Managed Data: Modular Strategies for Data Abstraction

Alex Loh
University of Texas at Austin
alexloh@cs.utexas.edu

Tijs van der Storm
Centrum Wiskunde & Informatica (CWI)
storm@cwi.nl

William R. Cook
University of Texas at Austin
wcook@cs.utexas.edu

Abstract

Managed Data is a two-level approach to data abstraction in which programmers first define data description and manipulation *mechanisms*, and then use these mechanisms to define *specific kinds of data*. Managed Data allows programmers to take control of many important aspects of data, including persistence, access/change control, reactivity, logging, bidirectional relationships, resource management, invariants and validation. These strategies for implementing these features are implemented once as reusable modules. Managed Data is a general concept that can be implemented in many ways, including reflection, metaclasses, and macros. In this paper we argue for the importance of Managed Data and present a novel implementation of Managed Data based on *interpretation* of *data models*. The paper also discusses strategies for implementing Managed Data. In object-oriented languages, Managed Data typically requires programmer control over the meaning of the field/method access operator (the "dot" operator). Bootstrapping poses difficulties, because implementing the first data manager requires some data. As a case study, we used Managed Data in a web development framework from the Ensō project to reuse database management and access control mechanisms across different data definitions.

Categories and Subject Descriptors D.3.3 [*Language Constructs and Features*]: Data management

General Terms Data management, Aspect-oriented programming, Model-based development

Keywords Schema, Interpretation, Composition

[Copyright notice will appear here once 'preprint' option is removed.]

1. Introduction

Mechanisms for organizing and managing data are a fundamental aspect of any programming model. Most programing models provide built-in mechanisms for organizing data. Well-known approaches include *data structure definitions* (as in Pascal, C [14], Haskell [11], ML [18]), *object/class models* (as in Java [1], Smalltalk [8], Ruby [26]), and *predefined data structures* (as in Lisp [21], Matlab). Languages may also support abstract data types (as in ML, Modula-2 [29], Ada [28]), or a combination of multiple approaches (e.g JavaScript, Scala [19]). A key characteristic of all these approaches is that the fundamental mechanisms for structuring and manipulating data are predefined. Predefined data structuring mechanisms allow programmers to create specific kinds of data, but they do not allow fundamental changes to the underlying data structuring and management mechanisms themselves.

Predefined data structuring mechanisms are insufficient to cleanly implement many important and common requirements for data management, including persistence, caching, serialization, transactions, change logging, access control, automated traversals, multi-object invariants, and bi-directional relationships. The difficulty with all these requirements is that they are pervasive features of the underlying data management mechanism, not properties of individual data types. It is possible to define such features individually for each particular kind of data in a program, but this invariably leads to large amounts of repeated code. To implement these kinds of crosscutting concerns, developers often resort to preprocessors, code generators, byte-code transformation, or modified runtimes or compilers. The resulting systems are typically ad-hoc, fragile, poorly integrated, and difficult to maintain.

This paper presents *Managed Data*, an approach to data abstraction that gives programmers control over the data structuring mechanisms. Managed Data has three essential components: (1) *data description languages* that describe the desired structure and properties of data, (2) *data man-*

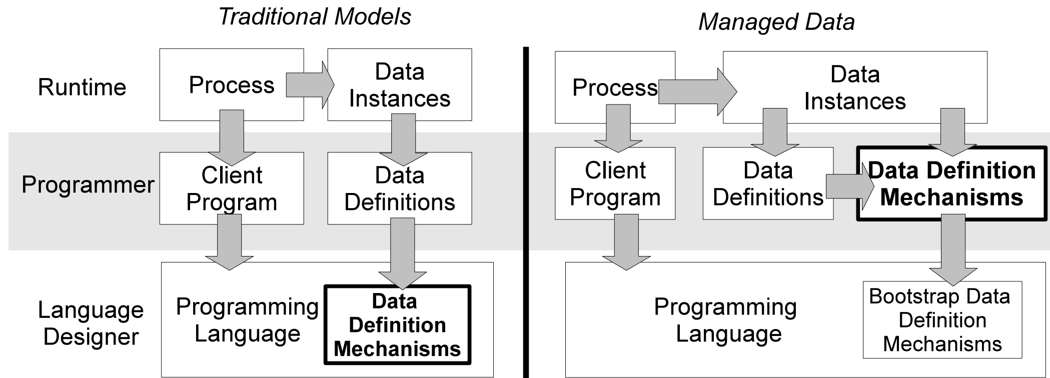

Figure 1. Traditional Data Mechanisms versus Managed Data

agers that enable creation and manipulation of instances of data that conform to the description, and (3) *integration* with a programming language, so that managed data instances can be used in the same way as ordinary data objects. Managed Data has a strong emphasis on modularity, allowing both data description languages and data managers to be modularly defined and reused. Additionally, Managed Data definitions may also themselves be Managed Data via a *bootstrapping* process, extending the benefits of programmable data structuring to their own implementation.

One way to understand Managed Data is as a *design pattern* that allows the programmer to define the behavior of data manipulation operations traditionally considered built-in primitives: initialization, type tests, casting, pointer equality, and field access.

This pattern can be implemented in many different ways in different languages, or in different programming styles, including object-oriented or functional. Some programming systems support a degree of control over the data structuring mechanisms. Meta-classes in Smalltalk define how classes are instantiated and compiled. The reflective features of Ruby, Python and Smalltalk can be used to trap and handle undefined methods and properties, which can be used to create proxies or implement virtual objects. Attributes and byte-code manipulation can be used to specify and implement pervasive data management behaviors in Java. Scheme macros are often used to create data structuring mechanisms. For example, the `defstruct` macro defines mutable structures with a functional interface. The Adaptive Object Model Architecture [30] provides an architecture for this approach, but does not discuss how it is bootstrapped or integrated with existing languages. In general, static languages are less able to support Managed Data directly, so they require the use of external code generators. Dynamic languages often provide reflective hooks that can be used to implement Managed Data.

Our implementation of Managed Data uses the meta-programming framework of Ruby. In Sections 2 and 3, we begin with a implementation based on dynamic proxies, which declares properties and methods on the fly with Ruby's `method_missing` mechanism. A second, more static implementation, introduced in Section 4, uses `define_method` to define instance methods as closures at run-time.

Section 5 demonstrates the use of Managed Data in EnsōWeb, a web development framework. Managed Data is used to configure pervasive data management concerns such as persistence, security and logging in a data-independent way, without boilerplating at the client code.

Finally, Section 6 compares and classifies related work and Section 7 concludes.

2. Simple Managed Data

Managed Data lifts data management up to the application level, by allowing a programmer to build *data managers* that handle the fundamental data manipulation primitives usually hardcoded into the programming language. The input to a data manager is a *schema*, which describes the structure and behavior of the data to be managed.

The three essential components of Managed Data are a *data description language* for writing schemas that describe the desired structure and properties of data, a *data manager* to interpret the data description language, and *integration* of the data manager with the programming language to allow data to be created and manipulated. Although there can be many flavors of schemas and data managers in a Managed Data system, it is useful to illustrate the concepts using a simple first example.

2.1 Schemas for Data Description

A basic record is a collection of named fields each of whose value is constrained to be a value of a primitive type. To define basic records as Managed Data, we need a simple data description language, in this case a mapping from field names to primitive type names. While we will eventually consider stand-alone data description languages, for now we can describe a basic record using a Ruby hash. For example, the following hash describes simple two-dimensional points:

```
Point = { x: Integer, y: Integer }
```

This defines a *hash* in Ruby 1.9 syntax. The hash is an object that represents a mapping from values to values. In this case the keys of the hash are the symbols `x` and `y`. Both these symbols are mapped to the class `Integer`. Classes are values in Ruby, as in Smalltalk. Although this definition appears to be a *type*, it is actually just a value. In what follows we will interpret this value as a kind of specification, or dynamically checked type, but it is important to keep in mind that `Point` is just a hash value.

`Point` describes records with `x` and `y` fields whose values are integers. `Point` is a simple example of a *schema*. A schema is a description of the structure, and possibly behavior, of data. In later sections we will consider more complex schema languages, but for now we will consider just simple schemas that describe basic records. It is easy to describe many different kinds of records using this simple notation.

```
Person = { name: String,
           birthday: Date,
           erdos: Integer }
```

Schemas are not complete specifications without a corresponding data manager. In this case our simple schema is a definition of a data type, and does not state whether the records are immutable or mutable, if they can be stored in a database, or transformed in other ways. Schemas can be interpreted in many different ways to create different kinds of records.

2.2 Using Managed Data

Before showing how to implement managed data, it is important to consider how managed data is used. This study will create a set of requirements, or a specification, that guides the implementation.

Our goal is to have objects that conform to the basic record schema provided and can be used within a programming language. At the same time, they should have additional behavior, such as logging and access control, as defined by the data manager. Assume we have a data manager, `BasicRecord`, which enables creating and updating basic records, given a schema. Here is an example of how we might use this data manager:

```
p = BasicRecord.new Point
p.x = 3
p.y = -10
print p.x + p.y
```

In this example, the `BasicRecord` manager is instantiated and given the `Point` schema as an input. The result is a new point object, that is, an object that conforms to the point schema. The object `p` has fields `x` and `y` which can be accessed and assigned.

Attempting to violate the schema results in an error. Some examples are given below:

```
print p.z # unknown field z
p.x = "top" # x must have type Integer
p.z = 3   # assigning unknown field z
```

The `BasicRecord` manager interprets the `Point` schema to manage points. We have specified that `BasicRecord` supports mutable fields, and that it checks types and checks the legality of field names. These simple checks are typical of how data is managed in most object-oriented programming languages. In later sections we will augment them with different data managers implementing a variety of features, including immutability, persistence, invariants, etc.

The fields of the managed data object are dereferenced using the “dot” operator, like in most object-oriented languages. For languages such as Java, C#, JavaScript, and PHP, this requires the class the object belongs to statically declare those fields or methods. However, since the schema that contains those fields is only known dynamically, the data manager must be able to determine the fields and methods of the managed data object dynamically.

2.3 Implementing a Data Manager

Our implementation of the data manager uses the *reflection* capabilities of the Ruby programming language. In Ruby, the form `o.f` can be redefined by implementing a `method_missing` method. This technique is illustrated in Figure 2, which defines a simple data manager for updatable records. The managed object stores the schema in the `@types` member variable. It then defines a `@values` record and initializes it with default values for each type (the way that default values are defined is omitted).

The key mechanism for implementing a dynamic object here is the `method_missing` method. It is called whenever a method (or field) is called on the object that is not defined. Since the class does not define any methods, the `method_missing` is called for all fields and methods. The arguments to `method_missing` are the method name and the arguments of the original call. There is one special case. In Ruby, an assignment to a field, of the form `obj.field = val` is converted into a call of the form `o.method_missing(:field=, val)`, where `:field=` is a symbol formed by appending `=` to the field name.

The `method_missing` method first checks if the call is a field assignment, in which case it sets `assign` to true. If it is an assignment, the field name is updated to remove the `=`. It then checks if the field is defined in the schema, and returns an error if it is not. For an assignment, it checks if the argument is of the right type, and raises an error if not. It then updates the `@values` map with the new value. For regular field access, it returns the current value for the named field in `@values`. One drawback of using `method_missing` is that it does not handle name clashes correctly. Specifically, if a field has the same name as an existing method, such as `_set` and `_get`, or one of the methods defined by Ruby’s `Object` class, like `clone`, `method_missing` will call that method in-

```
class BasicRecord
  def initialize(types)
    @types = types
    @values = {}
    # assign default values to all fields
    types.each do |name, type|
      @values[name] = type.default_value
    end
  end
  # internal methods for getting and setting fields
  def _set(name, value); @values[name] = value; end
  def _get(name); @values[name]; end
  # all properties and methods are handled here
  def method_missing(name, *args)
    if name =~ /(.*)=/
      name = $1.to_sym
      if !@types[name]
        raise "assigning unknown field #{name}"
      end
      if !(args[0].class <= @types[name])
        raise "#{name} must have type #{@types[name]}"
      end
      _set(name, args[0])
    else
      if !@types[name]
        raise "unknown field #{name}"
      end
      return _get(name)
    end
  end
end
```

Figure 2. A data manager for simple records

stead. To avoid such problems, programmers need to follow the convention of not prefixing field names with an underscore and ensuring that they do not share names with the default Ruby classes.

Our data manager loads and directly interprets the schema, and creates dynamic objects that act according to its specification. We do not stipulate *how* a data manager is implemented. It could work by code generation, reflection, bytecode manipulation, or other techniques. The advantage of our approach is that there is no need for an additional code generation step to translate the schema into code.

Figure 1 illustrates the difference between traditional built-in data structuring mechanisms and Managed Data. In the traditional approach, the programming language includes a process and data sublanguages, which are both predefined. With Managed Data, the data structuring mechanisms are defined by the programmer by interpretation of data definitions. Since a data definition model is also data, it requires a meta-definition mechanism. This infinite regress is terminated by a boot-strap data definition that is used to build the Managed Data system itself.

```
class LockableRecord < BasicRecord
  def _lock; @locked = true; end
  def _set(name, value)
    if @locked
      raise "cannot change {name} of locked object"
    end
    super
  end
end
```

Figure 3. A lockable data manager

```
class InitRecord < LockableRecord
  def initialize(fields, init)
    super(fields)
    # assign default values to all fields
    init.each do |name, value|
      _set(name, value)
    end
    _lock()
  end
end
```

Figure 4. A data manager with field initialization

3. Alternative Data Managers

Managed Data drops the idea of a single predefined data description language, and allows programmers to create or extend their own schema languages and data managers. Data managers naturally manage many schemas, but there can also be multiple data managers for a given schema to, for example, implement in-memory strategy versus a database strategy for managing the data. Data managers may be composed to form a stack of managers that has their combined behavior.

In this section we present a few alternative data managers that enhance the basic data structuring mechanism provided by object-oriented programming languages.

3.1 Immutability

The data manager in Figure 3 introduces a locking mechanism to protect a record from changes. This can be useful for certain types of optimization or to implement constant values. `LockableRecord` inherits from the default data manager and overrides its `_set` field. In this implementation locking is irrevocable, but we could also have as easily included an `_unlock` option.

```
class ObserverRecord < BasicRecord
  def initialize(fields)
    super
    @_observers = Set.new
  end

  def _observe(&block)
    @_observers.add(block)
  end

  def _set(name, value)
    super
    @_observers.each do |obs|
      obs.call(self, name, value)
    end
  end
end

class DataflowRecord < ObserverRecord
  def _get(name, dependent=nil)
    if not dependent.nil?
      _observe do |obj, field, value|
        if field == name
          # inform dependent
          dependent.set_dirty
        end
      end
    end
    super(name)
  end
end
```

Figure 5. A data manager for the Observer Pattern

3.2 Instance Initialization

Constant objects are initialized at the beginning of their lifespan and immutable henceforth. The data manager in Figure 4 extends `LockableRecord` with the option to initialize fields during construction. The constructor parameter `init` is a map from field names to initial values and can be used as follows:

```
p = InitRecord.new Point, x: 3, y: 5
```

We extended `InitRecord` from `LockableRecord` because we wanted to build constant objects, but in general immutability and initialization are orthogonal concepts that can apply independently.

3.3 Observers

Figure 5 presents a data manager that supports the OBSERVER PATTERN [7]. The following code snippet logs changes to the record by printing out a message whenever a point is changed.

```
p = ObserverRecord.new Point
p._observe do |obj, field, value|
  print "updating #{field} to #{value}\n"
end
p.x = 1; p.y = 6
p.x = p.x + p.y
```

Output:

```
updating x to 1
updating y to 6
updating x to 7
```

An observer is a useful pattern especially for event-driven tasks such as enforcing inverses and dataflow programming. `DataflowRecord` is a record that allows callers who access a particular field to register themselves as dependents. Dependents will be told to re-compute their cached values when the value in that field changes.

Compared to `LockableRecord` and `InitRecord`, `ObserverRecord` and `DataflowRecord` demonstrate another kind of dependency where one data manager rely on the services provided by another. These examples expose the need for a general strategy for combining data managers such that dependencies between data managers are respected and independent data managers can be selected modularly. Although outside the scope of this paper, possible approaches may include building the data manager into the type system, or using feature models [12].

4. Self-Describing Schemas

A self-describing schema is a schema that can be used to define schemas (including itself). Self-describing schemas are important because they allow schemas to be managed data. The concept of self-description is well known. The Meta-Object Facility (MOF) meta-metamodel [20] is self-describing. It is possible to write a BNF grammar for BNF grammars. Rather than starting from an existing metamodel, we will first develop what we believe to be a minimal self-describing schema, and then present a more complete and useful version based on this foundation.

4.1 A Minimalist Schema Schema

In the previous sections, the schema was a simple mapping from field names to primitive types, which can describe the structure of simple records. This simple schema format cannot be used to describe itself, because a simple schema is not a record. To model the structure of a schema, we need to be able to describe a record type as a collection of fields, each of which having a name and a type. This immediately requires a schema to have two concepts, namely “type” and “field”. A third concept arises, a “schema”, as a collection of types. Finally, we must recognize that some fields are single values, for example the name or type of a field, while other fields are many-valued, including the fields of a type and the

```
class Schema          class Field        class String
  classes: Class*       name: String     class Bool
class Class             type: Class
  name: String          many: Bool
  fields: Field*
```

Figure 6. A minimalist self-describing schema

```
class Schema
! types: Type*

class Type
# name: string
  schema: Schema / types

class Primitive < Type

class Class < Type
  supers: Class*
  subclasses: Class* / supers
! defined_fields: Field*
  fields: Field*
    = supers.flat_map(.fields) + defined_fields

class Field
# name: string
  owner: Class / defined_fields
  type: Type
  optional: bool
  many: bool
  key: bool
  traversal: bool
  inverse: Field? / inverse
  computed: Expr?

primitive string
primitive bool
```

Figure 7. An Ensō Schema Schema

types in a schema. These concepts can be represented in the minimalist self-describing schema shown in Figure 6.

In this notation, a *class* is introduced by the keyword **class** followed by the class name and then a list of *field definitions*. Each field definition has the form *name*:*type* giving the name and type of the field. A *type* is a class name optionally followed by `*`, which indicates that the field is many-valued, i.e. a collection of values of the given type.

In the example, the class `Schema` has a single field, `classes`, which is a collection of `Class` values. A `Class` has a `name` field of type `String` and a collection of `fields`. A `Field` has a `name`, a `type`, and a boolean flag indicating whether the field is single or many-valued. To be complete, it is necessary to define `String` and `Bool` as classes.

This explanation of the content of the schema also demonstrates why it is self-describing, because every concept used in the explanation is included in the definition. One small point is that the type of a field is a `Class` in the schema, but the type is written as a name in the figure. During parsing or interpretation of the textual presentation of the schema, the named must be looked up to find the corresponding `Class`.

While this minimalist schema could be used directly, we believe that it is more useful to work with a slightly more complex self-describing schema. There are two major problems with this schema: (1) it does not distinguish between primitive classes (e.g. `String`) and structured classes (e.g. `Field`), and (2) it does not support *inverse* relationships. Primitives could be distinguished by adding a boolean flag to the class `Class`, but we find it more natural to introduce a structural distinction between the concept of a `Class` and a `Primitive`. These are two specific cases of the general notation of a `Type`. Representing this concept in the schema requires introducing a type hierarchy, or inheritance, into the schema.

4.2 The Ensō Schema Schema

Figure 7 defines a condensed version of the Schema schema used in Ensō. It is similar to the minimalist schema, but Ensō introduces several new concepts:

1. A class definition can include a list of superclasses, written `<` *class*. In this schema, `Primitive` and `Class` are subclasses of `Type`, and `Type` is used in where `Class` was used in the minimalist schema. Multiple inheritance is allowed but no two ancestors may define a field with the same name.
2. A field type include an *inverse* field, written *type* `/` *field*. The inverse field must be a field in the class given by the type. The schema introduces three inverses: the `schema` for a class is the schema that it belongs to, the `owner` of a field is the class it belongs to, and the `inverse` of a field is the field that it is an inverse of.
3. A field can be *computed*, written `=` *expression*. A computed field cannot be assigned. A single computed field is used to implement inheritance! The `fields` of a class are computed as the union of the fields of all its superclasses, combined with the fields that are defined on the subclass. The `defined_fields` field is just the fields that are actually defined on a class. Note that fields could be overridden, depending on the interpretation of `+`.
4. A field can be *optional*, indicated by `?`, in addition to being multi-valued. Inverses and computed expressions are optional.
5. A field can be marked as a key with `#`, which forces its value to be unique within collections of the field's class. The `name` fields in `Class` and `Field` are both marked as keys, so the schema cannot have duplicate class names, and a class cannot have duplicate fields.

6. A field can be marked as a *traversal* with `!` in front of its name. Traversal fields delineate a distinguished minimum spanning tree called a *spine*. Spines provide a standardized way to view the object graph as a tree, avoiding inconveniences such as returning different result for different implementations of depth-first search. Additionally, the spine is used to a derive unique, canonical address for each object in the graph by tracing its path from the root. In practise, traversal fields often correspond to *composition*, or 'is a part of', relationships, but they do not necessarily have to be.
7. Orderedness in many-valued fields is implicitly defined by whether the type of the field is keyed. By default, collections containing keyed objects are unordered hash tables while unkeyed objects are indexed arrays.

Additional properties are added to `Field` and `Class` to represent these new capabilities. There are many possible self-describing schemas, and Ensō does not stipulate that one must be used over another.

4.3 The Ensō Data Manager

Figure 8 defines the data manager used in Ensō, called a *factory*, and Figure 9 contains the managed object it creates. `Factory` has only one method, `_make`, which it uses to build a `ManagedObject`. Managed data object are dynamically created based on the schema provided and generally indistinguishable from ordinary Ruby objects. If desired, the factory can completely replace class definitions without methods in Ruby.

`ManagedObject` uses its two methods, `_get` and `_set`, to manipulate the underlying data. These methods have additional checks and controls built in and can be overridden for even more. `ManagedObject` also provides convenience methods that aliases the familiar "dot" operator to `_get` and `_set`. This avoids the name clashes and scoping issues with `method_missing`. The snippets shown here are stripped down versions of our actual implementation, in particular we omitted code listings for `ManyField` and `ManyIndexedField`, which handle collections and can themselves be overridden by other data managers. We also left out the details for the `notify` method, which is a call to an Observer used to enforce inverse constraints.

In order to support stacking of data managers, the constructor of `Factory` needs to accept a dictionary of initialization parameters for overriding data managers that require different inputs, such as `InitRecord`. Note also that even though `Factory` refers to `ManagedObject` by name here, in the actual implementation it uses the PROTOTYPE PATTERN [7] so that data managers can modify their private copy of `ManagedObject` without polluting the shared copy.

Unlike in earlier examples, `Factory` and `ManagedObject` are built using modules rather than classes. This allows a stack of data managers to be constructed modularly without pre-defining the inheritance chain. Ruby modules implement mixin inheritance, a feature that is also present in languages like Smalltalk, Python and Scala. Normal inheritance do not work as the inheritance chain is not predefined in each data manager. Languages that do not support mixins, such as Java, can use the DECORATOR PATTERN [7] if all data managers share the same interface. Refering to our earlier examples from Section 3, `InitRecord` and `LockableRecord` can be re-written as decorators on the basic record, but that will not allow the `_observe` method in `ObserverRecord` to be overridden. Alternatively, data managers can also be implemented using some form of meta-programming or by explicitly passing around a `this` reference, depending on how much boilerplate is tolerated.

```
class Factory
  module FactoryBase
    def initialize(schema, params={})
      @schema = schema
      #create object methods
      schema.classes.each do |c|
        _create_methods(c.name)
      end
    end

    def _make(name, *inits)
      ManagedObject.new(@schema.classes[name], *inits)
    end
  end
  include FactoryBase

  def _create_methods(name)
    define_singleton_method(name) do |*inits|
      _make(name, *inits)
    end
  end
end
```

Figure 8. Data manager for any schema

4.3.1 Bootstrapping

The Schema schema is itself Managed Data with its own data manager checking for consistency and handling updates. Bootstrapping is needed to jumpstart this process.

Figure 10 summarizes the relationships between the different levels of schemas and data managers. At the lowest level, data objects such as points are described by a schema, the Point schema. This schema is managed by a data manager capable of initialization, allowing points to be created with starting values. The Point schema is in turn described by the Schema schema. The Schema schema is self-describing, following the spirit of modularity, we bootstrap the Schema schema from the minimal bootstrap schema that has only classes and fields. This minimal bootstrap schema is necessarily self-describing as it must manage itself, and

```
class ManagedObject
  module MObjectBase
    attr_reader :schema_class
    def initialize(schema_class, *initializers)
      @schema_class = schema_class; @values = {}
      schema_class.fields.each do |field|
        init = initializers.shift #shift pops the leftmost element of an array
        if !field.many
          @values[field.name] = init
        else
          #create the appropriate collection
          if (key = ClassKey(field.type))
            @values[field.name] = ManyIndexedField.new(key.name, self, field)
          else
            @values[field.name] = ManyField.new(self, field)
          end
          if !init.nil?
            init.each {|x| @values[field.name] << x} #initialize values
          end
        end

        #create convenience accessor methods for this field
        _create_methods(field.name)
      end
    end

    def _get(name)
      field = @schema_class.fields[name];
      raise "Accessing non-existant field '#{name}'" unless field
      return field.computed ? instance_eval(field.computed) : @values[name]
    end

    def _set(name, new)
      field = @schema_class.fields[name]
      raise "Assign to invalid field '#{name}'" unless field
      raise "Can't assign field '#{name}'" if field.computed || field.many
      check_type(field.type, new) #check_type not shown in this snippet
      if @values[name] != new
        notify(name, new) #notify observers that this field has been changed
        @values[name] = new
      end
    end
  end
  include MObjectBase

  def _create_methods(name)
    define_singleton_method(name) { _get(name) }
    define_singleton_method(name+"=") {|new| _set(name, new) }
  end
end
```

Figure 9. Managed object implements dynamic object conforming to schema class

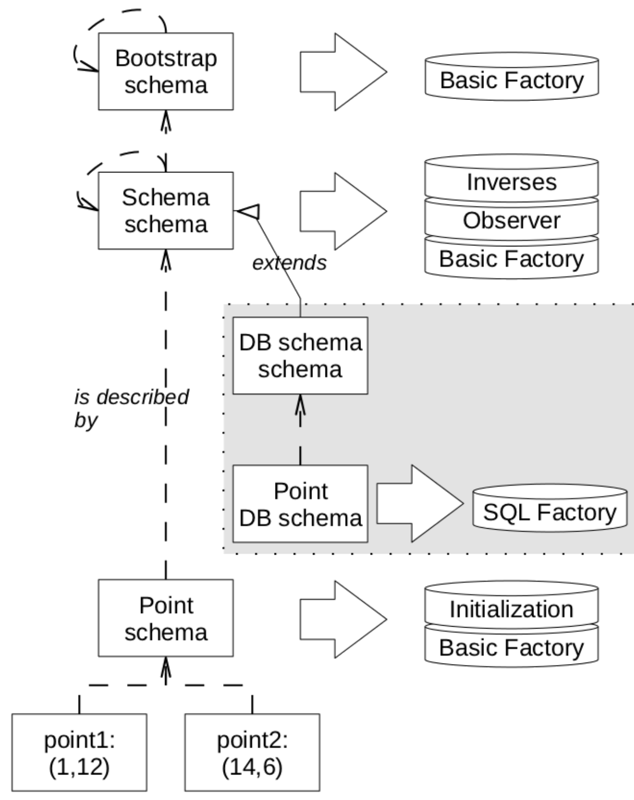

Figure 10. Bootstrapping in Ensō

it possesses a simplistic data manager that only allow updating. It is also hardcoded. Bootstrapping from a minimal schema allows us to customize even the Schema schema in very fundamental ways.

This is not the only path the schema can take, however. Some data managers may require additional information from their schema. For instance, a data manager with support for relational databases will require the schema to provide a mapping from classes and fields to table and column names, as well as additional information on indices and keys. In the diagram, Database schema extends Schema schema so that DB Point schema can provide it with the relevant fields. Note that even though it is different from Point schema, DB Point schema is still able to describe Point objects, so it is possible to migrate them from one data manager to another.

5. Case Study: EnsōWeb

We have used Managed Data to build EnsōWeb, a web development framework. EnsōWeb loosely follows the Model-View-Presenter architecture and comprises a number of DSLs for expressing data models, web interfaces, and business logic such as security policies. In this section we are primarily concerned with how data models are managed. This part of EnsōWeb is analogous to ActiveRecord in Ruby on Rails [26] or Java's Hibernate [2]. The data model manager can take on a few different roles:

- At its core, the data manager is expected to perform standard data modeling functions such as enforcing inverses and cardinality constraints. This basic manager is very similar to the Ensō data manager presented in Section 4.2.
- The data manager may optionally be required to implement low-level security. In turn, security failures may need to be logged.
- The data manager needs to support different backends for persistence depending on the requirements of the application. Possible options include in-memory, SQL, and XML databases.
- Versioning may be needed for backup and recovery.

For the most parts, the data manager in EnsōWeb provide similar services to Hibernate and Active Records, the chief difference being that Managed Data, as used in EnsōWeb, allow data management to be specifically tailored for an application (e.g. adding security constraints where none existed before), while ActiveRecord and Hibernate both only allow configuration within a fixed framework (e.g. defining a different set of validation rules). Managed Data also takes a modular approach to data management, allowing different predefined managers to be selected and composed.

Note that *all* of EnsoWeb is built around Managed Data, so models for web pages, logging, access control all have their own data managers performing generic tasks like versioning, merging, and validation, although they are not shown here.

5.1 Security

The optional security module extends the basic data manager and administers *policy-based access control* [5]. Given the state of the system, the current logged in user or his role, the operation requested (i.e. create, read, update, or delete) and object to be operated on, it returns a boolean permission and an optional error message. Internally, its security policy is a set of user-defined rules that either allow or deny an operation based on a predicate. This simple scheme removes the need to explicitly declare permissions for new objects.

The managed data object can be used in the web client code without additional boilerplate and with access control transparently enforced underneath, as shown below:

```
# Policy:
#  allow read(s:Student) if s.grade != 'F'
#  allow update(s:Student{grade}) if s.section == 1

# Data:
#  {name: 'Alice', grade: 'A', section: 1}
#  {name: 'Bob', grade: 'B', section: 2}
#  {name: 'Cathy', grade: 'F', section: 1}
```

```
module SecureObject
  def _set_user(user); @user=user; end
  def _check(op, obj, field=nil)
    # returns false if operation is not allowed
    # implementation omitted for brevity
  end
  def _get(name)
    if _check("Read", self, name)
      result = super
      _check("Read", result) ? result : nil
    else
      nil
    end
  end
  def _set(name, new)
    super if _check("Update", self, name)
  end
end
class ManagedObject
  include SecureObject
end

module LoggedSecureObject
  include SecureObject
  def _check(op, obj, field=nil)
    auth = super
    # write error to log file
    log.write "Security fail: #{op} #{obj}" if !auth
    auth
  end
end
```

Figure 11. Security extension for the data manager

```
for s in students
  print s.name
end
#Output: Alice, Bob

students['Alice'].grade = 'A+'
print students['Alice'].grade
#Output: 'A+'

students['Bob'].grade = 'A+'
print students['Bob'].grade
#Output: 'B'
```

In this snippet, records without read permissions (ie students with a failing grade) will not appear when iterating over the list. Likewise, any attempts to write to records without update permissions will be silently blocked.

Low-level security at the data model is a second line of defense to augment security at the user interface level, which is still needed to control the behavior of on-screen widgets and to provide meaningful error messages. The capability to filter out unpermitted records from the result of reads, either to return a shortened list of records or a null object, and to block write attempts to unpermitted records, is built into the data manager. Figure 11 provides a snippet from the secured data manager. Logging is implemented as a dependent data manager that overrides the `_check` method from the security data manager to record failed authorization attempts.

5.2 Persistence Layer

There are several alternatives for persisting data depending on the needs of the application: the relational databases (RDBMS) used for this example, in-memory, and XML-based options are the most common choices. Even among databases, the need for configurable data management exists [24]. Domain-specific SQL dialects for streams and sensor networks, different indexing schemes like B-trees and R-trees, OLAP cubes, query language features like stored procedures, recursive queries, views, biases between read and write operations, are some of the many decisions database designers need to make. Managed Data facilitates switching between these choices by making persistence a part of the data structuring mechanism and largely independent of the client code.

Figure 13 presents a data manager that persists the managed data object with a RDBMS backend. Classes are mapped to tables and fields to columns. Many-to-many relationships have to be transformed to use a junction table. Incidentally, many-to-one relationships without an inverse also use a junction table because the target table does not have a column to serve as a back pointer. Another thing this data manager needs to do is add a `key` field to every class, because unlike in-memory objects that are identified by their memory reference, database tuples can only be uniquely identified by their primary keys.

This RDBMS data manager is an example of coupling data managers with schema extensions. Because EnsōWeb is running on top of a relational SQL database, the schema of the data model needs some way of associating classes to tables and fields to columns and specifying a name for the schema and root table. Figure 12 the additions to the schema used by the database manager. This schema extension is merged into the schema schema via *overriding union* to produce the DB schema schema, which allows schemas to specify table and column names. Note that only the schema that describes the managed data is changed, the managed data objects themselves are the same regardless of which scheme is used.

5.3 A Family of Data Managers

By selecting the relevant data managers, the programmer can dynamically construct Managed Data with the desired set of behavior. Ideally, the different data managers should be orthogonal and not interact with each other. This is true

```
class Schema
  root: str?
  root_class: Class = "@classes[@root]"

class Class
  table: str?

class Field
  column: str?

primitive str
primitive bool
```

Figure 12. DB schema extension for specifying table and field names

in many cases such as the security and versioning data managers. However in practice, data managers have dependencies on other data managers whose service they rely on. There might also be unanticipated interactions between data managers. For instance, when using a secured manager backed by a relational database, the generated key field must be made readable by all users who have access to the object, as this is the only way to identify records in the database. This has to be resolved by applying a patch that adds security permissions for the new field.

Currently, interactions between data managers in Ensō must be manually coordinated, but we are investigating how to use a feature configuration [12] language to specify valid selections of data managers.

6. Related Work

One of our goals in this work is to provide a name and a better understanding for programming practices that have been used for many years, in many different forms. These existing approaches to managed data can be classified along several dimensions.

- The first, and most important, is *how the schema is defined*. The schema is the metadata that describes the structure and behavior of the data being managed. The schema can be defined as *external data*, *program data*, or *class attributes*. External data refers any data outside the program, whether an XML file or a special schema file. Program data is dynamic data that is instantiated within the program code, for example a Lisp S-expression. Class attributes are a form of metadata that is becoming increasingly popular. The schema can also be defined by a combination of these approaches.
- The second dimension is whether the data manager is implemented by *code generation* or *interpretation*. Code generators can either generate program source code or byte-code instructions. Interpretation is characterized by a lack of explicit manipulation of code syntax. Thus creating closures does not count as code generation.
- The final dimension is whether on not the client language, which manipulates the managed data, is *statically typed*. In some cases the typing is a *hybrid* of static checking with dynamically generated type information.

```
module DBFactory
  def initialize(schema, params={})
    # create a new database
    SQLexec("create database #{schema.name};")
    SQLexec("use #{schema.name};")
    @schema = schema
    #create object tables
    schema.classes.each do |c|
      str = "create table #{c._table_name} ("
      str += "DBKey int,"
      c.fields.each do |f|
        if !f.many
          if f.type.Primitive? #single−valued prim
            str += "#{f.name} #{convert(f.type)},"
          else #single−valued reference
            str += "#{f.name} int,"
          end
        else
          if f.inverse.nil? #1−to−many (no inv)
            #create a junction table
            junction(f._field_name, c._table_name,
                  f.type._table_name)
          elsif !f.inverse.many? #1−to−many
            #do nothing, resolved by inverse
          else #many−to−many
            #create a junction table
            junction(f._field_name, c._table_name,
                  f.inverse.ownder._table_name)
          end
        end
      end
      str += ")"
      SQLexec(str)
    end
  end
end
class Factory
  include DBFactory
end
```

Figure 13. Data manager for an RDBMS backend

The following table summarizes the related work relative to these criteria:

System:	Typing	Schema	Implementation
Ensō	Dynamic	External	Interpreted
AOM	Dynamic	External	Interpreted
Ruby on Rails	Dynamic	Reflection	Interpreted
Macros	Dynamic	Internal	Compiled
MOP	Dynamic	Internal	Compiled
RDBMS	Dynamic	External	Interpreted
MorphJ	Static	Reflection	Compiled
EMF	Static	External	Compiled
Type Providers	Hybrid	*	*
IDispatch	Hybrid	*	*

The Adaptive Object-Model (AOM) Architectural Style [30] is closely related to Managed Data. This architectural style is a "reflective architecture" or "meta-architecture" because the actual code of the system does not define the behavior and properties of the domain objects and business rules that are manipulated by the system. Instead these domain objects and rules are defined by explicit metadata, which is interpreted by the system to generate the desired behavior. There is generally no static type checking of the resulting system. In some ways the Adaptive Object-Model style is more general than Managed Data, because it is described at a very high level as a pattern language and it also covers business rules and user interfaces, in addition to data management. On the other hand, the Adaptive Object-Model does not discuss issues of integration with programming languages, the representation of data schemas, or of bootstrapping, which are central to Managed Data. The Adaptive Object-Model is also presented as a technique for implementing business systems, not as a general programming or data abstraction technique.

The languages in the Lisp family, including Scheme [13], have a long tradition of supporting user-defined data abstraction mechanisms. The `defstruct` macro is a widely used and has many variations and options, including mutability, serialization, pattern matching, and initialization [22]. A variant, the `define-type` macro, supports immutable sums-of-products [17]. Despite these prominent examples, it is not clear how widespread the practice of Manage Data is in Lisp-based languages. The `defstruct` macro is generally presented as a standard feature of the language or as part of the standard library, rather than an example of a general concept that users might practice themselves.

Object-oriented extensions of Lisp are often implemented within Lisp itself using macros or other encodings [4, 21, 31]. This approach was developed in a pure form in *The Art of the Metaobject Protocol* (MOP) [16]. The principles underlying MOP are the same as those underlying Managed Data: that the programmer should be able to control the interpretation of structure and and behavior in a program. In the case of MOP, this focus is on behavior of objects and classes, while in Managed Data the focus is more narrowly defined as data management. The MOP approach is general enough to include Managed Data as a special case.

F# 3.0 [25] has recently introduced *type providers* as a new mechanism for accessing and manipulating external data. Type providers are a form of Managed Data, because the type provider defines the structure and behavior of data values that appear as native data types in F#, but are in fact virtual values that can be drawn from any source. Type providers are a very interesting example of Managed Data, because they provide semi-static access to dynamic data from a strongly-typed functional language. The access is semi-static because the types provided by the provider can change between the time the program is compiled and when it is executed.

Relational Database Managed Systems (RDBMS) can be viewed as providing a form of Managed Data, where the SQL Data Definition Language (DDL) defines a schema. The RDBMS interprets the schema to create tables, which can then be accessed in SQL queries. Managed Data imports this approach from databases into the core of a programming language. One of the key issues with RDBMS has been the problem of integration with existing programming languages.

6.1 Class-based Metadata

One common approach to managed data is to extend an existing object-oriented language with attributes, which are interpreted to add additional behaviors to the class instances. In statically typed languages, the attributes are usually processed by a compiler or code generator, while in dynamic languages the attributes can be interpreted dynamically.

Ruby on Rails [27] implements a form of Managed Data. The schema information comes from metadata that is attached to class definitions. Ruby allows the information in a class definition to be extended easily, allowing many kinds of metadata to be included. These metadata attributes are interpreted at runtime. The Ruby on Rails engine can be viewed as a kind of data manager.

Hibernate [2] implements a restricted form of Managed Data for Java, but uses byte-code manipulation and code generation rather than dynamic interpretation. Hibernate is best understood as a specific data manager that supports binding to a relational database, rather than a general system for Managed Data.

MorphJ [10] is a system for compile-time transformation of class definitions. A source class provides metadata for a user-defined generation of a new class definition. One of the advantages of MorphJ is that the transformations are statically typed. However, using a class as metadata is more restrictive than other systems which allow arbitrary schema definitions.

6.2 Proxies

Many languages support a form of *dynamic proxy* that can be used to implement Managed Data. Proxies have long ex-

isted in dynamic languages. The Data Managers in Ruby defined in this paper use a form of dynamic proxy. A similar implementation is possible in Smalltalk. More recently, JavaScript 1.8.5 has proposed a Proxy API which provides similar functionality. Proxies are also possible in statically-typed languages, including Java [6] and C#, although they generally cannot implement the full range of Managed Data features as described in this paper. Rather than wrapping an existing object, a proxy can also be used to implement a new object dynamically. The main problem with this approach is that the interfaces for the managed objects must be predefined.

The `IDispatch` interface in Microsoft COM is a powerful tool for implementing Managed Data. It provides two main operations, `GetTypeInfo` and `Invoke`. The first operation returns a (possibly dynamically generated) description of the operations that are possible on the object. The `Invoke` operation allows operations to be invoked dynamically by passing an operation identifier and a list of parameters. Visual Basic, VBScript and JScript all provide special support for invoking COM objects that implement the `IDispatch` interface. The syntax `o.m(args)` is automatically converted to a call to `Invoke`.

6.3 Data Modeling

There has also been a long history of developing high-level data modeling languages and notations. While the concept of Managed Data is independent of the particular data description language being used, it is worth noting that the Ensō Schema schema is closely related to a large body of existing work. Examples include the Semantic Data Model [9], Entity-Relationship modeling [3], and Eclipse Modeling Framework's (EMF) [23] ECore model.

EMF is toolchain for model-driven engineering (MDE) in Java for Eclipse. MDE is a programming paradigm based on definable data descriptions specialized for a specific application domain. Their data models are managed, with native support for inverse and cardinality constraints. EMF uses the data description, called an ECore model, to generate code for commonly used scaffoldings, such as persistence and logging. Model-driven engineering share our goal of reusing data management services across different data descriptions, however, in most implementations, including EMF, data management is configurable only as part of the tool, whereas in Managed Data data managers can be configured programmatically.

6.4 Aspect-Oriented Programming

Aspect-oriented programming (AOP) [15] enables compiler support for injection of code at quantified join points. While AOP is not directly related to the idea of Managed Data, it is commonly used to modularize pervasive data management mechanisms such as logging and security, similar to how we define data managers. Like AOP, our approach allow functionality to be weaved in without the need for explicitly prepared variability hooks. However, our approach is coarser in granularity, operating only at method boundaries. Also, because we are dealing exclusively with data managers, we know beforehand the set of possible join points, and thus quantification become unnecessary. Nevertheless, the Ensō interpreter framework does support universally quantified modifications like aspects in the general context.

7. Conclusion

Managed Data is a powerful approach to data abstraction that gives programmers control over the fundamental mechanisms for creation and manipulation of data. A *schema* provides a description of the structure and behavior of desired data. The schema is interpreted by a *data manager* that implements the necessary strategies for managing data. Managed Data gives programmers more degrees of freedom, while separating concerns. Programmers can change specific schemas, or the schema language, or the data managers themselves, depending on what level of functionality is needed.

While the idea of Managed Data has appeared in various forms in the past, it has not been identified and studied as a fundamental programming concept. We have analyzed previous approaches to Managed Data, and proposed a new implementation based on interpretation of external schema languages. Our approach to Managed Data is the foundation of the Ensō system. As a case study, we demonstrated how Managed Data can be used in the context of a web development framework to reuse database management and access control services across different data definitions while hiding their implementation from the client code. Both our implementation of Managed Data and the example web framework EnsōWeb are available from `http://www.cs.utexas.edu/~alexloh/enso`.

In the future, we intend to explore static typing for Managed Data. Another promising direction is improving data manager performance through partial evaluation, especially since bootstrapping introduces a significant slowdown based on our current implementation.

References

[1] Ken Arnold, James Gosling, and David Holmes. *The Java Programming Language*. Addison-Wesley Professional, 4th edition, 2005.

[2] Christian Bauer and Gavin King. *Java persistence with hibernate*. Manning Publications Co., Greenwich, CT, USA, second edition, 2006.

[3] Peter Pin-Shan Chen. The entity-relationship model - towards a unified view of data. *ACM Trans. Database Syst.*, 1(1):9–36, March 1976.

[4] Pierre Cointe. Une extension de vlisp vers les objets. In *Science of Computer Programming*, volume 4, pages 291–322. North-Holland.

[5] William R. Cook. Policy-based authorization. 2003.

[6] Oracle Corporation. Proxy (Java 2 Platform SE v1.4.2). http://docs.oracle.com/javase/1.3/docs/guide/reflection/proxy.html.

[7] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.

[8] Adele Goldberg and David Robson. *Smalltalk-80: the language and its implementation*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1983.

[9] Michael Hammer and Dennis McLeod. The semantic data model: a modelling mechanism for data base applications. In *SIGMOD '78: Proceedings of the 1978 ACM SIGMOD international conference on management of data*, pages 26–36. ACM Press, 1978.

[10] Shan Huang, David Zook, and Yannis Smaragdakis. Morphing: Safely shaping a class in the image of others. In Erik Ernst, editor, *ECOOP 2007 âĂŞ Object-Oriented Programming*, volume 4609 of *Lecture Notes in Computer Science*, pages 399–424. Springer Berlin / Heidelberg, 2007. 10.1007/978-3-540-73589-2-19.

[11] Paul Hudak, Simon Peyton Jones, Philip Wadler, Brian Boutel, Jon Fairbairn, Joseph Fasel, María M. Guzmán, Kevin Hammond, John Hughes, Thomas Johnsson, Dick Kieburtz, Rishiyur Nikhil, Will Partain, and John Peterson. Report on the programming language haskell: a non-strict, purely functional language version 1.2. *SIGPLAN Not.*, 27(5):1–164, May 1992.

[12] K. Kang, S. Cohen, J. Hess, W. Novak, and S. Peterson. Feature–Oriented Domain Analysis (FODA) Feasibility Study. Technical Report CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie Mellon University, November 1990.

[13] Richard Kelsey, William Clinger, and Jonathan Rees. Revised 5 report on the algorithmic language Scheme. *ACM SIGPLAN Notices*, 33(9), 1998.

[14] Brian W. Kernighan and Dennis M. Ritchie. *C Programming Language*. Prentice Hall, 2nd edition, 1988.

[15] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean-Marc Loingtier, and John Irwin. Aspect-oriented programming. In Mehmet Aksit and Satoshi Matsuoka, editors, *ECOOP'97*, volume 1241 of *Lecture Notes in Computer Science*, pages 220–242. Springer Berlin / Heidelberg, 1997. 10.1007/BFb0053381.

[16] Gregor Kiczales and Jim Des Rivieres. *The Art of the Metaobject Protocol*. MIT Press, Cambridge, MA, USA, 1991.

[17] Shriram Krishnamurthi. *Programming Languages: Application and Interpretation*. April 2007.

[18] Robin Milner, Mads Tofte, and David Macqueen. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1997.

[19] Martin Odersky, Lex Spoon, and Bill Venners. *Programming in Scala: A Comprehensive Step-by-Step Guide*. Artima Incorporation, USA, 2nd edition, 2011.

[20] OMG. *Meta Object Facility (MOF) Specification*. Object Management Group, 2000.

[21] Guy Steele. *Common Lisp: The Language*. Digital Press, Newton, MA, USA, 1990.

[22] Guy L. Steele, Jr. *Common LISP: the language (2nd ed.)*. Digital Press, Newton, MA, USA, 1990.

[23] Dave Steinberg, Frank Budinsky, Marcelo Paternostro, and Ed Merks. *EMF: Eclipse Modeling Framework*. Addison-Wesley Professional, second edition, 2008.

[24] Michael Stonebraker and Ugur Cetintemel. "one size fits all": An idea whose time has come and gone. In *Proceedings of the 21st International Conference on Data Engineering*, ICDE '05, pages 2–11, Washington, DC, USA, 2005. IEEE Computer Society.

[25] Don Syme, Adam Granicz, and Antonio Cisternino. *Expert F# (Expert's Voice in .Net)*.

[26] Dave Thomas, Chad Fowler, and Andy Hunt. *Programming Ruby: The Pragmatic Programmers' Guide*. Addison-Wesley Professional, second edition, 2008.

[27] Dave Thomas, David Hansson, Leon Breedt, Mike Clark, James Duncan Davidson, Justin Gehtland, and Andreas Schwarz. *Agile Web Development with Rails*. Pragmatic Bookshelf, 2006.

[28] David A. Watt, Brian A. Wichmann, and William Findlay. *Ada language and methodology*. Prentice Hall International (UK) Ltd., Hertfordshire, UK, UK, 1987.

[29] Niklaus Wirth. *Programming in MODULA-2 (3rd corrected ed.)*. Springer-Verlag New York, Inc., New York, NY, USA, 1985.

[30] Joseph W. Yoder and Ralph E. Johnson. The adaptive object-model architectural style. In *Proceedings of the IFIP 17th World Computer Congress - TC2 Stream / 3rd IEEE/IFIP Conference on Software Architecture: System Design, Development and Maintenance*, WICSA 3, pages 3–27, Deventer, The Netherlands, The Netherlands, 2002. Kluwer, B.V.

[31] Čric Tanter. *Object-Oriented Programming Languages: Application and Interpretation*. 2010.