

L-DSMS – A Local Data Stream Management System

Christian Hänsel, Hans Jürgen Ohlbach, Edgar Stoffel

Department of Computer Science, University of Munich
{haensel,ohlbach,stoffel}@ifi.lmu.de

Abstract. L-DSMS is a Local Data Stream Management System. It is a Java Program which can read an XML-file with a description of a network of processing nodes for streaming data. L-DSMS automatically combines all the processing nodes into a single Java program which then processes the data. L-DSMS has a number of predefined nodes, together with an interface for implementing new processing nodes. The generated network can be remotely monitored and reconfigured by a client, Visu-L-DSMS. An example application of L-DSMS is the transformation of RDS-TMC traffic messages into KML, which, in turn, can be visualised by Google Earth.

1 Introduction

Computer programs can operate in quite different modes. The simplest mode is: they get started, read some data, compute something, output some results and terminate. Another mode is the server mode: they wait for some input from a user or a client, do something, and then wait for the next input. A further mode is the *streaming mode*: they get permanently fed with data, process it, and dump the results somewhere, while the next input is already waiting. The data may come from sensors, or, nowadays more typically, from sources on the Internet.

Data Stream Management Systems [1,2] have been developed to connect networks (grids) of computers in such a way that each computer can receive data, process them and forward the results to some other computer. The ideas behind L-DSMS are quite similar to these kind of Data Stream Management Systems. The main difference is that the processing takes place within a single computer, or, more precisely, within a single Java Virtual Machine (JVM). It realises a pipe-and-filter architecture within a JVM. Several instances of L-DSMS can of course run on different computes and be connected in the same way as Data Stream Management Systems.

Programs which process data streams on a single computer can be implemented in different ways.

- The easiest way is to implement a concrete application as a single monolithic program.
- A more flexible and comfortable way is to split the program into separate “processing nodes”. Processing nodes receive data from some standard input

interface, process them in some way, and deliver the results to some standard output interface. A particular application can then be realised by writing a program that loads the necessary processing nodes, and connects them such that the data can be forwarded from node to node. This seems to be the approach of the MeDICi Integration Framework [3]. Their system allows one to connect processing nodes which are even written in different programming languages.

- Alternatively to writing an application specific program that connects the necessary nodes, one can specify the network of processing nodes in an XML-file. A general network configurator can then read such a specification, load the necessary processing nodes, connects them and starts the processing. The network configurator is completely independent of the actual application. Therefore no programming is necessary any more for generating special applications. This is the approach of the L-DSMS system, presented in this paper.
- Instead of specifying the network configuration in an XML-file, one can write a user interface that allows one to specify a particular network by arranging icons on the screen. Yahoo pipes (http://www.jumpcut.com/pipes_team) is a nice example for this approach. It has a number of predefined nodes which process, for example, news feeds. A user can use it to specify his particular view on Internet messages.

The Visu-L-DSMS component of L-DSMS goes halfway this line. Visu-L-DSMS is a client program which can visualise and monitor L-DSMS networks running on some remote servers. It can also be used to change parameters of processing nodes in a running L-DSMS application, but so far it cannot be used to configure a new network.

The prototype application, which has been used to test L-DSMS is a system which receives traffic information via RDS-TMC radio signals and converts them into KML-files, which in turn, can be displayed by Google Earth. This way, Google Earth is able to integrate traffic information into its displayed road maps. The URL is <http://nihiru.pms.ifi.lmu.de/ge-tmc-server/GeoData.php?type=1&format=1>.

This paper describes the general ideas and concepts. The technical details and the code are available from the L-DSMS home page (<http://www.pms.ifi.lmu.de/reverse-wga1/ldsms/>) and in a deliverable of the EU Network of Excellence REVERSE [4].

2 Node Types

The L-DSMS system distinguishes three different node types: *sources*, *drains* and *general processing nodes*. Sources push data into the system (usually by reading them from some external source). Drains receive data from the system, and usually forward them to some destination outside the system. Finally, processing nodes receive data, process them and forward them to other nodes.

2.1 Source

A source is always at the head of a production path and produces the data that needs to be processed. Each source may produce data in a different way. The L-DSMS core package contains sources that read data from files, sockets or external hardware (cf. Sect. 7). In most of the cases, a source receives its data from outside L-DSMS (e.g. from a sensor or from some data source on the web). L-DSMS does not care about where this data originally came from.

To provide additional information about the produced data, a source can create optional meta information for each produced data package. A data packet together with its optional meta data forms the output of a source. A source needs to have at least one connected component that receives the output. There can be an arbitrary number of drains or processing nodes as connected components of a source node (cf. Fig. 1). Which drains or nodes are connected to a source, is specified in the configuration file (cf. Sect. 8). When connecting a node to a source node, one has to ensure that the data and meta data output types of the source are compatible to the data and meta data input types of all connected components. This means, if source S_1 produces data of type t_1 and meta data of type t_2 , all connected nodes have to accept data of type t_1 and meta data of type t_2 . If not, L-DSMS produces an error message and terminates. A drain or node always accepts data and meta data of any type if its own data input type is *java.lang.Object*.

The configuration file contains, beside the relation between sources and their connected components, some source specific properties. Two properties can be specified for all sources. The first property is the mandatory `class` attribute with value 'source', 'node' or 'drain'. The second property is an optional `name` attribute. If a name is specified, it has to be unique because this name is used by other nodes to identify sources as additional parents. All other properties are source specific and are listed in the detailed description for each source. The `ByteArrayFileSource` (cf. Sect. 7.1), for example, has the additional attributes `file`, `delay` and `repeat`.

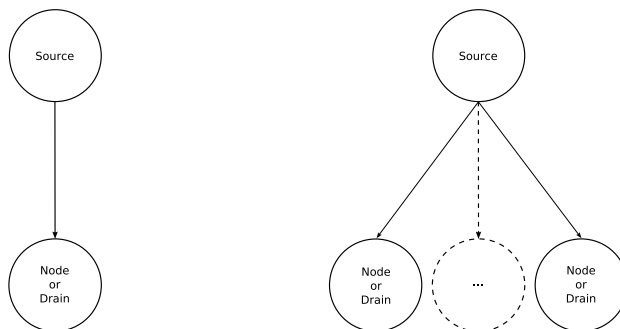


Fig. 1. Source with one or more Children

2.2 Drain

A drain is a component which consumes data: It is always the tail of a production path. A drain is the final receiver of data, because it can not have any child components inside of the system boundaries of L-DSMS. The data can of course be forwarded to some receivers outside L-DSMS. A `SocketDrain`, for example, forwards the incoming data to every process connected to this drain via a socket connection. These data leave the system boundaries of L-DSMS. The L-DSMS core package contains drains for writing the incoming data into files, sending them over socket connections, or simply printing them to the console screen.

Drains can receive their data and meta data from one ore more sources or processing nodes (cf. Fig. 2). The connections between a drain and its parent components are specified in the configuration file. The output types of the data and meta data of nodes which deliver data to a drain must of course also be compatible to the input types of the drain's data and meta data. A drain with input types `java.lang.Object` can accept data of all types.

The configuration file contains the relation between the drain and its parents together with drain specific properties. Two properties can be specified for all types. The first property is again the mandatory `class` attribute. The second property is the optional attribute `sourcerefs` and contains the names of additional sources. The names in `sourcerefs` specify source nodes or processing nodes as parent components for this drain. All other properties are drain specific and are listed in the detailed description for each drain. The `ByteArrayFileDrain` (cf. Sect. 5.2), for example, has the additional attribute `file`.




Fig. 2. Drain with one or more Parents

2.3 Processing Node

Processing nodes are a combination of sources and drains and can be positioned at every possible place within a production line. They receive data and meta data

from one or more components, process them and forward the results to an arbitrary number of consumer nodes. Since processing nodes are sources as well as drains, they inherit the properties of both of them. Therefore, they have the same type restrictions for connecting them to other components. Nodes are configured in the same way as sources and drains. Besides node specific properties, they all have the mandatory `class` attribute and an optional `name` attribute. In addition they have sources (like a drain), which are specified with the `sourcerefs` attribute.

Each node in the L-DSMS core package processes the data in a different way. Therefore there is no general assumption about the outcome of nodes. Some of them produce data of the same type as the input and others produces data of a totally different type. They even need not produce output at all if the incoming data does not satisfy certain conditions (e.g. in a filter).




Fig. 3. Node with an Arbitrary Number of Parents and Children

3 How the L-DSMS Network Operates

When it gets started, the L-DSMS system reads a configuration file and arranges the network of processing nodes. The network, hopefully, has some source nodes, possibly some further processing nodes and some drain nodes. Each node, except the drains, have a list of successor nodes, the *consumer nodes*, and each of them has a `consume` method. The source nodes have a `start` method.

After the network is ready for operation, a network broker calls the `start` method of each source node. It is the responsibility of the source node's `start` method to start a thread that does the actual work. If a source node does not start a thread, then the `start` method is just called only once. If this call termi-

nates, it is never called again. This makes sense if there is only one single source node.

After a source node has assembled a packet of data, it calls a `send` method for the data packet together with the meta data. The `send` method in turn calls the `consume` methods for all attached consumers. The `consume` method of a node `N` can do some processing and then call its `send` method, which in turn calls the `consume` methods of all consumers attached to `N`. This way, eventually the `consume` methods of the drain nodes get called and dump the data somewhere outside the system.


If there is a single source node and the network has a tree structure then this procedure passes the data through the tree in a depth first left to right order. The `consume` methods could, however, start their own threads when they get called first. Each time they get called next, they just forward the data to the corresponding thread and terminate. The threads can then process the data in parallel. This would result in a more breadth first like traversal through the net.

Another alternative is to implement node classes whose `consume` methods *synchronise* data from different sources. The first time, such a `consume` method is called from one source, it just stores the data locally and terminates. If it is called a second time, maybe from a different source, it can combine the new data with the previously stored data.

With this architecture the network broker need not manage any threads. It depends on the implementation of the `consume` methods to operate with or without threads. Even an agent architecture platform can be implemented this way.

4 Hello World

This example illustrates a very basic configuration, that prints “Hello World” onto the screen.



The following steps set up the ‘Hello World’ example.

- Create a file `<LDSMS_HOME>/examples/hello_world/input.txt` with a line of text “HELLO WORLD”.
- Create a file `config.xml` within this folder and insert the content from listing 1.1 (without the line numbers) into this file.
- Open a console and change to `<LDSMS_HOME>`.

- Type “java -jar ldsms.jar examples/hello_world/config.xml” and press ‘ENTER’ (cf. listing 1.2).
- Press ‘Strg’+‘C’ to stop LDSMS, after ‘HELLO WORLD’ has been printed.

Listing 1.1. examples/hello_world/config.xml

```

1 <?xml version="1.0" encoding="ISO-8859-1" ?>
2
3 <server>
4   <logging level="INFO" />
5   <services>
6     <network>
7       <source class="generics.StringFileSource" file="
8         examples/hello_world/input.txt">
9         <drain class="generics.ConsoleDrain" />
10      </source>
11    </network>
12  </services>
  </server>

```

Listing 1.2. Run the ‘Hello World’ example

```

java -jar ldsms.jar examples/hello_world/config.xml
1 [main] INFO de.lmu.ifi.pms.ldsms.network.Server -
  Configuring Server ...
572 [main] INFO de.lmu.ifi.pms.ldsms.network.Server -
  Initializing Server ...
573 [main] INFO de.lmu.ifi.pms.ldsms.network.Server -
  Starting Server ...
HELLO WORLD

```

Remarks: In line 7 of the configuration file, a `StringFileSource` (see Sect. 7.1) was specified. This `FileSource` reads every data from the text file, specified by the `file` attribute of the source element (‘examples/hello_world/input.txt’ in this case). In line 8, a `ConsoleDrain` was specified (see Sect. 5.2). A `ConsoleDrain` prints every input to the console. Because the `ConsoleDrain` is specified as a child element of `StringFileSource`, every output of the `StringFileSource` (here, the data from the text file) is passed to the input of the `ConsoleDrain`.

5 How to Extend L-DSMS

Although the L-DSMS library contains already a number of predefined node types, every new application will need its own specific processing nodes. L-DSMS supports adding new node types by providing corresponding interfaces and abstract classes. They specify exactly how the new classes have to be implemented.

5.1 Source Nodes

Every component that should be treated as a source node has to implement, either directly or indirectly, the interface `<ldsms core package>.network.Source`. A detailed description of its methods can be found in the JavaDoc documentation at the L-DSMS project page (<http://www.pms.ifi.lmu.de/reverse-wga1/ldsms/>). The abstract class `SourceImpl` already implements the interface ‘Source’. It can be used as the superclass for new Source classes. Fig. 1.3 shows an example for a source class. The `startSending` method keeps on sending ‘Hello’ and then sleeping for 5 seconds until the network broker stops it.

Listing 1.3. Example for a Source Class

```

package test;
import de.lmu.ifi.pms.ldsms.network.SourceImpl;
public class HelloSource extends SourceImpl<String, String>{
    private Thread helloThread ;

    public HelloSource() {
        super(String.class, String.class);}

    public void start() throws Exception {
        super.start();
        startSending();}

    public void startSending(){
        super.startSending();
        try{helloThread = new Thread() {
            public void run(){
                while(!isStopped()){
                    try{send("Hello", null);
                        Thread.sleep(5000);}
                    catch(InterruptedException ie){break;}}}};
            helloThread.start();}
        catch(Exception e){}}

    public String getDescription(){
        return "A simple Source that sends 'Hello' every 5 seconds
            .";}

    public String getShortDescription(){return "HelloSource";}

    public String toString(){return "HelloSource";}}

```

5.2 Drains Nodes

Every new drain class has to implement, either directly or indirectly, the interface `<ldsms core package>.network.Drain`. The abstract class `DrainImpl` already implements this interface and can be used as superclass of a new drain class. Drains

that need additional attributes from the configuration file, additionally have to implement the interface `org.apache.avalon.framework.configuration.Configurable`. This ensures, that the configuration information from the configuration file are passed to the Drain. Drains that need to be started or stopped, additionally have to implement the interface `org.apache.avalon.framework.activity.Startable`. This ensures, that the server starts the drain after the configuration has been finished and that the server stops the drain if the system is forced to terminate. This is useful, if additionally threads are used or streams have to be opened and closed.

5.3 Processing Nodes

Nodes are a combination of a drain and a source. Every node has therefore to implement both interfaces, for drains and for sources. The abstract class `Node` implements both interfaces already. It can therefore be inherited by a new node class. The most important method to implement for processing nodes and drain nodes is `consume(data, metadata)`. It is called by other nodes to pass data and meta data to the current node.

6 Managing L-DSMS with VISU-L-DSMS

VISU-L-DSMS is the graphic user interface for L-DSMS that was developed to ease the management of L-DSMS. It can manage instances of L-DSMS located at the same host as VISU-L-DSMS, as well as instances located on remote hosts. This allows one to manage running instances of L-DSMS from a single management workstation.

Linux as well as Windows users can start VISU-L-DSMS by simply executing a script that is shipped with VISU-L-DSMS. Executing the script opens the VISU-L-DSMS main window. After some operations it could look like in Fig. 4. This main window contains three areas and one menu panel. At the left hand side there is the **Network View**, a graphic representation of all components, together with their relationship to each other. Each component is represented as a node (coloured symbol) and their relationships to each other by edges (black lines).

At the right hand side there is the **Capturing View**. This area contains two frames. The upper frame is used to show the incoming data of a component selected in the network area. The lower frame is used to show the outgoing data of a component (the same or another one).

In the bottom area there are two tabs that provide additional information about the selected components of the network window (the **Overview Tab**) and their properties (the **Properties tab**).

These three areas are empty until VISU-L-DSMS is connected to a running L-DSMS instance.

To manage an instance of L-DSMS, a connection between VISU-L-DSMS and this instance has to be established first. This is done by selecting the **Connect**

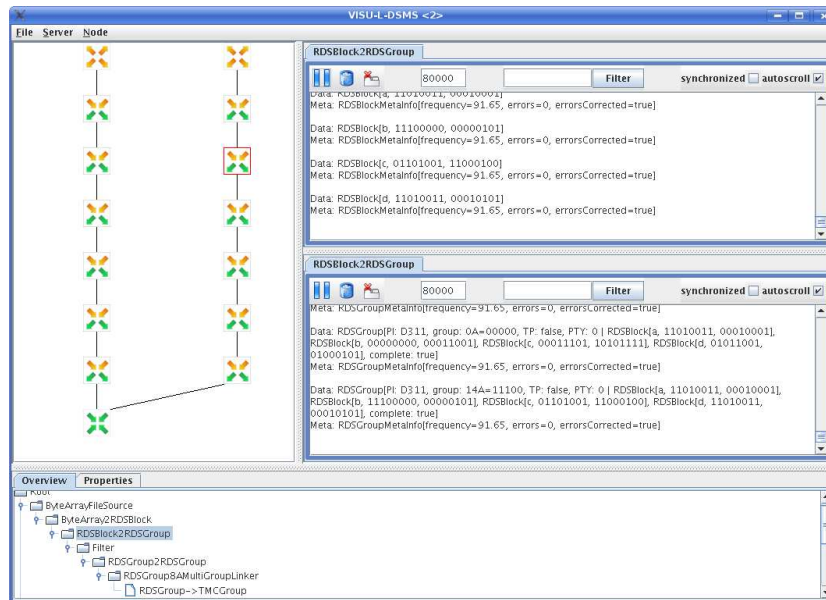


Fig. 4. VISU-L-DSMS Window

item from the **Server** menu. A connection dialog is opened where the necessary connection parameters, host name, port number and the name of a running instance of L-DSMS have to be specified.

The overview tab shows the names of all components in a tree structure with sources being the parents and their drains being the children (recursively). Each component is listed only once, even if it has more than one source. In this case, it is only listed once as a child of the first of its sources. If a component is selected in the network area, its representation in the overview tab is selected as well, and vice versa. This gives a compact overview about the network structure while the names of the components are listed in the overview tab.

The properties tab is used to *display* and *edit* the properties for the last component that has been selected. Each property is presented either as a text field, a list or a check box. Not all properties, however, can be edited. Editing values doesn't affect the L-DSMS instance, until the changes are saved. Saving properties causes the new values to be sent over the network to the observed L-DSMS instance and to change the corresponding node attributes. This way, a running L-DSMS system can be controlled remotely.

The capturing area is used for monitoring the incoming and outgoing data of the components. To capture data, select one or more components, either in the network or the overview area, open the **node** menu and select **start listening**. This opens a tab for the incoming data in the upper frame and a tab for the outgoing data in the lower frame for each selected component. If only selected

data is to be displayed, a regular expression can be defined in the *filter field*. Only the data matching the regular expression is then displayed.

The current version of VISU-L-DSMS can only change node properties of running L-DSMS instances. In principle it would be possible to extend VISU-L-DSMS and allow it to also change the network configuration. This way, a program for processing data streams could be generated just by arranging some graphic symbols in a suitable graphical editor. The only Java programming which would then still be necessary would be to extend the the L-DSMS library with new types of processing nodes.

7 Predefined Node Types

The L-DSMS core package contains a library of predefined node types. In this paper we list only the general purpose node types. A number of additional node types have been implemented for the test application, a system that feeds traffic information into Google Earth.

7.1 Sources

The following source node types are available:

ByteArrayFileSource produces binary data by reading it from a file stream.

The data is received in the form of byte arrays. The byte arrays are read from the stream as follows: the first four bytes describe the length of the data array as an integer, let it be n . The next n bytes are interpreted as the data array. The next four bytes describe the length of the meta data array as an integer, let it be m . The next m bytes are then interpreted as the meta data array.

ByteArraySocketSource reads binary data from a socket connection. The sequence of bytes is interpreted in exactly the same way as for **ByteArrayFileSources**.

ObjectSocketSource reads objects from a socket connection. The source starts reading as many bytes as are necessary to cast them into an object of the specified data type. Then it continues reading as many bytes as are necessary to cast the data into an objects of the meta data type. Bytes which cannot be casted this way are dropped.

StringFileSource reads a file as a list of strings. The content of the file is sent line by line as a sequence of strings (including '\n').

StringSocketSource reads strings from a socket connection and sends them in the same way as **StringFileSource**.

7.2 Drains

The following drain node types are available:

- ByteArrayFileDrain** writes the incoming data byte arrays and meta data byte arrays into the specified file such that **ByteArrayFileSource** can read them again.
- ByteArraySocketDrain** sends the incoming data byte arrays and meta data byte arrays into the specified socket such that **ByteArraySocketSource** can read them again.
- ConsoleDrain** prints every incoming data without any formatting to the console.
- ObjectSocketDrain** writes the incoming data and meta data into the output stream of a socket. The objects must implement the `java.io.Serializable` interface. If no client is connected, the data is dropped.
- StringFileDrain** writes the incoming data strings with UTF-8 encoding into the specified file. The meta data is only used, to indicate that the corresponding data is a Goodbye Message. This data is written to the specified file just before the **StringFileDrain** stops.
- StringSocketDrain** writes the incoming byte strings to every client, connected at the specified port.
- SpexNode** filters data from an XML stream, using the **SpexProcessor** and a XPath expression. The **SpexProcessor** extracts from streaming XML data the elements which are described by the given XPath expression [5].

7.3 Processing Nodes

The following processing node types are available:

- Buffer** caches information until it can be delivered to attached consumers. It can be used, to cache data until one or more drains are able to consume it.
- ByteArray2String** transforms incoming byte arrays into strings, using the specified encoding format. It can be used to connect a source with a drain, where the source produces byte arrays but the drain expects strings.
- Cast** casts the input to the specified type and filters out any incompatible data or meta data objects. It can be used to connect a source with a drain, where the output types of the source are different but compatible with the input types of the drain.
- Filter** tests whether the data and meta data meets a certain condition and passes it to its drains only if the condition is met. There is an elaborated language for specifying filters.
- String2ByteArray** takes as input Strings and forwards them as byte arrays. The parameter 'encoding' can be used to specify the encoding format (e.g. US-ASCII, UTF-8, UTF-16BE, UTF-16LE, UTF-16 etc.).

8 The Configuration File

The configuration file is an XML file. It contains the precise specification of all nodes, their parameters and the connections between the nodes. Listing 1.1 shows an example of such a file. The Document Type Definition (DTD) in Listing 1.4 illustrates the structure of the configuration file.

Listing 1.4. DTD for the Configuration File

```

<!ELEMENT server (logging , plugin-list?, services)>
<!ELEMENT logging (logger*)>
<!ATTLIST logging level (OFF|FATAL|ERROR|WARN|INFO|DEBUG|ALL) #
  #REQUIRED>
<!ELEMENT logger EMPTY>
<!ATTLIST logger name CDATA #REQUIRED>
<!ATTLIST logger level (OFF|FATAL|ERROR|WARN|INFO|DEBUG|ALL) #
  #REQUIRED>
<!ELEMENT plugin-list (plugin-prefix*, plugin*)>
<!ELEMENT plugin-prefix (#PCDATA)>
<!ELEMENT plugin ANY>
<!ATTLIST plugin class CDATA #REQUIRED>
<!ELEMENT services (network)>
<!ELEMENT network (source+,node*)>
<!ELEMENT source (node|drain)*>
<!ATTLIST source class CDATA #REQUIRED>
<!ATTLIST source name ID #IMPLIED>
<!-- other attributes depend on the value of 'class' -->
<!ELEMENT node ((and?|or?),node*)>
<!ATTLIST node class CDATA #REQUIRED>
<!ATTLIST node name ID #IMPLIED>
<!ATTLIST node sourcerefs IDREFS #IMPLIED>
<!-- other attributes depend on the value of 'class' -->
<!ELEMENT drain EMPTY>
<!ATTLIST drain class CDATA #REQUIRED>
<!ATTLIST drain sourcerefs IDREFS #IMPLIED>
<!-- other attributes depend on the value of 'class' -->
<!ELEMENT and (or?,condition+)>
<!ELEMENT or (and?,condition+)>
<!ELEMENT condition EMPTY>
<!ATTLIST condition class CDATA #REQUIRED>
<!-- other attributes depend on the value of 'class' -->

```

9 The Test Application

The L-DSMS system was developed because we wanted to implement a system that receives traffic information and feeds them into a system for displaying road maps. The traffic information comes via RDS-TMC (Radio Data System, Traffic Message Channel, see ISO standard 14819-3) as digitally encoded radio signals. It can be received by special receivers which deliver sequences of bytes to a computer. The messages are densely encoded. They consist of an event code, a location code and some time information. About 1460 different event types are encoded according to the Alert-C-Standard. The mapping between the location code and the actual road segments is done by a separate location table.

Our test implementation reads the TMC messages and uses the event code lists and location tables to enrich the data packets with information necessary

for incorporating them into road maps. The final result after a sequence of transformations are KML files (see <http://code.google.com/apis/kml/>), which in turn can be read by Google Earth. Google Earth integrates the messages into its road maps. The KML files are updated as soon as new TMC messages are received.

The version of L-DSMS which can be downloaded has therefore quite a number of node types which are specific for this application. They can serve as examples or templates for other applications.

10 Summary

L-DSMS is a local data stream management system. The configuration XML-files specify the structure of a processing network connecting source nodes with drain nodes via intermediate processing nodes. A network broker, implemented in Java, can read a configuration file and turn it into an executable Java program. A running L-DSMS instance can be monitored and, to a certain degree, modified remotely by Visu-L-DSMS. The system and its documentation is open source. It comes with a library of predefined node classes, together with interfaces for adding new node classes.

So far, only the parameters of node instances can be modified remotely by Visu-L-DSMS. A next step could be to specify and modify the whole network remotely, such that no XML editing would be necessary any more to configure a data stream processing system.

Acknowledgements

This research has been funded by the European Commission and by the Swiss Federal Office for Education and Science within the 6th Framework Programme project REWERSE number 506779 (cf. <http://rewerse.net>).

References

1. Brian Babcock, Shivnath Babu, Mayur Data, Rajeev Motwani, and Jennifer Widom. Models and issues in data stream systems. In *Proceedings of 21st ACM Symposium on Principles of Database Systems (PODS 2002)*, 2002.
2. Lukasz Golab and M. Tamer Özsu. Issues in data stream management. *ACM SIGMOD Record*, 2003.
3. Ian Gorton, Adam Wynne, Justin Almquist, and Jack Chatterton. The medici integration framework: A platform for high performance data streaming applications. In *WICSA '08: Proceedings of the Seventh Working IEEE/IFIP Conference on Software Architecture (WICSA 2008)*, pages 95–104, Washington, DC, USA, 2008. IEEE Computer Society.
4. Christian Hänsel. Implementation: L-DSMS - A Local Data Stream Management System. REWERSE deliverable A1-D10-3, Institute for Informatics, Ludwig-Maximilians-Universität München, 2008. URL: <http://idefix.pms.ifi.lmu.de:8080/rewerse/index.html#REWERSE-DEL-2008-A1-D10-3>.

5. Dan Olteanu. *Evaluation of XPath Queries against XML Streams*. Dissertation/Ph.D. thesis, Institute of Computer Science, LMU, Munich, 2005. PhD Thesis, Institute for Informatics, University of Munich, 2005.