

Better Construction with Factories

Tal Cohen and **Joseph (Yossi) Gil**
Department of Computer Science,
Technion—Israel Institute of Technology
Technion City, Haifa 32000, Israel

“The Factory-Owning Class Controls the Means of Production.”
K. Marx [14]

The polymorphic nature of object oriented programs means that client code expecting an instance of class C may use instead an instance of a class C' inheriting from C . But, in order to use such a different instance, one must create it, and in order to do so in current languages, must be familiar with the name of creating class. To break this coupling, we propose the novel notion of *factories*, which are class services (alongside methods and constructors) that manage the instance-creation step of object construction. In making the case for factories we propose a five-dimensional framework for understanding and analyzing the class notion in various programming languages.

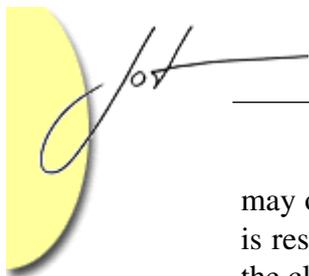
We show that factories can naturally replace the “creational” design patterns, and describe the design and implementation of a JAVA language extension supporting both *supplier-side* and *client-side* factories. Possible implementations in other languages are discussed as well.

1 INTRODUCTION

Good programming languages support, at the language level, the general principle of hiding implementation details from the client [19]. Indeed, most contemporary object oriented programming languages let, sometime even force, the programmer to hide the implementation details of methods that a class offers. An inspiring case in point is Meyer’s *principle of uniform access* [15, p.57], stating that

“All services offered by a module [i.e., a class] should be available through a uniform notation, which does not betray whether they are implemented through storage or through computation.”

This paper starts from the observation that despite the progress in language design, there is still a family of services which reveal more than they should of their implementation secrets. These services are what is known as *creation procedures* in some languages and *constructors* in others. Constructors are distinguished from the other services that a class



may offer in that the client cannot apply them to a polymorphic object; instead the client is responsible for creating such an object, and therefore must know the precise name of the class that creates it.

The polymorphic nature of classes is advertised as means for separating interface from implementation. Object oriented polymorphism means that a client may use instances of different subclasses to implement the same protocol. But, the trouble is that in order to be able to use such instances, one needs to create them somewhere, and the creation process is coupled with the name of the creating class. Breaking this coupling seems to be an intriguing chicken and egg riddle: Interface (or protocol) can be separated from implementation, but in order to select a particular implementation of a given protocol one must be familiar with at least one of these implementations.

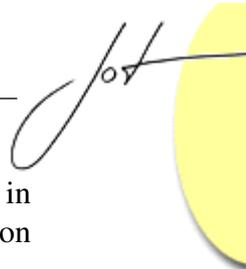
Our solution to this cyclic dilemma is by making the selection of an implementation part of the interface. In the object-oriented terminology, this means that we allow a class to offer a set of services, what we call *factories*, for generating instances of its various subclasses. Factories are first-class class members (alongside methods and constructors), but, unlike constructors, factories encapsulate instance management decisions without affecting the class's clients. Our contribution includes also a re-implementation of the JAVA compiler that supports factories; this implementation requires no changes to the JVM.

Factories directly attack the *change advertising problem*: Suppose that the implementation of a class (indeed, the internals of any software unit) is changed or specialized, but, as is the case with inheritance or dynamic aspects, that the original version still remains. Then, the fact that there was a change must be advertised to the clients that wish to enjoy its benefits. Specifically, an instance of a class C' inheriting from C can be used anywhere an instance of C is used; but clients must be aware that C' exists, and be familiar with its name and its particular repertoire of constructors, in order to create such instances.

Existing solutions to the change advertising dilemma can be found in several popular frameworks, which act outside of the programming language. This includes, for example, the J2EE [20] mechanism for obtaining instances of Enterprise JavaBeans (EJBs, [6]). Clients must not directly invoke constructors for EJBs; rather, special methods of “home objects” must be used, effectively encapsulating the creation process and providing the platform with the ability to decide an instance of which (sub)class will be generated.

Likewise, users of the Spring Application Framework¹ should only obtain instances (of any class) by using special “bean factory” objects. The need for factories is further evident from the popularity and usefulness of design patterns that strive to emulate their functionality, including ABSTRACT FACTORY, FACTORY METHOD, SINGLETON [10], and OBJECT POOL [12]. However, both the frameworks and the design patterns introduce certain restrictions that the developers must adhere to (such as never invoking constructors directly). Just like these design patterns, factories are not compelled to return a *new* class instance. In not betraying the secret whether a new instance was generated or an existing one was fetched, they can be thought as applying the principle of uniform notation to

¹<http://www.springframework.org>



instantiation. Much as with uniform access for “features” (attributes or functions) in Eiffel, factories prevent upheaval in client classes whenever an internal implementation decision of the class is changed.

More concretely, we describe the design and implementation of an extension to the Java programming language to support factories. In this extension, factories act as methods that overload the `new` operator. But, unlike `new` overloading in C++, factories are not concerned with memory allocation but rather with instance creation and specific subclass selection decisions. We offer two varieties of factories:

- *Client-side factories* help localize instantiation statements, whereby a re-implementation can be selectively injected to certain clients.
- *Supplier-side factories* provide classes with fine control over their instantiation, and help in a global advertising of a change in the implementation.

Factories enable the encapsulated implementation of the “creational” design patterns listed above, either for all clients (using supplier-side factories) or for specific ones (using client-side factories). They provide a language-level solution to the change advertising dilemma, without presenting developers with any restrictions or complications.

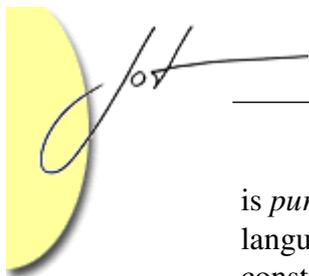
Outline Sec. 2 starts by setting forth a common terminology for the discussion, and tries to unify some of the different perspectives offered in the literature to the class concept. Using this terminology, Sections 3 and 4 expand on the motivation, by highlighting certain limitations of constructors. Factories are the subject of Sec. 5, which describes their Java syntax and some of the applications. This section also shows how factories support many classical design patterns. Sec. 6 describes how coupling between classes can be decreased using factories, and Sec. 7 describes the notion of client-side factories. Finally, Sec. 8 discusses the extension of the factories idea to other programming languages and concludes.

2 TERMINOLOGY

There are many ways in which people perceive the notion of class: as a “*repository for behavior associated with an object*” [2, p.13], a “*unit of software decomposition*” and a “*type*” [15, pp.170–171], a “*tool for creating new types*” [21, p.223], a “*group [of objects]*” [13, p.50]², a “*set of objects that share a common structure and a common behavior*” [1, p.93], etc. This section tries to unify these perspectives and propose a terminology (a conceptual framework if you will) for comparing and understanding the notion of a class in different programming languages.

We distinguish five, not entirely orthogonal, dimensions of class analysis: *commonality*, *encapsulation*, *morphability*, *binding*, and *purpose*. The most interesting dimension

²but also a “*template for several objects . . . [a description of] how these objects are structured internally*”



is *purpose*, by which we identify, for each syntactical element of a class, a programming-language purpose. In Sec. 3 we shall argue that, judged by these dimension of evaluation, constructors make a bit of weird bird. Let us now describe in greater detail each of the five dimensions in turn.

1. Commonality. This dimension makes the distinction between *common* elements of the class notion (e.g., class variables and methods in SMALLTALK) and *particular* such elements (e.g., instance variables and methods). More precisely, an element is common if its incarnation in different instances of the class is identical; otherwise, it is particular. Thus, particular elements may be used only in association with a specific class instance. Also, common elements cannot access particular elements.

2. Encapsulation. (Also known as *Visibility*.) A class may encapsulate (i.e., set the visibility of) its elements. C++'s three visibility levels, just as JAVA's four, are orthogonal to commonality.

3. Morphability. Morphability indicates the class element's ability to obtain a shape, or be re-shaped, in a subclass. In other words, morphability pertains to the kind of changes that a subclass may apply to components of the base class in the course of inheritance.

There is a great variety in the morphability capabilities in different programming languages. For example, C++ allows a subclass to decrease the visibility of inherited members, OBERON [22] forbids overriding, JAVA sports **final** members and allows data members to be hidden [11, Sect. 8.3.3], while EIFFEL allows re-implementation of a data member as a method, and method renaming. The analysis of this variety in full is beyond the scope of this paper.

4. Binding. As the name suggest, in this dimension we make the distinction between statically-bound and dynamically-bound elements. Of course, this distinction can be made only for class elements which can be replaced or altered in a subclass. Non-**virtual** methods in C++ are famous for being statically bound.

Observe that in most languages, commonality and binding are not orthogonal. Specifically, we find that *common elements are often statically bound*. The linkage between static binding and commonality is so entrenched that common methods and fields in languages such as JAVA, C# and C++ are marked with the **static** keyword.

The phenomena can be explained by the reliance of dynamic binding on dispatching information associated with individual objects. Common elements are statically bound since they may exist even when there are no instances to the class.

5. Purpose. Classes, being a unit of software decomposition, can be subjected to Parnas's [19] classical distinction between the *interface* and *materialization* (which Parnas calls "implementation") perspectives of a software component. We say that the interface and materialization are *purposes* that the class serves as a whole, and characterize its elements by this purpose.

But, unlike the software components of the seventies, classes are instantiable. Accordingly, we break the interface of a class into two facets: the *forge* and the *type*. Similarly, we distinguish between three facets in the materialization: the *implementation* of the type,



the *mill* behind the forge, and the *mold* into which instances are cast.

More specifically, the *forge* of the class is the collection of operations that can be used to create objects; the *type* is the set of messages that these instances may receive, along with their visibility specification; and, the *implementation* is the body of the methods executed in response to these messages. There is a subtle distinction between the mill and the mold, which together realize the class's forge: The mold is the memory layout which instances of this class follow. It consists solely of field definitions. The mill is the set of constructor bodies.

To understand these terms better, consider class `Vector` from the standard JAVA library. The forge of `Vector`, depicted in Fig. 2.1(a), includes the signature of the four

```
Vector:
  public Vector();
  public Vector(Collection);
  public Vector(int);
  public Vector(int, int);
  (a) The forge.
```

<code>int</code> capacityIncrement	32 bits
<code>int</code> elementCount	32 bits
<code>Object[]</code> elementData	32 bits
Fields inherited from superclasses	
Hidden fields added by the JVM	

(b) The mold.

```
Vector<E>:
  protected int capacityIncrement;
  protected int elementCount;
  protected Object[] elementData;
  public void addElement(E);
  public int capacity();
  // ... etc.

  // From AbstractList:
  public ListIterator listIterator();
  public List sublist(int, int);
  // ... etc.

  // From AbstractCollection:
  public int size();
  public void clear();
  // ... etc.

  // From Object:
  public Object clone();
  public void wait();
  public void notify();
  public boolean equals(Object);
  // ... etc.

  // Upcast operations:
  public (AbstractList) ();
  public (AbstractCollection) ();
  public (Object) ();
  public (Serializable) ();
  public (Iterable<E>) ();
  public (Collection<E>) ();
  public (List<E>) ();
```

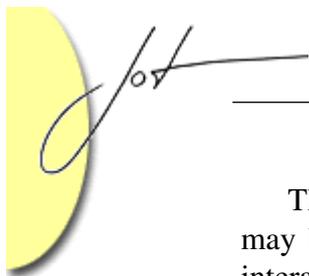
(c) The type.

Figure 2.1: The forge, type and mold of `java.lang.Vector`.

constructors provided by the class: the default constructor `Vector()`, the copy constructor `Vector(Collection)`, a variant that specifies the initial capacity (`Vector(int)`), and a variant that specifies both the initial capacity and the capacity growth increment (`Vector(int, int)`).

Fig. 2.1(c) shows the type of `Vector`, methods such as `addElement`, `capacity`, and others, as well as fields such as `capacityIncrement` and `elementData`. Superclasses also add to the type; in this case, the type of `Vector` includes methods and fields inherited from three superclasses. Each superclass and superinterface also adds an upcast operator.

We see that the type includes the signature of all non-`private` fields and methods of the class. Thus what we call type here is in fact the class's structural type, to which JAVA applies a name, making it a nominal type.



The type does not include details such as a specification of the order by which methods may be invoked, pre- and post-conditions, or other classes with which the class may interact while implementing each method. These may be thought of as the class *protocol*.

The *mold* for creating new objects is defined by the collection of all fields in this class and all of its supertypes. Specific languages or language implementations can include hidden fields in the mold, such as run-time type information, the Virtual Method Table [8] used in C++, etc. Fig. 2.1(b) presents the mold defined by class `Vector`. It includes fields defined in `Vector` as well as any fields inherited from superclasses, along with any hidden field added by the JVM.

Finally, the *implementation* is the body of the methods defined by the class or any of its superclasses, while the *mill* is the body of the constructors defined in this class.

3 CONSTRUCTOR ANOMALIES

Factories, the language extension proposed in this paper, are methods which return new class instances. Syntactically, a factory is a method which overloads the `new` operator with respect to a certain class.

In the terminology of the previous section, the signature of a factory belongs in the forge, while its body belongs in the mill. In this respect, factories are similar to constructors in mainstream object-oriented languages, the means by which a class' clients may obtain instances.

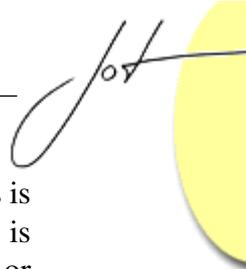
In analyzing constructors (in, e.g., JAVA or C++) with this terminology, we find that exhibit three fundamental anomalies, which underline the need for the alternative approach that factories offer:

1. Commonality. In JAVA, the syntax for creating an instance of class `MyClass` is `new MyClass()`, i.e., it refers to the class name. In contrast, in EIFFEL the syntax is `!!myInstance`, i.e., referring to a variable. The difference between the languages is not a coincidence. Constructors are anomalous in that they are *simultaneously common and particular*: common—since they are invocable without an instance; particular—since they work on an object.

This anomaly raises the dilemma of method binding inside constructor bodies. Method invocation from the mill follows a static binding scheme in C++³; in JAVA and C#, however, dynamic binding is used. Neither approach is without fault. Static binding can lead to illegal invocation of pure virtual methods. Dynamic binding prevents methods, invoked from within the mill, from assuming that all fields were properly initialized. Dynamic method binding in constructors leads, among other things, to difficulties in implementing non-nullable types, as described by Fähndrich and Leino [9]: during construction, fields of non-null types may contain null values.

2. Morphability. In examining the morphability of the five facets of a class purpose, we

³Even for **virtual** methods.



find that changes to four out of these are not arbitrary: The type definition of a subclass is an *extension* of the type definition of the superclass. Similarly, the mold of a subclass is an *extension* of the mold of the superclass. Also, the implementation can either *replace* or *extend* the implementation in the superclass, and the mill (constructor body of a subclass) must *extend* (i.e., invoke) the mill of the superclass.

In contrast, the forge of the subclass is *independent* of the forge of the superclass—the forge cannot be extended: it is not even inherited, and each class must define its own set of constructor signatures anew. The second constructor anomaly lies in the fact that the *the construction protocol is not inherited, yet, each constructor body must invoke a constructor of the base class*.

3. Binding. Third, it is mundane to see that a call to a constructor obeys a static binding scheme, and it takes just a bit of pondering to understand the difficulties that this scheme brings about. If a class C' inherits from C , then C' should be always substitutable for C . An annoying exception is made by constructor invocation sites in client code; these have to be manually fixed in switching from C to C' .

The Gang of Four [10, p.24] place this predicament first in their list of causes for redesign, saying: “Specifying a class name when you create an object commits you to a particular implementation instead of a particular interface”.

Interestingly, in EIFFEL, although it has a strict dynamic binding policy, and although creation methods can be overridden, and although creation syntax is similar to method invocation, it is still the case that creation instructions such as `!!x.make` are statically bound.

4 STAGES OF OBJECT CREATION

Fig. 4.1 demonstrates another issue with constructors. The figure depicts abstract class `Baby` whose constructor announces the baby’s birth, and concrete class `NamedBaby` inheriting from it. Method `announce` is refined in `NamedBaby`, extending the announcement with details about the newborn’s gender and name.

A client who has new baby boy named “John”, may then write

```
NamedBaby myBoy = new NamedBaby("John", true);
```

and be surprised by the printout “New baby: Her name is null”, which is explained by the announcement being made before the subclass’s fields are initialized. This lack of crisp separation between field initialization and the rest of the construction code, can even result in runtime exceptions, e.g., if `NamedBaby.announce` invokes `name.length()`.

C++ is not much better: The C++ equivalent of Fig. 4.1 would print partial (albeit more sensible) output, “New baby:”. Also, C++ would produce a runtime error if `announce()` is made abstract in class `Baby`.

This example motivates our distinction between three conceptual steps in an instance’s birth process (later we shall argue that the separation between these is better served by factories): (a) *Creation*, in which the object’s actual type is selected, memory is allocated and structured by the mold; (b) *Initialization*, in which fields are set to their initial values;

```

1 abstract class Baby {
2     public Baby() { announce(); }
3     public void announce() { System.out.print("New baby: "); }
4 }

6 class NamedBaby extends Baby {
7     String name; boolean isBoy;

9     public NamedBaby(String name, boolean isBoy) {
10        this.name = name; this.isBoy = isBoy;
11    }

13    public void announce() {
14        super.announce();
15        System.out.println((isBoy ? "His" : "Her") + " name is " + name);
16    }
17 }

```

Figure 4.1: Interwoven initialization and setup in JAVA constructors.

and (c) *Setup*, in which the mill is executed.

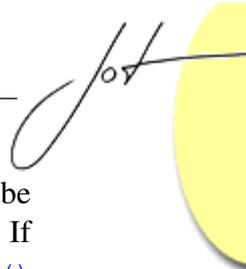
These three steps correspond exactly to steps C1, C2 and C4 in the effects of a creation instruction $!!x$ in EIFFEL [15, p.237]. The missing step, C3, is the attachment of the newly created object to the reference variable x ; however, in languages such as JAVA and C++ the invocation of a constructor is an *expression* rather than a statement, and can be performed without assigning the result to a variable. (EIFFEL also supports the invocation of a creation procedure as an expression [7, Sec. 8.20.18], in which case step C3 is absent.)

The initialization step is realized in C++ by what is called the initialization list (written just after the constructor's signature). In JAVA and C# it is expressed using initializer values (or defaults) for fields, whereas the instance initializer block and the constructor bodies perform the setup. In EIFFEL, it is the assignment of standard default values to fields. As the example shows, however, initialization with default values is insufficient. Developers should be able to initialize all fields, across all levels of inheritance (i.e., complete step (b)) before setup code is being executed (step (c), the announcement in our example); *initialization* and *setup* should be unwoven. We further note that none of these languages provides the developer with control over the *creation* step.

Note that overloading the **new** operator in C++ grants us control over memory allocation, but not over the kind of object to be created, nor the decision if a new object has to be created at all.

We argue that good design of elaborate software systems often requires intervention in the creation step. Indeed, there are a number of successful design patterns, including ABSTRACT FACTORY, FACTORY METHOD, SINGLETON, and OBJECT POOL, which address precisely this need. The control that these “creational patterns” grant the programmer over the creation step is achieved by replacing constructor signatures from the forge facet with a different, statically-bound, common method (e.g., `getInstance`)⁴.

⁴Such methods are sometimes called *factory methods*. While serving a similar purpose, they are different



Unfortunately, in contrast with most other patterns, the creational patterns cannot be implemented in OO languages without revealing implementation details to the client: If class `T` is implemented as a `SINGLETON`, then clients of this class cannot write `new T()` and expect the correct instance to be returned; rather, they must be aware of the nonstandard creation mechanism, in violation of the uniform access principle. As a result, if a class evolves during development so that the new version employs (e.g.) an instance pool, all clients must be updated to use the `getInstance` method rather than the constructors; the use of creational patterns cannot be encapsulated as part of the class implementation.

Creational patterns often collide with inheritance. To enforce the use of a `getInstance` method and prevent accidental direct access to the constructors, all constructors can be made `private`, with the undesired implication that the class cannot be subclassed. The alternative of defining the constructor as `protected`, is problematic in `JAVA`, since such constructors are visible to all classes in the same package.

Worse still, since method `getInstance` must be shared, it cannot be overridden in subclasses: If C' is a subclass of C , then the expression $C'.getInstance()$ is valid—but returns an instance of C ! This happens because `getInstance` is technically part of the type, while conceptually being part of the forge.

We shall see that factories enable a clear-cut separation between creation and initialization and setup, and allow for proper encapsulation of the creation step.

5 FACTORIES

Class `STemplate` in Fig. 5.1 shows how the `SINGLETON` design pattern can be realized by overriding `new` with the factory defined in lines 4–7. This factory is invoked whenever the

```
1 class STemplate {
2     private static STemplate instance = null;

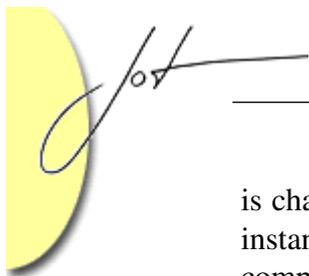
4     public static new() {
5         if (instance == null) instance = this();
6         return instance;
7     }

9     STemplate() { /* ... setup code ... */ }
10 }
```

Figure 5.1: A Singleton defined using a factory.

expression `new STemplate()` is evaluated, in class `STemplate` or any of its clients. Note that the factory is declared `static`, which stresses that it binds statically, and that (unlike constructors) it has no implicit `this` parameter. Examining the factory body we see that it always returns the same instance of the class. Thus, clients need not be explicitly aware of `STemplate` being a singleton, and will not be affected if this implementation decision

than our notion of *factories*.



is changed. (In the specific case of the SINGLETON design pattern, clients can compare instances to realize that only one exists. Other patterns, such as INSTANCE POOL, can be completely invisible to clients.)

A factory must either return a valid object of the class, or throw an exception. (Should the factory's return value be `null`, a `NullPointerException` results.)

Suppose that C' is a subclass of C . Then, a factory of C can return an instance of C' . This can be done by invoking any method which returns an instance of C' , including a factory of C' —e.g., by a statement such as `return new C'(...)`. If the factory however chooses to create an instance of class C , then it should invoke the constructor; yet writing `new C(...)` (e.g., `new STemplate()` in the example) would recurse infinitely. Instead, the factory invokes the class constructor directly with the expression `this(...)` (line 5 in the example).

We chose to overload the keyword `this`, particularly, its use for invoking a constructor. No ambiguity arises: In constructors, the function call `this(...)` occurring in the first line can substitute the mandated call to `super` with a call to a different constructor in the same class (as in standard JAVA). Such a call does not create an instance, nor does it return a value, and it must appear only as the very first step in the constructor body.

In a factory, `this(...)` stands for a call to a constructor of the class. The call creates a new instance and returns a value; it may occur multiple times (or not at all), and in any location inside the factory body. The factory can choose to return the value generated by such a call. (In the case of the `STemplate` class, the value is cached to a static field, which is then returned.)

The constructor can only be called from a factory in the same class; any use of `new C(...)`, either from outside class C or from inside it, will invoke a factory rather than a constructor.

While there are many different solutions to the specific issue of singletons, (e.g., declaring an object—rather than a class—in SCALA [18], or using prototype-based languages, such as CECIL [5]), the factory solution is not specific to singletons, and can be used for any creational design pattern. More examples will be presented in the sequel.

As usual with overloading, a factory may have parameters, which are matched against the actual parameters in the creation expression. A parameterized factory could be used for, e.g., implementing the FLYWEIGHT pattern: To do so, the factory returns, if possible, an existing object from its pool, and only creates an instance if no such object exists.

Like constructors, factories are not inherited. Had class C' inherited a factory `new()` from its superclass C , then the expression `new C'()` might yield an instance of C , contrary to common sense. Thus, the problem of `C'.getInstance()` yielding an instance of C , described in Sec. 3, does not occur with factories.

In contrast, when factories are employed, the expression `new C()` can yield an instance of C' , since a subclass is always substitutable for its superclass.

Factories also allow developers to separate the initialization and setup stages of object construction. The mixup of Fig. 4.1 is resolved by the factory based implementation in Fig. 5.2, in which the call `new NamedBaby("John", true)` yields the expected



“New baby: His name is John” output.

```
1 abstract class Baby {
2     public Baby() { } //No fields that require initialization
3     public void announce() { System.out.print("New baby: "); }
4 }

6 class NamedBaby extends Baby {
7     String name; boolean isBoy;

9     public NamedBaby(String name, boolean isBoy) {
10        this.name = name; this.isBoy = isBoy; //Initialization
11    }

13    public void announce() {
14        super.announce();
15        System.out.println((isBoy ? "His" : "Her") + " name is " + name);
16    }

18    public new(String name, boolean isBoy) {
19        NamedBaby result = this(name, isBoy); //Construction
20        result.announce(); //Setup
21        return result;
22    }
23 }
```

Figure 5.2: Re-implementation of Fig. 4.1 with factories

The implementation in the figure adheres to the simple rule that fields are initialized in constructors, and other setup is carried out by the factory. In particular the announcement of the birth is made in the factory of `NamedBaby` (line 20).

Automatically Generated Factories

A definition of a factory with a certain signature hides the constructor with the same signature. Such hidden constructors can only be invoked from the factory of a class, regardless of their access level. Let us now deal with the dual situation, i.e., a constructor without a factory. Backward compatibility of our extension is achieved by the following perspective: An expression of the form `new S(...)` is always implemented by a factory whose signature matches the actual parameters. This can be either a user-defined factory, or an *automatically generated factory* (AGFa). The automatic generation of factories is governed by:

The AGFa Rule: Let c be a constructor with a signature σ in a non-abstract class S . Then, either (a) S has an explicit factory with signature σ , or (b) it has static AGFa with signature σ , which invokes c .

Fig. 5.3 shows an example of the AGFa rule. The class defined in Fig. 5.3(a) has a factory with no parameters. It also has a two-parameters constructor, with no matching factory. Fig. 5.3(b) shows the AGFa that the compiler (internally) injects into the class.

```
class Complex {
    public static final Complex origin = new Complex(0,0);
    public Complex(double x, double y) { /* instance setup ... */ }
    public static new() { return origin; }
}
```

(a) A class in which the no-args factory returns a fixed instance.

```
public static new(double x, double y) { return this(x,y); }
```

(b) The factory added to the class by the AGFa rule.

Figure 5.3: A class (a) with a constructor and its AGFa (b).

Recall that in plain JAVA, instances of abstract classes cannot be created, even though such classes have constructors. The following argument uses the AGFa rule to explain this: *Instances can only be created by a `new` expression, which must have a matching factory. However, by the AGFa rule, abstract classes in plain JAVA do not have factories.*

Conversely, if an abstract class S_a does define factories, then you can write `new S_a (...)` in your code. Fig. 5.4 shows an abstract class, `ScrollBar`, with a factory. This example is modelled after the famous example [10, p.87] of the ABSTRACT FACTORY design pattern. The code in the figure improves on the original implementation of the design pattern, in that the client is not aware that an abstract factory stands behind the scenes of the simple call `new ScrollBar()`. (As we shall see later, the internal implementation of the widget factory class itself can also be improved with factories.)

```
public abstract class ScrollBar {
    public static new() {
        WidgetFactory f = WidgetFactory.currentFactory();
        return f.CreateScrollBar(); // Select concrete subclass
    }
    // ... rest of the class omitted
}
```

Figure 5.4: An abstract class with a factory.

As shown in Fig. 5.5, interfaces may also have factories. The figure shows an interface, `DirectoryEntry`, whose factory makes it possible to obtain an instance of either of two implementing classes, `Folder` and `File`, depending on the parameter value.

6 BETTER DECOUPLING WITH FACTORIES

The use of factories in interfaces can eliminate coupling between client code and library code. Consider, for example, the JAVA collection libraries. The standard library designers



```
public interface DirectoryEntry {
    public static new(String name) {
        if (FileSystem.isDirectory(name)) return new Folder(name);
        return new File(name);
    }
    // ... rest of the interface omitted
}
```

Figure 5.5: An interface with a factory.

require, in very strong words, that interface types (like `List` and `Set`) will be used for method arguments:

“... it is of paramount importance that you declare the relevant parameter type to be one of the collection interface types. Never use an implementation type.”

– [3, p.526]; emphasis in the original.

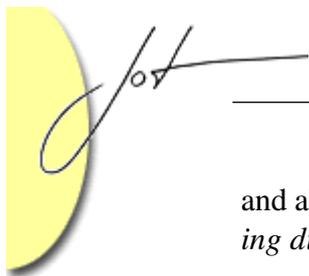
Similar recommendations apply to return types, field types, etc., all in spirit of Canning et al.’s original suggestions for separating the type and class notions using interfaces [4]. The coupling of client code to concrete implementation is indeed reduced by following this recommendation. But, such a coupling still remains, particularly at the point where a client is required to create an object.

Interfaces with factories can eliminate this coupling. In the case of the `List` interface, clients can generate instances of some default implementation by writing (say) `new List()`. The factory can choose the proper concrete implementation, possibly based on hints provided by the client. Fig. 6.1 provides an example factory that can be used by the `List` class in JAVA’s collections framework. Should new and improved implementa-

```
public interface List {
    /**
     * @param synch indicates if a thread-safe list is needed
     * @param randomAccess indicates if O(1) element access is needed
     */
    public static new(boolean synch, boolean randomAccess) {
        if (synch) {
            if (randomAccess) return new Vector();
            return Collections.synchronizedList(new LinkedList());
        }
        // Else, synchronization is not needed.
        if (randomAccess) return new ArrayList();
        return new LinkedList();
    }
    // ... rest of the interface omitted.
}
```

Figure 6.1: One possible factory for the `List` interface.

tions appear in future versions of the JAVA class libraries, this factory can be upgraded,



and all clients will immediately benefit from the change. This solves the *change advertising dilemma* for new implementations of interfaces.

We would like to draw attention to the fact that following the recommendation of using interfaces rather than classes as method parameters, may in some situations *increase the burden* on clients rather than reducing it. Consider the learning effort of a user in search of a specific service in a software library. Suppose that this service is provided by a method m in an interface I . Then, before m can be invoked, the user must search for all the different implementation of I , say classes C_1, C_2, C_3, \dots , study them, and choose which of these to instantiate in order to generate an instance of I . Further, suppose that m takes a parameter of type interface I' . Then, the user must also search for all implementations of I' , say classes C'_1, C'_2, C'_3, \dots , study them all and choose the one appropriate for instantiation prior to invoking method m . If the constructor of the chosen class expects a third interface parameter I'' , then, the user must further search for implementations $C''_1, C''_2, C''_3, \dots$ of I'' , etc.

A small example is method `Security.getProviders` in the JAVA standard library taking a `Map` as a parameter. In this parameter, the user can provide a set of selection criteria. Before the method may be used, even for testing or experimentation, the programmer must create an object representing such a test, and to do so, choose an implementation of the `Map` interface—but there are no less than seventeen such implementations in version 1.5 of the JDK.

Another example is method `JPanel.setBorder()` from the Swing GUI libraries, which expects a parameter of the `Border` interface. In order to use this method, the client must spend time in studying the different implementations of this class, only to realize that yet a third class, `BorderFactory`, should be used to generate instances. With factories, the functionality of `BorderFactory` can be embedded in `Border`.

Searches for implementations of a given interface is usually not easy: implementations may be done by various different vendors, the list may change over time, and the selection between these may require a hefty learning effort. Interfaces (and abstract classes) with factories can therefore simplify the adoption of new, unfamiliar classes. Sometimes such a search is inevitable, but in many cases, it can be saved if the interface itself provides a reasonable implementation.

Writing a unit test code for a class whose methods take interface parameters is greatly simplified if these interfaces give ready-made instantiations. It is even conceivable that interfaces provide a stub implementation just for this purpose. For example, the standard JAVA interface `Runnable` can provide a stub implementation (perhaps defined as an inner class) in which the `run()` method does nothing.

7 CLIENT-SIDE FACTORIES

All examples so far defined factories in the same class on which the overload takes place. Factories of this sort are called *supplier-side factories*. It is also possible to define *client-*



side factories, as demonstrated in Fig. 7.1.

```

1 class Bank {
2   public static new Account(Customer c) {
3     if (c.hasBadHistory()) return new LimitedAccount(c);
4                               // LimitedAccount is a subclass of Account
5     return Account.new Account(c);
6   }
7   // ... rest of the class omitted
8 }

```

Figure 7.1: A client-side factory for `Accounts` in class `Bank`.

Line 2 in the figure starts the definition of a factory. Unlike the previous examples, this definition specifies the returned type. The semantics is that the definition overloads `new` when used for creating `Account` objects from within class `Bank`. It is invoked in the evaluation of an expression of the form `new Account(c)` (where `c` is of type `Customer` or any of its subclasses) in this context. This factory chooses an appropriate kind of `Account` depending on the particular business rules used by the enclosing class.

Unlike supplier-side factories, client-side factories *are* inherited by subclasses. Therefore, the factory from Fig. 7.1 will also be used for evaluating expressions of the form `new Account(c)` in subclasses of `Bank`.

This client-side factory can be used by other classes as well, by writing `Bank.new Account(...)`, or, after making a static `import` of class `Bank`, by simply writing `new Account(...)`.

Fig. 7.2 shows an implementation of the ABSTRACT FACTORY pattern with static binding. Classes `MotifWidgetFactory` and `PMWidgetFactory` each overload the `new`

```

class MotifWidgetFactory {
  public new ScrollBar() { return new MotifScrollBar(); }
  public new Window() { return new MotifWindow(); }
  // ... factories for other widget classes ...
}

class PMWidgetFactory {
  public new ScrollBar() { return new PMScrollBar(); }
  public new Window() { return new PMWindow(); }
  // ... factories for other widget classes ...
}

```

Figure 7.2: Widget-factory classes defined using client-side factories.

operator of all the GUI widgets. A client wishing to use Motif, may write `import static MotifWidgetFactory.*`. This may be changed later to `import static PMWidgetFactory.*`, should the GUI library need replacing.

The full semantics of a `new` call can now be explained as follows: Whenever a class is used in a `new` expression, its supplier-side factories enjoy an implicit `import static`. A client-side factory in scope can override this import.

The abstract widget factory example we have just described suffers from the problem that switching from Motif to PM requires a change to the client's `import static` statements. But there may be many such statements, in many source files. The remedy is to simply define an empty class,

```
class WidgetFactory extends PMWidgetFactory {}
```

and statically import it in all clients. This will direct all widget factory calls to `PMWidgetFactory`. The GUI can now be globally replaced with a single change, specifically replacing `WidgetFactory`'s superclass.

Dynamically Bound Factories

The above `WidgetFactory` can be thought of as a statically-bound implementation of the ABSTRACT FACTORY pattern, in that the decision on the concrete implementation is made at compile time. To make a dynamically-bound widget factory, we need *dynamically-bound factories*. These are defined, as the name suggests, without the `static` keyword. Fig. 7.3 shows how such factories can be used in the classical implementation of the ABSTRACT FACTORY design pattern.

```
public abstract class WidgetFactory {
    public abstract new ScrollBar();
    public abstract new Window();
    // ... and other widgets.

    private static WidgetFactory f;
    public static new() {
        if (f != null) return f;
        if (GUI.isMotif()) return f = new MotifFactory();
        if (GUI.isPM()) return f = new PMFactory();
        //... etc.
    }
}
```

(a) The abstract widget factory class

```
class MotifWidgetFactory extends WidgetFactory {
    public new ScrollBar() { return new MotifScrollBar(); }
    public new Window() { return new MotifWindow(); }
    //...
}

class PMWidgetFactory extends WidgetFactory {
    public new ScrollBar() { return new PMScrollBar(); }
    public new Window() { return new PMWindow(); }
    //...
}
```

(b) Two concrete widget factory subclasses

Figure 7.3: Using non-`static` factories to implement a dynamically bound abstract factory class.

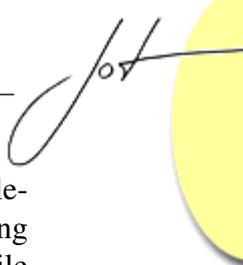


Fig. 7.3(a) shows the abstract factory, while Fig. 7.3(b) shows two concrete implementations. The factories of the widgets are all non-**static** and obey a dynamic binding scheme. Also worthy of note is the factory of this abstract class itself, which (while realizing the SINGLETON design pattern) determines at runtime the correct GUI library.

Fig. 7.4 shows how dynamically-bound factories can be used to implement the FACTORY METHOD pattern (also known as VIRTUAL CONSTRUCTOR). The code in this figure implements the classic example (from [10, p.107]) of an abstract `Application` class, bound to an abstract `Document` class. Each concrete subclass of `Application` can bind itself to a concrete subclass of `Document`, by overriding the dynamically-bound factories. The resulting code is very similar to the original GoF example, except that the

```
abstract class Application {
    List<Document> docs;
    protected abstract new Document ();

    public void newDocument () {
        // Handles the File|New menu option
        doc = new Document (); docs.add(doc); doc.open ();
    }
    // ... rest of the class omitted
}
```

(a) The abstract `Application` class

```
class MyApplication extends Application {
    protected new Document () {
        return new MyDocumentType (); // A concrete subtype
    }
    // ... rest of the class omitted
}
```

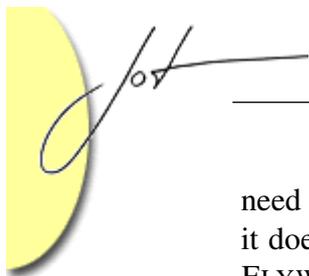
(b) One possible concrete subclass

Figure 7.4: Implementing pattern FACTORY METHOD with dynamically bound factories.

`newDocument` method uses ordinary construction syntax (implemented using our notion of a factory) rather than the nonstandard “factory method” dictated by the pattern.

Syntactically, the invocation of a dynamically-bound factory defined in class C for objects of class S is written as $c.\text{new } S(\dots)$, where c is an instance of class C . The prefix “ $c.$ ” can be dropped for code inside class C (so it is replaced with **this**).

It is not a coincidence that this looks very much like the JAVA syntax for creating an instance of a dynamic inner class: $c.\text{new } I(\dots)$, where c is an instance of the containing class (possibly **this**) and I is the inner class’s name. The constructor of a (non-**static**) inner class in JAVA is a method of the containing class, and not of the class it constructs—just like a client-side factory is a member of the containing class, and not of its target class. In fact, Nystrom, Chong and Myers [16] have shown that if the concept of inner classes is extended (using *nested inheritance*), most of the need for the FACTORY METHOD design pattern disappears. But while nested inheritance has many distinct advantages with regard to code modularity and the creation of extensible software systems, it only solves the



need for factory methods for classes defined inside the same module as their client. Also, it does not remove the need for instance-management patterns like INSTANCE POOL or FLYWEIGHT.

8 DISCUSSION

Factories may be a minor perturbation to language syntax, but they are of benefit to language designers and programmers alike. We implemented factories as a JAVA extension using the Polyglot [17] extensible compiler framework (v. 2.0a4). This took approximately two workdays of a single programmer.

In our implementation, supplier-side factories (both explicit and AGFa) are realized as methods named `new` in the container class. The return type of `new` is the containing class itself.

Client-side factories are stored in the client, and are named `newclassname`, where *classname* is the fully-qualified target class name, with every dot replaced by `dot`. For example, the factory for `Accounts` in class `Bank` (Fig. 7.1) is realized as a method called `newcomdotbankdotAccount` (assuming `Account`'s fully qualified name is `com.bank.Account`). The return type of client-side factories is the target type (e.g., `Account`).

Any use of `new` is replaced by the proper method invocation, wrapped in a test that ensures a non-`null` value is returned (and throws an exception otherwise).

The addition of factories to interfaces is less straightforward, since interfaces in JAVA cannot contain any concrete methods. Instead, our extension, synthesizes an `static` inner class (called `$NewHolder$`) for the interface, and places factories in this class.

The implementation generates bytecode that can be used on any JAVA virtual machine. As discussed in Sec. 5, the introduction of AGFas implies that JAVA-with-factories is fully compatible with existing JAVA source code. However, the code generated by our compiler assumes that all instantiated classes have been compiled using the same compiler, and thus have supplier-side factories (either explicit or AGFa). If factories are integrated into the JAVA language, full backwards compatibility with existing, pre-compiled classes can be achieved by having the class-loader (rather than the compiler itself) add any required AGFa to each class. This will work equally well for old and newly-compiled classes.

Clearly, the notion of factories is not limited to JAVA alone. It is not so difficult to approximate supplier-side—but not client-side—factories in SMALLTALK, by overriding the `new` class method. Adding factories to C# seems rather straightforward, but it might take some cunning to add them to C++, since the language introduces two obstacles: First, C++ intrinsic overloading of the `new` operator, is focused on the memory allocation problem rather than on instance generation. One possible solution is to introduce a new keyword, such as `factory`, to the language. Declarations for `factory new` can then exist alongside those for `operator new`. Such definitions can include both supplier-side factories (no explicit return type) and client-side ones (with a specific return type). Client



calls to `new` will then be redirected to the factory, and should the factory decide to create a new instance, the `new` operator will be used for memory allocation (as before).

The second obstacle is due to C++ value semantics. The compiler must know the space requirements of class instances allocated e.g., on the stack, but this is not possible with factories. A simple solution is that classes with factories are restricted to reference semantics representation only.

The Eiffel language presents a different challenge for introducing factories. Unlike constructors in C++ or Java, creation procedures in Eiffel are named. The advantage of this approach is that the distinction between the different kinds of objects that may be created is not by the kind of arguments, but rather through a meaningful name.

In terms of syntax design, the problem is that we must find a way, other than a special name, to distinguish between factory methods (which have no object to work on), and methods and creation procedures (which start their work with a system-supplied object.)

We propose to the integration of factories into Eiffel by introducing a new part to the Eiffel class declaration, alongside `features`, `creates`, etc. The part is called `factory`, and it may be included only in non-`expanded` types. Following Eiffel's accessibility rules, a class may provide different factories to different client classes by qualifying the `factory` part with a type list. Supplier-side factories have the return type "`like Current`"; any other return type indicates a client-side factory.

A subclass may re-classify a creation procedure as a factory (or vice-versa) when overriding it, and in particular, the default creation procedure, `default_create` (defined in the root class `ANY`) may be changed to a factory by any class that so desires. Following the principle of uniform access, clients that include a creation instruction (or a creation expression) employ the exact same syntax regardless of whether a creation procedure or a factory is being used. The syntax `!!x.make` is used by clients to obtain an instance, regardless of whether `make` is a creation procedure or a factory. Interestingly, the distinct name for each factory and creation method implies that this extension maintains backwards compatibility with existing code, without resorting to automatically-generated factories (AGFas).

Fig. 8.1 shows an Eiffel version of the singleton class from Fig. 5.1. This class re-classifies `default_create` as a factory, so clients can use the creation instruction `!!x` (for a variable `x` of type `S_TEMPLATE`) to obtain the shared instance.

As we can see from the figure (line 7), no special syntax is needed to create an instance from inside the factory (the equivalent of the special `this()` call in the Java version): Since a class may include both creation procedures and factories, each with distinct names, there is no risk of undesired recursion. Whenever a new instance is required, the factory simply calls a (possibly private) creation procedure.

```

1 class
2   S_TEMPLATE

4 factory -- obtain an instance
5   default_create: like Current is
6     once
7       !!Result.instance
8     end

10 create {NONE} -- private instance creation mechanism
11   instance is
12     do
13       -- initialize fields, etc.
14     end

16 end -- class S_TEMPLATE

```

Figure 8.1: A singleton defined in Eiffel using a factory.

REFERENCES

- [1] G. Booch. *Object Oriented Design with Applications*. The Benjamin/Cummings Publishing Company, Inc., 1991.
- [2] T. A. Budd. *An Introduction to Object-Oriented Programming*. Addison-Wesley Publishing Company, first ed., 1991.
- [3] M. Campione, K. Walrath, and A. Huml. *The Java Tutorial: A Short Course on the Basics*. Addison-Wesley Publishing Company, 2000.
- [4] P. S. Canning, W. R. Cook, W. L. Hill, and W. G. Olthoff. Interfaces for strongly-typed object-oriented programming. In *OOPSLA'89*.
- [5] C. Chambers. The Cecil language, specification and rationale. Technical Report TR-93-03-05, University of Washington, Seattle, 1993.
- [6] L. G. DeMichiel, L. U. Yalçinalp, and S. Krishnan. Enterprise JavaBeans specification, version 2.0. <http://java.sun.com/j2ee/>, 2001.
- [7] ECMA International. *Standard ECMA-367: Eiffel Analysis, Design, and Programming Language*. ECMA International, 2005.
- [8] M. A. Ellis and B. Stroustrup. *The Annotated C++ Reference Manual*. Addison-Wesley Publishing Company, 1994.
- [9] M. Fähndrich and K. R. M. Leino. Declaring and checking non-null types in an object-oriented language. In *OOPSLA'03*.
- [10] E. Gamma, R. Helm, R. E. Johnson, and J. M. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Publishing Company, 1995.



- [11] J. Gosling, B. Joy, G. L. J. Steele, and G. Bracha. *The Java Language Specification*. Addison-Wesley Publishing Company, third ed., 2005.
- [12] M. Grand. *Patterns in Java, Volume 1*. John Wiley & Sons, 1998.
- [13] I. Jacobson. *Object-Oriented Software Engineering - A Use Case Driven Approach*. Addison-Wesley Publishing Company, first ed., 1992.
- [14] K. Marx. *Das Kapital: Kritik der politischen Oekonomie*. Otto Meissner, 1867.
- [15] B. Meyer. *Object-Oriented Software Construction*. Prentice-Hall, Englewood Cliffs, New Jersey 07632, second ed., 1997.
- [16] N. Nystrom, S. Chong, and A. C. Myers. Scalable extensibility via nested inheritance. In *OOPSLA'04*.
- [17] N. Nystrom, M. R. Clarkson, and A. C. Myers. Polyglot: An extensible compiler framework for Java. In *CC'03*.
- [18] M. Odersky, P. Altherr, V. Cremet, B. Emir, S. Maneth, S. Micheloud, N. Mihaylov, M. Schinz, E. Stenman, and M. Zenger. An overview of the Scala programming language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland, 2004.
- [19] D. L. Parnas. Information distribution aspects of design methodology. In *IFIP'71*.
- [20] B. Shannon. *Java 2 Platform Enterprise Edition Specification, v1.4*. Sun Microsystems Inc., 2003. http://java.sun.com/j2ee/j2ee-1_4-fr-spec.pdf.
- [21] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley Publishing Company, third ed., 1997.
- [22] N. Wirth and M. Reiser. *Programming in Oberon—Steps Beyond Pascal and Modula*. Addison-Wesley Publishing Company, 1992.

ABOUT THE AUTHORS

Tal Cohen is a computer science Ph.D., currently employed in Google's engineering center in Haifa. This research was done during his Ph.D. studies in the Technion in Haifa, Israel. He can be reached at ctal@cs.technion.ac.il. See also <http://tal.forum2.org/cv>.

Yossi Gil is on the faculty of the department of computer science at the Technion and head of the software and systems development laboratory there. His research interests include software engineering, programming languages and database systems. All of Gil's academic titles were conferred by the Hebrew University of Jerusalem: B.Sc. (summa cum laude) in mathematics and physics, M.Sc. (summa cum laude) in computer science, and Ph.D. in computer science. He can be reached at yogi@cs.technion.ac.il.