

**OBJECTIVE: A Benchmark for Object-Oriented  
Active Database Systems**

Uğur Çetintemel, Jorgen Zimmermann, Özgür Ulusoy, and Alejandro Buchmann

**BILKENT UNIVERSITY**

**Department of Computer Engineering  
and  
Information Science**

Technical Report BU-CEIS-96-10

# OBJECTIVE: A Benchmark for Object-Oriented Active Database Systems

U. Çetintemel<sup>1</sup>, J. Zimmermann<sup>2</sup>, Ö. Ulusoy<sup>1</sup>, and A. P. Buchmann<sup>2</sup>

<sup>1</sup> Bilkent University, 06533 Ankara, Turkey

<sup>2</sup> University of Darmstadt, 64293 Darmstadt, Germany

## Abstract

Although much work in the area of Active Database Management Systems (ADBMSs) has been done, it is not yet clear how the performance of an active DBMS can be evaluated systematically. In this paper, we describe the OBJECTIVE Benchmark for object-oriented ADBMSs, and present experimental results from its implementation in an active database system prototype. OBJECTIVE can be used to identify performance bottlenecks and active functionalities of an ADBMS, and compare the performance of multiple ADBMSs.

## 1 Introduction

An Active Database Management System (ADBMS) detects certain situations and performs corresponding user defined actions in the form of Event-Condition-Action (ECA) rules [DBB<sup>+</sup>88]. ADBMSs have received great attention lately, and several prototypes of object-oriented ADBMSs are already available (e.g., ACOOD [Ber91], NAOS [CCS94], Ode [AG89], REACH [BZBW95], SAMOS [GGD94], SENTINEL [CVKB94]). We are currently in a position to evaluate the performance of ADBMSs by concentrating on

- the performance requirements of different architectural approaches; i.e., integrated versus layered,
- different techniques used for standard tasks of an ADBMS; i.e., rule maintenance, event detection, and
- a variety of functionalities provided by an ADBMS; e.g., garbage collection and parameter passing.

Benchmarking is a very important process in the sense that database users base their purchasing decisions partially relying on benchmark results, and database designers measure the performance of their systems by using an appropriate benchmark. There has been much work in the area of database benchmarking; e.g., the Wisconsin Benchmark [BD84], the OO1 Benchmark [CS92], and the OO7 Benchmark [CDN93]. However, there have been only a few attempts to evaluate the performance of ADBMSs, the most important of which are the BEAST Benchmark [GGD95], and the ACT-1 Benchmark [ZBD95].

In this paper, we describe the OBJECTIVE<sup>1</sup> Benchmark which is a simple but comprehensive test of active functionalities provided by an *object-oriented* ADBMS, and give performance results of its implementation in an ADBMS prototype. OBJECTIVE can be used to identify performance bottlenecks and active functionalities of an ADBMS, and compare the performance of multiple ADBMSs. The philosophy of OBJECTIVE is to isolate components providing active functionalities, and concentrate only on the performance of these components while attempting to minimize the effects of other factors (e.g., underlying platform). OBJECTIVE operates on a very simple database structure consisting of completely synthetic classes, events, and rules. Although the design is very simple (for ease of reproducibility and portability), this simplicity does not contribute negatively to the benchmark in any manner.

The OBJECTIVE Benchmark addresses the following issues with respect to object-oriented ADBMS performance and functionality:

- method wrapping penalty,
- detection of primitive and composite events,
- rule firing,
- event-parameter passing,
- treatment of semi-composed events, and
- rule administration tasks.

The OBJECTIVE Benchmark comprises a number of operations that evaluate the issues stated above, and those operations were first run on REACH [BZBW95]. REACH is a full-fledged operational object-oriented ADBMS which is tightly integrated in Texas Instruments' Open OODB [WBT92]. The results reported in this paper reveal that REACH combines the most advanced features of current ADBMS proposals from the functionality point of view. As for its performance, a single bottleneck operation is identified.

The remainder of the paper is organized as follows. Section 2 discusses the main features of ADBMSs and Section 3 discusses previous related work. The OBJECTIVE Benchmark is introduced along with performance results of its implementation in REACH in Section 4. Finally, Section 5 concludes and gives directions for future work. In addition, Appendix A gives simplified codes which illustrate the implementation of the OBJECTIVE operations, and Appendix B presents the complete set of REACH results.

## 2 Active Database Management Systems

ADBMSs extend passive DBMSs with the ability to specify and implement reactive behavior which is typically specified in term of ECA rules. The general form of an ECA rule is: on *event* if *condition* do *action*. The semantics of such a rule is that *when* the *event* occurs, the *condition* is checked, and *if* it is satisfied *then* the *action* is executed. Therefore, an ADBMS has to monitor events (of interest) and detect their occurrences. After an event is detected, it is *signalled*. This signalling is a

---

<sup>1</sup>OBJECT-oriented actIVE database systems benchmark

notification that an interesting event has occurred, and rule execution should take place. ECA rules require, at least, the operations *insert*, *delete*, and *fire*. These operations are used to insert a new rule into the database, delete an existing rule from the database, and trigger a rule, respectively. For some applications it may be useful to *disable* rules temporarily, which can afterwards be *enabled* when necessary [DBM88].

This section discusses main issues in ADBMSs to an extent which is necessary for the comprehension of the rest of the paper.

## 2.1 Events

ECA rules are triggered on the occurrence of particular events. An event can be either *primitive* or *composite*. Primitive events are atomic events which can be associated with a point in time. The most commonly referred primitive event types are [BZBW95, GD93, CM93]:

- *method events*  
A method invocation can be defined as an event of interest. In such a case, an event occurs when its corresponding method is executed. Since a method execution corresponds to an interval rather than a point in time, usage of time modifiers like BEFORE or AFTER is mandatory. The semantics of BEFORE and AFTER modifiers, respectively, is that the method event is to be raised just before the invocation of the method, and immediately after the execution of the method.
- *state transition events*  
A change in the state of the object space can be an event; e.g., modification of an object attribute. It is then necessary to define operators to access old and new values of relevant entities.
- *temporal events*  
Basically, two types of temporal events exist; *absolute* and *relative*. Absolute temporal events are defined by giving a particular point in time (e.g., 01.10.1996, 11:23), whereas relative temporal events are defined relative to other events (e.g., 10 minutes after commit of a particular transaction). The latter type can also include events which occur periodically (e.g., every day at 17:30).
- *transaction events*  
Transaction events correspond to standard transaction operations like *begin of transaction* (BOT), *end of transaction* (EOT), *abort of transaction* (ABORT), and *commit of transaction* (COMMIT).
- *abstract events*  
Abstract events are user-defined events whose occurrences are directly signalled. Therefore, the underlying system does not need to monitor abstract events; i.e., they are explicitly raised by the user and associated with a point in time.

Different techniques are used for the detection of method events. A straightforward approach is to modify the body of the method for which an event is to be defined with an explicit raise of an event [GD93]. Another technique, *method wrapping*, is to replace the original method with a *method wrapper* that contains an explicit event raise operation and a call to the original method [BZBW95].

Unlike primitive events which are atomic, composite events are defined as a combination of primitive (and possibly other composite) events. The meaningful ways to build composite events from its constituent events are usually specified through an *event algebra* that defines certain *event constructors*. Some useful event constructors are [DBM88, GGD94]:

- The *disjunction* of two events, event1 and event2, is raised when either of event1 or event2 occurs.
- The *conjunction* of two events, event1 and event2, is raised when both event1 and event2 occur.
- The *sequence* of two events, event1 and event2, is raised when event1 and event2 occur in that order.
- The *closure* of an event, event1, is raised exactly once regardless of the number of times event1 occurs (provided that event1 occurs at least once).
- The *negation* of an event, event1, is raised if event1 does not occur in a given time interval.
- The *history* of an event, event1, is raised if event1 occurs a given number of times.

For the last three event constructors, it is appropriate to define time intervals in which composition of events should take place. The definition of a time interval is mandatory for negation, and optional for history and closure.

Composite events can further be grouped into *aggregating* composite events and *non-aggregating* composite events [ZB95]. The former group contains composite events that are constructed with the operators sequence, disjunction, and conjunction, whereas the latter group comprises composite events constructed with history, negation, and closure.

Several different approaches are used for composite event detection including syntax graphs [Deu94, CKAK93], Petri nets [GD94], finite state automata [GJS92], and arrays [Eri93].

## 2.2 Conditions

The condition part of a rule is usually a boolean expression, a predicate, or a set of queries, and it is satisfied if the expression evaluates to true, or all the queries return non-empty results, respectively. In addition to the current state of the database, the condition may access the state of the database at the time of event occurrence by the use of event parameters.

## 2.3 Actions

The action part of a rule is executed when the condition is satisfied. In general, an action can be database operations, transaction commands (e.g., abort transaction), or arbitrary executable routines. Therefore, during the execution of an action some events may also occur. This may lead to the triggering of other rules which is called *cascaded rule triggering*. The action may access, besides the current database state, the database state at the time of event occurrence and the time of condition evaluation which can be accomplished by parameter passing.

## 2.4 Execution Model

An *execution model* specifies the semantics of rule execution in a transaction framework. A transaction which triggers rules is called a *triggering transaction*, and the (sub-)transaction which executes the triggered rule is called the *triggered (sub-)transaction*. An important issue which is determined by an execution model is the coupling between the triggered transaction and the triggering transaction. Additionally, an execution model also describes concurrency control and recovery mechanisms used to achieve a correct and reliable rule execution. These two issues are discussed in more detail in the rest of this subsection.

Coupling modes determine the execution of rules with respect to the transaction which triggers them. The Event-Condition (EC) and Condition-Action (CA) coupling modes, respectively, determine when the rule's condition is evaluated with respect to the triggering event, and when the rule's action is executed with respect to the condition evaluation. Three basic coupling modes are introduced in [Day88]: *immediate*, *deferred*, and *decoupled*.

For EC coupling, the intended meaning of each mode is:

- In *immediate* EC coupling mode, the condition is evaluated in the triggering transaction, immediately after the detection of the triggering event.
- In *deferred* EC coupling mode, the condition is evaluated at the end but before the commit of the triggering transaction.
- In *detached* EC coupling mode, the condition is evaluated in a separate transaction which is independent from the triggering transaction.

For CA coupling, the semantics of each mode can be given analogously.

If several triggered rules have to be executed at the same point in time, they form a *conflict set* [HW92]. In this case, some sort of *conflict resolution* (e.g. priorities) must be employed to control their execution order. The ability to do such a resolution is especially desirable if we want to impose a particular serial order of execution.

Since condition and action parts of a rule may act on database objects, the execution of rules must be done in a transaction framework. The nested transaction model [Mos85] is the most prevalent approach for rule execution in ADBMSs, primarily due to the fact that it captures the semantics of (cascaded) rule triggering well. In this model, the triggered rules are either executed as subtransactions of the triggering transaction, in case of immediate and deferred coupling modes, or as an independent transaction in case of detached coupling mode.

## 3 Related Work

Although much work in the area of ADBMSs has been done, it is not yet clear how the performance of an ADBMS can be evaluated systematically. In fact there have been very few attempts including [GGD95, ZBD95, BM93, Ker95]. In this section, we discuss these efforts in some detail.

### 3.1 The BEAST Benchmark

The BEAST is the first benchmark proposed for testing the performance of object-oriented ADBMSs [GGD95]. It was presented as a designer's benchmark; i.e., the designers of an ADBMS can use it

to determine performance bottlenecks of their systems. It uses the database and schema of the OO7 Benchmark [CDN93]. The BEAST Benchmark focuses on event detection, rule management, and rule execution aspects of an ADBMS. These three aspects represent the whole active behavior and should be covered by any ADBMS.

The BEAST Benchmark runs a series of tests to determine the functionality of each component. It consists of:

- *Tests for event detection* concentrate on the time to detect particular events. A set of primitive and composite events are tested. Tests for primitive event detection consist of the detection of value modification, the detection of message sending, the detection of transaction events, and the detection of a set of primitive events. The BEAST tests for composite event detection comprise the detection of a sequence of primitive events, the detection of a non-occurrence of an event within a transaction, the detection of a repeated occurrence of a primitive event, the detection of a sequence of composite events, the detection of a conjunction of method events sent to the same object, and the detection of a conjunction of events belonging to the same transaction.
- *Tests for rule management* evaluate the rule management component of an ADBMS by measuring the retrieval time of rules. It does not include any tests for rule definition and storage.
- *Tests for rule execution*  
The tests for rule execution consider both the execution of single and multiple rules. For the execution of single rules, a rule is executed with different coupling modes. In the case of multiple rule execution, the tests concentrate on the overhead of enforcing an ordering on the triggered rules, optimization of condition evaluation and raw rule execution power of the underlying system.

In all these tests response time was accepted as the sole performance metric. In the experiments, the number of defined events (primitive and composite), and the number of rules were used as benchmark parameters, and a set of quantitative results were obtained for each particular setting of these parameters. To date, BEAST has been run on four object-oriented ADBMS prototypes, and the performance results are presented in [GBLZ96].

### 3.2 The ACT-1 Benchmark

The ACT-1 Benchmark [ZBD95] concentrates on the minimal features of object-oriented ADBMSs. Four basic issues are addressed in this benchmark:

1. *Method wrapping penalty* measures the useless overhead of method wrapping for the detection of method events.
2. *Rule firing cost* measures the cost of raising an event and firing the corresponding rule.
3. *Minimal event composition cost* aims to assess the cost of a simple event composition (the sequence of two events).
4. *Sequential rule firing cost* concentrates on the overhead of serialization of a set of rules that have to be executed at the same time (two rules that are triggered by the same event at the same coupling mode).

ACT-1 uses a simple database with objects and rules modeling the operation of a power plant. Four operations, WRAPPING PENALTY, FIRING COST, BUILD UP, and SEQ EXEC, are implemented in REACH and some preliminary results based on response times of these operations are presented.

### 3.3 Other ADBMS Benchmarking Related Work

There are several other performance evaluation studies on ADBMSs. Actually, these are not devoted performance evaluation works; rather, they present a rule (sub)system and then evaluate its performance. For instance [BM93] mainly addresses the problem of handling large rule sets. It argues that the techniques used in current active database prototypes are not appropriate for handling large rulebases. It proposes a novel indexing technique for rule activation and gives performance results of DATEX, a database rule system, which uses this particular technique. Storage size and number of disk accesses are used as the cost metrics in this evaluation.

[Ker95] presents another performance study on active functionality in DBMSs. It gives a performance evaluation of the rule system of MONET—a parallel DBMS kernel aimed to be used as a database back-end—by using a simple benchmark. This simple *core* benchmark is designed mainly for testing the implementation of MONET, and it consists of three basic experiments. The *countdown* experiment tries to assess the cost of handling a single event and subsequent firing of a single rule (i.e., an abstract event is signalled and a rule is fired by this event. This fired rule notifies the same event which further leads to the triggering of the same rule. This is repeated a predetermined number of times). The *dominoes* experiment is aimed to determine the cost of isolating a firable rule instance. The *pyramid* experiment has the purpose of investigating the performance of the system under high active workloads.

## 4 The OBJECTIVE Benchmark

The aim of the OBJECTIVE Benchmark is to identify the bottlenecks and functionalities of an object-oriented ADBMS, and to create a level-playing field for comparison of multiple object-oriented ADBMSs. The BEAST Benchmark is a very good initial step towards a benchmark which will cover a bigger set of functionalities of an ADBMS. However, we require a more generic benchmark to be able to test *both* performance and functionality. Typically, a system with little functionality can be implemented more efficiently than a system with more functionality. As an example, consider the (useless) overhead of method wrapping. At one extreme, there are systems that hand-wrap only those methods on which a rule is defined, and at the other extreme there are systems that do automatic wrapping of all the methods. The latter systems allow the definition of new rules without requiring the recompilation of classes, but pay for the wrapping when a method that is not an event type for any rule is invoked. Likewise, a system that allows event parameters to be passed to condition and action parts of rules will be much more flexible than the one which does not support such a functionality, but at the same time it will face an overhead in event composition and rule execution in non-immediate coupling modes. Therefore, in order not to skew results in favor of sys

After introducing the operations of the OBJECTIVE Benchmark along with a requirements analysis in Section 4.1, we describe the synthetic database of OBJECTIVE in Section 4.2. In Section 4.3, we describe the implementation of the OBJECTIVE operations in detail while presenting experimental results of their implementation in REACH.



TEST	DESCRIPTION
MW1	Method wrapping <i>penalty</i>
PED1	Detection of a <i>method invocation</i> event
PED2	Detection of a <i>BOT</i> event
PED3	Detection of a <i>COMMIT</i> event
CED1	Detection of a <i>sequence</i> of primitive events
CED2	Detection of a <i>conjunction</i> of primitive events
CED3	Detection of a <i>negation</i> of a primitive event
CED4	Detection of a <i>history</i> of a primitive event
CED5	Detection of a <i>closure</i> of a primitive event
RF1	<i>Retrieval</i> of a rule
RF2	Rule firing in <i>deferred</i> coupling mode
RF3	Rule firing in <i>decoupled</i> coupling mode
RF4	Rule <i>execution</i>
RF5	<i>Conflict resolution</i> of triggered rules
RF6	<i>Cascaded</i> rule triggering
EPP1	The passing of event parameters in <i>immediate</i> coupling mode
EPP2	The passing of event parameters in <i>deferred</i> coupling mode
EPP3	The passing of event parameters in <i>decoupled</i> coupling mode
GC1	The <i>garbage collection</i> of semi-composed events
RA1	<i>Creating</i> a rule
RA2	<i>Deleting</i> a rule
RA3	<i>Enabling</i> a rule
RA4	<i>Disabling</i> a rule
RA5	<i>Modifying</i> a rule

Table 1: The OBJECTIVE operations

#### 4.1 The OBJECTIVE Operations

The OBJECTIVE Benchmark addresses the following issues [ZB95] by the operations which are described briefly in Table 1:

1. *Method wrapping penalty*

In an object-oriented database system where method wrapping is used for method event detection, there is a *useless* overhead which is generated when a method which does not generate any event or which generates an event that does not contribute to the triggering of any rule is invoked (i.e., such an event is neither a primitive event for a rule, nor a part of a composite event for a rule). Ideally, the introduction of active capabilities should not deteriorate the performance when they are not in effect. In other words, ADBMS users should not pay for active functionality when they do not use it. Therefore, an ADBMS must keep such a (useless) overhead minimal.

2. *Event detection*

An ADBMS should support primitive and composite events and response times for event detection, both primitive and composite, are crucial for the performance of an ADBMS. The

primitive event types should minimally include method events and transaction events. For composite events, at least, the detection time for an aggregating event and a non-aggregating event should be measured.

### 3. *Rule firing*

Rules typically reside in secondary storage and have to be fetched into main memory for execution. Therefore, efficient retrieval of rules whose events are signalled is indispensable for an ADBMS. As well as for capturing the semantics of some applications, (non-immediate) coupling modes are introduced primarily for increased performance with respect to execution of rules. If different coupling modes cannot be supported effectively, then there will hardly be any point in keeping them. Therefore, efficient firing of rules in different coupling modes is a crucial issue. Different approaches can be taken in the storage of condition/action parts of a rule (e.g., compiled code). Regardless of their internal representation, efficient access and execution of these parts is mandatory. Another pragmatic issue is the conflict resolution of a set of rules that are to be executed at the same point in execution flow. In addition, the ability to treat application/program execution and rule execution uniformly is also significant. Extra overhead should not be introduced for detection of events and firing of rules during rule execution.

### 4. *The handling of event parameters*

For some applications, e.g., consistency-constraint checking and rule-based access control, event parameters must be passed to the condition-action part of the rule. Otherwise, expressing conditions and actions with proper bindings is not possible. This requires the usage of some intermediate storage in case the rule is executed in either deferred or detached coupling mode. In immediate coupling mode it may be sufficient to pass a pointer to the parameters instead of passing the parameters themselves. However, this approach may not be applicable in deferred and detached coupling modes, because the parameters to be passed might be transient objects rather than persistent ones. The way event parameters are handled, thus, has a great impact on the performance of the system.

### 5. *Garbage collection of semi-composed events*

The problem of garbage collection exists for some composite events that are not fully composed, and whose extents have expired [BZBW95]. If no garbage collection is done for such semi-composed events, the database size will increase unnecessarily which will lead to a further increase in response time. So, an efficient mechanism for garbage collection of semi-composed events must be employed from the performance point of view.

### 6. *Rule administration*

An ADBMS should be able to create, destroy, enable and disable rules on-line. The ability to maintain rules dynamically is very important because of well-known reasons of availability and extensibility. Although the execution speeds of these tasks are not of great importance, a comprehensive benchmark should take them into account.

## 4.2 Description of the OBJECTIVE Database

Generation of a synthetic database is an important issue in all benchmarks for database systems [Gra91]. In a benchmark for active database systems, the most interesting part of database specifica-

```

class Name{
    int attribute;
    double data;
public:
    void doNothing()                {; }
    void setAttribute(int i)        {attribute = i;}
    int getAttribute()              {return attribute;}
    void setData(double d1, double d2) {data = d1 - d2;}
    void setMinData()               {data = 0.0;};

```

Figure 1: A class example

tion is the specification of events and rules, because tests of the benchmark will typically concentrate more on rules and events than particular objects in the database.

The database for the OBJECTIVE Benchmark consists of completely synthetic object classes with the same methods and attributes (see Figure 1 for a generic class definition<sup>2</sup>), and it has a very simple schema. The rationale for this decision is twofold: First, a benchmark should be easily reproducible and portable, and second OBJECTIVE is designed to be a generic benchmark, not a domain-specific benchmark; i.e., the aim of OBJECTIVE is to test important aspects of system performance and functionality, not to model a particular application. Thus, we do not want to add extra complexity which will not contribute to the benchmark in any manner, but which will make the implementation more difficult.

Several events and rules are defined (Figure 2 and Figure 3) to be used in benchmark operations. The rules are defined in the rule programming language REAL<sup>3</sup> (REAch rule Language) [ZBB<sup>+</sup>96]. The events, however, are defined in a hypothetical language based on the event definition notation of REAL<sup>4</sup>. These events and rules are discussed in detail in Section 4.3 where we describe the implementation of the benchmark operations. The naming convention used for objects, events, and rules are based the name of the relevant operation; e.g., the objects, events, and rules of name EPP1 are the ones that will be utilized in operation EPP1.

In addition to these events, rules, and classes which are used in the benchmark tests, we also utilize dummy event, rule, and class types. By changing the number of instances of these dummy types, we can run our operations for different database configurations, and see their effects on system performance.

We define dummy classes with the same methods and attributes, and the instances of these dummy classes form the (dummy) database objects. The methods of these dummy classes are used to generate before/after (dummy) method events. The event constructors *sequence* and *history* are used to generate non-aggregating and aggregating (dummy) composite event types, respectively. The number of

---

<sup>2</sup>We use a notation for our class definitions and test routines which is the de facto standard for object-oriented languages, namely the notation of C++ programming language.

<sup>3</sup>Rules in REAL consist of parts for defining a rule's event, condition and action along with EC and CA coupling modes and priorities. The default value for a coupling mode is *imm*(ediate) and the default values for method event modifiers and priorities (priority range is {1, 2, ..., 10}) are *after* and *5*, respectively. In addition, there is a *decl*(laration) section in which variables are specified in a C++ manner.

<sup>4</sup>REAL does not consider the definition of stand-alone event types

```

event PED1 { PED1::doNothing(); };
event PED2 { BOT('PED2'); };
event PED3 { COMMIT('PED2'); };
event CED1 { ABSTRACT(CED1_1) then ABSTRACT(CED1_2); };
event CED2 { ABSTRACT(CED2_1) and ABSTRACT(CED2_2); };
event CED3 { not ABSTRACT(CED3_2) in (ABSTRACT(CED3_1) , ABSTRACT(CED3_3)); };
event CED4 { 1 times ABSTRACT(CED4_2) in (ABSTRACT(CED4_1) , ABSTRACT(CED4_3)); };
event CED5 { all ABSTRACT(CED5_2) in (ABSTRACT(CED5_1) , ABSTRACT(CED5_3)); };

```

Figure 2: The events related to event detection operations

component events to form a composite event is selected at random<sup>5</sup> from range  $\{2, 3, \dots, 10\}$ . Likewise, the component event types are selected randomly from the already generated method event types. The dummy rules choose their event types at random from the existing dummy primitive and composite event types. Both the condition and action parts of dummy rules are defined as empty.

We include four parameters for the OBJECTIVE Benchmark; *NumEvents*, *FracCompEvents*, *NumRules*, and *NumObjects*, which define the number of (dummy) events, the fraction of composite events, the number of (dummy) rules, and the number of (dummy) data objects, respectively. The database configurations based on these parameters are summarized in Table 2.

### 4.3 Implementation and Results

In this section, we discuss the implementation of the benchmark operations which are described briefly in Section 4, and present the results of their application to REACH.

In all the operations described in this section, we assume that access to the internals of an ADBMS is not possible. This assumption is made due to two primary reasons: First, this is generally the case in reality, and second we want our benchmark to be a general one so that it can be applied to different ADBMSs through their external interfaces. Although this assumption makes accurate time measurement impossible for certain tests, we can circumvent it to a certain extent by keeping all the other non-interesting phases as small as possible by using appropriate events and rules. Actually, we assume that we can run our tests by just using the application programming interface of an ADBMS.

We make use of two time measures for the OBJECTIVE operations (whenever appropriate); *cold* and *hot* times representing the elapsed times when a measurement is done beginning with empty buffer, and beginning with completely initialized buffer, respectively. However, we do not present both cold and hot time results for all operations. Instead, we prefer to present the more meaningful and informative time measure for a given operation according to the focus of that operation. As a case in point, it is more meaningful to concentrate on the cold times for an operation concerned with rule retrieval, while one should emphasize the hot time results for conflict resolution of triggered rules.

We consider the *CPU time* used by the process running an operation instead of wall-clock time, because we ran the benchmark in a normal operating environment (i.e., not in an isolated machine), and we do not want to include the effects of certain operating system tasks in our results. Another important point to note is that we always use *transient* objects rather than persistent ones in order

---

<sup>5</sup>Uniform distribution is used in all random selections.

```

rule RF1{
  decl ;
  event ABSTRACT(RF1);
  cond FALSE;
  action ;
};
rule RF4{
  decl ;
  event ABSTRACT(RF4);
  cond TRUE;
  action ;
};
rule RF6{
  decl RF6 *obj;
  int i;
  event obj->setAttribute(i)
  cond obj->getAttribute() > 0;
  action obj->setAttribute(i-1);
};
rule EPP3{
  decl EPP3 *obj;
  double d1;
  double d2;
  event obj->setData(d1,d2);
  cond dep d1 < d2;
  action obj->setMinData();
};

rule RF2{
  decl ;
  event ABSTRACT(RF2);
  cond def FALSE;
  action ;
};
rule RF5_1{
  decl ;
  event ABSTRACT(RF5);
  cond FALSE;
  action ;
  prio 1;
};
rule EPP1{
  decl EPP1 *obj;
  double d1;
  double d2;
  event obj->setData(d1,d2);
  cond d1 < d2;
  action obj->setMinData();
};
rule GC1{
  decl ;
  event 1000 times ABSTRACT(GC1) in
    BOT('GC1') ,
    COMMIT('GC1')
  cond FALSE;
  action ;
};

rule RF3{
  decl ;
  event ABSTRACT(RF3);
  cond dep FALSE;
  action ;
};
rule RF5_2{
  decl ;
  event ABSTRACT(RF5);
  cond FALSE;
  action ;
  prio 2;
};
rule EPP2{
  decl EPP2 *obj;
  double d1;
  double d2;
  event obj->setData(d1,d2);
  cond def d1 < d2;
  action obj->setMinData();
};

```

Figure 3: The OBJECTIVE Benchmark rules

to exclude any database overhead<sup>6</sup>.

The following general order of execution is used for the implementation of each operation:

1. clear the system buffer,
2. open the database,
3. perform cold and hot time measurements, and
4. close the database.

The environment in which we benchmark REACH is a SUN-SPARC 10/512 with 112 MB of RAM under Solaris 2.5, Open OODB 0.2.1, and the EXODUS Storage Manager<sup>7</sup> 2.2. Each operation has been run about 50 times for the same setting of database parameters. Table 3 depicts the mean values of the REACH results. All results along with 90% confidence intervals and standard deviations are presented in Appendix B.

<sup>6</sup>Only exceptions are the operations that require the passing of objects as event parameters in non-immediate modes. In such a case, it only makes sense to pass persistent objects, not transient ones.

<sup>7</sup>Open OODB uses EXODUS as its storage manager.

Parameter	Empty	Small	Medium	Large
<i>NumEvents</i>	0	100	500	1000
<i>FracCompEvents</i>		0.3,0.6,0.9	0.3,0.6,0.9	0.3,0.6,0.9
<i>NumRules</i>	0	100	500	1000
<i>NumObjects</i>	0	5000	25000	50000

Table 2: The OBJECTIVE database configurations

#### 4.3.1 Method Wrapping

The purpose of operation MW1 (Figure 4) is to assess the cost of the (useless) overhead generated by the invocation of a method which is wrapped to provide active functionality whenever required. In operation MW1, a method is invoked which does not generate any event.

The useless overhead paid by REACH is quite acceptable, because REACH minimizes this overhead by assigning a global variable to each method indicating the presence/absence of a detector for that method in the database; thereby reducing it to a memory look-up rather than a database access. Nevertheless, the database must be scanned for the relevant event detectors and the corresponding variables must be set in the memory before the start of an application program.

#### 4.3.2 Event Detection

The primitive event detection operations (Figure 5) examine how efficiently an ADBMS detects primitive events of interest. The aim of operation PED1 is to measure the time it takes to detect a method event. We invoke a method which generates a primitive event which is not an event type for any rule. By this way, we try to discard the time for rule execution, and concentrate on event detection only. Operation PED2 tries to measure the time it takes to detect a transaction event. Unfortunately, in any transaction operation the underlying system does certain bookkeeping operations which is not interesting to us. We chose the BOT operation since it seems to contain minimum uninteresting operations when compared with the other transaction operations. This transaction operation generates an event which does not trigger any rule. On the other hand, operation PED3 considers the COMMIT operation which is, we believe, the most representative of all the transaction operations. The primary focus of this operation is, unlike that of PED2, not only on the detection of a transaction event, but also on getting an insight about the influence of the support for some active functionalities (e.g., event history management).

The composite event detection operations (Figure 6) examine the event composition of an ADBMS. In order to concentrate on composition costs only, we used minimum (meaningful) number of component primitive events for testing different composition types. It would be just as easy to use a larger number of component events, but then it would be very hard to justify a particular number, and more importantly there would be a relatively high risk that the composition costs be overshadowed. To stress the composition costs even more, abstract events are used as component events to exclude event detection and parameter passing time. As in the case of primitive event detection operations,

---

<sup>8</sup>The subscripts <sub>c</sub> and <sub>h</sub> represent cold and hot time results, respectively.

TEST	CONFIGURATION									
	EMPTY	SMALL			MEDIUM			LARGE		
		0.3	0.6	0.9	0.3	0.6	0.9	0.3	0.6	0.9
<i>MW1<sub>h</sub></i>	0.03	0.04	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.03
<i>PED1<sub>h</sub></i>	2.04	2.19	2.27	2.15	2.31	2.57	3.07	3.50	3.80	3.70
<i>PED2<sub>c</sub></i>	12.72	13.79	14.67	14.19	13.37	13.22	15.03	16.97	17.70	17.37
<i>PED3<sub>c</sub></i>	318	1005	1062	5447	10069	20921	42758	35436	46321	74865
<i>CED1<sub>h</sub></i>	3.50	3.72	3.77	3.76	4.01	3.82	4.45	5.42	5.49	5.35
<i>CED2<sub>h</sub></i>	4.16	4.31	4.30	4.37	4.48	4.51	5.31	6.43	6.80	6.39
<i>CED3<sub>h</sub></i>	3.60	3.68	3.69	3.56	3.97	3.91	4.53	5.52	5.81	5.50
<i>CED4<sub>h</sub></i>	4.69	4.86	4.84	4.84	5.17	5.21	6.15	7.47	7.50	7.49
<i>CED5<sub>h</sub></i>	4.73	4.87	4.88	4.79	5.11	5.12	6.13	7.02	7.36	7.68
<i>RF1<sub>c</sub></i>	10.58	12.21	12.38	12.79	13.37	13.77	14.64	16.48	16.54	16.84
<i>RF2<sub>h</sub></i>	1.68	1.92	1.94	1.92	2.48	2.47	2.60	3.31	3.20	3.26
<i>RF3<sub>h</sub></i>	2.38	2.61	2.65	2.66	2.54	2.75	3.08	3.95	3.71	4.26
<i>RF4<sub>h</sub></i>	1.50	2.04	2.02	1.91	2.11	2.53	2.70	2.48	2.59	2.53
<i>RF5<sub>h</sub></i>	1.46	2.44	2.44	2.33	2.21	2.71	3.17	3.03	3.84	3.48
<i>RF6<sub>h</sub></i>	2.40	3.02	3.04	2.96	3.28	3.84	4.05	4.09	4.58	4.37
<i>EPP1<sub>h</sub></i>	2.12	2.86	2.84	2.75	2.89	3.05	3.58	3.81	3.86	3.78
<i>EPP2<sub>h</sub></i>	2.84	3.07	3.05	2.96	3.44	3.73	4.06	5.16	5.18	5.90
<i>EPP3<sub>h</sub></i>	3.40	3.44	3.84	3.57	3.66	3.96	4.53	5.81	5.93	6.32
<i>GC1<sub>c</sub></i>	19712	19423	19785	26483	18981	26010	48674	102149	171610	272112
<i>RA1<sub>c</sub></i>	4.48	4.53	4.60	4.57	4.63	4.71	5.40	7.64	7.92	8.85
<i>RA2<sub>c</sub></i>	2.18	2.13	2.25	2.27	2.22	2.21	2.42	2.23	2.34	3.71
<i>RA3<sub>c</sub></i>	2.07	2.06	2.17	2.08	2.17	2.16	2.40	2.52	2.39	2.55
<i>RA4<sub>c</sub></i>	2.22	2.14	2.22	2.07	2.07	2.48	2.58	2.46	2.51	2.66

Table 3: The OBJECTIVE results<sup>8</sup>for REACH (in milliseconds)

the composite event detection operations generate composite events which are not event types for any rules. It is important to note here that, in all the event detection operations, there is also an overhead for looking up rules to be fired. The primitive and composite event types<sup>9</sup> relevant to event detection operations are defined in Figure 2.

REACH optimizes useful overhead of method event detection as well as useless overhead. This is accomplished by using a prefetching mechanism. This mechanism, by examining the relevant application programs and header files, prefetches the necessary primitive method event detectors, composite event detectors containing those primitive event types as constituents, and the rules to be triggered by the occurrences of these event types. However, this prefetching is done only for method event types, not for transaction or abstract events. This explains why the PED2 results are worse

<sup>9</sup>Since REACH does not allow the creation of events without any associated rules through its rule definition interface, we had to create these events manually. REACH takes a different approach in this respect, because almost all object-oriented ADBMSs (e.g., SAMOS, NAOS, SENTINEL, and ACOOD) encourage the stand-alone definition of events for reusability reasons.

than the PED1 results. The comparison of results of operation PED2 and those of operation PED3 reveals that the COMMIT operation itself, not its detection, shows very poor performance. A more thorough investigation leads us to the fact that this behavior is primarily a consequence of REACH's poor event history<sup>10</sup> maintenance; i.e., at commit time REACH updates the event history with the events occurred in that transaction. However, it is also evident (from the dependency of the results on database configuration) that this update is implemented inefficiently. In addition, it can be inferred from the large standard deviations of PED3 results (see Table 7, Table 8, and Table 9) that, duration of the event history maintenance task (thus duration of the commit operation) depends on the size of the event history which increases at each run of the benchmark operations. Another contributing factor is the underlying platform, Open OODB, which always writes back the whole buffer at commit time.

Results for the composite event detection operations show almost no dependency on database configuration. This is a direct consequence of the use of extended syntactic trees for event composition. For each composite event type, a specialized event detector object is constructed; hence, the overhead of using more generic models (e.g., Petri Nets) is eliminated; making the event composition process very fast. The results for operations CED1 and CED2 are slightly better, as they do not require the confirmation of a validity interval as is done in operations CED3, CED4, and CED5. In general, composite event detection process scales very well; even the most crucial parameters for this test, *NumEvents* and *FracCompEvents*, do not have a notable effect on the results.

### 4.3.3 Rule Firing

The rule firing operations (Figure 7) of the OBJECTIVE Benchmark focus on different aspects of rule firing in an ADBMS. Operation RF1 measures the cost of fetching a rule from the rulebase by triggering a rule in immediate coupling mode. In order to keep the elapsed time for rule execution minimal, which is not interesting to us in this operation, the triggered rule has a FALSE condition part, so that condition evaluation is relatively cheap, and no action is executed. Operations RF2 and RF3 trigger rules in deferred and decoupled coupling modes, respectively. These operations do not measure the time to fire and execute rules in different coupling modes; rather, they examine the cost of storing the information that the triggered rule will be fired just before commit, and in a new transaction, respectively. Although the task measured by RF3 is similar to that measured by RF2 (i.e., abstract event signalling and notification of the current transaction to store a particular bit of information), the contribution of operation RF3 is mainly with respect to functionality (i.e., is decoupled mode supported?). The focus of operation RF4 is on determining how efficiently a rule's condition/action parts are accessed and executed (or interpreted). This operation triggers a rule with a TRUE condition part, so that its action part (though empty) is executed. Operation RF5 reveals the overhead when an event occurs and two rules have to be fired. Different priorities are assigned to these rules to force a particular serialization order. Operation RF6 invokes a method event which triggers a rule that generates the same event in its action part. Therefore the same rule is triggered a second time, but with a condition which evaluates to FALSE; stopping this cascading rule firing. The rules which are triggered by the rule firing operations are defined in Figure 3.

As REACH treats rules as first-class objects, rules are fetched just like ordinary objects by using their names. The (cold time) results for operation RF1 suggest a dependency of rule retrieval time

---

<sup>10</sup>event history is the log of all event occurrences since system startup.



on database configuration. It is important to emphasize here that cold times make sense for this operation as no prefetching mechanism is used for abstract events. Results for operations RF2 and RF3 show that it is slower to initialize the triggering of a rule in decoupled mode than to initialize it in deferred mode. Such a behavior is not surprising at all, since operation RF3 contains the initialization of a new transaction to execute the rule. The results for operation RF4 indicate mainly the time for accessing and executing the action part of the rule. These results are almost constant for all database configurations, because the condition and action parts of a rule are stored as compiled code in shared library allowing very fast access and execution independent of database parameters. The figures for operation RF6 are slightly worse than those for operation RF5. Although both operations contain two rule triggerings, RF6 generates two method event occurrences, whereas in RF5 rules are triggered by a single abstract event.

#### 4.3.4 Event Parameter Passing

The event parameter passing operations (Figure 8) test how efficiently an ADBMS passes event parameters to the condition and action parts of the rules in different coupling modes. The operation EPP1 measures the cost of parameter passing as well as rule execution in immediate coupling mode, whereas the operations EPP2 and EPP3 measure just the cost of using an intermediate storage for passing event parameters. From the point of view of the triggered rules, there is a similar overhead due to the retrieval of the event parameters from the storage where they reside temporarily; but this overhead is not measured by our operations. The rules triggered by the event parameter passing operations are defined in Figure 3.

REACH supports the ability to pass all arguments of a method invocation that triggers a rule to condition and action parts of that rule. In immediate mode all arguments are stored in a bag (i.e., bytestring) and access to the arguments is accomplished by using an array of pointers that store addresses of the arguments. The same mechanism is used in deferred mode, but the dereferenced value of a pointer argument is stored in the array instead of the pointer itself. In detached mode, as the execution of the rule will take place in a different address space, the bag and the pointer array are written in a file. Different requirements for the implementation of these approaches show their effects in the results, making parameter passing somewhat expensive in detached mode due to the inevitable use of an intermediate secondary storage.

#### 4.3.5 Garbage Collection of Semi-Composed Events

The purpose of operation GC1 (Figure 9) is to examine the overhead of flushing an event composition structure that is used in the detection of a composite event. In this operation, we first produce garbage (i.e., create a semi-composed event), and then try to measure the time for collecting the garbage. Such a garbage collection can typically be accomplished at two different points (from a black-box point of view):

- immediately after the monitoring interval is finished, or
- at commit time.

In the former case, the time for the operation generating the end-of-interval event, and in the latter case, the time for commit operation should be measured. For generality of the test, we take COMMIT

to be also the end-of-interval event so that garbage collection can only be accomplished during commit for this operation. Unfortunately, isolation of garbage collection inside commit is not possible by using the results of this operation only. However, we can circumvent this problem to a certain extent by using the difference of the results of this operation and those of operation PED3 (i.e., detection of COMMIT) in which no time for garbage collection is involved. In this manner, we may have an *estimation* of the times indicating the duration of the garbage collection task, which is the best we can do with our black-box view of the system.

As in the case of operation PED3, we encounter very poor results for operation GC1. It is argued above that results of operation PED3 be used in the interpretation of the results of GC1. Unfortunately, it is out of question to get an understanding of the performance of the system under the intended task even by using results of PED3. The reason is that, as mentioned in Section 4.3.2, commit time is dependent on the size of the event history in REACH, and the size of the event history is not the same in respective runs of operation PED3 and operation GC1; making it impossible to interpolate the time for garbage collection by using the results of these two operations.

#### 4.3.6 Rule Administration

The rule administration operations (Figure 10) are somewhat different from the other OBJECTIVE operations in the sense that they are more likely to be included in a *feature* benchmark. However, we deem the functionalities examined by these operations so important from the functionality point of view that they must be included in a comprehensive benchmark for ADBMSs.

Operation RA1 creates a new rule and stores it in the rulebase, and operation RA2 deletes an existing rule from the rulebase. Operation RA3 and RA4, enables and disables a rule, respectively. Operation RA5 changes the action part of a rule. In all these operations, the relevant rules are kept very simple in order to focus on the efficiency of the provided rule administration facility.

All of the rule administration operations are implemented using the rule management commands of REACH from its command line interface [ZBB<sup>+</sup>96]. The implementation of operation RA1 in REACH consists of the creation of a rule and compilation of the shared library containing the condition/action parts of rules in the form of two C functions by using the REACH command `rl_cc`. The other operations, RA2, RA3, and RA4, are implemented using REACH commands `r_delete`, `r_enable`, and `r_disable`, respectively. Unfortunately, we were not able to get results for operation RA5 (although it is possible modify rules dynamically in REACH) because of a bug in the system. The results for the presented rule administration operations, except RA1, show a constant behavior under all database configurations. The exceptional results for operation RA1 are possibly due to the compilation time of the shared library whose size is directly proportional to the number of rules.

## 5 Conclusions and Future Work

We presented the OBJECTIVE Benchmark for object-oriented ADBMSs, and illustrated it with the results obtained from its implementation in REACH. Although OBJECTIVE is designed to be very simple in nature, it is also very comprehensive in its coverage of active functionalities.

The results obtained from the implementation of OBJECTIVE on REACH reveal that REACH supports a high level of active functionality efficiently. Almost all components of REACH perform and scale well. The only exception we encountered is the problematic commit operation of REACH.

This operation is a real bottleneck as it is a *must* operation for all applications running inside a transaction framework, and this bottleneck must be surmounted to achieve acceptable overall system performance. The implementation phase also helped to disclose a number of bugs in the system. The results of REACH alone are sufficient to identify its bottleneck components. However, results to be taken from different systems (with possibly different approaches and architectures for supporting ADBMS tasks) would be highly welcome to make an objective judgment about the degree of efficiency with which these tasks are supported by a particular ADBMS.

We believe that the OBJECTIVE operations cover an important subset of issues with respect to ADBMS performance and functionality. The remaining issues are mainly the ones related to event consumption policies, condition optimization, and parallel rule execution.

An open related research area is the evaluation of ADBMS performance in multi-user environments. There is considerable performance difference between single-user and multi-user environments which results from issues of optimal system resource utilization. Therefore, the results obtained from a single-user benchmark do not necessarily represent the real performance of the system. It is especially interesting to investigate the effects of the number of concurrently running transactions to event detection and rule execution.

An interesting thing to note here is that all the benchmarks that have been proposed so far for ADBMSs, including OBJECTIVE, are generic benchmarks. This is, we think, a consequence of the lack of adequate information about the characteristics of ADBMS tasks (even the notion of an ADBMS task is elusive for now). As the application areas for ADBMSs mature, we expect to see the development of domain-specific benchmarks to evaluate end-to-end performance in order to have a better understanding of ADBMS performance.

As a final remark, we hope that the OBJECTIVE Benchmark finds acceptance as a useful yardstick for evaluating ADBMS performance and functionality.

## References

- [AG89] R. Agrawal and N. H. Gehani. ODE (object database and environment): The language and the data model. In *Proc. 1989 ACM-SIGMOD Conference on Management of Data*, Portland, Oregon, June 1989.
- [BD84] H. Boral and D. J. DeWitt. A methodology for database system performance evaluation. In *Proc. of the 1984 SIGMOD Conference*, pages 176–185, Boston, June 1984.
- [Ber91] Mikael Berndtsson. ACOOD: An approach to an active object oriented DBMS. Master’s thesis, University of Skövde, Department of Computer Science, Skövde, Sweden, September 1991.
- [BM93] D. A. Brant and D. P. Miranker. Index support for rule activation. In *ACM SIGMOD Conference on Management of Data*, pages 42–48, Washington D.C., May 1993.
- [BZBW95] A. P. Buchmann, J. Zimmermann, J. Blakeley, and D. L. Wells. Building an integrated active OODBMS: Requirements, architecture and design decisions. In *Proceedings of Data Engineering Conference*, pages 117–128, 1995.

- [CCS94] C. Collet, T. Coupaye, and T. Svensen. NAOS: Efficient and modular reactive capabilities in an object-oriented database system. In *Proceedings of the 20th International Conference on Very Large Databases*, Santiago, Chile, September 1994.
- [CDN93] M. J. Carey, D. J. DeWitt, and J. F. Naughton. The OO7 benchmark. In *ACM SIGMOD Conference*, pages 12–21, May 1993.
- [CKAK93] S. Chakravarthy, V. Krishnaprasad, E. Abwar, and S. K. Kim. Anatomy of a composite event detector. Technical Report UF-CIS-TR-93-039, CIS Department, University of Florida, December 1993.
- [CM93] S. Chakravarthy and D. Mishra. SNOOP: an expressive event specification language for active databases. Technical Report UF-CIS-TR-93-007, CIS Department, University of Florida, March 1993.
- [CS92] R. Catell and J. Skeen. Object operations benchmark. *ACM Transactions on Database Systems*, 17(1):1–31, March 1992.
- [CVKB94] S. Chakravarthy, Z. Tamizuddin V. Krishnaprasad, and R. H. Badani. ECA rule integration into an OODBMS: Architecture and implementation. Technical Report 94-023, CIS Department, University of Florida, 1994.
- [Day88] U. Dayal. Active database management systems. In *Proc. 3rd International Conference on Data and Knowledge Bases*, pages 150–169, Jerusalem, June 1988.
- [DBB<sup>+</sup>88] U. Dayal, B. Blaustein, A. Buchmann, U. Chakravarthy, M. Hsu, R. Ladin, D. McCarthy, and A. Rosenthal. The HiPAC project: Combining active databases and timing constraints. *ACM SIGMOD Record*, 17(1):51–70, March 1988.
- [DBM88] U. Dayal, A. Buchmann, and D. McCarthy. Rules are objects too: a knowledge model for an active, object-oriented database system. In *Proceedings of 2nd International Workshop on Object-Oriented Database Systems, Lecture Notes in Computer Science 334*. Springer, 1988.
- [Deu94] Alin Deutsch. Detection of method and composite events in the active DBMS REACH. Master’s thesis, Technical University Darmstadt, July 1994.
- [Eri93] J. Eriksson. Cede: Composite event detector in an active object-oriented database. Master’s thesis, Department of Computer Science, University of Skövde, 1993.
- [GBLZ96] A. Geppert, M. Berndtsson, D. Lieuwen, and J. Zimmermann. Performance evaluation of active database management systems using the BEAST benchmark. Technical Report 96.01, CS Department, University of Zurich, February 1996.
- [GD93] S. Gatzia and K. R. Dittrich. Events in an active object-oriented database system. In *Proc. of the 1st International Workshop on Rules in Database Systems*. Springer, September 1993.

- [GD94] S. Gatziau and K. R. Dittrich. Detecting composite events in active database systems using Petri nets. In *IEEE RIDE Proc. 4th International Workshop on Research Issues in Data Engineering*, Houston, Texas, USA, 1994.
- [GGD94] S. Gatziau, A. Geppert, and K. R. Dittrich. The SAMOS active DBMS prototype. Technical Report 94.16, CS Department, University of Zurich, 1994.
- [GGD95] A. Geppert, S. Gatziau, and K. R. Dittrich. A designers benchmark for active database management systems: OO7 meets the BEAST. Technical Report 95.18, Dept. of Computer Science, University of Zurich, 1995.
- [GJS92] N. Gehani, H. V. Jagadish, and O. Shumeli. Composite event specification in active databases: Model and implementation. In *Proc. 18th International Conference on Very Large Data Bases*, Vancouver, Canada, August 1992.
- [Gra91] Jim Gray. *The Benchmark Handbook for Database and Transaction Processing*. Morgan Kaufmann, 1991.
- [HW92] E. Hanson and J. Widom. An overview of production rules in database systems. Technical Report 92-031, CIS Department, University of Florida, 1992.
- [Ker95] M. L. Kersten. An active component for a parallel database kernel. In *Rules in Database Systems*, Workshops in Computing, pages 277–291. Springer, September 1995.
- [Mos85] J. E. B. Moss. *Nested Transactions: An Approach to Reliable Distributed Computing*. MIT Press, 1985.
- [WBT92] D. L. Wells, J. A. Blakeley, and C. W. Thompson. Architecture of an open object-oriented database management system. *IEEE Computer*, 25(10):74–81, October 1992.
- [ZB95] J. Zimmermann and A. P. Buchmann. Benchmarking active database systems: A requirements analysis. In *OOPSLA '95 Workshop on Object Database Behavior, Benchmarks, and Performance*, Austin, Texas, 1995.
- [ZBB<sup>+</sup>96] J. Zimmermann, H. Branding, A. P. Buchmann, A. Deutsch, and A. Geppert. Design, implementation and management of rules in an active database system. In *Proceedings of Database and Expert System Application*, 1996. To be published.
- [ZBD95] J. Zimmermann, A. Buchmann, and A. Deutsch. The ACT-1 benchmark. Technical report, Technical University Darmstadt, Germany, 1995.

## A Illustration of the OBJECTIVE Operations

We do not include every step of the implementation in our codes which illustrate the benchmark operations; as they may vary widely from system to system, and will not contribute much to the understanding of the benchmark. In particular, we skipped the code for handling the database, flushing the system buffer and measuring the elapsed time. For ease of illustration, all the tests of a particular group are shown in the same main program, although each is implemented as independent operations comprising all the steps outlined in Section 4.3. Each block of statements written in bold font indicates the interesting parts of a particular test and is wrapped around by the code for time measurement. In the implementation, each such block is executed eleven times; we take the cold time to be the time elapsed for the first iteration and the hot time to be average elapsed time for the last ten iterations.

```
void main(){
    MW1 *objectMW1=new MW1;

    Transaction::begin();
    objectMW1->doNothing();
    \\ no event generation
    Transaction::commit(); }
```

Figure 4: The Method Wrapping program

```
void main(){
    PED1 *objectPED1=new PED1;

    Transaction::begin();
    objectPED1->doNothing(); \\ generate event PED1
    Transaction::commit();

    \\ generate event PED2
    Transaction::begin('PED2');
    Transaction::commit();
    Transaction::begin('PED3');
    \\ generate event PED3
    Transaction::commit(); }
```

Figure 5: The Primitive Event Detection program

```

void main(){
    Transaction::begin();

    AbstractEvent::raise('CED1_1');
    AbstractEvent::raise('CED1_2');
    AbstractEvent::raise('CED2_1');
    AbstractEvent::raise('CED2_2');
    AbstractEvent::raise('CED3_1');
    AbstractEvent::raise('CED3_3');
    AbstractEvent::raise('CED4_1');
    AbstractEvent::raise('CED4_2');
    AbstractEvent::raise('CED4_3');
    AbstractEvent::raise('CED5_1');
    AbstractEvent::raise('CED5_2');
    AbstractEvent::raise('CED5_3');
    Transaction::commit(); }

```

Figure 6: The Composite Event Detection program

```

void main(){
    RF6 *objectRF6=new RF6;

    Transaction::begin();
                                \\ trigger rule RF1
    AbstractEvent::raise('RF1');
                                \\ trigger rule RF2
    AbstractEvent::raise('RF2');
                                \\ trigger rule RF3
    AbstractEvent::raise('RF3');
                                \\ trigger rule RF4
    AbstractEvent::raise('RF4');
                                \\ trigger rules RF5_1
                                \\ and RF5_2
    AbstractEvent::raise('RF5');
                                \\ trigger rule RF6 twice
    objectRF6->setAttribute(1);
    Transaction::commit(); }

```

Figure 7: The Rule Firing program

```

void main(){
    EPP1 *objectEPP1=new EPP1;
    EPP2 *objectEPP2=new EPP2;
    EPP3 *objectEPP3=new EPP3;

    Transaction::begin();
                                \\ trigger rule EPP1
    objectEPP1->setData(1.0, 2.0);
                                \\ trigger rule EPP2
    objectEPP2->setData(1.0, 2.0);
                                \\ trigger rule EPP3
    objectEPP3->setData(1.0, 2.0);
    Transaction::commit(); }

```

Figure 8: The Event Parameter Passing program



```

void main(){
    Transaction::begin('GC1');
                                \\ create a semi-composed event
    for(int i=0; i < 999; i++)
        AbstractEvent::raise('GC1');
    Transaction::commit(); }

```

Figure 9: The Garbage Collection program

```

void main(){
                                \\ create rule RA1 and compile shared library
    system('createRule.RA1');
    system('rl_cc');
                                \\ delete rule RA2
    system('r_rm /RA2');
                                \\ enable rule RA2
    system('r_enable /RA3');
                                \\ disable rule RA2
    system('r_disable /RA4'); }

```

Figure 10: The Rule Administration program

## B The Complete Results for REACH

TEST		CONFIGURATION			
		EMPTY	SMALL		
			0.3	0.6	0.9
$MW1_h$	average	0.03	0.04	0.03	0.04
	std. dev.	0.01	0.02	0.01	0.02
	conf. int.	[0.02,0.05]	[0.02,0.06]	[0.02,0.05]	[0.03,0.06]

Table 4: The Method Wrapping results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$MW1_h$	average	0.03	0.03	0.03
	std. dev.	0.01	0.02	0.01
	conf. int.	[0.02,0.04]	[0.02,0.05]	[0.02,0.04]

Table 5: The Method Wrapping results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$MW1_h$	average	0.03	0.03	0.03
	std. dev.	0.01	0.02	0.01
	conf. int.	[0.03,0.05]	[0.02,0.05]	[0.03,0.04]

Table 6: The Method Wrapping results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			0.3	0.6	0.9
$PED1_h$	average	2.04	2.19	2.27	2.15
	std. dev.	0.05	0.14	0.05	0.06
	conf. int.	[1.97,2.15]	[1.87,2.40]	[2.19,2.38]	[2.08,2.27]
$PED2_c$	average	12.72	13.79	14.67	14.19
	std. dev.	0.7	1.01	1.18	1.07
	conf. int.	[11.67,13.67]	[11.32,15.44]	[12.75,17.37]	[12.12,15.31]
$PED3_c$	average	318	1005	1062	5447
	std. dev.	4	203	1907	13219
	conf. int.	[313,325]	[333,1097]	[1036,1093]	[993,45106]

Table 7: The Primitive Event Detection results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$PED1_h$	average	2.31	2.57	3.07
	std. dev.	0.87	0.16	0.10
	conf. int.	[0.37,3.31]	[2.44,2.63]	[2.98,3.24]
$PED2_c$	average	13.37	13.22	15.03
	std. dev.	1.81	0.75	1.17
	conf. int.	[11.17,15.98]	[12.02,13.85]	[13.75,17.34]
$PED3_c$	average	10069	20921	42758
	std. dev.	1268	23190	70541
	conf. int.	[329,23601]	[4484,70043]	[4615,207086]

Table 8: The Primitive Event Detection results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$PED1_h$	average	3.50	3.80	3.70
	std. dev.	0.20	0.03	0.12
	conf. int.	[3.32,3.87]	[3.77,3.84]	[3.61,3.99]
$PED2_c$	average	16.97	17.70	17.37
	std. dev.	0.82	0.65	0.58
	conf. int.	[15.88,18.35]	[17.28,19.00]	[16.71,18.52]
$PED3_c$	average	35436	46321	74865
	std. dev.	32020	16039	77360
	conf. int.	[10572,97230]	[30302,73728]	[9013,222630]

Table 9: The Primitive Event Detection results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			<b>0.3</b>	<b>0.6</b>	<b>0.9</b>
$CED1_h$	average	3.50	3.72	3.77	3.56
	std. dev.	0.10	0.21	0.21	0.12
	conf. int.	[3.35,3.73]	[3.51,4.17]	[3.65,4.40]	[3.40,3.82]
$CED2_h$	average	4.16	4.31	4.30	4.37
	std. dev.	0.14	0.18	0.09	0.27
	conf. int.	[3.99,4.51]	[4.08,4.71]	[4.17,4.46]	[4.11,5.05]
$CED3_h$	average	3.60	3.68	3.69	3.56
	std. dev.	0.10	0.12	0.14	0.11
	conf. int.	[3.46,3.76]	[3.61,3.80]	[3.50,3.95]	[3.38,3.82]
$CED4_h$	average	4.69	4.86	4.84	4.84
	std. dev.	0.20	0.10	0.09	0.10
	conf. int.	[4.50,5.24]	[4.62,4.99]	[4.74,4.99]	[4.70,5.03]
$CED5_h$	average	4.73	4.87	4.88	4.79
	std. dev.	0.10	0.15	0.19	0.20
	conf. int.	[4.59,4.94]	[4.62,5.24]	[4.59,5.25]	[4.56,5.33]

Table 10: The Composite Event Detection results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$CED1_h$	average	4.01	3.82	4.45
	std. dev.	0.22	0.16	0.14
	conf. int.	[3.70,4.23]	[3.77,3.94]	[4.25,4.67]
$CED2_h$	average	4.48	4.51	5.31
	std. dev.	0.44	0.15	0.11
	conf. int.	[3.95,5.24]	[4.44,4.58]	[5.14,5.44]
$CED3_h$	average	3.97	3.91	4.53
	std. dev.	0.36	0.12	0.15
	conf. int.	[3.39,4.24]	[3.75,4.08]	[4.33,4.71]
$CED4_h$	average	5.17	5.21	6.15
	std. dev.	0.45	0.16	0.14
	conf. int.	[4.52,5.62]	[5.04,5.51]	[5.96,6.25]
$CED5_h$	average	5.11	5.12	6.13
	std. dev.	0.43	0.13	0.10
	conf. int.	[4.51,5.49]	[5.02,5.40]	[5.98,6.25]

Table 11: The Composite Event Detection results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$CED1_h$	average	5.42	5.49	5.35
	std. dev.	0.32	0.16	0.11
	conf. int.	[5.00,5.97]	[5.22,5.71]	[5.22,5.49]
$CED2_h$	average	6.43	6.80	6.39
	std. dev.	0.21	0.39	0.39
	conf. int.	[6.16,6.71]	[6.08,7.28]	[6.02,7.08]
$CED3_h$	average	5.52	5.81	5.50
	std. dev.	0.19	0.36	0.26
	conf. int.	[5.36,5.88]	[5.12,6.11]	[5.26,5.92]
$CED4_h$	average	7.47	7.50	7.49
	std. dev.	0.44	0.42	0.24
	conf. int.	[6.84,8.05]	[6.75,7.99]	[7.19,7.81]
$CED5_h$	average	7.02	7.36	7.68
	std. dev.	0.31	0.25	0.2
	conf. int.	[6.61,7.42]	[6.98,7.74]	[7.38,8.07]

Table 12: The Composite Event Detection results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			0.3	0.6	0.9
$RF1_c$	average	10.58	12.21	12.38	12.79
	std. dev.	0.24	0.10	0.38	0.64
	conf. int.	[10.24,10.94]	[10.35,14.24]	[11.80,12.96]	[11.72,14.02]
$RF2_h$	average	1.68	1.92	1.94	1.92
	std. dev.	0.01	0.03	0.01	0.01
	conf. int.	[1.63,1.72]	[1.82,1.94]	[1.91,1.98]	[1.91,1.99]
$RF3_h$	average	2.38	2.61	2.65	2.66
	std. dev.	0.15	0.28	0.18	0.22
	conf. int.	[2.00,2.50]	[2.01,2.89]	[2.16,2.82]	[2.11,2.88]
$RF4_h$	average	1.50	2.04	2.02	1.91
	std. dev.	0.04	0.24	0.08	0.02
	conf. int.	[1.46,1.55]	[1.48,2.56]	[1.96,2.10]	[1.89,1.95]
$RF5_h$	average	1.46	2.44	2.44	2.33
	std. dev.	0.05	0.32	0.07	0.15
	conf. int.	[1.39,1.54]	[2.01,2.97]	[2.33,2.55]	[2.21,2.78]
$RF6_h$	average	2.40	3.02	3.04	2.96
	std. dev.	0.03	0.20	0.04	0.19
	conf. int.	[2.34,2.44]	[2.39,3.21]	[2.97,3.12]	[2.82,3.51]

Table 13: The Rule Firing results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$RF1_c$	average	13.37	13.77	14.64
	std. dev.	1.12	0.70	0.39
	conf. int.	[11.67,14.55]	[13.24,15.18]	[14.11,15.06]
$RF2_h$	average	2.48	2.47	2.60
	std. dev.	0.09	0.08	0.04
	conf. int.	[2.41,2.62]	[2.33,2.54]	[2.52,2.64]
$RF3_h$	average	2.54	2.75	3.08
	std. dev.	0.31	0.05	0.10
	conf. int.	[2.42,2.68]	[2.72,2.78]	[3.01,3.20]
$RF4_h$	average	2.11	2.53	2.70
	std. dev.	0.11	0.18	0.05
	conf. int.	[2.01,2.22]	[2.31,2.76]	[2.68,2.74]
$RF5_h$	average	2.21	2.71	3.17
	std. dev.	0.63	0.03	0.02
	conf. int.	[1.45,2.86]	[2.67,2.75]	[3.11,3.31]
$RF6_h$	average	3.28	3.84	4.05
	std. dev.	0.54	0.04	0.08
	conf. int.	[2.43,3.79]	[3.81,3.86]	[3.97,4.21]

Table 14: The Rule Firing results (MEDIUM Database Configuration)



TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$RF1_c$	average	16.48	16.54	16.84
	std. dev.	0.46	0.75	0.35
	conf. int.	[15.86,16.95]	[15.51,17.53]	[16.59,17.44]
$RF2_h$	average	3.31	3.20	3.26
	std. dev.	0.18	0.09	0.16
	conf. int.	[3.15,3.66]	[3.10,3.30]	[3.08,3.47]
$RF3_h$	average	3.95	3.71	4.26
	std. dev.	0.16	0.02	0.07
	conf. int.	[3.82,4.19]	[3.69,3.73]	[4.12,4.31]
$RF4_h$	average	2.48	2.59	2.53
	std. dev.	0.07	0.12	0.18
	conf. int.	[2.41,2.59]	[2.47,2.70]	[2.32,2.78]
$RF5_h$	average	3.03	3.84	3.48
	std. dev.	0.18	0.54	0.06
	conf. int.	[2.87,3.28]	[3.51,4.93]	[3.41,3.55]
$RF6_h$	average	4.09	4.58	4.37
	std. dev.	0.09	0.28	0.01
	conf. int.	[3.99,4.21]	[4.36,5.12]	[4.36,4.38]

Table 15: The Rule Firing results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			<b>0.3</b>	<b>0.6</b>	<b>0.9</b>
$EPP1_h$	average	2.12	2.86	2.84	2.75
	std. dev.	0.14	0.37	0.18	0.05
	conf. int.	[2.03,2.50]	[2.17,3.47]	[2.70,3.35]	[2.69,2.84]
$EPP2_h$	average	2.84	3.07	3.05	2.96
	std. dev.	0.19	0.10	0.07	0.06
	conf. int.	[2.75,3.40]	[2.81,3.17]	[2.95,3.19]	[2.91,3.10]
$EPP3_h$	average	3.40	3.44	3.84	3.57
	std. dev.	0.19	0.02	0.10	0.08
	conf. int.	[3.22,3.93]	[3.38,3.44]	[3.80,3.91]	[3.51,3.61]

Table 16: The Event Parameter Passing results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		<b>0.3</b>	<b>0.6</b>	<b>0.9</b>
$EPP1_h$	average	2.89	3.05	3.58
	std. dev.	0.48	0.05	0.15
	conf. int.	[2.10,3.39]	[3.00,3.13]	[3.49,3.90]
$EPP2_h$	average	3.44	3.73	4.06
	std. dev.	0.42	0.10	0.07
	conf. int.	[2.82,3.87]	[3.65,3.87]	[3.99,4.19]
$EPP3_h$	average	3.66	3.96	4.53
	std. dev.	0.37	0.13	0.20
	conf. int.	[3.33,4.17]	[3.88,4.11]	[4.35,4.67]

Table 17: The Event Parameter Passing results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		<b>0.3</b>	<b>0.6</b>	<b>0.9</b>
$EPP1_h$	average	3.81	3.86	3.78
	std. dev.	0.28	0.06	0.05
	conf. int.	[3.39,4.17]	[3.80,3.95]	[3.73,3.85]
$EPP2_h$	average	5.16	5.18	5.90
	std. dev.	0.14	0.12	0.06
	conf. int.	[5.03,5.43]	[5.10,5.38]	[5.82,5.98]
$EPP3_h$	average	5.81	5.93	6.32
	std. dev.	0.10	0.32	0.17
	conf. int.	[5.66,5.93]	[5.59,6.45]	[6.12,6.49]

Table 18: The Event Parameter Passing results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			0.3	0.6	0.9
$GC1_c$	average	19712	19423	19785	26483
	std. dev.	1780	1238	1032	22577
	conf. int.	[16953,22963]	[17829,21327]	[17926,72951]	[17677,92111]

Table 19: The Garbage Collection results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$GC1_c$	average	18981	26010	48674
	std. dev.	2582	34177	36241
	conf. int.	[16579,24199]	[22206,28280]	[31481,129709]

Table 20: The Garbage Collection results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$GC1_c$	average	102149	171610	272112
	std. dev.	62925	75462	89994
	conf. int.	[43224,169073]	[37115,228039]	[40254,449263]

Table 21: The Garbage Collection results (LARGE Database Configuration)

TEST		CONFIGURATION			
		EMPTY	SMALL		
			0.3	0.6	0.9
$RA1_c$	average	4.48	4.53	4.60	4.57
	std. dev.	1.03	1.22	1.76	0.37
	conf. int.	[3.79,6.88]	[3.83,6.95]	[3.22,7.08]	[4.22,5.04]
$RA2_c$	average	2.18	2.13	2.25	2.27
	std. dev.	0.05	0.07	0.02	0.10
	conf. int.	[2.15,2.29]	[2.02,2.21]	[2.22,2.27]	[2.17,2.42]
$RA3_c$	average	2.07	2.05	2.17	2.08
	std. dev.	0.07	0.06	0.18	0.04
	conf. int.	[1.97,2.19]	[1.99,2.16]	[2.00,2.27]	[2.03,2.16]
$RA4_c$	average	2.24	2.14	2.22	2.07
	std. dev.	0.15	0.12	0.02	0.05
	conf. int.	[2.09,2.50]	[2.03,2.33]	[2.18,2.23]	[2.02,2.15]

Table 22: The Rule Administration results (EMPTY and SMALL Database Configurations)

TEST		CONFIGURATION		
		MEDIUM		
		0.3	0.6	0.9
$RA1_c$	average	4.63	4.71	5.39
	std. dev.	0.33	0.05	0.17
	conf. int.	[4.20,5.01]	[4.67,4.78]	[5.23,5.57]
$RA2_c$	average	2.22	2.21	2.41
	std. dev.	0.13	0.15	0.16
	conf. int.	[2.07,2.39]	[2.01,2.64]	[2.26,2.58]
$RA3_c$	average	2.17	2.16	2.40
	std. dev.	0.04	0.05	0.01
	conf. int.	[2.12,2.21]	[2.10,2.21]	[2.39,2.41]
$RA4_c$	average	2.07	2.48	2.58
	std. dev.	0.10	0.06	0.03
	conf. int.	[1.97,2.20]	[2.41,2.58]	[2.56,2.61]

Table 23: The Rule Administration results (MEDIUM Database Configuration)

TEST		CONFIGURATION		
		LARGE		
		0.3	0.6	0.9
$RA1_c$	average	7.64	7.92	8.85
	std. dev.	0.12	0.08	0.21
	conf. int.	[7.51,7.76]	[7.83,8.02]	[8.61,8.99]
$RA2_c$	average	2.23	2.34	3.71
	std. dev.	0.02	0.13	0.17
	conf. int.	[2.21,2.25]	[2.23,2.45]	[3.52,3.90]
$RA3_c$	average	2.52	2.39	2.55
	std. dev.	0.43	0.20	0.35
	conf. int.	[2.08,2.95]	[2.27,2.56]	[2.21,2.81]
$RA4_c$	average	2.46	2.51	2.66
	std. dev.	0.02	0.22	0.14
	conf. int.	[2.44,2.48]	[2.33,2.69]	[2.42,2.79]

Table 24: The Rule Administration results (LARGE Database Configuration)