

A Framework Based on Design Patterns for Providing Persistence in Object-Oriented Programming Languages

Jörg Kienzle

*Software Engineering Laboratory
Swiss Federal Institute of Technology
CH - 1015 Lausanne Ecublens
Switzerland
email: Joerg.Kienzle@epfl.ch*

Alexander Romanovsky

*Department of Computing Science
University of Newcastle
NE1 7RU, Newcastle upon Tyne
United Kingdom
email: Alexander.Romanovsky@newcastle.ac.uk*

Abstract. This paper describes an approach to providing object persistence in object-oriented programming languages without modifying the run-time system or the language itself. By successively applying design patterns such as the *Serializer*, *Factory Method*, and *Strategy* patterns we develop an object-oriented framework for providing object persistence. The advantages of object-orientation are highlighted: structured classification through class-hierarchies, extensibility and promotion of reuse. The framework clearly separates persistence control from storage control. A hierarchy of different storage types, useful in different application domains, is introduced. The framework does not rely on any kind of special programming language features. It only uses basic object-oriented programming techniques, and is therefore implementable in any object-oriented programming language. An experimental implementation in Ada 95 is presented.

Keywords. Persistence, Stable Storage, Object-Oriented Framework, Design Patterns.

1 Introduction

Research into persistent programming languages and systems in recent years has shown that this technology is useful for developing complex software in many problem domains. Persistence is used whenever data values from a program execution are saved so that they can be used in a later execution. Software fault tolerance mechanisms based on backward error recovery use persistence to provide state restoration in case of computer crashes. Transaction durability is often achieved using persistence techniques. How the data is saved and what kind of storage medium is used for that purpose depends on the applications demands and can vary considerably from one application to another. Unfortunately, widely used object-oriented programming languages still do not offer support for persistence.

This paper describes the design of a framework for providing object persistence in object-oriented programming languages. The outline of the paper is as follows: section 2 introduces persistency and gives a brief overview of some programming languages that address persistence; section 3 shows how parts of the problem can be solved by applying design patterns; section 4 presents the full framework, obtained by putting together the partial solutions obtained in section 3; section 5 shows an experimental implementation of the framework in Ada 95; section 6 looks at other related work in this area, and section 7 draws some conclusions.

2 Persistence and Programming Languages

Persistence is a general term for mechanisms that allow application data from a program's execution space to somehow survive the execution of the program, so that in a later execution it can be used again. There are many possible schemes for supporting persistence. For a complete survey, the reader should refer to [1].

The most sophisticated and desired form of persistence is *orthogonal persistence* [2]. It is the provision of persistence for *all* data irrespective of their type. In a programming language providing orthogonal persistence, persistent data is created and used in the same way as non-persistent data. Loading and saving of values does not alter their semantics, and the process is transparent to the application program. Whether or not data should be made persistent is often determined using a technique called *persistence by reachability*. The persistence support designates an object as a *persistent root* and provides applications with a built-in function for locating it. Any object that is "reachable" from the persistent root, for instance by following pointers, is automatically made persistent.

The first language providing orthogonal persistence, *PS-Algol* [3], was conceived in order to add persistence to an existing language with minimal perturbation to its initial semantics and implementation. There are persistent versions of functional programming languages such as *Persistent Poly* and *Poly ML* [4]. There has also been work on adding orthogonal persistence to widely used programming languages. Probably the most interesting project nowadays is *PJava* [5], a project that aims at providing orthogonal persistence to the *Java* [6] programming language.

Due to the demanding requirements of orthogonal persistence, all these implementations had to slightly modify the programming language and / or modify the run-time system. The paper [7] for instance investigates adding orthogonal persistence to the Ada 95 [8] language. The authors identified the following problems:

- Orthogonal persistence requires that both data and types can have indefinite lifetimes. If a persistent application is to evolve, structural equivalence and dynamic type checking are necessary when a program binds to an object from the persistent store. When introducing orthogonal persistence, type compatibility within an execution extends to type compatibility across different executions. This may conflict with the typing rules of the programming language.
- Often programming languages allow the use of static variables inside classes or even as standalone global variables. It is possible that a programmer uses such static variables to link objects together, such as for instance a static table that links *key* values to some other data. Now if the *key* values are made persistent, the table should also persist, or else the *key* values are useless. It might be tricky to provide automatic persistence for such static variables without breaking orthogonal persistence.
- Orthogonal persistence also requires that elaborate types such as task types / threads and subprogram pointers values persist. This can raise severe implementation problems.
- A program might evolve and change the definition of types and classes, but still try and work with values saved in previous executions. To make this work, some form of version control must be provided, and additional dynamic checking is required. The problem can be even more complicated when considering inheritance.
- Another important problem when providing orthogonal persistence is storage management. Persistent data that will not be used anymore must be deleted, for storage leaks will result in permanent loss of storage capacity. This basically requires some form of automatic garbage collection, at least for all persistent data.

Finally the authors conclude that adding orthogonal persistence to the Ada 95 language would require major changes, making the new language backwards incompatible. It is interesting to note here that even in the case of the Java language, a modern object-oriented language that already provides automatic garbage collection and a powerful reflection mechanism, the virtual machine executing the Java byte code had to be modified in order to support orthogonal persistence [5].

As soon as we do not require orthogonal persistence, persistence support for conventional programming languages can be provided in multiple ways. Many languages have been extended or provide standard libraries that allow data to be made persistent for instance by saving it to disk. Avalon [9] for instance is an extension to C++ that provides persistence and transactions. The authors have extended the C++ language, providing additional keywords such as *stable* used to designate class attributes that are to be made persistent.

Persistence support in object-oriented programming languages must provide a mechanism that allows the state of an object to persist between different executions of an application. It can be quite challenging to find means for taking the in-memory representation of the objects state and writing it to some storage device. Fortunately, object-oriented programming languages often provide some kind of streaming functionality that allows transforming the state of an object into a flat stream of bytes. Some languages go even further and provide streams that allow a user to write objects into files or other storage devices (Ada `Stream_IO` [8 A.12.1], Java [6] `FileOutputStreams`). Unfortunately, the facilities provided by the programming language are not always sufficient, or they lack modularity and extensibility, making the definition of new persistent objects or the addition of new storage devices difficult or even impossible.

We believe that when designing persistence support for object-oriented programming languages one should strive to achieve the following:

- *Clear separation of concerns*: The persistent object should not know about storage devices or about the data format that is used when writing the state of the object onto the storage device and vice versa.
- *Modularity and extensibility*: It should be straightforward to define new persistent objects or add new storage devices.
- *Safe storage management*: Storage leaks leading to wasted space on the storage device must be prevented.

In order to help the programmer we propose a general framework for providing object persistence that maximizes these goals. It can be used by two different types of programmers: *persistence support programmers* and *application programmers*.

Persistence support programmers will add support for new storage devices to the framework using object-oriented programming techniques. *Application programmers* will use the framework to declare persistent objects. When instantiating a new persistent object, the application programmer specifies where the state of the object will be saved by choosing among the existing implementation of storage devices. The operations defined for persistent objects allow the application programmer to save or restore the state of the object at any time.

The framework does not rely on any kind of special programming language features. It only uses basic object-oriented programming techniques and is therefore implementable in any object-oriented programming language. Most of its structure is based on well-known design patterns. The remainder of this paper documents the construction of the framework.

3 Applying Design Patterns Successively

3.1 Classification of Storage Types

At some point, persistent objects must save their state to some kind of storage, so that it can be retrieved again in a later execution. The term storage is used in a wider sense here. Sending the state of the object over a network and storing it in the memory of some other computer would for instance also make sense, as long as the data survives program termination. Nevertheless, all storage types do not have the same properties, and therefore must not all provide the same set of operations. An abstract class hierarchy is the most natural way to represent the structure of such storage types. A concrete implementation of a storage type must derive from one of the storage classes and implement the required operations. We propose to classify the different storage types in the way presented in figure 1.

The `Storage` class represents the interface common to all storage types. The operations `Read` and `Write` represent the operations that allow the user to read and write data from and to the storage device. What kind of value types they must support will be discussed in more detail in the next subsection.

The storage hierarchy is split into *volatile* storage and *non-volatile* storage. Data stored in volatile storage will not survive program termination. An example of a volatile storage is conventional computer memory. Once an application terminates, its memory is usually freed by the operating system, and therefore any data still remaining in it is lost. Data stored in non-volatile storage on the other hand remains on the storage device even when a program terminates. Databases, disk storage, or even remote memory are common examples of non-volatile storage. Since the data will not be lost when the program terminates, additional housekeeping operations are needed to establish connections between the object and the actual storage device, to cut off existing connections, and to delete data that will not be used anymore. These operations are `Open`, `Close` and `Delete`.

The kind of storage to be used for saving application data depends heavily on the application requirements. Properties such as performance, capacity of the storage media and particularities of usage (for instance *write-once* devices like CD writers) may affect the choice. Persistence can be implemented in a stronger form to support fault tolerance of different sorts, including tolerating software design faults (bugs) for instance by using the recovery block scheme [10], or tolerating faults of the underlying hardware for instance by using checkpoints [11]. To apply persistence properly and to choose the suitable storage type, the application programmer has to identify the fault assumptions and to know the reliability of the storage devices which can be used.

This is why among the different non-volatile storage devices, we distinguish *stable* and *non-stable* devices. Data written to non-stable storage may get corrupted, if the system fails in some way, for instance by crashing during the write operation. Stable storage ensures that stored data will never be corrupted, even in the presence of application crashes and other failures.

Stable storage has been first introduced in [12]. The paper describes how conventional disk storage that shows imperfections such as bad writes and decay can be transformed into stable storage, an ideal disk storage with no failures, using a technique called *mirroring*. When using this technique, data is stored *twice* on the disk (often two different physical disks are used to store the two copies of the data to increase reliability even




Figure 1: Storage Class Diagram

more). If a crash occurs during the write operation of the first copy, the previously valid state can still be retrieved using the second copy. If the crash happens during the write operation of the second copy, the new state has already been saved in the first copy. When the system restarts later on, the state stored in the first copy must be duplicated and saved over the second copy.

Using this mirroring technique, any non-volatile, non-stable storage can be transformed into stable storage. It is therefore possible to write an implementation of the mirroring algorithm that is independent of the actual non-volatile storage that will effectively be used to store the data. To achieve this decoupling, the *Strategy* design pattern described in [13] has been used. The *Strategy* design pattern has three types of participants: the *Strategy*, the *Concrete Strategy* and the *Context*.

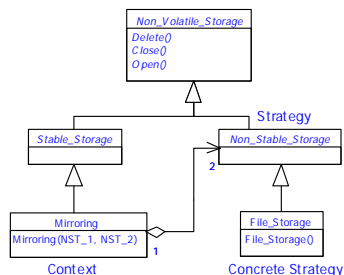


Figure 2: Stable Storage Using Mirroring

At instantiation time, two non-volatile storage objects must be passed as a parameter to the constructor of the mirroring class. That way, a variety of stable storages can be created reusing concrete implementations of non-volatile storage. What kind of non-volatile storage will be chosen depends on the needs of the application. To help the user make his choice, a concrete non-volatile storage must document the assumptions under which the storage is considered non-volatile and other information that might be useful for the application programmer such as for instance performance.

The mirroring technique is not the only one that can be used to create stable storage. Database systems for instance have their own mechanism to guarantee atomic updates of data. Typically this is done by structuring updates of data as transactions. A transaction can be committed, in which case the updates will be made permanent, or aborted, in which case the system remains unchanged. If any kind of failure occurs during the transaction, the data also remains unchanged. It is possible to write a concrete stable storage class that provides a bridge between an object-oriented programming language and a database.

Yet another form of providing stable storage is replication. The state of a persistent object can be broadcasted over the network and stored for instance in remote memory. Although memory is usually not seen as non-volatile, from the application point of view it is, since it survives program termination. The group of replicas as a whole can be considered stable, for as long as at least one of the remote machines remains accessible, the data can always be retrieved on a later execution.

Just as in the mirroring example, the replicated solution can be implemented in a generic way using the *Strategy* pattern. This time, the relationship between the context and the strategy is one to many as depicted in figure 3.

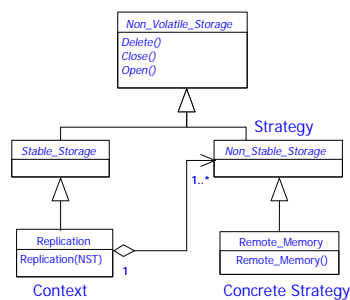


Figure 3: Stable Storage Using Replication

The replication class implements broadcasting and other replica management algorithms that handle failures of replicas during program execution.

3.2 Object Serialization

When storing the state of a persistent object on some kind of storage device, the data must first be transformed from its representation in memory into some form that can be stored by the device. Most of the time the most convenient form will be a flat stream of bytes e.g. for storing data in flat files or sending data through network transport buffers. Interfaces to ODBMs can be more elaborate.

The *Serializer* design pattern described in [14] is an ideal solution for this kind of problem. It provides a mechanism to efficiently stream objects into data structures of any form as well as create objects from such data structures. The participants of the *Serializer* pattern are the *Reader / Writer*, the *Concrete Reader / Concrete Writer*, the *Serializable* interface, *Concrete Elements* that implement the *Serializable* interface and different *Backends*. The structure of the *Serializer* pattern is shown in figure 4.

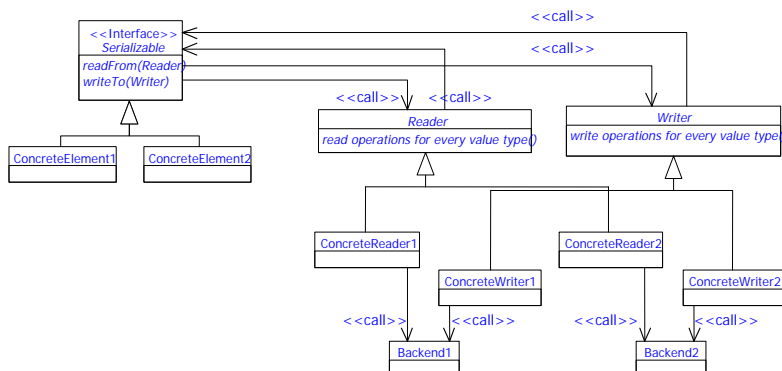


Figure 4: The Serializer Pattern

The *Reader* and *Writer* parts declare protocols for reading and writing objects. These protocols consist of read respectively write operations for every value type, including composite types, array types and object references. The *Reader* and *Writer* hide the *Backend* and the external representation format from the serializable objects. *Concrete Reader* and *Concrete Writer* implement the *Reader* and *Writer* protocols for a particular backend and external representation format.

The *Serializable* interface defines operations that accept a *Reader* for reading and a *Writer* for writing. It also should provide a *Create* operation that takes a class identifier as an argument and creates an instance of the denoted class. *Concrete Element* is an object implementing the *Serializable* interface, which allows it to read and write its attributes to a *Concrete Reader / Concrete Writer*.

The *Backend* is a particular backend, and corresponds to our storage class shown in the previous subsection. A *ConcreteReader/ConcreteWriter* reads from/writes to its backend using a backend specific interface. Relational database front-ends, flat files or a network buffers are examples of concrete backends.

When invoked by a client, the *Reader / Writer* hands itself over to the serializable object. The serializable object makes use of its protocol to read / write its attributes by calling the read / write operations provided by the *Reader / Writer*. For certain value types such as composite types, the *Reader / Writer* might call back to the serializable object or forward the call to other objects that implement the *Serializable* interface. This results in a recursive back-and-forth interplay between the two parties.

The bigger the set of supported value types of the *Reader / Writer* interface is, the more type information can be used by the *Concrete Reader / Concrete Writer* to efficiently store the data on the backend. On the other hand, there are backends that support only a small set of value types. Flat files for instance only support byte transfer. For these kinds of backends the *Concrete Reader / Concrete Writer* must contain implementation code that maps the `read / write` operations of unsupported value types to the ones that are supported.

The big advantage of the *Serializer* pattern is that the application class itself has no knowledge about the external representation format which is used to represent their instances. If this were not the case, introducing a new representation format or changing an old one would require to change almost every class in the system.

In some object-oriented programming languages, such a serialization mechanism is already provided, which means that the `readFrom / writeTo` operations defined in the *Serializable* interface have predefined implementations for all value types of the programming language that are not covered by the *Reader / Writer* interface. The Java Serialization package [15] or Ada streams [8, 13.13] are examples of such predefined language support. If no language support is available, the `readFrom / writeTo` operations of the *Serializable* interface must be implemented for every *Concrete Element*.

3.3 Creation of Persistent Objects

When creating an instance of a persistent object, the user must be able to specify on what kind of storage he wants the state of the object to be saved. The object can then create an instance of the corresponding storage class and thus establish a connection to the storage device.

The information needed to create an instance of a concrete storage class is device dependent. To create a new file, a user must typically provide a file name that follows certain conventions, and maybe also a path that specifies in which directory the file should be created. To access remote memory, an IP number or machine name must be provided. To solve this problem, a parallel hierarchy of storage parameters has been introduced. It has the same structure as the storage hierarchy (see figure 5). This allows each storage class to define its own storage parameter type containing all the information it needs to uniquely identify data stored on the device.

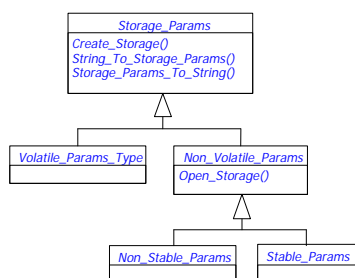


Figure 5: The Storage Parameter Class Hierarchy

`Create_Storage`. A *Concrete Creator*, in our case a concrete storage parameter class, must provide an implementation for this method: the corresponding creator function of the storage class must be called, passing as a parameter the information stored inside the concrete storage parameter instance. Non-volatile storage needs a second creator function, `Open_Storage`, that will instantiate the non-volatile class

At the same time, the storage parameter class allows a user to create instances of storage classes. This is done using the well-known *Factory Method* pattern described in [13]. The participants of this design pattern are the *Product*, the *Concrete Product*, the *Creator* and the *Concrete Creator*.

The *Product* and *Concrete Product* are in our case the storage class and its descendants, as they define and implement the interface of the objects the factory will create.

The *Creator* is the storage parameter class, for it declares the abstract factory method

without creating a new storage on the device. Instead a connection between already existing data and the storage object will be established.

The `Create_Storage` and the `Open_Storage` operations define the connection between the two parallel class hierarchies.

3.4 Identification of Persistent Objects

Since the state of a persistent object survives program termination, there must be a unique way to identify a persistent object that remains valid during several executions of the same program. The storage parameter that has been introduced in the previous subsection uniquely identifies a location on the storage device, and can therefore also be used as a means for object identification.

Sometimes it can be convenient for a user to treat persistent objects in a uniform way. An object name in the form of a string has proven to be an elegant solution for uniform object identification [6]. The two functions `Storage_Params_To_String` and `String_To_Storage_Params` provide a mapping between the two identification means.

3.5 Storage Management

Once a persistent object has been created and its state saved to a non-volatile storage, the data will theoretically remain on the storage forever. The only way to remove the data and free the associated storage space is to explicitly delete the object. This situation can lead to permanent storage leaks, if the user forgets to store the parameters that allow him to identify the object on subsequent application runs.

A simple solution to this problem is to provide some sort of reliable persistent directory. The parameters of every persistent object created so far are automatically stored in it. At any time, the user can consult the list of existing persistent objects to determine which of them he still needs and which of them he wants to delete.

Since the objects persist, the state of the directory should also survive program termination. The directory itself therefore is just another persistent object. When writing the state of the directory to the storage, the storage parameters of all persistent objects that have been created in the system must be written to the storage. It is therefore important that the storage parameter class also implements the *Serializable* interface.

The directory must be reliable. Even a crash during the update of the directory should not corrupt the data. This can be achieved by storing the directory on a stable storage. But this is not enough. Storage leaks can occur if a crash occurs after a new persistent object has been created, but before the creation has been registered in the directory. To prevent this problem, the creation and deletion of objects and the updating of the directory to reflect the change must be executed atomically.

4 Putting Everything Together

With the previous solutions in mind, we can now put together the overall system. Its structure is shown in figure 6. For simplicity, the *Reader* and *Writer* parts of the *Serializer* pattern are shown as one class.

Before using any part of the framework, the user must initialize the persistence support. He has to choose where to store the persistent directory by instantiating the appropriate storage parameters and passing them to the `Initialize` operation of the persistence support class. In order to make the directory reliable a stable storage must be used. A good idea is to hard-code the storage parameters of the persistent directory in the application code, for on subsequent runs the same parameters must be used again. During the `Initialize` operation, the persistent support class will try and

restore a previously valid directory, or if this fails, create a new, empty directory instead. Once the initialization has been performed, the user can create, restore, save and delete persistent objects.

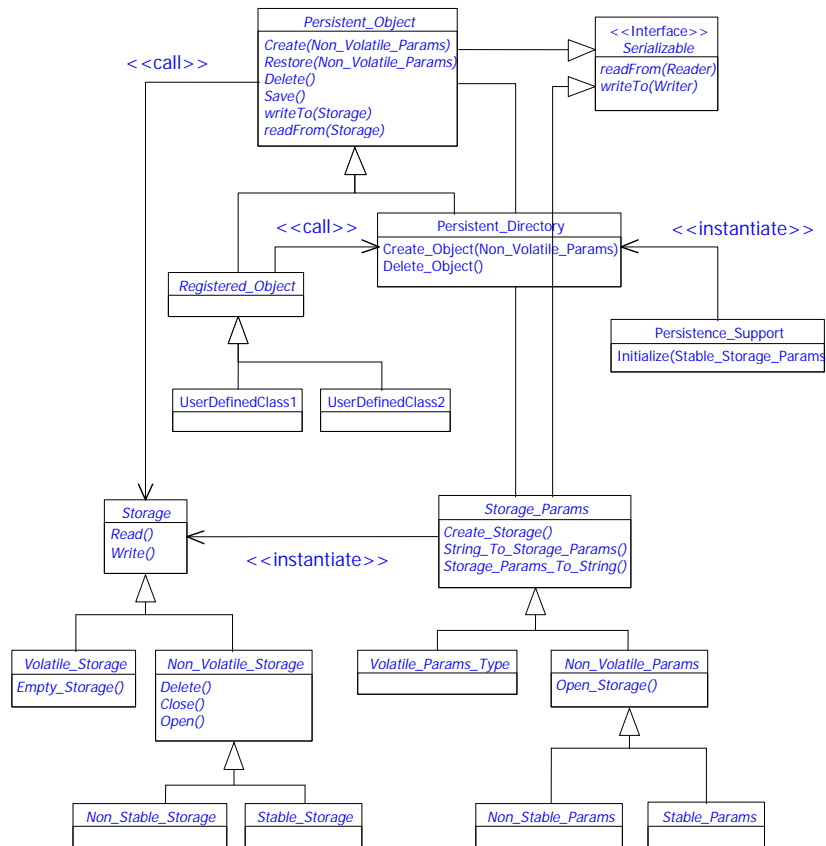


Figure 6: Framework Overview

The root class of the framework is the persistent object class. It implements the *Serializable* interface described in the previous section and the operations *Create*, *Restore*, *Save* and *Delete*.

The registered object class derives from the persistent object class. It is responsible for registering any new persistent object instances with the persistent directory. In order to define a new persistent object, the user must derive from the registered object class and add any application dependent state using new class attributes. He must also implement the *readFrom* and *writeTo* operations of the *Serializable* interface, if the underlying programming language does not provided them automatically.

The user-defined class must also provide two constructors, *Create*, to create a new instance of a persistent object, and *Restore*, to restore the state of an already existing persistent object. They must perform initialization of the application dependent object state if needed and then up-call the corresponding constructor of the registered object class, which will perform the necessary operation for registering the object with the persistent directory.

The *Create* operation of the registered object class will call the *Create_Object* operation of the persistent directory. Inside *Create_Object*, the operation *Cre-*

ate of the persistent class is called. This will actually create the object on the specified store as explained below. At the same time, the storage parameters of the newly created class are stored inside the persistent directory. As mentioned before, this must be done atomically. A simple way of implementing this is to use a simple form of logging.

The implementation of the constructors `Create` and `Restore` in the persistent object class will create an instance of the storage type identified by the storage parameter using the `Create_Storage` or `Open_Storage` factory methods. If `Restore` has been called, the state of the object is read from the storage using the operations of the *Serializable* interface.

This is again an application of the *Strategy* design pattern. The persistent object, the *Context*, is configured with a storage, the *Strategy*, at instantiation time. The user chooses the storage medium for his object by passing the corresponding storage patterns to the `Create` or `Restore` constructors. Once this association has been set up, the user can write the state of the persistent object to the associated storage using the `Save` operation. The implementation will then write an object identifier and successively all object attributes to the associated storage.

5 Experimental Implementation

An experimental implementation of the framework [16] has been realized using the object-oriented programming language Ada 95 [8]. The fact that Ada 95 supports streaming of objects has simplified the implementation, but also narrowed down the read / write operations of the storage type to support byte reads and writes only. This fact is reflected in the specification of the abstract `Storage_Type`. There is only one pair of read / write operations, and it operates on stream element arrays (arrays of bytes):

```
with Ada.Streams; use Ada.Streams;
package Storage_Types is
  type Storage_Type (<>) is abstract tagged limited private;
  type Storage_Ref is access all Storage_Type'Class;
  procedure Read (Storage : in out Storage_Type;
                 Item     : out Stream_Element_Array;
                 Last     : out Stream_Element_Offset) is abstract;
  procedure Write (Storage : in out Storage_Type;
                  Item     : in Stream_Element_Array) is abstract;
private
  ...
end Storage_Types;
```

A persistence support programmer writing a new interface for a storage device must derive from this class (or more precisely from a subclass such as `Non_Stable_Storage_Type` whose properties correspond to the properties of the device) and implement the `Read` and `Write` operations.

The following sample code shows how an application programmer uses the framework. First, he must declare his own persistent type by deriving from the registered object class, here called `Registered_Object_Type`, adding additional attributes that will contain the application data. He must also provide a constructor that allocates the new object, performs initialization and then up-calls the constructor of the registered object class. The following example shows how to declare a persistent integer type:

```
with Persistent_Objects.Registered; use Persistent_Objects.Registered;
package Persistent_Integers is
  type Persistent_Int_Type is new Registered_Object_Type with record
    Value : Integer;
  end record;
end Persistent_Integers;
```

```

    type Persistent_Int_Ref is access all Persistent_Int_Type'Class;
    function Create (Storage_Params : Non_Volatile_Params_Type'Class)
        return Persistent_Int_Ref;
end Persistent_Integers;

```

With these declarations, instances of the persistent integer class can now be created and used.

```

-- include the necessary files
with File_Storage_Params;
with Persistent_Integers; use Persistent_Integers;
-- Create a new persistent integer and save it's contents
declare
    I : Persistent_Integer_Ref :=
        Create (File_Storage_Params.String_To_Storage_Params ("foo"));
begin
    I.Value := 123;
    Save (I.all);
end;
-- Restore the contents of the previously created persistent integer
declare
    I : Persistent_Integer_Ref := Persistent_Integer_Ref (
        Restore (File_Storage_Params.String_To_Storage_Params ("foo")));
begin
    -- work with I
end;

```

When creating the persistent integer, the user chooses on which storage the state shall be stored by calling the `String_To_Storage_Params` function of the chosen storage class. In the example above, the persistent integer is stored in a file, since `String_To_Storage_Params` of the `File_Storage_Params` class is called.

6 Related Work

The PJava project [17] aims to provide orthogonal persistence for the Java language without modifying the language. Roots of persistence have been defined, where individual objects can be registered during run-time. All objects reachable from a persistent root are made persistent (*persistence by reachability*). This is achieved by modifying the Java Virtual Machine.

The CORBA Persistent Object Service [18] is a standardized CORBA service that allows CORBA Objects to make all or part of their state persistent. Whether or not the client of such a persistent object is aware of the persistent state is a choice the object has. By supporting special interfaces, and describing the persistent data using an interface definition language, a persistent object can delegate the management of its persistent state to other objects. The persistent data finally get stored in so-called data-stores. This is the CORBA way to abstract the real storage devices, such as file systems or databases. Each persistent object can dynamically be connected to a data-store. From then on it is possible to save and restore the persistent state.

The PerDiS project [19] takes a very different approach to persistence. They address the issue of providing support for distributed collaborative engineering applications such as CAD programs, where large volumes of fine-grain, complex objects must be shared across wide-area networks. They present the user with some form of a persistent, distributed memory. The application accesses this memory transactionally. The memory is divided into clusters containing objects. Named roots provide the entry points. Objects are connected by pointers. Reachable objects are stored persistently in clusters on disk; unreachable objects are garbage-collected automatically.

Transactional systems often incorporate a persistence support for guaranteeing the *durability* property. Several object-oriented transactional systems use class libraries to

declare transactional objects and to manipulate them. For example, *Arjuna* [20], one of the best known object-oriented transactional systems, is delivered as a set of C++ classes. It provides two types of object stores: *persistent* (the stable storage is implemented as a set of UNIX files) and *volatile* (implemented in the main memory). These systems usually hide all details of the ways they implement persistence and do not offer programmers any object-oriented extendable approaches to manipulate the stable storage (e.g. to include a new storage class to support a new media, or, to change the storage dynamically). They incorporate object storage control into the transactional control.

7 Conclusions

We have shown the construction of an object-oriented framework providing persistence support for object-oriented programming languages. It does only rely on basic object-oriented programming techniques, and can therefore be implemented in any object-oriented programming language.

The advantage of using object-oriented programming is obvious. Class hierarchies have allowed us to clearly classify the different storage devices and show the dependencies among them. Abstract classes provide the interface for the different types of storage.

The structure of the framework is based on well-known design patterns. The advantages of using design patterns in this context are substantial:

- Design patterns enhance the modularity and flexibility of object-oriented programming. They offer solutions that have been proven to be successful.
- The *Serializer* pattern makes it easy to add new data representation formats for objects that must be written to new storage devices by introducing a new *Reader / Writer* pair. It also moves the knowledge about the external data representation format out of the persistent object itself.
- Encapsulating the storage devices in separate *Strategy* classes lets the user change and replace particular storage implementations or even extend the storage device hierarchy without modifying the persistent object class. It also promotes reuse when implementing stable storage on top of non-volatile storage.
- Using the *Factory Method* pattern in combination with a parallel class hierarchy representing storage parameters allowed us to provide a uniform way for persistent objects to create storage devices. At the same time the storage parameter provides unique identification of persistent objects.
- People familiar with the used design patterns will be able to understand the structure of the framework faster.

8 References

- [1] Atkinson, M. P.; Morrison, R.: “Orthogonally Persistent Object Systems”. *VLDB Journal* **4**(3), pp. 319 – 401, 1995.
- [2] Atkinson, M. P.; Bailey, P. J.; Chisholm, K. J.; Cockshott, W. P.; Morrison, R.: “An Approach to Persistent Programming”. *Computer Journal* **26**(4), pp. 360 – 365, 1983.
- [3] Atkinson, M. P.; Chisholm, K.; Cockshott, W.: “PS-Algol: An Algol with a Persistent Heap”. *ACM SIGPLAN Notices* **17**(7), pp. 24 – 31, July 1981.
- [4] Matthews, D. C. J.: “A persistent storage system for Poly and ML”. *Technical Report TR-102*, Computer Laboratory, University of Cambridge, January 1987.

- [5] Atkinson, M. P.; Jordan, M. J.; Daynès, L.; Spence, S.: “Design Issues for Persistent Java: a Type-Safe, Object-Oriented, Orthogonally Persistent System”. In *Proc. of the 6th Int. Workshop on Persistent Object Systems*, Cape May NJ (USA), May 1996.
- [6] Gosling, J.; Joy, B.; Steele, G. L.: *The Java Language Specification*. The Java Series, Addison-Wesley, Reading, MA, USA, 1996.
- [7] Oudshoorn, M. J.; Crawley, S. C.: “Beyond Ada 95: The Addition of Persistence and its Consequences”. In *Reliable Software Technologies - Ada-Europe’96*, volume 1088 of *Lecture Notes in Computer Science*, pp. 342 – 356, Springer Verlag, 1996.
- [8] ISO: *International Standard ISO/IEC 8652:1995(E): Ada Reference Manual*, Lecture Notes in Computer Science **1246**, Springer Verlag, 1997; ISO, 1995.
- [9] Eppinger, J. L.; Mummert, L. B.; Spector, A. Z.: *Camelot and Avalon - A Distributed Transaction Facility*. Morgan Kaufmann Publishers, San Mateo, CA, 1991.
- [10] Randell, B.: “System structure for software fault tolerance”. *IEEE Transactions on Software Engineering* **1**(2), pp. 220 – 232, 1975.
- [11] Lee, P. A.; Anderson, T.: “Fault Tolerance - Principles and Practice”. In *Dependable Computing and Fault-Tolerant Systems*, volume 3, Springer Verlag, 2nd ed., 1990.
- [12] Lampson, B. W.; Sturgis, H. E.: “Crash Recovery in a Distributed Data Storage System”. *Technical report*, XEROX Research, Palo Alto, June 1979.
- [13] Gamma, E.; Helm, R.; Johnson, R.; Vlissides, J.: *Design Patterns*. Addison Wesley, Reading, MA, 1995.
- [14] Riehle, D.; Siberski, W.; Bäumer, D.; Megert, D.; Züllighoven, H.: “Serializer”. In *Pattern Languages of Program Design 3*, pp. 293 – 312, Addison Wesley, 1998.
- [15] Sun Microsystems: *Java Object Serialization Specification*, November 1998.
- [16] Kienzle, J.; Romanovsky, A.: “A Flexible Approach for Streaming”. *Technical Report 99/323*, Swiss Federal Institute of Technology, November 1999.
- [17] Atkinson, M. P.; Daynes, L.; Jordan, M. J.; Printezis, T.; Spence, S.: “An orthogonally persistent Java”. *ACM SIGMOD Record* **25**(4), pp. 68 – 75, December 1996.
- [18] Object Management Group, Inc.: *Externalization Service Specification*, December 1998.
- [19] Ferreira, P.; Shapiro, M.; Blondel, X.; Fambon, O.; Garcia, J.; Kloosterman, S.; Richer, N.; Roberts, M.; Sandakly, F.; Coulouris, G.; Dollimore, J.; Guedes, P.; Hagimont, D.; Krakowiak, S.: “PerDiS: design, implementation, and use of a PERSistent Distributed Store”. *Technical report*, QMW TR 752, CSTB ILC/98-1392, INRIA RR 3525, INESC RT/5/98, October 1998.
- [20] Parrington, G. D.; Shrivastava, S. K.; Wheeler, S. M.; Little, M. C.: “The Design and Implementation of Arjuna”. In USENIX (Ed.), *Computing Systems, Summer, 1995.*, volume 8, pp. 255 – 308, Berkeley, CA, USA, Summer 1995, USENIX.