Accessing Databases from Esterel[*]

David White & Gerald Lüttgen

Department of Computer Science, University of York
Heslington, York YO10 5DD, UK

E–mail: {dhw100, luettgen}@cs.york.ac.uk

Technical Report, December 2004

Abstract

A current limitation in reactive systems design is the lack of database support in tools such as Esterel Studio. This report proposes a way of integrating databases and Esterel by providing an Application Programming Interface (API) for relational database use inside Esterel.

As databases and Esterel programs are often executed on different machines, result sets from database queries may be processed either locally and according to the synchrony hypothesis, or remotely along several reactive cycles. These different scenarios are reflected in the development of two APIs detailed in this report. Their utility is demonstrated by means of a case study modelling a warehouse storage system as might be used by a direct order company. The system employs a robot whose task it is to collect items from a customer's order and assemble them in one place. In addition to customer and order data, the underlying database stores spatial data detailing the position of items in the warehouse. Both robot and warehouse are implemented using the Lego Mindstorms robotics system.

1 Introduction

One of the current limitations in reactive systems programming is the lack of *database support* available within *synchronous languages*, such as Esterel [3, 4] and Lustre [7], and their development environments, *Esterel Studio* and *SCADE* [9], respectively. As is, a system designer needs to modify auto–generated code by hand in order to interface to databases, which is both a difficult and error–prone business. This is a problem very much relevant in industry since some reactive systems programmed in synchronous languages would

[*]This research has been funded by an Undergraduate Research Bursary of the Nuffield Foundation under grant ref. URB/01528/6.

Dagstuhl Seminar Proceedings 04491
Synchronous Programming - SYNCHRON'04
http://drops.dagstuhl.de/opus/volltexte/2005/161

benefit from an easy model of database interaction. For example, synchronous languages are often used to build the flight software for aeroplanes. Adding database interaction would enable mapping data to be retrieved and processed directly by the reactive kernel implementing the flight software, so that maps of the currently overflown area could be displayed to pilots in real time.

This report addresses that limitation by providing an *Application Programming Interface* (API) for using relational databases within the Esterel programming language. We choose MySQL [14] as the database and, since reactive kernels are produced as C programs by the academic Esterel compiler [2, 8], the APIs are implemented using the MySQL C interface [15]. MySQL is selected here simply for its convenience, since it is widely used in academia. However, our work can as easily be applied to other relational databases. To the best of our knowledge, no articles on database integration within Esterel have been published before in the academic literature. Indeed, other reactive systems design tools seem to be very limited in this respect as well, including Mathworks' *Simulink/Stateflow* [5] and iLogix' *Statemate* [11].

Because database transactions are relatively complex when compared to responses of reactive kernels, databases and reactive programs must be considered as running asynchronously to each other. This is true regardless of whether they reside on the same machine or on different machines. In the former case, however, *result sets* to database queries may reasonably be assumed to be processed within a single synchronous step of the reactive kernel. In the latter case, result sets are necessarily processed asynchronously to the reactive kernel. For these reasons, an API for each situation is provided: a *Local Result Set API* and a *Remote Result Set API*. The realisation of both relies on Esterel's formidable support for extending the language via external data types and external functions and procedures. While the local result set API makes heavy use of external functions and procedures, the remote result set API only employs valued signals. This is because external functions and procedures are expected to be instantaneous in Esterel.

We demonstrate the utility of our APIs by means of a case study involving a *warehouse storage system*. The idea behind this is that of a direct order company, i.e., orders must be picked from items stored in a warehouse. Accordingly, a *robot* is built and programmed to pick up and drop off items, therefore making it possible for a complete order to be collected and stored. The orders, as well as information about the items, are provided by a database running remotely to the robot on a PC. Part of the information stored on an item is its position in the warehouse, thereby providing mapping data for the robot. The case study is implemented using *Lego Mindstorms* robotics kits [1, 13], which provides a small programmable brick, called *RCX*, that houses a microcomputer capable of running an Esterel reactive kernel. Sensors and actuators connected to the RCX, such as *touch*, *light* and *rotation sensors* and *motors*, respectively, permit interaction with the RCX's environment. The RCX also has a built–in infrared port, which we use to communicate with the warehouse database on the PC.

The remainder of this report is structured as follows. The next section describes both our APIs, emphasising the general model of interaction between

Esterel reactive kernels and databases. Our case study is presented in Sec. 3 which provides an example of the usage of the Local Result Set API. Sec. 4 contains our conclusions and suggestions for future work. Details of the realisation of our APIs and the case study can be found in a technical report [18].

2 Database APIs for Esterel

Due to the variety of computing architectures in which an Esterel reactive kernel together with a database may be used, two different APIs for enabling database access within Esterel are devised. The APIs exploit the built–in extensibility of the Esterel language and provide dedicated signals, external data types and external functions and procedures for use inside Esterel.

Both APIs allow local and remote databases to be queried. They consider databases as part of the system environment and as running asynchronously to the reactive kernel. This is because database transactions are typically more complex to process than ordinary reactions. The APIs differ in the storage location considered for the result set returned by a query. Result sets are database tables presented as sets of rows, from which the reactive kernel extracts desired information according to its needs. The *Local Result Set API* views the result set as being local to the reactive kernel, whence operations on the result set can conceptually be considered to take zero time, satisfying the synchrony hypothesis [10]. This API can thus make heavy use of user–defined external functions and procedures, which are required to be instantaneous in the Esterel language [3]. The *remote result set API* allows the result set to be stored remotely to the reactive kernel. As a consequence, this API cannot profit from the elegance and simplicity of employing external functions and procedures, but must solely rely on signals for processing result sets.

2.1 Local Result Set API

The most logical way to view the interaction between a reactive kernel programmed in Esterel and a database is to regard the database simply as an extension of the reactive kernel's environment. For this reason all interactions with the database from within Esterel are modelled using input and output signals, as these are Esterel's facilities for communicating with the environment. Therefore, to perform an operation on the database, a dedicated output signal is emitted, parameterised in a string that formulates a query in SQL syntax. The database's response is awaited via a dedicated input signal whose parameter carries an identifier pointing to the result set. During the time between the emitted query and the results returning, the database is queried and the whole result set is transferred back to the site that also runs the reactive kernel. Note that multiple databases can simply be supported by declaring a dedicated output and input signal for each database.

Once a database has been queried and a result set returned, data can be extracted from the result set using dedicated operations implemented using

Esterel's external function and external procedure facility. Note that this is possible since both the result set and the reactive kernel reside at the same site, which implies that accesses of the result set by the kernel may be considered as instantaneous. If the query's SQL command is one that does not return results, such as the command for the deletion of data items, then the only operation provided is one to check the number of affected rows. If the SQL command did return a result set however, the set may be accessed by successively reading the result set row–by–row, extracting the specific data items from each row and coercing them into native Esterel types. Once all rows have been processed, an operation shall be called to free the memory occupied by the result set.

Table 1: Services offered by the Local Result Set API

```
type MYSQL_RES_ptr;
type MYSQL_ROW;

procedure appstr() (string, string);
procedure appint() (string, integer);
procedure appbol() (string, boolean);
procedure appflt() (string, float);
procedure appdou() (string, double);

output <Signal name for emitting query> : string;
input <Signal name for returning results> : MYSQL_RES_ptr;

function check_result(MYSQL_RES_ptr) : boolean;
function get_next_row(MYSQL_RES_ptr) : MYSQL_ROW;
function num_rows(MYSQL_RES_ptr) : integer;

function getint(MYSQL_ROW, integer) : integer;
function getbol(MYSQL_ROW, integer) : boolean;
function getdou(MYSQL_ROW, integer) : double;
function getflt(MYSQL_ROW, integer) : float;
function getstr(MYSQL_ROW, integer) : string;

function num_affected_rows(MYSQL_RES_ptr) : integer;
procedure clear_results() (MYSQL_RES_ptr);
```

Esterel's interaction with a remote database and its local processing of the result sets thus leads to the API displayed in Table 1. In the remainder of this section we explain the API's services in more detail.

We begin with the formation of a query string containing SQL commands. Since Esterel does not provide any facilities for building strings, the string must be generated using a series of append operations. The API offers an append operation for each of Esterel's native data types and implements these operations using Esterel's external procedure call function. For example, the following

generates a query which uses an integer variable `order_id`:

```
var query_string : string in
    query_string := "select * from orders where order_id = ";
    call appint() (query_str, order_id);
end var
```

As mentioned before, the database is interacted with via an output query signal and an input result signal. For each database used, they should be declared as such:

```
output item_db_query : string;
input item_db_results : MYSQL_RES_ptr;
```

Each pair of signal names can be chosen by the user. The mapping between chosen signal names and the actual databases is defined elsewhere (cf. Sec. 2.3). The data returned on a result signal is simply an identifier of the external type `MYSQL_RES_ptr` defined in our API's implementation. Note that the identifier value should never be copied since this would not result in a full copy being performed. Our framework effectively allows only one result set per database connection; however, if more result sets are required simultaneously within some Esterel application, additional connections to the same database may be declared.

To ensure that the results are received correctly from the database, the results signal should be *awaited* immediately following the emission of the query:

```
emit item_db_query("select * from item");
await item_db_results;
```

If two queries are issued simultaneously, then the result signals must be awaited in parallel or using *immediate await* statements.

In order to check the success of the SQL command, the boolean function `check_result` should be called and passed the identifier of the result set, i.e., the value of the input result signal. If it returns true, the query succeeded and the operations described below may be used to access the data inside the result set. If it returns false, the data in the result set is not valid.

For working with a result set that contains data — as opposed to an empty one returned by, e.g., an SQL insert statement — rows must be declared inside Esterel. Rows are declared to be of external type `MYSQL_ROW` which is defined in our API's implementation. The lifetime of any data loaded into a row from a result set lasts only as long as the result set itself, i.e., up to the time the `clear_results` operation is called. We recommend that rows are only declared locally and that their scope finishes before the call to `clear_results` occurs.

The functions provided to operate on the result set and rows will now be described. Most of the operations mirror the equivalent MySQL function from the MySQL C API [15]. This interface was chosen for two reasons: the MySQL C functions are widely and well known and, at a later date, additional functions can easily be included, if desired. Function `get_next_row` is required to load

data into a row from a result set. It is passed a result set identifier and, each time it is called, it will return the next row in the result set. Generally, the program will need to know how many rows there are in the result set and, therefore, how many times to call `get_next_row`. This is accomplished by a call to function `num_rows` which, when passed a result set identifier, returns the number of rows in the result set. Once a row has been loaded from a result set, data can be extracted using a `get<type>` function which is provided for each of the native Esterel data types. In addition to a row, an integer is also passed indicating the index of the column from where the data is to be retrieved. The following is a simple example of data extraction using our API:

```
var row_holder : MYSQL_ROW,
    item_name : string,
    item_location : integer in
  row_holder := get_next_row(?item_db_results);
  item_name := getstr(row_holder, 1);
  item_location := getint(row_holder, 2);
end var;
```

In this case, the results are identified with the valued signal `item_db_results`, and the item's name and location are in the second and third column of the row, respectively. Note that indexing is as in the C programming language and thus starts with index 0.

One function that remains to be described for accessing the result set is `num_affected_rows`. This is used when the result set contains no data but a user wants to know how many rows were affected by the SQL command. As such, `num_affected_rows` can be employed to test the success of an SQL query, e.g., to check whether a delete query has had the desired effect.

The final operation provided by the API, which has already been referred to above, clears the memory occupied by the result set:

```
procedure clear_results()(MYSQL_RES_ptr);
call clear_results()(?item_db_results);
```

It is essential here that there are no rows loaded from the result set after it is cleared, since the data within these rows is cleared with the result set as well.

2.2 Remote Result Set API

The Remote Result Set API should be used in situations where it is not feasible to transfer the entire result set to the system running the reactive kernel, i.e., when both the database *and* the result set must be viewed as part of the environment. Since remote communication must be taken into account, the API is quite different to that of the Local Result Set API. This is because external functions and procedures can only be used in Esterel if their operations may be considered as instantaneous [3]. Consequently, one must either employ Esterel's task concept or must solely rely on signals. In both cases and as a consequence

of operations on the result set not being instantaneous, the Esterel kernel must be informed of when an operation is complete. This is accomplished by awaiting an "acknowledge" signal after every operation. In the remainder we focus on the solution via signals rather than tasks, as the authors were unfamiliar with Esterel's intricate task concept at the time the research was carried out.

An important aim of the Remote Result Set API is to minimise the amount of data that must be transfered. To achieve this, rows are not handled by the reactive kernel but are moved database–side instead. As a consequence, rows are part of the environment, like the result set. To prevent the complexity of handling multiple rows in the kernel, each database is limited to only one row. This is a reasonable restriction since systems that use this API are unlikely to be performing complex database manipulations that require multiple rows. If multiple row access should indeed be necessary, additional connections can be specified to achieve that. This work–around can also be employed to support multiple databases, similar to what is suggested for the Local Result Set API.

Table 2: Services offered by the Remote Result Set API

```
input <db_id>_ackstr : string;
input <db_id>_ackint : integer;
input <db_id>_ackbol : boolean;
input <db_id>_ackflt : float;
input <db_id>_ackdou : double;

procedure appstr() (string, string);
procedure appint() (string, integer);
procedure appbol() (string, boolean);
procedure appflt() (string, float);
procedure appdou() (string, double);

output <db_id>_query_out : string;   after emission, await <db_id>_ackbol
output <db_id>_fetch_next_row;       after emission, await <db_id>_ackbol
output <db_id>_num_rows;             after emission, await <db_id>_ackint
output <db_id>_num_affected_rows;    after emission, await <db_id>_ackint
output <db_id>_clear_results;        after emission, await <db_id>_ackbol

output <db_id>_getint : integer;     after emission, await <db_id>_ackint
output <db_id>_getbol : integer;     after emission, await <db_id>_ackbol
output <db_id>_getflt : integer;     after emission, await <db_id>_ackflt
output <db_id>_getdou : integer;     after emission, await <db_id>_ackdou
output <db_id>_getstr : integer;     after emission, await <db_id>_ackstr
```

Our API for remote result set access is displayed in Table 2. The remainder of this section explains the API's services in more detail. Similar to the naming of the query and result signals in the Local Result Set API, each signal `signal_name` is prefixed with a textual database identifier `db_id`, which

we denote by `<db_id>_signal_name`. Moreover, the offered string generation functions are identical to those in the Local Result Set API.

The main difference to the Local Result Set API is the way in which the results to a query are accessed. There is now one result set and one row per database defined, and since each signal is prefixed by a unique string, there is no need for a result set identifier to be returned. The only data returned after issuing a query is the success of that query. Therefore, after a query has been issued, the boolean acknowledge input signal for that database must be immediately awaited:

```
emit item_db_query("select * from item");
await <db_id>_ackbol;
```

Its carried value is the same as the one returned by the `check_result` operation in the Local Result Set API.

The value of the `<db_id>_ackbol` input signal should then be tested to determine whether the query has succeeded or not. If the query has succeeded and the result set is not empty, then the first row can be loaded by emitting the `<db id>_fetch_next_row` signal. It is again necessary to await `<db_id>_ackbol` after this to determine when the operation has completed. True is returned if there exists a valid row to load, and false if there are no more rows available.

Now that a row has been loaded, the signals for accessing it may be used. Similar to the Local Result Set API, there is a `<db_id>_get<type>` signal for each of the Esterel native types. In the following example, it is shown how to extract an integer from the first column of the row:

```
emit <db_id>_getint(0);
await <db_id>_ackint;
order_id := ?<db_id>_ackint;
```

All the operations contained in the Local Result Set API are also available in the Remote Result Set API, but implemented as signals instead of external functions. However, in order to cope with the additional remote nature of result sets and thus rows, it is necessary to await an acknowledge signal after each operation.

2.3 Implementation of the APIs

Both APIs are implemented in a combination of Perl and C. The implementations heavily rely on the MySQL C API [15] and are, for most parts, rather straightforward. The complete source code for both APIs is included in [18].

The involvement of the Perl scripting language [17] in the realisation of the APIs may be surprising at first. The reason is our desire to support *multiple* databases with *user–declared* signal names for emitting SQL queries and awaiting result sets. As reactive systems typically interact with an arbitrary but fixed number of databases, it is unnecessary to provide API services for dynamically

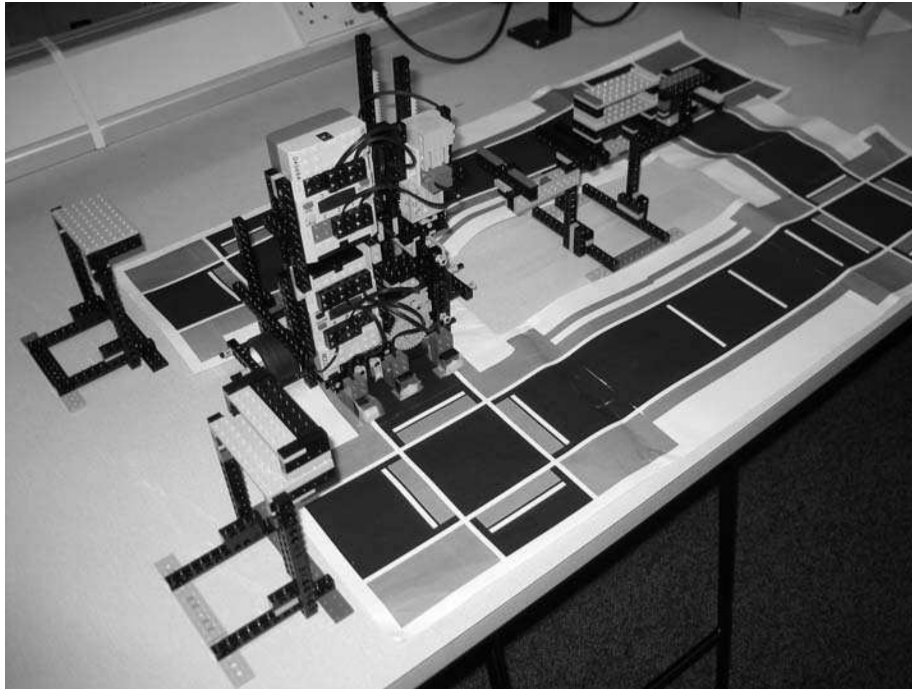

Figure 1: Aerial photograph of our warehouse system.

binding signal names to databases; even a static binding should not be defined within an Esterel program at all, as it is not part of reactive behaviour. Instead, we opt to provide such a binding as a parameter to our Perl script for each API, which appropriately combines the C–code generated by the academic Esterel compiler [8] for some reactive program with C–code implementing the APIs services used in the program.

Since database connections are generally permanent throughout the time a reactive system is running, our APIs provide no explicit facilities for connecting and disconnecting from a database. Instead, connection and disconnection is handled implicitly by the APIs. However, if explicit connect and disconnect services would be required from within an Esterel program, the APIs could be extended by defining and implementing according dedicated signals.

3 Case Study

In this section we present a case study demonstrating the utility of our APIs: an automated warehouse storage system modelling a direct order company, where items from orders are picked, stored and finally removed from the warehouse. This requires producing a warehouse containing various items and a robot capable of moving the items within the warehouse, for which we use *Lego Mindstorms' robotics kits* [13] (cf. Fig. 1). Lego Mindstorms provides both a

construction tool with sensors and actuators and a microcontroller, called the RCX, which is capable of running a reactive kernel programmed in Esterel. The database behind our warehouse model is that of a standard order system but which also includes mapping data about the location of the items. As this case study is meant to exemplify the use of our database APIs, only the part of the solution employing the APIs is focused on below.

3.1 Lego Mindstorms, the RCX and BrickOS

Lego Mindstorms is a platform for building computer–controlled robots within the *Lego system* [13]. At the heart of Lego Mindstorms is the RCX. This “brick” is a small battery–powered computer system capable of controlling up to three *actuators* and reading up to three *sensors*. In Lego, actuators are normally motors, and sensors can be light, rotation or touch sensors. Each RCX also provides an *infrared transmitter and receiver* used for both downloading programs from a PC and for inter–RCX communication. The infrared download device used on the PC can also participate in communications with RCXs.

The RCX provides great flexibility though its re–programmable firmware. BrickOS [6], formerly known as LegOS, is an open source replacement firmware for the RCX. It boasts a number of features that make it considerably more complex than the standard Lego firmware. Foremost, it allows programs written in the C programming language to be executed on the RCX. Obviously this is especially important for this project since the Esterel compiler [8] generates C code as a target language.

BrickOS also provides infrared communication through the *Lego Network Protocol* (LNP) [6]. This protocol has two layers, an integrity layer and an addressing layer. The integrity layer guarantees that, if a message is received, it will be the same message that was sent — similar to the Internet protocol UDP. However, it differs from UDP in that it implements a broadcast mode, i.e., any RCXs in the receiving area will pick up the message. To provide *unicast* messages, an addressing layer is placed on top of the network protocol stack. In our Lego Mindstorms’ setting, each RCX has a unique identifier which serves as its address and is specified when the BrickOS firmware is downloaded.

3.2 Hardware

The hardware requirements of our warehouse storage system are high in Lego Mindstorms’ terms, requiring more sensors and actuators than one RCX can control. Therefore, it is necessary to use two RCXs, one to control the movement of the robot and the second to control the forklift, hereafter referred to as the *Movement RCX* and *Forklift RCX*, respectively. Again, communication between the two RCXs is handled using the infrared link provided on each RCX. Because the Forklift RCX does not need to communicate with the PC, the infrared download tower is set up to allow the following communication to take place: host computer to Movement RCX and Movement RCX to Forklift RCX, as shown in Fig. 6 below.

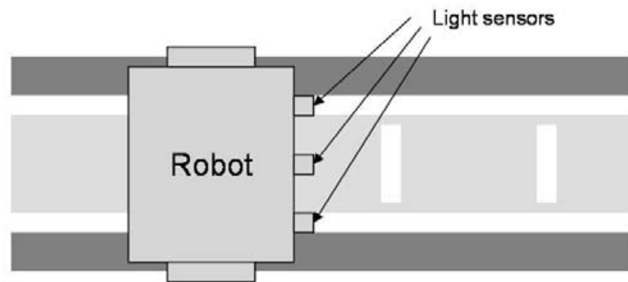

Figure 2: Track section.

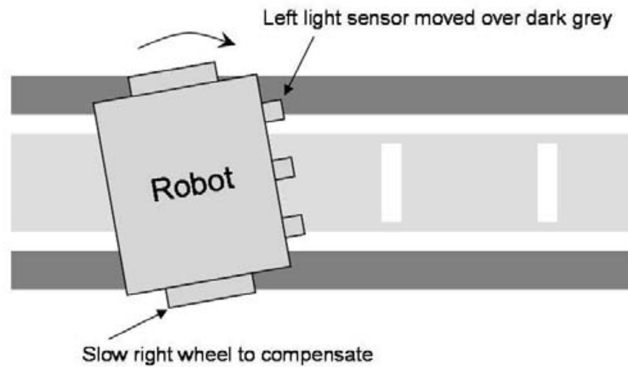

Figure 3: Track enabling straight robot movement.

3.2.1 Movement Subsystem and Track Description

Since the robot is required to pick up and place items in the warehouse with high precision and since the movement of the robot greatly influences its precision, it is necessary to make the movement subsystem as accurate as possible. Our solution to this challenge involves a combination of a drive system that guarantees straight line motion and facilities for the robot to detect an error in its direction and correct it. The underlying drive that moves the robot is implemented as a *dual differential drive*. This provides mechanically guaranteed straight motion and precise control of each wheel's speed. Each section of track that the robot moves on has two widely separated thin white lines over which two light sensors rest (cf. Fig. 2, outer light sensors). By using *thin* widely spaced tracks, the number of positions the robot can be in and that the light sensors still register as straight are limited, and therefore accuracy is enhanced. Any error in the starting direction will be detected by the light sensors moving over a non–white coloured part of the track. The areas inside the track and outside the track are different colours, so the error's direction can be inferred and corrected for by slightly slowing one wheel (cf. Fig. 3). This is the kind of precise wheel control that the dual differential drive excels at.

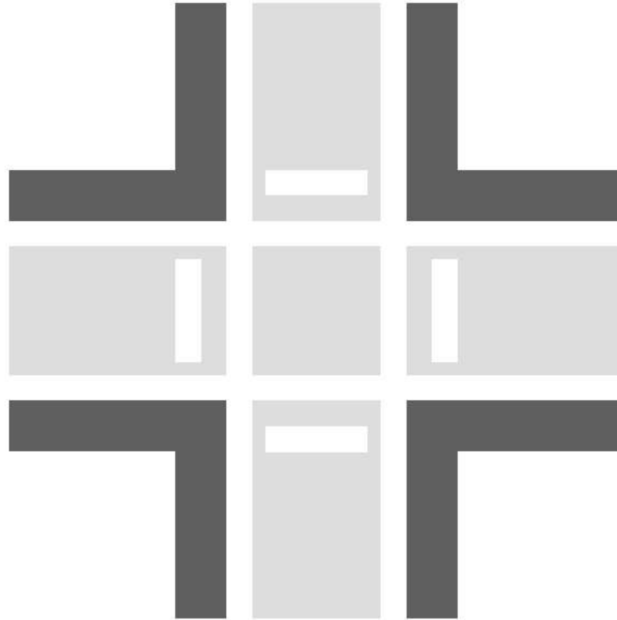

Figure 4: Track intersection.

Distance travelled by the robot is determined using a third light sensor which rests over the middle of the track. Thin white lines in the middle of the track, perpendicular to the direction of travel, are counted by this sensor, therefore enabling the robot to determine what distance it has travelled. Rather than placing the white lines a regular distance apart, they are placed beside objects that the robot might need to perform an operation at, such as to turn or to pick up an item. This enables the robot to stop at precisely the correct place and prevents error in the measured distance from accumulating.

At corners where sections of track intersect, the white lines used for distance measuring are replaced by the movement lines of the perpendicular track (cf. Fig. 4). The movement lines of the perpendicular track are designed to be identical in width to the distance measuring lines, thus providing a system that works over a track intersection.

At intersections, the same movement lines help the robot to perform precise 90 degree turns. After beginning the turn, the condition of both light sensors over a white part of the track is awaited. This indicates that the robot has turned 90 degrees. Again, because the lines are widely spaced, at the end of a turn the robot will be very close to straight for the next section of track.

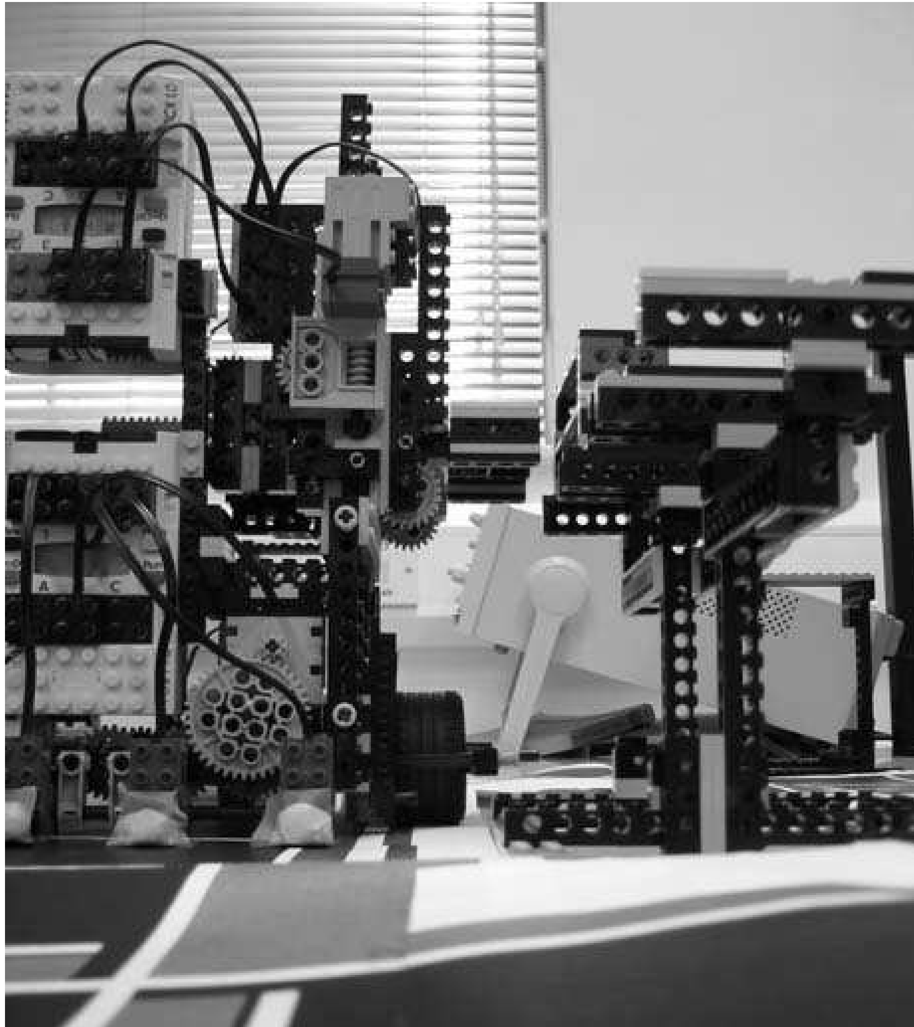

Figure 5: Forklift and palates.

3.2.2 Forklift Subsystem and Item Storage Design

To manipulate the warehouse items, a *forklift* design was chosen. To be compatible with the forklift design, all items must be placed on *palates* that are the right size for the forklift. Placing the items on palates also allows items to be stacked. We assume that all items of the same type, i.e., the same product, will be stored in one stack (cf. Fig. 5).

The forklift subsystem differs slightly from a normal forklift: instead of moving the whole robot forward to get the lift under a palate, the forklift arm is simply extended. This means that the forklift is situated on one side of the

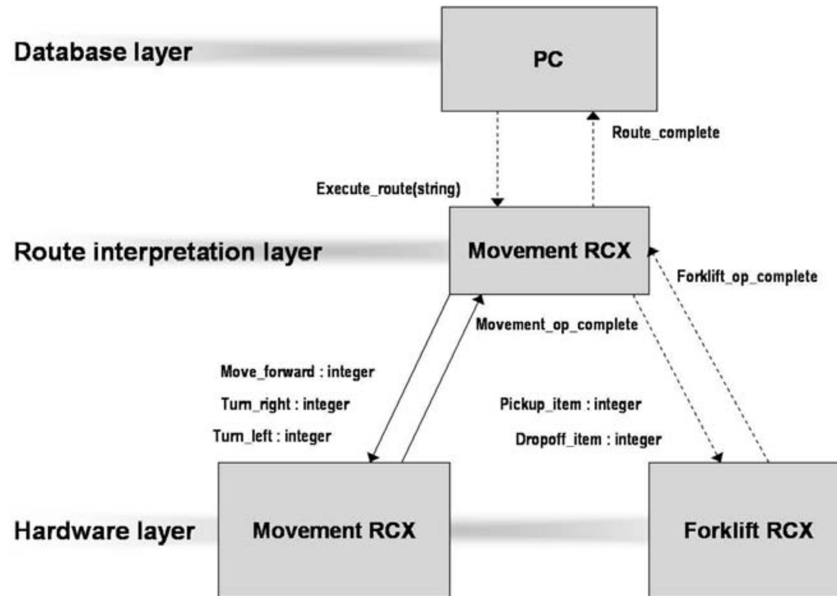

Figure 6: Diagram showing the inter–layer signals. (A normal line indicates internal communication and a dashed line communication over infrared link.)

robot rather than the front. The raising of the forklift and the extending of the arm are both controlled by motors and feedback is provided by two rotation sensors. Also, because the forklift subsystem is physically built on top of the movement subsystem, it is not possible to retrieve items from ground height. Therefore, all items are raised on top of a shelf system.

3.3 Software

The software for the warehouse system is structured in three layers: *database access layer*, *route interpretation layer* and *hardware layer*. Each device used in the system, the two RCXs and the PC, executes a reactive kernel programmed in Esterel. The database layer runs on the PC and is responsible for accessing the warehouse's database and for generating routes for the robot to execute. The second layer runs on the Movement RCX and interprets the route sent from the PC. The third layer is responsible for interacting with the Lego hardware and performs operations such as *move the robot* and *pick up the item*. This layer is present on the Movement RCX and the Forklift RCX. The signals used for communicating between the various layers are shown in Fig. 6.

3.3.1 Database

Our warehouse's database models a simple ordering system. Since the main emphasis within this case study is on retrieving spatial data, the aspects of the database concerning ordering information are kept as simple as possible: a customer may make many orders, each order is identified by an order id and must contain one or more order lines, and each order line must contain exactly one item. Stored separately is a table describing the drop–off bins in the warehouse. A drop–off bin is used by the robot to place parts of an order before it is complete. When the order is complete, the items the bin contains are removed from it, and the bin is then ready for another order. For each bin, its location is stored, as is the order id of the order in the bin when it is in use.

3.3.2 Database Access Layer

The database access layer uses our Local Result Set API. It maintains two connections to the same database since, at one point in the program, it is necessary to manipulate the database while retaining the result set of an earlier operation. The input and output signals for the main connection to the database are called `orders_query_out` and `orders_results`, respectively, and for the additional connection `stock_results` and `stock_query_out`, respectively, since those are only used to update stock levels:

```
output orders_query_out : string;
input orders_results : MYSQL_RES_ptr;

output stock_query_out : string;
input stock_results : MYSQL_RES_ptr;
```

Since the database access layer's only function is to wait for an order that needs to be picked and then to instruct the robot how to pick it, the main module is constructed as a loop. Inside the loop is a *trap* statement which handles all the ways in which the database and the system can fail. Therefore, all failure modes are dealt with in one place, and the error handling process is simplified.

The operation of the database access layer is roughly as follows (cf. Fig. 7). First, an order is retrieved from the database. One of the lines of this order is then extracted and the item details are stored locally. The robot is then sent to retrieve all items, one by one, and to deliver them in an available drop off bin using a pre–generated route. After each item has been collected, it is necessary to update the stock level of the item's type. However, at this point the database result set still contains uncollected order lines which are required for later in the program's execution. Therefore, the second database connection is used to perform the update without overwriting the result set containing the order's details.

Since the operation of the database access layer is quite simple and most database operations are effectively the same, only the database operation that

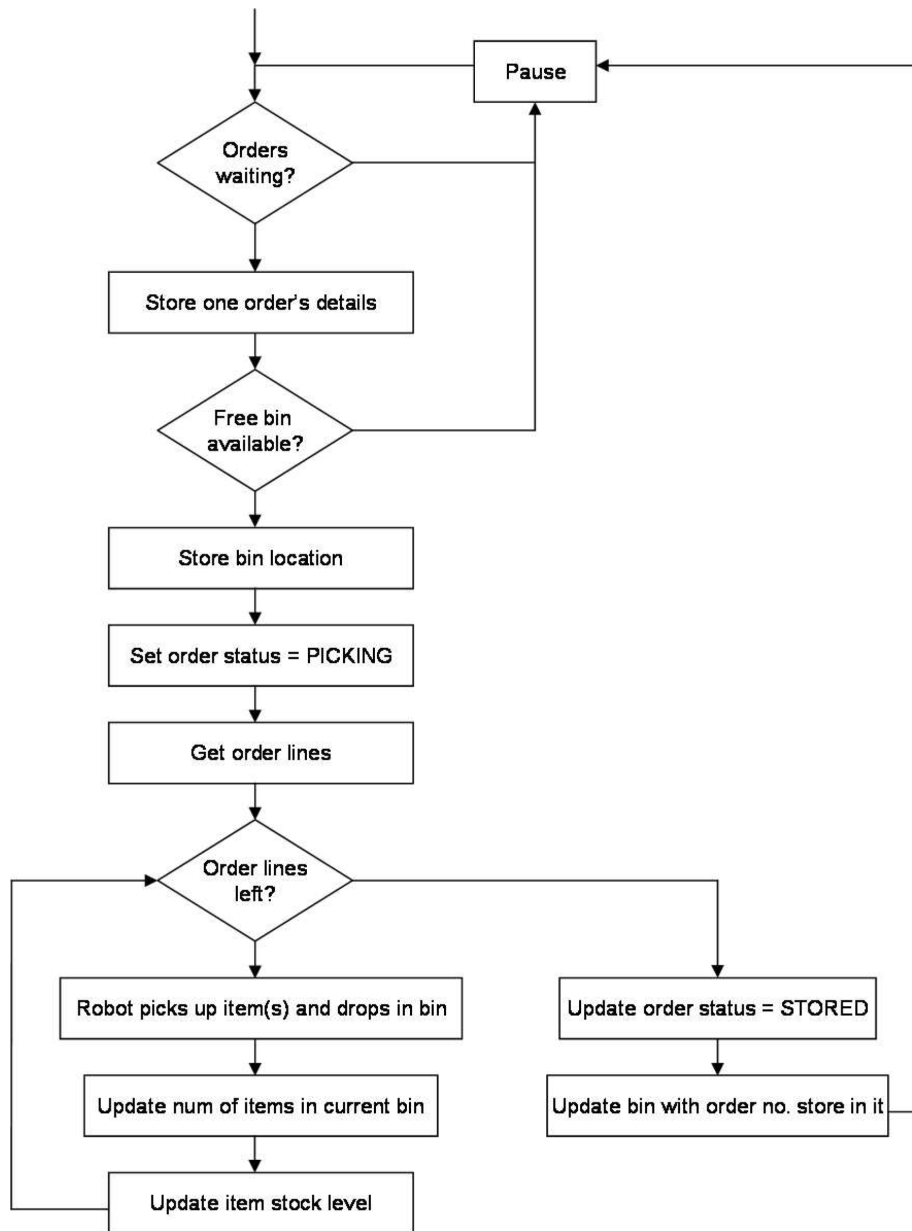

Figure 7: Flowchart describing the database access layer's behaviour.

retrieves waiting orders will be given here. The complete Esterel program can be found in [18]. The operation starts by emitting the query on output signal `orders_query_out` and then awaits the results:

```
emit orders_query_out("select order_id, customer.customer_id,
    name, address from customer, orders where
    customer.customer_id = orders.customer_id and
    orders.status = 'AWAITING_PICKING' order by order_id");
await orders_results;
```

The returned results are then checked for validity using our API function `check_result`. If the query has succeeded, then `num_rows` is called on the result set to see if any rows were returned. If rows have been returned, then there are orders waiting and, consequently, the first row is loaded into the local variable `row`. The details of the row are then retrieved using the `get<type>` functions and emitted on the corresponding local signals: `order_id`, `customer_id`, `customer_name` and `customer_address`. Now the result set is no longer needed, and the `clear_results` procedure is called. The following Esterel code captures this algorithm:

```
if (check_result(?orders_results)) then
  if (num_rows(?orders_results) > 0) then
    var row : MYSQL_ROW in
      row := get_next_row(?orders_results);
      emit order_id(getint(row,0));
      emit customer_id(getint(row,1));
      emit customer_name(getstr(row,2));
      emit customer_address(getstr(row,3));
    end var;
    call clear_results()(?orders_results);
  else ...
```

There are two ways in which this operation for retrieving orders can fail: firstly, if the query fails and, secondly, if there are no waiting orders. Each is catered for by an *exit* statement which corresponds to the trap mentioned above:

```
if (check_result(?orders_results)) then
  if (num_rows(?orders_results) > 0) then
    %Code snipped
  else
    call clear_results()(?orders_results);
    exit no_waiting_orders;
  end if;
else
  exit bad_query;
end if;
```

In each case, the program flow jumps to the end of the loop and emits an appropriate error message before pausing and then repeating the main loop. Note that function `clear_results` must be called after it has been determined that there are no waiting orders, freeing the memory occupied by the result set.

The generation of the pick–up and drop–off routes is the only non–database related function of the database access layer. To generate the route, first the item's height is computed and passed to a function called `generate_route_to_pickup_item()` along with the item's x and y coordinates. The function generates a string consisting of *op codes* that represent the operations the robot must perform to pick up that item and to return to the communication point in the warehouse. The string is sent to the robot via a signal, and then another signal indicating the completion of executing the route is awaited. After the pick–up route is complete, a drop–off route is computed by calling function `generate_route_to_drop_item()` and passing it the order's drop–off bin location. The route is emitted and the completion awaited on the same signals used by the pick–up operation. Again, details can be found in [18].

By using external functions to generate the route, the warehouse can be redesigned in any way consistent with the item locations stored in the database, and only the two route generating functions will have to be rewritten.

3.3.3 Route Interpretation Layer

To interpret a route string of op codes, a number of external C functions are provided that extract parts of the string. First, function `get_num_ops()` is called to determine how many separate operations the route string contains. A loop is then repeated this number of times. On each iteration, a function `get_op()` is invoked to extract the type of operation and `get_param()` to extract the parameters to the operation. Once the operation type has been determined, an emission is made on the appropriate signal with the value obtained from `get_param()`. When the robot has completed the operation, either the `movement_op_complete` or `forklift_op_complete` signal is emitted, informing the route interpretation layer that the robot is ready for the next operation. When all operations have been performed in this manner, the signal `route_complete` is emitted which lets the database access layer know that the robot has finished its route.

3.3.4 Hardware Layer

The Movement RCX is required to run both the route interpretation layer and part of the hardware layer. To accomplish this, both layers are run in parallel and local signals are used to communicate between them. To communicate with the hardware layer running on the Forklift RCX, the output signals `pickup_item` and `drop_item` and the input signals `forklift_op_complete` are used (cf. Fig. 6). These combined with the signals of the hardware layer running on the Movement RCX, i.e., signals `move_forward`, `turn_left`, `turn_right` and `movement_op_complete`, give the complete range of commands provided by the hardware layer. A full listing of all the layers' implementations in Esterel is contained in [18].

4 Conclusions

This report presented two APIs for interfacing the synchronous programming language Esterel to the relational database MySQL. The *Local Result Set API* assumes the result set to a database query to be stored locally to a reactive Esterel kernel and largely relies on the external function concept of Esterel. The *Remote Result Set API* considers the result set to be stored remotely and is realised via signals.

Both database APIs worked well in testing. In particular, the Local Result Set API is heavily used in our case study and, although none of the database operations are particularly complex, they are representative of the kind of database operations performed by reactive systems. In contrast to the elegance exhibited by the Local Result Set API, the Remote Result Set API appears to be slightly convoluted. This is due to modelling all database operations as signals, which became necessary since remoteness implies that one cannot expect instantaneous responses and, hence, cannot use external functions. It remains to be explored whether an implementation based on Esterel's task concept would be more elegant.

It should be emphasised that the introduction of a database using either API in an Esterel reactive system does not undermine Esterel's synchrony hypothesis. However, since the response times for returning query results or for accessing remote result sets cannot be guaranteed, the system can end up waiting for a signal that may never arrive. If the database is one that can provide guaranteed response times, such as a real–time database [12], the problem is elevated. Otherwise, the problem must be solved via timeouts in system design. In our case study, all database operations are performed at non–time–critical points, whence any unexpected delay from the database simply results in the system pausing, not malfunctioning.

Future Work

Future work is proposed to proceed along three directions. Firstly, some restrictions on the usage of our Local Result Set API may be removed, thereby making the API safer for use. In particular, our implementation implies that row variables become undefined once the result set is cleared. In a similar vein, checks could be implemented in our APIs to detect whether they are used in any way other than that intended.

Secondly, although MySQL is a popular database, our APIs would be applicable to others if they supported the *Open Database Connectivity* (ODBC) API [16]. This would allow any database to be used with no increase in complexity to our present approach.

Thirdly, database APIs for Lustre/SCADE are envisioned along similar lines than ours for Esterel/Esterel Studio.

Acknowledgements

We would like to thank the participants of the SYNCHRON '04 workshop for their constructive comments, especially for pointing out the appropriateness of Esterel's task concept for the Remote Result Set API.

References

[1] D. Baum, M. Gasperi, R. Hempel, and L. Villa. *Extreme Mindstorms – An advanced guide to Lego Mindstorms*. Apress, 2000.

[2] G. Berry. The constructive semantics of pure Esterel. CMA, Ecole des Mines, INRIA, 1999. Draft version 3.0.

[3] G. Berry. The Esterel v5 language primer. CMA, Ecole des Mines, INRIA, 2000.

[4] G. Berry. The foundations of Esterel. In *Essays in Honour of Robin Milner*. MIT Press, 2000.

[5] R.H. Bishop. *Modern Control Systems: Analysis and Design Using MATLAB and SIMULINK*. Addison Wesley, 1997.

[6] BrickOS. *brickos.sourceforge.net*, last visited in 2004.

[7] P. Caspi, D. Pilaud, N. Halbwachs, and J.A. Plaice. LUSTRE : A declarative language for programming synchronous systems. In *14th ACM Symposium on Principles of Programming Languages*. ACM, 1987.

[8] Esterel compiler. *www-sop.inria.fr/meije/esterel/esterel-eng.html*, last visited in 2004.

[9] Esterel Technologies. Esterel Studio and SCADE Suite & Drive design tools. *www.esterel-technologies.com*, last visited in 2004.

[10] N. Halbwachs. *Synchronous Programming of Reactive Systems*. Kluwer Academic Publishers, 1993.

[11] D. Harel and M. Politi. *Modeling Reactive Systems with Statecharts: The STATEMATE Approach*. McGraw Hill, 1998.

[12] K.-Y. Lam and T.-W. Kuo. *Real-time Database Systems: Architecture and Techniques*. Kluwer Academic Publishers, 2000.

[13] Lego Mindstorms robotics kit. *mindstorms.lego.com*, last visited in 2004.

[14] The MySQL open source database. *www.mysql.com*, last visited in 2004.

[15] The MySQL C API. *dev.mysql.com/doc/mysql/en/C.html*, last visited in 2004.

[16] R. Signore, J. Creamer, and M.O. Stegman. *The ODBC Solution: Open Database Connectivity in Distributed Environments*. McGraw Hill, 1995.

[17] L. Wall, T. Christiansen, and J. Orwant. *Programming Perl*. O'Reilly, 2000.

[18] D. White and G. Lüttgen. Accessing databases from Esterel. Technical report, Department of Computer Science, University of York, UK, 2005. To appear. Available electronically at *www.cs.york.ac.uk/luettgen/publications.html#YCS-2004* .