 OmiEd) to choose a set of components and assign each input exchange item
144 of a component to an output exchange item of another component. These
145 assignments are called *links* and there may be multiple links between two
146 components and may be in the same or opposite directions. At runtime a
147 component requests data from other components along each input link, typ-
148 ically before performing each time step. The request is made by calling the
149 GetValues function of each linked component specifying a date and time at
150 which the data is needed, as illustrated in Fig. 1. The GetValues function
151 returns a list of real numbers called a *value set* where each number represents
152 the state of the quantity at the requested point in time at a different spatial
153 location. As such, each call to GetValues may be considered to be a request
154 for the state of a quantity at a point in time for a list of spatial locations and
155 the response to be the list of numbers returned.

156 In addition to facilitating the exchange of data between components, the
157 GetValues function provides implicit coordinated execution of components
158 at runtime. The execution of a linked model is initiated when one of the

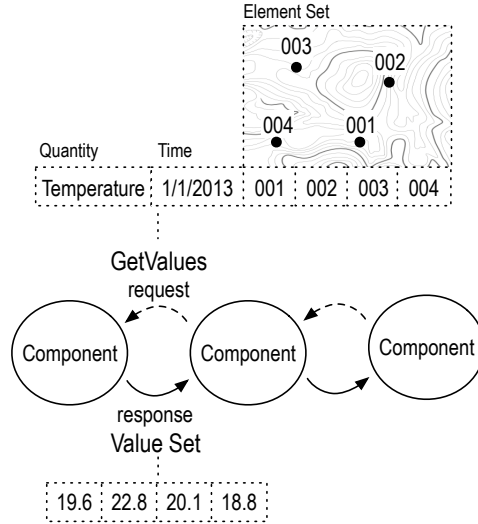


Figure 1: Lists of real numbers called *value sets* are exchanged between model components.

159 components begins executing. On each time step the component invokes
 160 GetValues on each component linked to it to obtain all the necessary input
 161 value sets for the time step, pausing its execution during each invocation.
 162 When GetValues is invoked on a component, it executes as many time steps
 163 as necessary to advance to the requested point in simulation time and returns
 164 a value set corresponding to that time. Thus a component only executes time
 165 steps on-demand in response to the invocation of its GetValues function by
 166 another component and may itself invoke GetValues on other components
 167 prior to performing each of its time steps. In this way components take turns
 168 executing and pull data from one another until the initiating component's
 169 simulation is completed.

170 Components are typically model programs that consume input data and
 171 produce simulated output data, but they can serve other purposes as well.

172 Examples include data conversion or transformation, data visualization, ac-
173 cess to databases, and access to online data services as in the case of our
174 data component.

175 **2. Methods**

176 *2.1. Overview*

177 The purpose of the data component is to serve as an intermediary between
178 online data services and model components, both providing model input data
179 retrieved from web services and delivering model output data to web services.
180 The design of the data component was guided by the following requirements:

- 181 1. To be general-purpose
- 182 2. To minimize the runtime of a simulation
- 183 3. To be scalable

184 Our design balances these three competing objectives making the data com-
185 ponent broadly applicable and suitable for use on both desktop computers
186 and compute clusters.

187 The first requirement of the data component is that it is general-purpose
188 such that its inputs and outputs can be defined, and redefined, by a scientist
189 as necessary for different sets of model components. The input and output
190 exchange items of the data component reflect the quantities exposed by a
191 web service: any quantities that a web service can provide or accept can
192 be configured as exchange items of the data component. This is possible
193 because the OpenMI defines the way in which data is exchanged between
194 software components and web service standards define the way in which data

195 is exchanged with online services. Together these standards make it possible
196 for the data component to serve as a data relay between model components
197 and web services.

198 The data component is configured (via a file) by specifying the list of
199 input and output quantities that a web service can provide and accept, along
200 with the element set definition of each and the web service URL and type.
201 These quantities become available as input and output exchange items when
202 the data component is added to a composition in the OmiEd application and
203 can be linked to model components in the same way that links are added
204 between model components.

205 The second requirement of the data component is that it minimizes its
206 impact on the runtime of a simulation, ideally causing no increase. If a data
207 component was to call a web service after each request received from a model
208 component to either obtain input data or send output data, the simulation
209 would be paused during the web service call (due to the synchronous execu-
210 tion of components) and increase the runtime of a simulation. This increase
211 in runtime can be reduced or eliminated by decoupling the calls to the web
212 services from the requests made by the model components. In order to de-
213 couple the web service calls from the model component requests, the data
214 component must have the ability to temporarily store model input and out-
215 put data in a *data store*. Rather than the data component call a web service
216 in response to each request for input data from a model component, it first
217 checks to see if the data is already available in the data store. If it is, then it
218 can be returned to the model component immediately, and if not, it can then
219 be requested from a web service. There are two cases in which the data may

220 already be available in the data store: (1) the data was previously requested
221 by a model component, and (2) the data was retrieved from a web service
222 ahead-of-time. We refer to the prior as *caching* and the latter as *prefetching*
223 and these techniques can reduce, and in some cases eliminate, the increase in
224 runtime due to the web service calls. In addition to minimizing the runtime,
225 caching also minimizes the amount of data downloaded from the web services
226 because each input is only retrieved once. The data store is shared among all
227 simulations executing across a compute cluster to maximize the reusability
228 of the cached data. With respect to sending output data, rather than call a
229 web service in response to each request from a model component, the data
230 component immediately stores the output data in the data store and sends
231 it at a later time. We refer to this as *buffering* and it eliminates the increase
232 in runtime otherwise due to sending output data to web services.

233 The third requirement of the data component is that it is scalable such
234 that many simulations, each containing an instance of the data component,
235 may execute concurrently across a compute cluster with minimal impact to
236 the runtime of the simulations. To these ends we employed two strategies:
237 (1) maximize network efficiency when sending data to web services, and (2)
238 separate the data component software into two tiers.

239 Network utilization is inefficient when the amount of data being sent is
240 small enough that the network latency is comparable to the transmission time
241 of the data (i.e. the duration of time and amount of data exchanged at the
242 network transport layer for establishing the connection and for sending the
243 data are similar). To ensure that the network bandwidth is used efficiently
244 when sending model output data to web services, the data component sends

245 a sufficient amount of data in each web service call. With respect to retriev-
246 ing data, which consists of values that each represent a quantity at a point
247 in time for a location, the data component could request groups of values in
248 each web service call for spans along any of these three dimensions in each
249 web service call. At one extreme it could make a web service call for each
250 individual value, and at the other extreme it could make a single web service
251 call to obtain all the input values required for a complete simulation. In the
252 prior case the network utilization may be inefficient due to the small data
253 size of a single value, and in the latter case the execution of a simulation
254 would be delayed until the data is retrieved and may require storing a large
255 amount of data for the lifetime of the simulation (in addition it would pro-
256 hibit both real-time online data access during the simulation and the ability
257 to utilize multi-threaded and multi-hosted web services). Efficient network
258 utilization can be balanced with real-time data access by requesting groups
259 of values in each web service call (essentially coalescing what would otherwise
260 be multiple requests into a single request). Values could be grouped by time,
261 quantity, and/or location, depending on the capabilities of a web service. In
262 addition, grouping by time would require the data component to be capable
263 of predicting the simulation times at which model components will request
264 data and grouping by quantity would only be possible in cases in which the
265 data component is providing multiple quantities to one or more model com-
266 ponents that are sourced from a single web service and requested for the
267 same points in simulation time. We designed the data component such that
268 requests are grouped by location (when supported by the web service) and
269 left grouping by time and quantity to be addressed in future work due to the

270 additional complexity.

271 The data component software is organized into two tiers that separate
272 the management of the data store and communication with web services
273 from the interactions between the components within a composition. This is
274 a more scalable design than our previous work (in which there was a single
275 tier) because the management of the data store requires considerable com-
276 puter resources (memory, processor, and network) yet accessing the data
277 for providing input data to model components and collecting output data
278 requires few resources. Without separating them, the resource demands of
279 the data store are imposed on each data component thus increasing the re-
280 source demands of every simulation. By separating them, the amount of
281 resources dedicated to the management of the data store can be managed
282 separately from those required by the individual simulations. The number of
283 *data managers* that manage the data store can be increased or decreased in-
284 dependently from, and as necessary to support, the number of simultaneous
285 simulations.

286 An overview of the system is illustrated in Fig. 2. Compositions of linked
287 components perform simulations on the nodes of a cluster. Each composition
288 includes a data component (labeled DC in the figure) whose input and/or
289 output exchange items are linked to model components. Model components
290 request input from data components (by invoking `GetValues`) for a quantity
291 at a specific time and element set in the same way as from other model com-
292 ponents. The data component in turn requests the input data from a data
293 manager which may obtain the data from the data store or retrieve it from
294 a web service to fulfill the request. Each time a model component produces

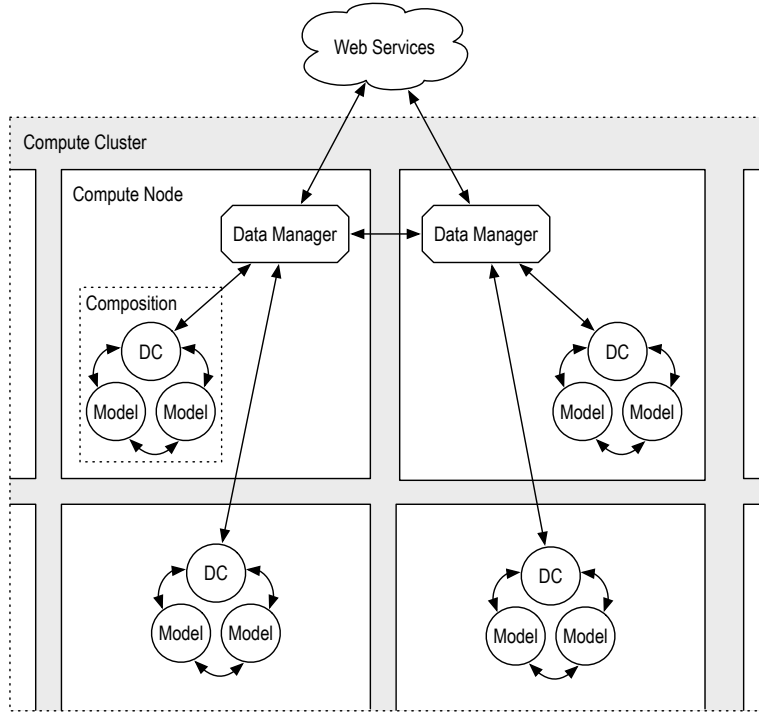


Figure 2: System overview. Arrows indicate the movement of data.

295 output data (in response to a `GetValues` request from another model com-
 296 ponent) the data component is notified. When notified, the data component
 297 obtains a copy of the output data (by invoking `GetValues` on the model com-
 298 ponent) and sends them to a data manager which stores the data for eventual
 299 delivery to a web service.

300 2.2. The data store

301 Data managers are responsible for both communicating with web services
 302 and managing the storage of model input and output data in the data store.
 303 We implemented a set of software modules that provide the functionality

304 to communicate with web services and utilized an existing software for the
305 data store functionality. The data store is a *key-value store*, which is a non-
306 relational database in which related data is aggregated together and stored
307 as an entry that is accessed via a unique identifier. We chose to utilize a
308 key-value store because storing data in this way achieves high performance
309 when scaling horizontally (i.e. increasing the number of compute nodes to
310 allow for higher capacity) because the data can be efficiently sharded and
311 replicated across compute nodes (i.e. each node stores a subset and/or copy
312 of the entries).

313 The data operations that may be performed on a key-value store include
314 inserting entries, accessing entries, and removing entries, typically referred
315 to as *put*, *get*, and *remove*. These operations rely on a unique *key* to be
316 associated with each entry when inserted into the store and is subsequently
317 used to locate the entry for access or removal. Locating entries based on
318 their key is very efficient, while iterating or searching through all the entries
319 is not, thus the way in which data is aggregated into entries dictates the
320 way in which it may be efficiently accessed and thus the overall performance
321 of a key-value store. The structure of the data exchanged between both
322 components and between the data component and web services is a value set
323 that consists of a list of real numbers that represent the state of a quantity
324 at a point in time over a set of locations. As the value set is the unit of
325 aggregation of data exchanged, storing each value set as an entry in the key-
326 value store aligns with the way in which the data is accessed by the data
327 component.

328 Aggregating data as value sets is not the only possibility, as it would

also be possible to aggregate data into larger units such as groups of value sets, or into smaller units such as the individual values that make up a value set (as in Bulatewicz and Andresen (2011)). Storing individual values as entries in the key-value store simplifies the process of assembling value sets ad-hoc from entries in the key-value store as they are requested by model components (to avoid the need to call a web service to obtain them) thus maximizing the reusability and the effectiveness of the cache and resulting in no storage of duplicate data. This also results in higher memory usage per entry as each entry incurs a constant overhead (approximately 260 B) that is approximately the data size of a single value resulting in 50% of memory usage being overhead, and greater processor and network usage as each entry must be inserted and removed from the key-value store individually. Storing value sets as entries in the key-value store (as in Bulatewicz and Andresen (2012)) minimizes overhead in terms of memory, processor, and network, but introduces the possibility of storing duplicate data in the key-value store in the case that the values stored in two value sets intersect and requires a more complex process to assemble value sets ad-hoc (see Section 2.3.1). In our earlier work we found that the overhead of storing individual values as entries limited the scalability of the system and thus in this work we designed the data component to store value sets as entries in the key-value store.

Each entry in the key-value store is a variably-sized object consisting of a quantity identifier (string), timestamp (string), element set identifier (string), scenario identifier (string), a delivery flag (boolean), array of values (double precision), and value count (long), that are serialized into an array of bytes. The keys used to access the entries in the store are strings formed by the

concatenation of the entry’s quantity identifier, element set identifier, times-
tamp, and scenario identifier, for example: `TemperatureSewardCounty2013-01-01T12:00:00S01`. Using keys of this form guarantees uniqueness and makes it possible to efficiently lookup a value set from the key-value store for a specific quantity, time, and element set, for a particular scenario identifier. The scenario identifier provides a way to partition, version, and identify value sets that are created by different linked models or instances thereof. For example, when executing several instances of a linked model, each instance may be assigned a unique scenario identifier so that the input and output value sets of each are distinct. The delivery flag indicates whether the value set is pending delivery to a web service.

When a value set is delivered to a web service, additional information must be provided that indicates the locations the the values represent. This information is not stored inside the entries in the key-value store because all the value sets for a particular element set would result in the storage of duplicate data. As element sets are static during a simulation run there is typically a high ratio of value sets to element sets, so the entries only store the element set identifier and the actual element set information is stored separately in the data store. In this way a data store can lookup the complete element set information for any value set before delivering it to a web service.

A number of different key-value store database systems could be utilized as the data store, such as Memcached (Memcached, 2013) or Cassandra (Cassandra, 2013). We chose to utilize the Hazelcast distributed data platform (Ozturk, 2010) because in our previous work we found it to be highly effi-

379 cient and require minimal configuration. Hazelcast is a clustering, scalable,
380 in-memory data platform that is implemented in Java and distributed as
381 a shared library that we compiled into the data manager program. When
382 the data manager is started it creates an instance of the Hazelcast platform
383 peer that runs as a set of threads inside the data manager process. Instances
384 within different data manager processes dynamically form a cluster by discov-
385 ering one another via multicast and communicating via TCP/IP. Instances
386 thus join and leave the cluster as data manager processes are started and
387 stopped. Each instance has a local memory that is logically organized into
388 one or more global hashmap data structures whose entries are distributed
389 across the instances of a cluster and it is these distributed hashmaps that
390 make up the data store. The instance running within a data manager is
391 self-contained and the software modules within the data manager may only
392 **put**, **get**, and **remove** entries (i.e. value sets) to and from the data store as
393 illustrated in Fig. 3.

394 The instances balance the entries in the data store such that they are
395 evenly distributed among the instances executing on a cluster and each in-
396 stance has approximately the same number of entries in its local memory. For
397 each entry stored in an instance there is a backup copy of the entry stored in
398 a different instance somewhere in the cluster in case an instance fails. When
399 instances leave a cluster its entries are migrated to and distributed among
400 the remaining instances. Each instance optionally persists the entries of its
401 local memory to a file between executions.

402 The Hazelcast platform supports *native clients* that may access the data
403 store managed by the cluster of instances. A client connects to an instance

415 threads maintain a direct and persistent network connection to the instance
416 threads. The data component communicates with the data manager through
417 the Hazelcast client-instance connection using two request queues managed
418 by the instance. The component inserts both requests to retrieve value sets
419 from web services and requests to store value sets in the data store into these
420 queues and the data manager and its software modules process the requests.

421 *2.3. Providing input data to models*

422 *2.3.1. Caching*

423 During the execution of a composition, several model components within
424 a single composition may request identical value sets from a data component.
425 In addition, model components in independently executing compositions on
426 different cluster nodes may request the same value sets from different data
427 components. In both cases it is advantageous for the data components to
428 cache the value sets that they retrieve from the web services and to share
429 those value sets across all the data components that are executing simulta-
430 neously in different compositions across a cluster. It is also advantageous for
431 the cached value sets to be persisted between executions as the same value
432 sets may be needed on subsequent executions of a composition.

433 When `GetValues` is invoked by a model component on a data component,
434 the data component checks to see if the requested value set exists in the data
435 store by creating the appropriate key and then performing a `get` operation
436 on the data store using the key. If the data component successfully retrieves
437 the value set from the data store then it is returned to the model component
438 and the execution of the composition continues. If the value set is not in the
439 data store then the data component inserts the key into the request queue.

440 After the insertion is completed, the data component periodically checks the
441 data store until the value set is available (during which the execution of the
442 composition is paused). The data component relies on the retrieval module
443 inside the data manager to obtain the requested value set from a web service
444 and insert it into the data store.

445 The retrieval module waits for a request to be inserted into the request
446 queue. When a request is inserted by a data component, it is removed by
447 the data manager provided that the amount of data in the local data store
448 has not reached the maximum limit (as configured in the data component).
449 The request queue may only hold a single request at-a-time and causes data
450 components to wait if they attempt to insert a request when there is already
451 a request in the queue. This prevents the data manager from becoming
452 overwhelmed with requests. The data store is checked for the requested value
453 set in case it was already retrieved while the data component was waiting to
454 insert the request. If it is not, the retrieval module attempts to assemble the
455 requested value set from other value sets that are already in the data store.

456 The element set of a requested value set may intersect with the element
457 sets of other value sets in the data store. As such, it may be possible to
458 assemble the requested value set by extracting the necessary values from
459 other value sets already in the data store whose element sets intersect with
460 the element set of the requested value set. This maximizes the reusability of
461 the cached data and minimizes the number of web service calls.

462 The algorithm given in Fig. 4 is utilized by the retrieval module to as-
463 semble value sets in such a way as to minimize the number of `get` operations
464 performed on the data store. Each element of each element set is compared to

```

for each ( value v in request_value_set )
  for each ( element_set s )
    for each ( element e in s )
      if ( v.element = e )
        list.add( e, s )

for each ( element_set s in list )
  key = create_key(request_quantity, request_time, s)
  value_set = data_store.get( key )
  if ( value_set is not null )
    for each ( value v in value_set )
      for each ( element e in request_element_set )
        if ( v.element = e )
          result_value_set.add( v )

return result

```

Figure 4: Algorithm for assembling value sets.

465 the requested element set to determine whether the elements in the requested
 466 element set exist in other element sets. If all the elements in the requested
 467 element set can be found in other element sets, then the value set map is
 468 checked for each source element set to see if a value set for the requested
 469 time exists. If it does then the required values are collected from it. If all the
 470 values in the requested value set are found then the assembled value set is
 471 inserted as a new entry into the data store. This requires one `get` operation
 472 per source element set. In the case of a requested value set whose element
 473 set is a subset of another element set whose data is in the value set map, it
 474 would require one `get` operation to obtain the necessary data to assemble the
 475 value set. The maximum number of `get` operations that may be necessary is
 476 equal to the size of the value set being requested, which occurs in the case
 477 that each value is sourced from a different element set.

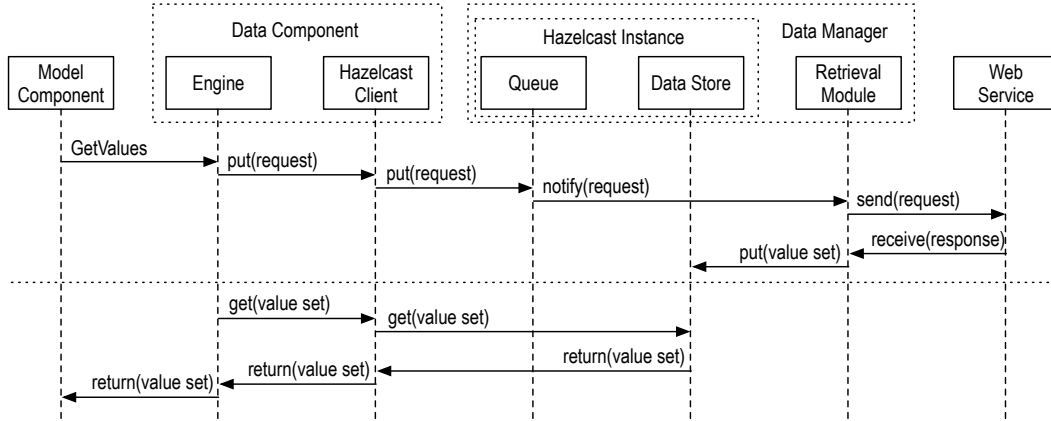


Figure 5: Sequence diagram of interactions involved in providing data from web services.

478 If a value set cannot be assembled from the values already in the data
 479 store, a web service call task is created for the request and added to a thread
 480 pool. Each task generates the appropriate web service request XML, calls
 481 the web service, and then parses the response into a value set that is inserted
 482 into the data store, as shown in Fig. 5. Multiple web service calls are issued
 483 simultaneously in a pipelined fashion to take advantage of multi-core and
 484 multi-host web services. The retrieval module limits the number of simulta-
 485 neous web service calls to the number of connected data components. This
 486 limit is necessary because data components may request value sets ahead-of-
 487 time (prefetch) which could result in the creation of so many threads that
 488 the system resources become exhausted.

489 2.3.2. Prefetching

490 The simulation of physical processes (especially those for which the OpenMI
 491 was initially designed) typically involve the calculation of output quantities

492 over a simulation time period. A component typically steps forward through
493 simulation time requesting value sets from the data component on each step.
494 To avoid causing a model component to wait for a value set while the data
495 component is retrieving it from a web service, the data component retrieves
496 value sets before they are requested, a technique called *prefetching*.

497 Throughout the execution of a composition the components are at approx-
498 imately the same point in simulation time. This is because each component
499 typically requires input data from the other components that reflect its cur-
500 rent simulation time, causing those components to advance to the same point
501 in simulation time. For this reason, all components should be prefetched to
502 the same future point in simulation time.

503 Prefetching relies on knowledge of what data will be needed before it is
504 requested. It is not possible for the data component to obtain this informa-
505 tion directly from model components, as the OpenMI does not support this
506 functionality. The data component predicts what value sets will be requested
507 in the future by observing what value sets have been requested in the past.
508 Components that use a fixed-length time step request data from the data
509 component at fixed intervals making it possible to identify these components
510 and determine the length of their time steps. In such cases the data com-
511 ponent can accurately predict the value sets that will be requested in the
512 future. It is more difficult for the data component to predict the data needs
513 of components that use a variable-length time step and is not addressed in
514 this work. The data component prefetches all links to a common future point
515 in simulation time (number of Julian days) given by: $t = \min\{p + i, e\}$ where
516 p is the earliest time to which all links have been prefetched, i is the longest

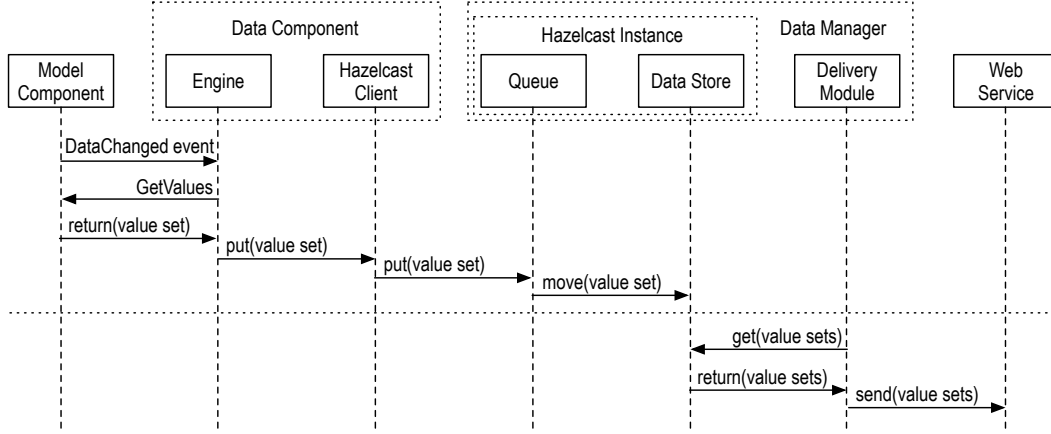


Figure 6: Sequence diagram of interactions involved in delivering data to web services.

request interval (i.e. longest time step) across all links, and e is the ending time of the composition.

2.4. Delivering output data to web services

The input exchange items of the data component may be linked to one or more model components within a composition. At initialization, the data component registers to be notified via a *DataChanged* event whenever a model component produces an output value set along any of its input links, which is typically raised after each time step. When the data component receives this notification it invokes the *GetValues* function on the model component to obtain a copy of the value set as shown in Fig. 6. The data component instructs the Hazelcast client to insert the value set into data queue within the Hazelcast instance that the client is connected to. This queue can only hold a single value set at-a-time so if a value set is in the queue, then additional insert attempts will wait until the value set is re-

531 moved, causing the data component to wait, and in turn causing the model
532 component to wait. The queue serves as a gate to prevent too much data
533 from being inserted into the data store, which would be possible if the client
534 inserted value sets directly into the data store. Whenever a value set is added
535 to the queue, the data manager checks if there is available space in the local
536 data store and if so moves the value set into the data store and sets a flag
537 within the value set that indicates it is pending delivery to a web service.
538 The amount of memory dedicated to the local data store is configurable via
539 the data store configuration file and must be equivalent among all connected
540 data stores (as required by Hazelcast).

541 The delivery module periodically searches the local data store for value
542 sets pending delivery and if there is a sufficient amount of data to be sent
543 such that network resources will be utilized efficiently then the value sets are
544 sent to the appropriate web service. The amount of data that is sent in each
545 web service call is configured in the data store as a number of bytes, called
546 the *delivery size*. The data component estimates the number of value sets
547 to include in each web service call by estimating the the size of an encoded
548 value set (as XML) via a constant per-value multiplier specific to each web
549 service.

550 The following algorithm is used by the delivery manager. The delivery
551 thread periodically iterates over the entries in the local data store and checks
552 whether each entry is pending delivery. If an entry is pending delivery it is
553 copied into a priority queue and flagged as no longer needing delivery in
554 the data store. The priority queue orders the value sets by earliest creation
555 date first. After iterating through all the entries in the local data store

556 and updating the priority queue, the priority queue for each web service
557 is checked to determine whether there are a sufficient number of value sets
558 whose encoded size is greater than the delivery size. If so, a sufficient number
559 of value sets are removed from the priority queue to meet the delivery size
560 and a thread pool task is created that serializes the value sets into the XML
561 encoding used by the web service and calls the web service. The process
562 repeats until both the simulation is completed and the number of entries
563 delivered is equal to or greater than the number of entries inserted into the
564 local buffer. The latter ensures that each data component delivers a fair
565 share of the entries and that only data components with excess capacity
566 deliver more entries than they collect.

567 The delivery size provides a means for both the regulation of network
568 efficiency and the control of the delay between the collection of a value set
569 and its delivery to a web service. The delivery manager attempts to remove
570 enough value sets from the buffer to meet the delivery size before sending
571 them in a single web service call. This may cause entries to remain in the
572 buffer for extended periods of time. This may be acceptable in cases in which
573 the data is being archived, but in cases where the data is consumed as the
574 simulation is being carried out it may be necessary to minimize the duration
575 that an entry may reside in the buffer before it is delivered, at the expense
576 of efficient network utilization. By setting the delivery size to 0, value sets
577 are sent individually as quickly as possible. Note that the delivery module
578 does not attempt to send more value sets than necessary to meet the delivery
579 size because this would require additional parameterization of a maximum,
580 since data may be inserted into the data store at a rate that is faster than

581 the delivery module is able to remove it, preventing the delivery of data.

582 The amount of memory available to the data store is finite and once
583 exhausted the operation of the data store will pause, since it relies on the
584 ability to store data. The effectiveness of the cache increases with the amount
585 of data in the cache, so it is beneficial to allow the data store to be filled as
586 much as possible without impeding the basic operation of the data store. For
587 these reasons a thread within the data manager periodically checks the local
588 buffer and if the data size is greater than 90% of the maximum size, then
589 an eviction is performed. The least accessed 15% of the entries are removed
590 from the buffer that (1) have been sent, or (2) been accessed at least once,
591 or (3) have been downloaded and cached. This prevents cached data, which
592 may or may not be used in the future, from preventing data retrieval from
593 web services or collection from components. In the case that the data store
594 is mostly full of unsent data, downloaded data is purged leaving only the
595 unsent data and if no memory is available, a request will remain in the data
596 queue blocking data components from adding more data requests. As data is
597 delivered to the web services sent data will be purged and data components
598 will again be able to insert requests into the data queue.

599 **3. Performance study**

600 To evaluate the scalability and efficiency of the data component we mea-
601 sured a set of performance metrics for varying numbers of linked models
602 simultaneously executing across a compute cluster.

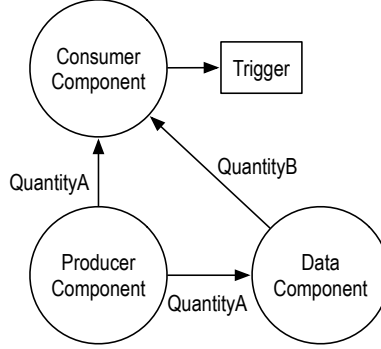


Figure 7: The composition used in the performance study. Arrows indicate the direction of data transfer between the components.

3.1. Baseline configuration

We created a composition that includes a data component linked to two model components such that the data component provides input to a *consumer* component and collects output from a *producer* component as illustrated in Fig. 7. The producer and consumer components serve as placeholders for model components and although they are capable of accepting and providing exchange items and advancing through simulation time, they do not perform any calculations but rather pause for 1 second on each time step, which we refer to as the *time step calculation time*. They discard the data they receive as input and produce constant-valued data as output. We configured the components to advance their simulation time by one day on each time step and configured the composition for a time horizon of 7 months, thus each component performs 212 time steps in each simulation. We empirically determined that using a higher number of time steps does not impact the performance results. The components exchange value sets corresponding

618 to an element set of 1000 elements as this value is a reasonable number of
619 grid cells for environmental models and is small enough to allow for a large
620 number of value sets to be stored in the data store for performance eval-
621 uation. We configured the data component to deliver the output from the
622 producer component to a web service hosted within the compute cluster and
623 provide input to the consumer from the service. The data component thus
624 provides 212 value sets to the consumer component and collects 212 value
625 sets from the producer component during the execution of a single instance
626 of the composition, which we refer to as a *simulation*.

627 A simulation begins when the trigger invokes GetValues on the consumer
628 component for its starting simulation time. The consumer component in turn
629 calls GetValues on both the producer component (which advances its time)
630 and data component which return the requested value sets to the consumer
631 component which then advances its time. The generation of the value set by
632 the producer raises a DataChanged event which causes the data component
633 to call GetValues on the producer component to obtain the new data. The
634 trigger repeatedly invokes GetValues on the consumer until the simulation
635 time of the consumer reaches the configured end time. We refer to the du-
636 ration of time that an instance of the composition spends executing as the
637 *simulation runtime*. As the components perform time steps sequentially and
638 pause for 1 second on each step the simulation runtime of the baseline config-
639 uration would be 424 seconds if the data component and exchanges between
640 the components incur zero time.

641 We executed the simulations on an onsite Linux-based Beowulf compute
642 cluster for the performance study. To limit variability in the results due

643 to differences between the hardware specifications of the compute nodes we
644 utilized a single class of machines that had dual 8-core Intel Xeon E5-2690
645 processors with 64 GB of memory and were connected via gigabit Ethernet.
646 A virtualized Windows-based server with a 4-core 2.7 GHz processor and 8
647 GB of memory hosted the web services and was connected to the compute
648 nodes via gigabit Ethernet.

649 Access to the cluster nodes is provided via a job scheduler (Sun Grid En-
650 gine) to which requests are made for resources (number of processor cores,
651 amount of memory, and maximum runtime) and when they become available
652 a set of scripts provision the nodes as necessary and then execute a set of
653 simulations. The job scheduler executed each set of simulations on multiple
654 nodes utilizing an average of approximately 5 cores on each node. We config-
655 ured the job scheduler to reserve one core for each data manager and one core
656 for every 4 simulations. Scheduling several simulations on each core made it
657 possible to execute a greater number of simulations than there were cores.
658 We verified that collocating several simulations on a single core did not affect
659 the performance results in our experimental configuration (likely because the
660 producer and consumer components require few computer resources).

661 The components are implemented in the C# programming language based
662 on the OpenMI 1.4 software development kit and were executed using the
663 OmiEd application (via the command line). We chose version 1.4 of the
664 OpenMI because the model components in our case study rely on libraries
665 based on this version (Bulatewicz et al., 2013; Castronova and Goodall, 2010)
666 although we are actively developing an implementation of the data compo-
667 nent for version 2.0 of the OpenMI as well. The data manager is imple-

668 mented in the Java programming language because this is the language that
669 Hazelcast is implemented in. The web service is implemented in the PHP
670 programming language and is hosted by the Apache HTTP server. We de-
671 veloped a custom REST-based web service and minimal XML data encoding
672 to avoid any bias that a more complex encoding may have on the results.
673 The web service parses the XML in each request and returns constant-valued
674 data in its response.

675 The delivery size that maximizes throughput when data is sent from the
676 data store to the web service is dependent on several factors including net-
677 work latency, available bandwidth, and software performance. We conducted
678 a series of measurements to empirically determine an appropriate delivery size
679 for our experimental configuration. We found that the maximum throughput
680 between a benchmark application and the web service was 47 MB/s when at
681 least 50 MB of data was sent. As the XML serialization of a 1000-element
682 value set requires 49 KB, the data store would have to send 1498 value sets
683 in each web service call to achieve maximum throughput. If value sets were
684 collected at a rate of 1 per second then they would be delivered every 25
685 minutes. To increase the rate at which value sets were delivered to the web
686 service while still maintaining good network efficiency we decided to use a
687 delivery size of 11 MB which achieves 50% of the maximum network through-
688 put.

689 3.2. Scalability

690 To verify that the design of the data component and data manager are ef-
691 ficient and capable of high performance when there are large numbers of data
692 components we measured the average simulation runtime for varying numbers

693 of simulations and data managers. Each simulation used a unique scenario
 694 identifier so that its input and output value sets were distinct. The results are
 695 presented in Fig. 8 (top). For a given number of data managers, the average
 696 simulation runtime increased as the number of simulations increased. This is
 697 because the data component pauses a simulation while it is waiting for the
 698 data manager to process its requests. When all simulations were supported
 699 by a single data manager the rate at which the average simulation runtime
 700 increased was most closely described by the function $0.0003 \times (n^{1.5})$ where n
 701 is the number of simulations ($R^2 = 0.996$). In the case of 4 data managers
 702 the rate at which the runtime increased was most closely described by the
 703 function $0.4 \times e^{0.001 \times n}$ ($R^2 = 0.994$). Based on these functions we expect the
 704 average simulation runtime to increase at a greater rate when there are more
 705 than 1000 simulations.

706 The number of data managers supporting the simulations had a varying
 707 impact on the average simulation runtime. Increasing the number of data
 708 managers from 1 to 4 significantly reduced the average simulation runtime,
 709 while increasing further to 8 only resulted in a small reduction for higher
 710 numbers of simulations. Increasing the number of data managers further to
 711 16 slightly increased the average simulation runtime due to the additional
 712 overhead incurred by the management of the distributed data store. We
 713 therefore estimate that the ideal ratio for our experimental configuration
 714 was approximately 1 data manager per 250 simulations.

715 The duration of time between when a value set is collected by a data
 716 component and when it is delivered to the web service, the *delivery time*,
 717 is a function of (1) the rate at which value sets are collected by the data

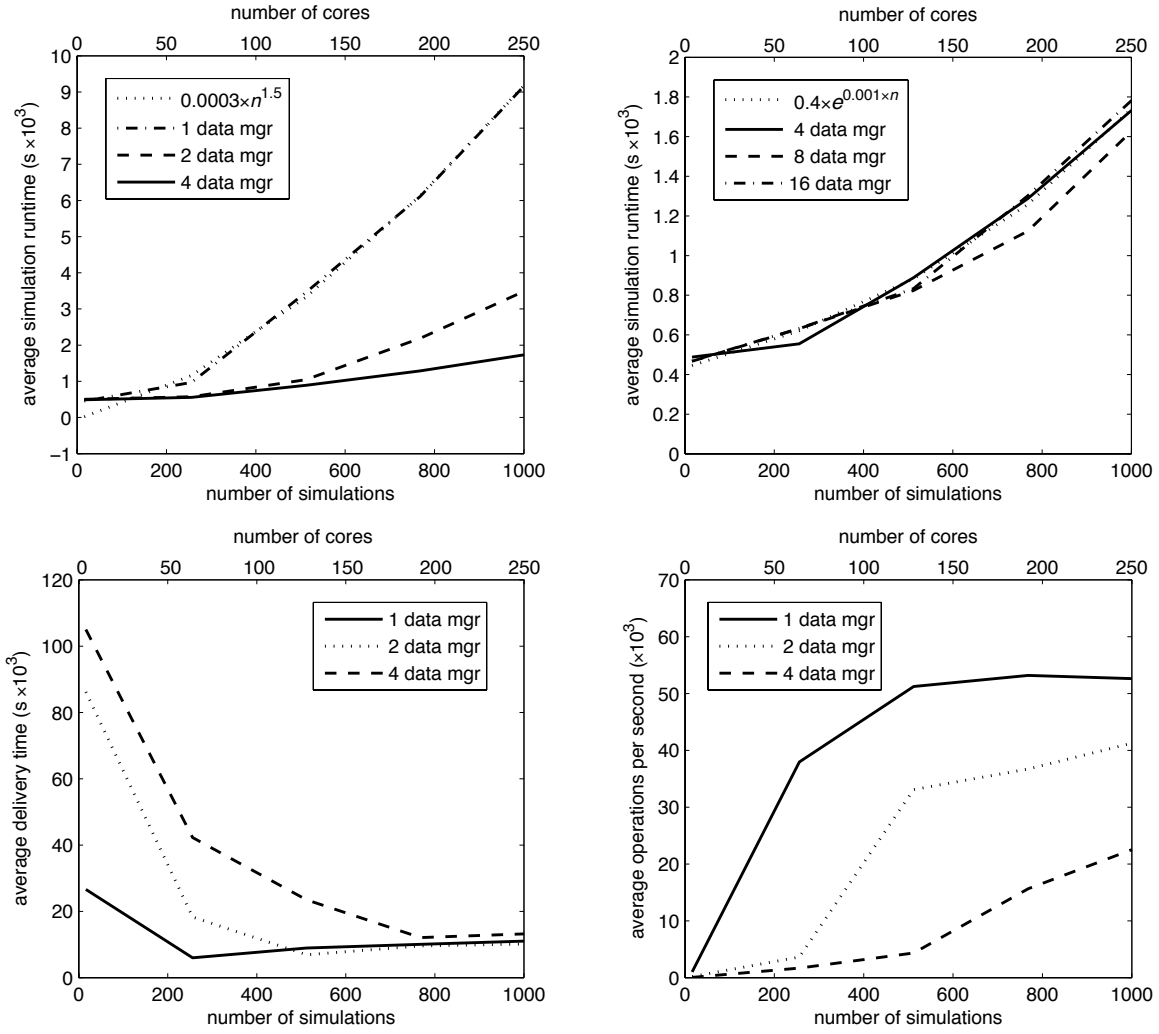


Figure 8: Scalability results.

718 components connected to a data manager, (2) the size of the value sets,
719 and (3) the delivery size that the data manager is configured to use. For our
720 experimental configuration the size of a value set with 1000 elements encoded
721 in XML was 49 KB and the delivery size was 11 MB, so the data manager only
722 sent data to the web service when it found 230 unsent value sets in its local
723 data store. For low numbers of simulations the rate at which value sets were
724 collected and added to the data store was low, resulting in the data manager
725 delaying the sending of the data, as shown in Fig. 8 (bottom-left). Higher
726 numbers of data managers amplified this effect, further reducing the rate at
727 which value sets were added to each data manager and further increasing the
728 delivery time. For high numbers of simulations, the delivery time was low due
729 to the faster rate at which value sets were collected by data components and
730 added to the data store which ensured there was always a sufficient number
731 of unsent value sets. In this case the delivery time increased slightly as the
732 number of simulations increased because greater numbers of value sets in the
733 data store increased the search time for locating unsent value sets.

734 The average number of operations per second (`put`, `get`, and `remove`) per-
735 formed by Hazelcast on the hashmap that stores the value sets increased
736 with the number of compositions as shown in Fig. 8 (bottom-right). The
737 maximum number of operations per second reached in our experimental con-
738 figuration was approximately 50000 when a single data manager was used
739 and was significantly less for greater numbers of data managers.

740 3.3. *Caching*

741 To investigate the effect of caching on the performance of the data compo-
742 nent we executed 16 simulations connected to a single data manager and ad-

743 justified the baseline configuration in two ways that facilitate caching. Caching
744 is most effective in reducing the average simulation runtime when (1) the in-
745 put value sets needed by the simulations are in the data store prior to when
746 they are requested, and (2) the time required to retrieve a value set from a
747 web service is non-zero. We therefore configured the simulations such that
748 one simulation was executed before the others and the web service waited
749 3 s before responding to each request (the *web service response time*) for
750 a value set. We refer to the time required to retrieve a value set from a
751 web service, including the time spent generating the request, calling the web
752 service, and processing the response, as the *retrieval time*. We measured
753 the average simulation runtime and amount of data retrieved from the web
754 service for both the baseline and alternate configurations with and without
755 caching (Fig. 9). For the “caching” scenario we assigned a common scenario
756 identifier to all the simulations causing them to all request identical input
757 data and in the “no caching” scenario we assigned different identifiers caus-
758 ing each to request distinct input data. The average simulation runtime for
759 the alternate configuration in the “no caching” scenario was approximately
760 2.5 times higher than in the baseline configuration due to the additional 3
761 s delay incurred for each of the 212 web service requests to retrieve data.
762 For the baseline configuration the average simulation runtime was similar in
763 both the “caching” and “no caching” scenarios because the retrieval time
764 was very low. For the alternate configuration, however, the average simu-
765 lation runtime in the “caching” scenario was 59.8% lower than in the “no
766 caching” scenario because the retrieval time was higher (approximately 3 s).
767 In both configurations the amount of data retrieved from the web service in




Figure 9: Caching results.

the “caching” scenario was 93.8% lower than in the “no caching” scenario because the simulations requested identical input data and thus only 1/16th the amount of data had to be retrieved from the web service. Thus the reduction in the average simulation runtime afforded by caching is dependent upon the magnitude of the retrieval time, while the reduction in the amount of data transferred is a function of the size of the value set. In general, the amount by which the average simulation runtime can be reduced is the percentage of the runtime that is due to the retrieval of the data (i.e. the retrieval time).

3.4. Prefetching

To investigate the effect that prefetching has on the performance of the data component we enabled the prefetching feature and measured the average runtime for 16 and 256 simulations connected to a single data manager where the simulations were configured such that the time step calculation time and web service response time were the same (2 s). This allows for the retrieval




Figure 10: Prefetching results.

of a value set (which involves generating the request, calling the web service, and processing the response) to take place while the components (producer and consumer) are calculating their time steps. The results of the alternate configuration as compared to the baseline configuration are presented in Fig. 10. In all cases, prefetching reduced the average simulation runtime because some portion of the data retrieval was performed during the time step calculation rather than waiting until each time step was calculated before requesting the data. In the baseline configuration with 16 simulations, prefetching resulted in a small decrease in the average simulation runtime (3.4%) because the retrieval time was small and hence only a small amount of time was saved by performing the retrieval during the time step calculation. In the alternate configuration with 16 simulations, prefetching resulted in a large decrease in the runtime (43.0%) because the retrieval time was high (due to the increased web service response time) and thus a significant

796 amount of time was saved by performing the retrieval during the time step
797 calculation. In the case of 256 simulations, the reduction in the average sim-
798 ulation runtime was less in both configurations due to the number of prefetch
799 failures that were a result of the data component’s prioritization of requests
800 for data over requests for prefetching data.

801 The reduction in the average simulation runtime possible by prefetching
802 is thus a function of the relative difference between the retrieval time and
803 the length of time between subsequent requests made to the data component
804 (i.e. the sum of the calculation times of all components for a time step). In
805 cases in which the retrieval time is less than the total amount of time the
806 components spend calculating a time step, the runtime of the simulation is
807 not affected by the web service calls and is masked by the time step calcula-
808 tion time (assuming the data manager does not reject the prefetch requests).
809 In general, the amount by which the average simulation runtime can be re-
810 duced is the percentage of the runtime that is due to the retrieval of the data
811 (i.e. the retrieval time). In cases in which the retrieval time is greater than
812 the time spent calculating time steps, the average simulation runtime will
813 increase proportionally according on the relative difference between them.

814 **4. Case study: A groundwater sustainability challenge**

815 To demonstrate how the data component may be incorporated into a
816 modeling study we present how we are utilizing it in an ongoing case study
817 to provide input data from an online database to the components of a linked
818 model executing on a desktop computer.




Figure 11: Conceptualization of an irrigated agricultural system.

819 4.1. Study area

820 The communities of western Kansas in the Central Plains of the United
 821 States have relied upon the availability of groundwater for irrigated agricul-
 822 ture for 50 years (Fig. 11). The rate at which water is extracted from the
 823 Ogallala Aquifer underlying the region has exceeded the rate at which it nat-
 824 urally recharges resulting in a gradual decline in the volume of water stored
 825 in the aquifer. In some areas it is no longer possible to extract water from
 826 the aquifer due to the decreased saturated thickness, a trend that will con-
 827 tinue to spread throughout the region unless agricultural practices transition
 828 to sustainable rates of water consumption. As this transition impacts the
 829 closely intertwined economy and ecology of the region it is essential that it
 830 be guided by multidisciplinary integrated assessment.

831 We consider the relevant natural and human processes in this system
 832 to be (1) the movement and volume of groundwater, (2) the choice of crop
 833 planted, and (3) the growth of the plants. Building on previous experience in
 834 integrated modeling for irrigated agriculture (Bulatewicz et al., 2010, 2013)
 835 we have developed three new model components that simulate these processes

836 and have created a prototype linked model integrating them.

837 *4.2. Model components*

838 The crop choice component is an iterative Positive Mathematical Pro-
839 gramming (PMP) model (Howitt, 1995) that simulates farmers allocation of
840 arable land to different crops. The model operates on an annual time step,
841 with each execution predicting farmers choices in a single growing season.
842 In addition to harvested crop prices and crop-specific costs of production,
843 the model accepts as inputs the current (county average) depth to water and
844 saturated thickness of the aquifer. Depth to water affects water extraction
845 costs, while saturated thickness affects the pumping rate of wells, which in
846 turn creates an upper bound on the annual extraction of irrigation water.
847 The model simulates land allocations as the solution to a constrained opti-
848 mization problem that represents farmers profit-maximizing mix of land uses,
849 given price conditions, water extraction costs, and the constraints on water
850 and land availability. The component has input exchange items for satu-
851 rated thickness and depth to water and output items for the predicted acres
852 planted to each crop (wheat, corn, sorghum, soybeans, and alfalfa). Details
853 on the model development, calibration, and data sources are in Clark (2008)
854 and Garay et al. (July 2010). The model is implemented in MATLAB and
855 interoperability with the OpenMI is provided by the Simple Script Wrapper
856 (Bulatewicz et al., 2013).

857 The groundwater model provides the groundwater elevation (head) as
858 a function of space and time. For this application, we have developed an
859 OpenMI component for the Hydrologic Response Function (HRF) approach
860 from Steward et al. (2009). Briefly, the aquifer is treated as a sloping base

861 with rectangular cells used to gather pumped water-use within cells that
862 contain uniform aquifer properties Steward (2007). Our OpenMI code fully
863 implements the HRF equations and enables the drawdown associated with
864 pumping to be communicated with neighboring cells. This approach was
865 chosen as it has been shown to accurately reproduce the cones of depression
866 formed by groups of wells in the study area (Steward et al., 2009), and the
867 code executes much faster than other approaches based upon the Analytic
868 Element Method (Steward et al., 2008) and finite gridded domain approaches
869 (Steward and Allen, 2013). We also incorporated the groundwater added to
870 the domain through leakage from surface water identified by Ahring and
871 Steward (2012). This was accomplished by adding recharge to cells that
872 coincide with rivers and adjusting the recharge rates until groundwater sur-
873 faces matched observations (see Steward et al. (2009) for a discussion of these
874 recharge volumes). The component has an input exchange item for irrigated
875 water-use and output exchange items for saturated thickness and depth to
876 water. The model is implemented in Scilab and interoperability with the
877 OpenMI is provided by the Simple Script Wrapper.

878 The crop production component provides crop yield and irrigated water
879 use data as simulated by the Erosion-Productivity Impact Calculator (EPIC)
880 model (Williams, 1995). EPIC is a process-based generalized crop model that
881 simulates daily crop growth by predicting plant biomass through the simu-
882 lation of carbon fixation by photosynthesis, maintenance respiration, and
883 growth respiration. In a previous work (Bulatewicz et al., 2009) we enabled
884 this legacy model to work with OpenMI by creating a wrapper component
885 that executed the unmodified model program on-demand. For this new com-




Figure 12: Component linkages. Data is transferred in the direction of the arrows.

886 ponent we took an alternative approach to model reuse in which we executed
 887 the original model program for all combinations (2500) of the primary model
 888 inputs of interest (soil, crop, management, weather) and created an index of
 889 the model output data that the component utilizes to lookup and provide
 890 the data to other model components. The input exchange items of the com-
 891 ponent are acreage per crop, soil ID, and weather ID. The output items are
 892 crop yield and irrigated water-use. The model operates on an annual time
 893 step over a 2-dimensional grid and is implemented in C# using the Simple
 894 Model Wrapper (Castronova and Goodall, 2010). We calibrated the model
 895 for use in western Kansas in an earlier work (Bulatewicz et al., 2009).

896 4.3. Linked model design

897 There are a total of 14 links between the models as illustrated in Fig. 12.
 898 The linked model prototype uses an element set consisting of a single ele-

ment that represents Seward County in southwestern Kansas. At the start of each year the crop production model requests the planted acreage of each crop from the crop choice model and the soil and weather information from the data component. The data component retrieves the data from an online database and provides it to the crop production model. The crop choice model requests the saturated thickness and depth to water from the groundwater model for the previous year which in turn requests the irrigated water use from the crop production model for that year. After receiving the response the groundwater model calculates the new saturated thickness and depth to water and provides them to the crop choice model which in turn predicts the crop acreages and provides them to the crop production model. The crop production model then calculates the crop yield and irrigated water-use for the current year.

4.4. *Using the data component*

To create the linked model we began by adding the 3 models to a new composition using the OmiEd application and then added the appropriate links between them. We then configured the data component by (1) defining the necessary output exchange items, and (2) specifying the information about the web service from which they should be retrieved. The exchange items and web service information are defined within the data component's configuration file as shown in Fig. 13. The format of the configuration file is based on that of the Simple Model Wrapper and was extended to include web service information. The element set and quantity of each exchange item (as well as the units information, not shown in the figure) is listed in the configuration file as well as the type, URL, and list of quantities provided

```

<Configuration>
  <ExchangeItems>
    <OutputExchangeItem>
      <ElementSetID>Seward</ElementSetID>
      <Quantity><ID>WeatherStationID</ID></Quantity>
    </OutputExchangeItem>
    <OutputExchangeItem>
      <ElementSetID>Seward</ElementSetID>
      <Quantity><ID>WeatherDataID</ID></Quantity>
    </OutputExchangeItem>
    <OutputExchangeItem>
      <ElementSetID>Seward</ElementSetID>
      <Quantity><ID>SoilTypeID</ID></Quantity>
    </OutputExchangeItem>
  </ExchangeItems>
  <TimeHorizon>
    <StartDateTime>2012-01-01T00:00:00</StartDateTime>
    <EndDateTime>2040-08-01T00:00:00</EndDateTime>
    <TimeStepInSeconds>86400</TimeStepInSeconds>
  </TimeHorizon>
  <WebServices>
    <WebService>
      <Type>WaterOneFlow</Type>
      <Url>http://host/Baseline/Service_10.asmx</Url>
      <RetrievableQuantities>
        <QuantityID>WeatherStationID</QuantityID>
        <QuantityID>WeatherDataID</QuantityID>
        <QuantityID>SoilTypeID</QuantityID>
      </RetrievableQuantities>
    </WebService>
  </WebServices>
</Configuration>

```

Figure 13: The data component configuration file (partial).

924 and accepted by each web service. The quantity ID specified in each output
925 exchange item must appear in the list of `RetrievableQuantities` for one of the
926 web services and each input item must appear in the `DeliverableQuantities`
927 After creating the configuration file we added the data component to the
928 composition and added links from each of its output exchange items to the
929 appropriate input of the crop production component.

930 The URL specified in the configuration file is that of a CUAHSI HIS
931 WaterOneFlow web service that was hosted on a virtualized server that we
932 setup on the cluster network and was publicly accessible via the Internet.
933 The web service was connected to a SQL Server database that was also
934 hosted on the server and used the Observations Data Model (Horsburgh et al.,
935 2008), which is a relational data model for the storage and retrieval of time
936 series hydrologic observations and associated metadata. The data component
937 provides interoperability between the ODM/WaterOneFlow web service and
938 the OpenMI by mapping their respective data models to one another in a
939 similar way as Castronova et al. (2013b) (e.g. mapping quantities to variables
940 and sites to elements). Thus, the IDs of the elements within the element sets
941 of the input and output exchange items specified in the configuration file
942 must exist as sites in the database (mapped to `SiteCode`) and the quantity
943 IDs of the exchange items must exist as variables in the database (mapped
944 to `VariableName`). The WaterOneFlow web service returns data as time
945 series whereas exchanges between OpenMI components require space series,
946 so had there been multiple elements in the element set the data component
947 would have made multiple web service calls for each time step (one for each
948 element).

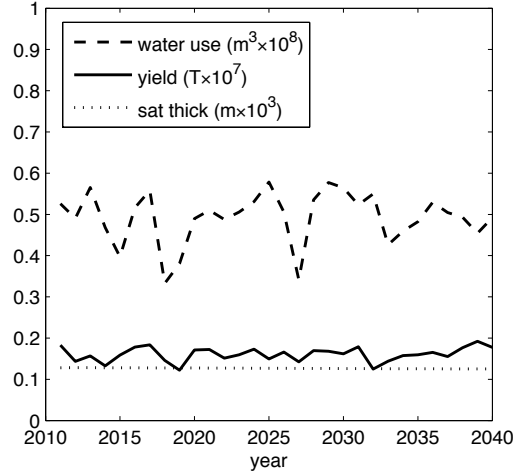


Figure 14: Output from the linked model for 3 indicators.

949 The output of the linked model simulation for 3 indicators is shown in
 950 Fig. 14 (does not include wheat data as it is not currently available in the crop
 951 production component). The county-wide total crop yield and irrigated water
 952 use varied from year to year according to the weather while the saturated
 953 thickness of the aquifer decreased at a constant rate.

954 5. Conclusions

955 We have presented the design of a general-purpose data component for
 956 the OpenMI, evaluated its performance, and demonstrated its application
 957 in a modeling study. The data component can mitigate data management
 958 challenges in modeling and simulation by serving as a bridge between model
 959 components and online services minimizing the reliance on data files and
 960 ad-hoc scripting. We adapted three techniques to the unique design of the
 961 OpenMI to enable efficient operation: caching, prefetching, and buffering,

Table 1: Summary of performance study results. (*estimated)

Technique	Improvement	Dimension	Cost Per Time Step
Caching	1% - 29%	Time	0.14 s - 3.78 s
Prefetching	3% - 43%	Data	131 KB
Buffering*	5% - 33%		

962 making it suitable for use on both desktop computers and high-performance
 963 compute clusters.

964 The data component is added to a composition and linked to model com-
 965 ponents in same way that model components are linked to one another. The
 966 scientist configures, and re-configures the data component for the input and
 967 output exchange items necessary for any given set of model components based
 968 on the data available via web services. It relies on a data manager program
 969 that communicates with web services and manages a distributed data store
 970 shared across all the data managers executing on a compute cluster. The
 971 data retrieved from web services is cached in the data store and the data
 972 collected from model components is buffered in the data store before being
 973 delivered to web services.

974 We evaluated the performance of the data component in terms of scal-
 975 ability and the effectiveness of caching and prefetching in minimizing the
 976 simulation runtime. The results are summarized in Table 1. The increase in
 977 simulation runtime incurred by the data component (as compared to using
 978 local data files) ranged from 0.14 s for 16 simulations to 3.78 s for 1000 simu-
 979 lations for each time step. The data transferred to and from the web service

980 was 131 KB per time step for a value set of 1000 values.

981 Caching can have a significant impact on the runtime of simulation in
982 some cases and little or no impact in other cases. We demonstrated this via
983 two configurations which resulted in a 1% to 29% reduction in the average
984 simulation runtime. This range only serves as an example of possible perfor-
985 mance, as the actual impact is a direct result of the retrieval time and the
986 number of times model components request identical data.

987 Prefetching can also have a significant impact on the runtime of a simula-
988 tion, but through different means than caching. Prefetching is only effective
989 when the time step processing time of a model component is comparable to
990 the retrieval time thus making it possible to overlap the model execution
991 with the retrieval of data. We demonstrated this via two configurations in
992 which the runtime was reduced by only 3% when there was no overlap and
993 43% when there was full overlap. In addition, prefetching is less effective
994 when a data manager is under high utilization.

995 Buffering always reduces the runtime of a simulation where the reduction
996 is directly proportional to the web service response time. Although the im-
997 pact of buffering on the simulation runtime cannot be measured empirically
998 (because buffering is inherent in the design of the data manager) its impact
999 can be estimated by adding the time spent sending the data on each time
1000 step. For the experimental configuration, in which the model components
1001 spend 2 s processing each time step, if the time spent sending data on each
1002 time step was 0.1 s then the reduction in runtime due to buffering would
1003 be 5% whereas if the time spent sending data was 1.0 s then the reduction
1004 would be 33%.

1005 Based on the results of the performance study, it can be expected that the
1006 simulation runtime will increase as the number of simulations is increased,
1007 and that buffering always results in improved runtimes while caching and
1008 prefetching may result in improvements depending upon the situation. Over-
1009 all, the runtime overhead of the data component is primarily determined by
1010 the web service response time and to a lesser degree the time step processing
1011 time of the model components and the value set size (as the data transfer size
1012 and parsing time are influenced by it). As the web service response time in-
1013 creases, the runtime increase incurred by the data component becomes larger
1014 while at the same time the benefit of buffering and the potential benefit of
1015 caching and prefetching increase as well. In general, the percentage of the
1016 runtime that is due to the web service calls is equivalent to the reduction that
1017 would be achieved in cases in which caching and prefetching are effective.

1018 We therefore conclude that the design of the data component meets the
1019 three requirements identified in Section 2. Standards for web services make
1020 it possible for the component to be configured and reconfigured as necessary
1021 to meet the needs of different linked model configurations and different web
1022 services. The increase in simulation runtime incurred by the data component
1023 (as compared to using local data files) is reasonable and in some cases can be
1024 eliminated by caching and prefetching data. The overall performance of the
1025 data component is reasonable for large numbers of simultaneous simulations.

1026 As the importance of data availability, interoperability, and transparency
1027 continue to rise, so too does the need for software tools that facilitate these.
1028 General-purpose tools that intelligently and efficiently provision, collect, and
1029 deliver data will become an essential part of OpenMI linked models on desk-

1030 top computers and compute clusters alike and this work provides a starting
1031 point for such tools.

1032 **Acknowledgements**

1033 This work was supported by the National Science Foundation (grants
1034 GEO0909515, EPS0919443, EPS1006860, CNS1126709) and the Ogallala
1035 Aquifer Project of the USDA/ARS. Access to the Beocat compute cluster at
1036 the Dept. of Computing and Information Sciences at Kansas State University
1037 was appreciated. Any findings, opinions, conclusions, or recommendations
1038 expressed herein are those of the authors and do not necessarily reflect the
1039 views of any funding units.

1040 **Appendix A. Software availability**

Software name: DataComponent

Developer: GRowE/Kansas State University

Contact address: 234 Nichols Hall, Kansas State University,
Manhattan, KS, 66502, 785-532-6350

E-mail: tombz@ksu.edu

Year first available: 2013

1041 Hardware required: Architecture independent

Required software: Windows/Linux

Program language: C#

Program size: 2 MB

Availability: Download available under MIT License at:

<http://code.google.com/p/data-component>

Cost: Free

1042 **References**

1043 Ahring, T.S., Steward, D.R., 2012. Groundwater surface water interactions
1044 and the role of phreatophytes in identifying recharge zones. *Hydrology and*
1045 *Earth System Sciences* 16, 4133–4142.

1046 Bavoil, L., Callahan, S.P., Scheidegger, C.E., Vo, H.T., Crossno, P.J., Silva,
1047 C.T., Freire, J., 2005. Vistrails: Enabling interactive multiple-view visu-
1048 alizations. *Visualization Conference, IEEE* 0, 18.

1049 Bulatewicz, T., Allen, A., Peterson, J.M., Staggenborg, S., Welch, S.M.,
1050 Steward, D.R., 2013. The simple script wrapper for openmi: Enabling

1051 interdisciplinary modeling studies. *Environmental Modelling & Software*
1052 39(2013), 283–294.

1053 Bulatewicz, T., Andresen, D., 2011. Efficient data access for open model-
1054 ing interface (openmi) components, in: *Proceedings of the International*
1055 *Conference on Parallel and Distributed Processing Techniques and Appli-*
1056 *cations (PDPTA) Volume 1*, ed. H. R. Arabnia, CSREA Press, Las Vegas,
1057 Nevada, USA, July 18-21, pp. 822–828.

1058 Bulatewicz, T., Andresen, D., 2012. Efficient data collection from Open
1059 Modeling Interface (OpenMI) components, in: *Proceedings of the Interna-*
1060 *tional Conference on Parallel and Distributed Processing Techniques and*
1061 *Applications (PDPTA) Volume 1*, ed. H. R. Arabnia, CSREA Press, Las
1062 Vegas, Nevada, USA, July 16-19, pp. 53–59.

1063 Bulatewicz, T., Jin, W., Staggenborg, S., Lauwo, S.Y., Miller, M., Das, S.,
1064 Andresen, D., Peterson, J., Steward, D.R., Welch, S.M., 2009. Calibration
1065 of a crop model to irrigated water use using a genetic algorithm. *Hydrol*
1066 *Earth Syst Sci* 13, 1467–1483.

1067 Bulatewicz, T., Yang, X., Peterson, J.M., Staggenborg, S., Welch, S.M.,
1068 Steward, D.R., 2010. Accessible integration of agriculture, groundwater,
1069 and economic models using the open modeling interface (openmi): method-
1070 ology and initial results. *Hydrology and Earth System Sciences* 14(3),
1071 521–534.

1072 Cassandra, 2013. Cassandra. [Http://cassandra.apache.org](http://cassandra.apache.org).

- 1073 Castronova, A.M., Goodall, J.L., 2010. A generic approach for developing
1074 process-level hydrologic modeling components. *Environmental Modelling
1075 & Software* 25(2010), 819–825.
- 1076 Castronova, A.M., Goodall, J.L., Elag, M.M., 2013a. Models as web services
1077 using the Open Geospatial Consortium (OGC) Web Processing Service
1078 (WPS) standard. *Environmental Modelling & Software* 41(0), 72–83.
- 1079 Castronova, A.M., Goodall, J.L., Ercan, M.B., 2013b. Integrated modeling
1080 within a Hydrologic Information System: An OpenMI based approach.
1081 *Environmental Modelling & Software* 39(0), 263–273. Thematic Issue on
1082 the Future of Integrated Modeling Science and Technology.
- 1083 Chandrasekaran, S., Silver, G., Miller, J., Cardoso, J., Sheth, A., 2002. Web
1084 service technologies and their synergy with simulation. *Winter Simulation
1085 Conference* 1, 606–615.
- 1086 Clark, M.K., 2008. Effects of high commodity prices on western kansas crop
1087 patterns and the ogallala aquifer. Unpublished M.S. Thesis. Department
1088 of Agricultural Economics, Kansas State University.
- 1089 Garay, P.V., Peterson, J.M., Golden, B.B., Smith, C.M., July 2010. Dis-
1090 aggregated spatial modeling of irrigated land and water use, in: *Selected
1091 Presentation at the Agricultural and Applied Economics Association An-
1092 nual Meeting*, Denver.
- 1093 Globus Online, 2013. Globus online. <http://www.globusonline.org>.
- 1094 Globus Toolkit, 2013. GridFTP user’s guide. <http://www.globus.org>.

1095 Goodall, J.L., Robinson, B.F., Castronova, A.M., 2011. Modeling water
1096 resource systems using a service-oriented computing paradigm. Environ-
1097 mental Modelling & Software 26(5), 573–582.

1098 Gregersen, J.B., Gijsbers, P.J.A., Westen, S.J.P., 2007. OpenMI: Open mod-
1099 eling interface. J Hydroinform 9(3), 175–191.

1100 Horak, J., Orlik, A., Stromsky, J., 2008. Web services for distributed and
1101 interoperable hydro-information systems. Hydrol. Earth Syst. Sci. 12, 635–
1102 644.

1103 Horsburgh, J.S., Tarboton, D.G., Maidment, D.R., Zaslavsky, I., 2008. A
1104 relational model for environmental and water resources data. Water Re-
1105 sources Research 44, W05406.

1106 Howitt, R.E., 1995. Positive mathematical programming. American Journal
1107 of Agricultural Economics 77, 329–342.

1108 Hull, D., Wolstencroft, K., Stevens, R., Goble, C., Pocock, M., Li, P., Oinn,
1109 T., 2006. Taverna: a tool for building and running workflows of services.
1110 Nucleic Acids Research 34(Web Server issue), 729–732.

1111 J. S. Horsburgh, D. G. Tarboton, M.P.D.R.M.I.Z.D.V., Whitenack, T., 2009.
1112 An integrated system for publishing environmental observations data. En-
1113 vironmental Modelling & Software 24(8), 879–888.

1114 Maidment, D.R., 2008. Bringing water data together. Journal of Water
1115 Resources Planning and Management 134(2), 95–96.

1116 Memcached, 2013. Memcached. [Http://memcached.org](http://memcached.org).

- 1117 OpenMI Association, 2010. News and announcements.
 1118 <http://www.openmi.org/archives/archived-news-and-announcements/2010>.
- 1119 Ozturk, T., 2010. Scalable data structures for java, in: Devovx, Metropolis
 1120 Antwerp Belgium.
- 1121 Portele, C., 2007. Opengis geography markup language (GML) encoding
 1122 standard, open geospatial consortium. OGC 07-036.
- 1123 Pullen, J.M., Brunton, R., Brutzman, D., Drake, D., Hieb, M., Morse, K.L.,
 1124 Tolk, A., 2005. Using web services to integrate heterogeneous simulations
 1125 in a grid environment. *Future Gener. Comput. Syst.* 21, 97–106.
- 1126 Rajasekar, A., Wan, M., Moore, R., Schroeder, W., 2006. A prototype rule-
 1127 based distributed data management system, in: HPDC workshop on Next
 1128 Generation Distributed Data Management, Paris, France.
- 1129 Shasharina, S., Li, C., Pundaleeka, R., Wang, N., Wade-Stein, D., Schissel,
 1130 D., Peng, Q., 2006. HDF5WS – web service for remote access of simulation
 1131 data. *APS Meeting Abstracts*, 2014.
- 1132 Steward, D.R., 2007. Groundwater response to changing water-use practices
 1133 in sloping aquifers. *Water Resources Research* 43, W05408:1–12.
- 1134 Steward, D.R., Allen, A.J., 2013. The analytic element method for rectan-
 1135 gular gridded domains and application to the ogallala aquifer. *Advances*
 1136 *in Water Resources* in review.
- 1137 Steward, D.R., Le Grand, P., Janković, I., Strack, O.D.L., 2008. Analytic
 1138 formulation of Cauchy integrals for boundaries with curvilinear geome-

- 1139 try. Proceedings of the Royal Society of London, Series A, Mathematical,
1140 Physical and Engineering Sciences 464, 223–248.
- 1141 Steward, D.R., Yang, X., Chacon, S., 2009. Groundwater response to chang-
1142 ing water-use practices in sloping aquifers using convolution of transient
1143 response functions. Water Resources Research 45, W02412:1–13.
- 1144 Tarboton, D.G., Horsburgh, J.S., Maidment, D.R., Whiteaker, T., Zaslavsky,
1145 I., Piasecki, M., Goodall, J., Valentine, D., Whitenack, T., 2009. Develop-
1146 ment of a community hydrologic information system, pp. 988–994.
- 1147 Vretanos, P.A., 2010. OpenGIS Web Feature Service 2.0 Interface Standard
1148 - OGC 09-025r1 and ISO/DIS 19142. Open Geospatial Consortium Inc. .
- 1149 Williams, J.R., 1995. The epic model. Computer Models of Watershed
1150 Hydrology , Chapter 25, 909–1000.
- 1151 Zaslavsky, I., Valentine, D., Whiteaker, T., 2007. CUAHSI WaterML, OGC
1152 07-041r1. Open Geospatial Consortium Inc. .