

Amnon H. Eden, Tom Mens. "Measuring Software Flexibility." *IEE Software*, Vol. 153, No. 3 (Jun. 2006), pp. 113–126. London, UK: The Institution of Engineering and Technology.

Copyright notice: This paper is a postprint of a paper submitted to and accepted for publication in *IEE Software* and is subject to IET copyright. The copy of record is available at IET Digital Library.

Also appeared as: Amnon H. Eden, Tom Mens. "Evolution Complexity." Department of Computer Science, University of Essex, Technical report CSM-424, ISSN 1744-8050, 2005.

Measuring Software Flexibility

Amnon H. Eden ⁽¹⁾, Tom Mens ⁽²⁾

Abstract. *Flexibility* has been recognized as a desirable quality of software since the earliest days of software engineering. Classic and contemporary software design literature suggest that particular implementations are more flexible than others but stop short of suggesting objective criteria for quantifying such claims. To measure software flexibility in precise terms, we introduce the notion of evolution complexity and demonstrate how it can be used to measure and compare the flexibility of—

1. programming paradigms (object-oriented vs. procedural programs)
2. architectural styles (Shared Data, Pipes and Filters, and Abstract Data Type)
3. design patterns (Visitor and the Abstract Factory)

We also demonstrate how evolution complexity can be used to choose the most flexible design policy. We conclude with experimental results corroborating our claims.

Key terms: Software evolution, software flexibility, science of software design, software metrics.

Related terms: Design patterns, software architecture, computational complexity, software complexity.

(1) Department of Computer Science, University of Essex, Colchester, Essex CO4 3SQ, United Kingdom, and Center For Inquiry, Amherst, New York, USA.

(2) Software Engineering Lab, Université de Mons-Hainaut, Av. du champ de Mars 6, 7000 Mons, Belgium.

1 Introduction

Rapid technological developments pervade every aspect of daily life, having a direct effect on the software we use. Every element of the software's operational environment [15] is in a state of constant flux: Frequent changes in the hardware, operating system, cooperating software, and client's expectations are motivated by performance improvements, bug-fixes, security breaches, and attempts to assemble synergistically ever more sophisticated software systems. Classic and contemporary literature in software design recognize the central role of flexibility in software design and implementation. Structured design, modular design, object-oriented design, software architecture, design patterns, and component-based software engineering, among others, seek to maximize flexibility. Textbooks about software design emphasize the flexibility of particular choices, thereby implying the superiority of the design policy they advocate. But despite the progress made since the earliest days of software engineering, from the 'software crisis' [21] through 'software's chronic crisis' [10], evolution (formerly 'maintenance') of industrial software systems has remained unpredictable and notoriously expensive, often exceeding the cost of the development phase [1]. Flexibility has therefore become a central concern in software design and in many related aspects in software engineering research. The purpose of this paper is to contribute to our understanding of this quality and to examine ways in which it can be quantified ('metrics').

We begin by observing two problems in the current notion of 'software flexibility'. The first is the absence of reliable metrics thereof. No formal criteria for flexibility have so far been offered ⁽³⁾ and no metrics for quantifying it are known to us.

The second problem we observe in the current notion of flexibility is that it is misconceived as an absolute quality. Such misconception is reflected for example in IEEE's definition of software flexibility:

Flexibility: The ease with which a system or component can be modified for use in applications or environments other than those for which it was specifically designed [12]

We find this surprising because hardly any artefact is 'flexible' in absolute terms. The RAM capacity of a desktop computer, for example, can be expanded only if the hardware and the operating system were specifically designed to accommodate for such changes ⁽⁴⁾. The same applies to articles of clothing (such as trousers and skirts) which cannot be expanded in size unless specific provisions were made for this explicit

⁽³⁾ with the exception of the works discussed in Section 5.

⁽⁴⁾ and even then, memory can only be extended to the extent to which it was specifically designed for.

purpose. Other examples can be drawn from consumer appliances to urban architecture, each of which is only flexible towards a particular class of changes. However, claims on the flexibility of particular programming paradigms, architectural styles, and design patterns (‘design policies’) are rarely qualified. For example, in his seminal paper on modular decomposition, Parnas [24] claimed that the Abstract Data Type architecture is ‘flexible’. Twenty years later Garlan, Kaiser and Notkin [9] qualified this claim by demonstrating a specific class of changes towards which the same architectural style is not flexible. (In Section 3.1, we corroborate and make precise Garlan et. al’s observation.) More recently, Parnas observes that predicting the class of changes is the key to understanding software flexibility ⁽⁵⁾:

To apply this principle [“design for change”], one begins by trying to characterise the changes that are likely to occur over the “lifetime” of the product. Since we cannot predict the actual changes, the predictions will be about classes of changes. ([25], p. 281.)

Similar reservations regarding the flexibility of design patterns were also made by Gamma et. al:

Each design pattern lets some aspect of system structure vary independently of other aspects, thereby making a system more robust **to a particular kind of change**. ([8], p. 24, our emphasis)

1.1 Contributions

Computational complexity [30] allows us to measure how the cost of computation grows as a function of the size of the input. The big Oh notation indicates an approximation to the number of steps taken by the algorithm. For example, an algorithm traversing a matrix of size $Columns \times Rows$, which can also be expressed in pseudocode as follows:

```

procedure traverse(matrix: aMatrix, integer: columns, integer: rows) is
  for (integer c := 1) until c == columns
    for (integer r := 1) until r == rows
      traverse(aMatrix[c,r]);

```

has the complexity $\mathcal{O}(Columns \times Rows)$. This means that the number of steps in traversing a matrix of size $Columns \times Rows$ grows proportionately to the product of the number of columns and the number of rows. In our example, this means that traversing a matrix of size 4×5 will take about twice as much as traversing a matrix of size 2×5 ⁽⁶⁾.

⁽⁵⁾ Or as Parnas puts it, the key to understanding “software aging”.

⁽⁶⁾ It is customary to omit constants in the big Oh notation.

We borrow this notion of measurement for estimating software flexibility. We say that a is more flexible than b towards a particular evolution step because the number of changes required for a is smaller than the number of changes required for b . For example, Shaw and Garlan [27] say that the Shared Data architectural style is inflexible towards changes in the data storage format because such a change “will affect almost all of the modules”. In contrast, the Abstract Data architectural style is flexible towards the same change because it “can be changed in individual modules without affecting others”. Evolution complexity allows us to establish this claim (Section 3.1), showing that the complexity of evolving Shared Data is linear in the number of modules in the implementation $\mathcal{O}(|SharedData|)$ whereas the complexity of same evolution step in Abstract Data Type is constant $\mathcal{O}(1)$. Thus, the complexity of an evolution step measures how inflexible is the implementation towards a particular class of changes: The less changes are required, the more flexible it is.

In this paper, we formulate the notion of evolution complexity to achieve the following:

- (a) To provide means for quantifying flexibility.
- (b) To corroborate and make precise informal claims on the flexibility of particular programming paradigms, architectural styles and design patterns.
- (c) To provide means for choosing the most flexible design policy.

We define ‘evolution step’ as the unit of evolution with relation to a particular change in the implementation, and demonstrate the following:

- ◆ When a particular evolution step is evidently ‘easy’, we demonstrate that the complexity of executing this particular evolution step is fixed (constant complexity);
- ◆ When a different evolution step is evidently ‘difficult’, we demonstrate that the complexity of executing this evolution step grows as a function of the size of the implementation (linear complexity).

1.2 Outline

The remainder of this paper is organized as follows: In Section 2, we formulate our terminology and illustrate it with an example. In Section 3, we examine metrics for quantifying the flexibility of four recognized design policies. In Section 4 we describe a small-scale experiment that corroborates the predictions evolution complexity makes. In Section 5 we discuss related work, including the notion of ‘software complexity’. In Section 6 we summarize the conclusions that can be drawn from our discussion.

2 Definitions

In this section, we define the terminology used in our discussion and motivate our definitions with an example. We formulate the term ‘evolution step’ as a unit of measuring the cost of the evolution process using the notion of ‘evolution cost metric’.

We use the term **design policy** with reference to any sensible collection of software design decisions, such as programming paradigms, architectural styles and design patterns. Design policies allow us to expand the scope of our discussion from a particular implementation into a class of such. We also use the terms ‘implementation’ and ‘program’ interchangeably with reference to a unit of source code.

2.1 Evolution step

Consider the operational environment [16] of a particular program: It consists of the hardware, the operating system, other programs, and other sources of input to the program. We abstract the operational environment as a set of functional and non-functional requirements, which we call the **problem**. Software evolution can be described as the process during which changes occur in the problem, which entail changes in the **implementation**. To distinguish between changes in the problem and changes in the implementation we refer to the former as **shifts** and to the latter as **adjustments**, jointly represented as an ‘evolution step’. This terminology is illustrated in Figure 1.

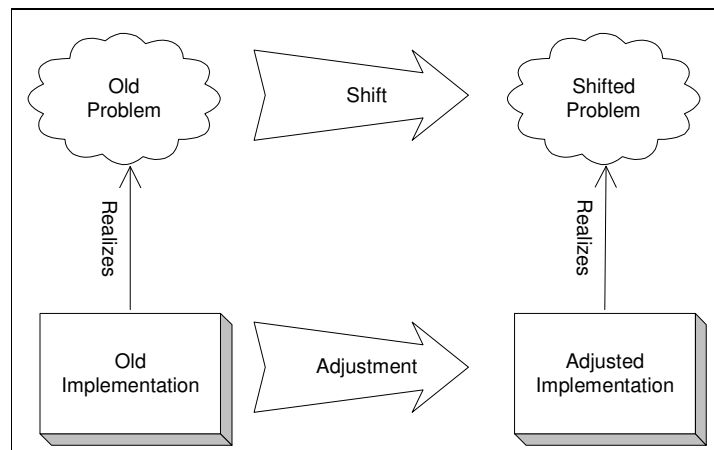


Figure 1. An evolution step

Let us represent the set of problems as \mathbb{P} , the set of implementations as \mathbb{I} (7). A step in the process of software evolution can be represented as a mapping of the combination of the old problem $p_{old} \in \mathbb{P}$, the shifted problem $p_{shifted} \in \mathbb{P}$, and the old implementation $i_{old} \in \mathbb{I}$ into the adjusted implementation $i_{adjusted} \in \mathbb{I}$. This mapping can thus be represented as the *evolution function*, a mathematical function \mathcal{E} which maps each tuple $\langle p_{old}, p_{shifted}, i_{old} \rangle$ to the adjusted implementation $i_{adjusted}$:

Definition 1. **Evolution function** is a functional relation \mathcal{E} such that

$$\mathcal{E} : \mathbb{P} \times \mathbb{P} \times \mathbb{I} \rightarrow \mathbb{I}$$

and $\mathcal{E}(p_{old}, p_{shifted}, i_{old}) = i_{adjusted}$ only if i_{old} realizes p_{old} and $i_{adjusted}$ realizes $p_{shifted}$.

Evolution step is a pair $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, \mathcal{E}(p_{old}, p_{shifted}, i_{old}) \rangle$

2.2 Example: Java's Collection interface

Let us illustrate our terminology using the problem of representing data structures and their implementations using the `Collection` interface in the `java.util` package (8). The design policy that guided the authors of `java.util` is also known as (part of) the Iterator design pattern [8]. These definitions are given to demonstrate how to make precise claims on the flexibility of particular implementations, but the reader may skip the technical detail included in this subsection without adversely affecting the readability of the remainder of this paper.

Problem

Provide several concrete data structures and operations thereon. We may unpack a concrete instance of this problem as follows:

(7) in this context, the constitution of which in terms of Zermelo-Fraenkel's set theory is of no relevant consequence.

(8) Accompanying the Java™ class library, version 1.4.2.

- ♦ $p_{DS/OP} \triangleq$ Provide a client with data structures

$$DS = \{LinkedList, ArrayList\}$$

and a collection of operations shared by them

$$OP = \{add, contains, size\}$$

To make this example concrete, we use the data structures and their respective operations to represent a movie cast and the audition process.

Implementation

The package `java.util` provides the uniform interface to all data structures, the `Collection` interface. The purpose of this interface is to hide the particulars of each data structure and thereby provide a fixed point of reference for clients of the data structures. Flexibility is supported by allowing the programmer to replace one data structure with another without requiring any adjustments to the remainder of the implementation. Table 1 depicts such an implementation, designated *Collection*, which uses the data structures to represent the collection of actors (`movieCast`) as determined by the director. Figure 2 depicts the UML class diagram of *Collection*.

Table 1. A Java program illustrating the use of `Collection`, which makes class `Director` flexible towards adjustments in `movieCast`.

```
public class Director {           // Flexible towards adjustments in movieCast:
    private Collection movieCast = new ArrayList;
    public void audition(Actor candidate) {
        if (CanAct(candidate))
            movieCast.add(candidate);
    }
    ...
}
```

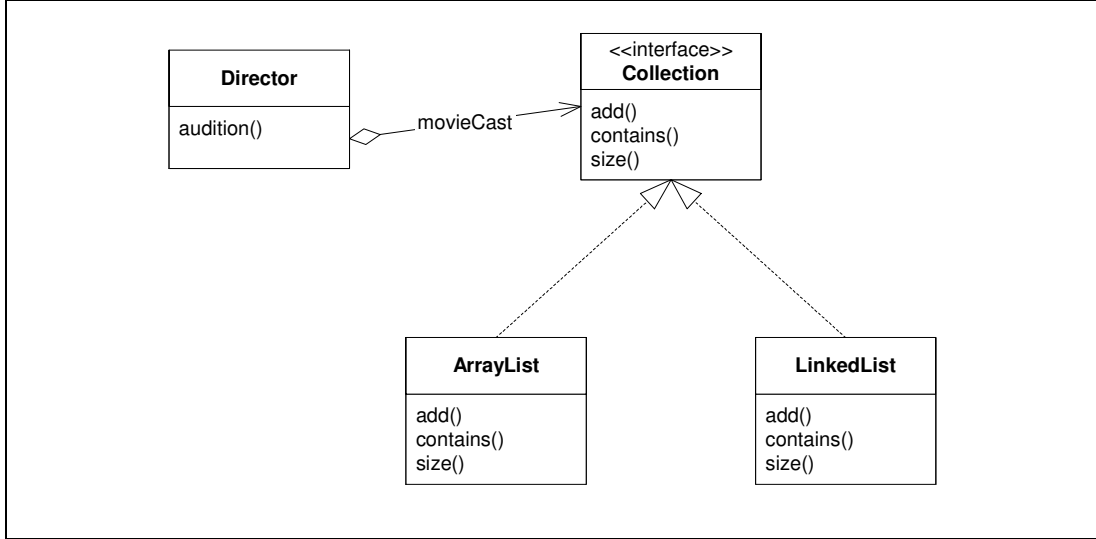


Figure 2. UML class diagram of the program in Table 1 and related modules.

Encoding evolution steps

Consider two possible shifts to the problem $p_{DS/OP}$:

- ◆ `movieCast` should be represented as a `LinkedList`
- ◆ Operation `remove` over `movieCast` is also required (for example because an actor has resigned)

We may encode the new ('shifted') problems as follows:

- ◆ $p_{Shifted-DS} \triangleq$ Use `LinkedList` where `ArrayList` was used;
- ◆ $p_{Shifted-OP} \triangleq$ Add the operation `remove` over the movie's cast.

We may encode the two revised implementations respectively:

- ◆ $Collection_{Adjusted-DS} \triangleq$ `Collection` with `LinkedList` where `ArrayList` was used
- ◆ $Collection_{Adjusted-OP} \triangleq$ `Collection` with method `remove` in `Collection` and interface and in all data structures implementing it

The scenarios following each one of these combinations of shifts and adjustments can be represented as an evolution step as follows:

- ◆ Evolve the data structure:
 $\varepsilon_{DS} \triangleq \langle \langle p_{DS/OP}, p_{Shifted-DS}, Collection \rangle, Collection_{Adjusted-DS} \rangle$
- ◆ Evolve the operation:
 $\varepsilon_{OP} \triangleq \langle \langle p_{DS/OP}, p_{Shifted-OP}, Collection \rangle, Collection_{Adjusted-OP} \rangle$

These evolution steps can also be expressed in terms of the evolution function (Definition 1):

- ♦ $\mathcal{E}(p_{DS/OP}, p_{Shift-DS}, Collection) = Collection_{Adjusted-DS}$
- ♦ $\mathcal{E}(p_{DS/OP}, p_{Shift-OP}, Collection) = Collection_{Adjusted-OP}$

Detailed encoding of shifts, adjustments, and evolution steps can be used to define the class of changes concerned but is not essential for the understanding of notion evolution complexity.

2.3 Evolution cost metrics

We may summarize at this point our intuition on the flexibility of the Java `Collection` interface as follows:

- ♦ It is flexible towards changing the data structure (ε_{DS}) in the sense that the cost of executing ε_{DS} does not depend on the number of data structures;
- ♦ It is inflexible towards adding an operation (ε_{OP}) in the sense that cost of executing ε_{OP} grows with the number of the data structures, $|DS|$.

In other words, the flexibility of `Collection` interface towards each evolution step ε can be quantified in terms of the ‘cost’ of the evolution process. In reality, however, the cost of an evolution process, whether measured in terms of time, equipment, labour, and other resource required for executing the evolution step depends on the size of the programming team, the relevant experience of its members and the software development tools at their disposal, the clarity and currency of the documentation, and on a complex range of other cognitive, social and economic factors that are difficult to measure. Given the complexity of this problem we ask, How can we quantify the cost of the evolution process?

A hint as to the solution was given in a talk by Parnas on the subject of Software Aging:

“...one organises the software so that the items that are most likely to change are ‘confined’ to a small amount of code, so that if those things do change, only a small amount of code would be affected.” ([25], p. 281)

In other words, flexibility (measured in terms of the cost of the evolution process) is directly linked to the amount of code that is affected. Thus, a first approximation to measuring the cost of executing an evolution step ε is given by the evolution cost metric which counts the number of modules that are affected by ε .

Let us also make the simplifying assumption that the costs of adding, removing, or changing each modular unit commensurate. Thus, the evolution cost metric we propose is obtained by calculating the number of modules that were added, removed,

or adjusted ⁽⁹⁾ as a result of the evolution. This number is obtained by calculating the symmetric set difference between the sets of classes in the old (i_{old}) vs. the adjusted ($i_{adjusted}$) implementations. Formally:

Definition 2. The **evolution cost metric** $\mathcal{C}_{Classes}^1$ measures the cost of executing evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$ in terms of the number of classes affected by the adjustment:

$$\mathcal{C}_{Classes}^1(\varepsilon) \triangleq |\Delta Classes(i_{old}, i_{adjusted})|$$

where $\Delta Classes(i_{old}, i_{adjusted})$ designates the symmetric set difference between the sets of classes in each program, namely: $(Classes(i_{old}) \setminus Classes(i_{shifted})) \cup (Classes(i_{shifted}) \setminus Classes(i_{old}))$.

Let us demonstrate how $\mathcal{C}_{Classes}^1$ can be used to compare the cost of changing the data structure (ε_{DS}) with the cost of adding an operation (ε_{OP}). For this purpose, imagine that a computer program is carrying out the evolution process: such a (meta-)program adjusts the old implementation. The complexity of this (meta-)program can be calculated using the evolution cost metric $\mathcal{C}_{Classes}^1$ (Definition 2) as follows:

$$\begin{aligned} \mathcal{O}(\mathcal{C}_{Classes}^1(\varepsilon_{DS})) &= \\ \mathcal{O}(\# \text{classes affected by replacing ArrayList with LinkedList}) &= \\ \mathcal{O}(1) \end{aligned}$$

and

$$\begin{aligned} \mathcal{O}(\mathcal{C}_{Classes}^1(\varepsilon_{OP})) &= \\ \mathcal{O}(\# \text{classes affected by adding operation remove to Collection}) &= \\ \mathcal{O}(|DS|) \end{aligned}$$

The complexity of each evolution step is summarized in Table 2.

Table 2. $\mathcal{C}_{Classes}^1$ -complexity of evolving the Collection interface

| Evolution step Design policy | Change data structure (ε_{DS}) | Add operation (ε_{OP}) |
|---|---|---|
| Collection interface | $\mathcal{O}(1)$ | $\mathcal{O}(DS)$ ⁽¹⁰⁾ |

⁽⁹⁾ The notion “a module that was changed” can be fixed in a number of ways, the simplest of which is to count the number of modules that require recompilation after the change.

This establishes the intuitions expressed in Section 2.3 on the flexibility of the Collection design policy towards each evolution step as follows:

- ◆ We say that the `Collection` interface is flexible towards changing the data structure because the complexity of executing ε_{DS} is fixed and independent of the size of the implementation: $\mathcal{O}(\mathcal{C}_{Classes}^I(\varepsilon_{DS}))=1$. This means, for example, that executing this evolution step in an implementation with 20 data structures cost about as much as executing the same in an implementation with only 10 data structures.
- ◆ We also say that the `Collection` interface is inflexible towards adding an operation because the complexity of executing ε_{OP} grows with the number of data structures in the implementation: $\mathcal{O}(\mathcal{C}_{Classes}^I(\varepsilon_{OP}))=|DS|$ (linear complexity). This means, for example, that the cost of executing ε_{OP} in an implementation with 20 data structures is twice as hard much as the cost of as executing the same in an implementation with only 10 data structures.

This result also corroborates the claim that the Collection design policy is not flexible in absolute terms. Instead, we demonstrated that it is flexible towards one class of evolution steps (adding data structures) and inflexible towards another (adding an operation.)

2.4 Evolution complexity

We suggested that the flexibility of the `Collection` implementation towards each evolution step ε can be quantified by measuring the complexity of the metaprogramming process that executes ε . This approach can be best described as an application of the principle of ‘software evolution is software too’.

In his award-winning paper “Software Processes are Software Too” [22], Osterweil suggests to treat human and executable processes uniformly. This is helpful because “manual and automated processes are both executed, they both address requirements that need to be understood, both benefit from being modelled by a variety of ... models, both must evolve guided by measurement, and so forth.” [23]. Evolution complexity can thus be defined as the application of Osterweil’s principle to software evolution: The evolution process is conceptualized as a computational process executing a particular evolution step (Definition 1), specifically as a metaprogram ⁽¹¹⁾ which simulates the respective evolution process. In particular, we can use computa-

⁽¹⁰⁾ The actual number of steps is $|DS| + 1$, whose complexity is equivalent to linear complexity by the abstraction conveyed by the conventional big Oh notation [30].

⁽¹¹⁾ A metaprogram is a program that manipulates other programs. The canonical example is a compiler. Metaprogramming functions are built into many programming languages.

tional complexity [30] to measure the complexity of the evolution process with relation to a particular evolution step. Formally:

Definition 3. The complexity of the evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$, also the **evolution complexity** of ε , is the complexity of a process executing ε .

Specifically, this means that the complexity of evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$ is the complexity of the process which adjusts i_{old} into $i_{adjusted}$ such that $i_{adjusted} = \mathcal{E}(p_{old}, p_{shifted}, i_{old})$.

Computational complexity [30] is concerned with estimating how the ‘cost’ of the computational processes grows as a function of the size of the problem. Similarly, evolution complexity is concerned with estimating how the cost of the evolution process grows as a function of the size of the implementation. The complexity of a computational process is measured by breaking it into commensurable sub-steps. Similarly, the complexity of the evolution process can be measured by breaking it into commensurable sub-steps. The metric $\mathcal{C}_{Classes}^I$ demonstrates how this number can be calculated: the complexity of adjusting the implementation towards shifts in the data structure is constant $\mathcal{O}(1)$ whereas the complexity of adjusting the implementation towards shifts in the operations is linear in the number of data structures $\mathcal{O}(|DS|)$.

Caveat

We identify the following possible misconceptions of Definition 3:

- ◆ We do not claim that software evolution can, should, or eventually will be fully automated. Rather, we argue that conceptualizing the evolution process *as if* it were automated is a metaphor that is useful for the purpose of quantifying the complexity of the evolution process. ⁽¹²⁾
- ◆ *Complexity* is a measure of growth, not an absolute value. Computational complexity does not measure the *actual* time that computational processes require but how it *grows*. Likewise, evolution complexity does not measure the *actual* cost of the evolution processes requires but how it *grows*.

⁽¹²⁾ Similar misconceptions of Osterweil’s general principle are discussed in [23].

2.5 More evolution cost metrics

The evolution cost metric $\mathcal{C}_{Classes}^1$ is inadequate in at least three situations:

- (a) *When the evolution of different modules do not commensurate.* During late phases in the software lifecycle, it may become evident that the cost of evolving one module does not commensurate with the evolution of another. For example, the cost of changing a ‘small’ class (for example, a class defined in two lines of code) does not commensurate with the evolution of a ‘large’ class (for example, a class defined in 200 lines of code).
- (b) *When modules were not yet implemented.* During even earlier phases of architectural design, the implementation has not been completed yet.
- (c) *When the programming language does not support classes.* Evolution cost metrics should also be defined for programming languages other than class-based.

We conclude that the metric must accommodate for varying degrees of modular granularity, as well as for varying degrees of information on each module. This leads us to define the generalized evolution cost metric:

Definition 4. The **generalized evolution cost metric** $\mathcal{C}_{Modules}^\mu$ measures the cost of executing evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$ in terms of the *software complexity* $\mu(m)$ of each module m affected by the adjustment:

$$\mathcal{C}_{Modules}^\mu(\varepsilon) \triangleq \sum_{m \in \Delta Modules(i_{old}, i_{adjusted})} \mu(m)$$

where μ is any *software complexity* metric and $\Delta Modules(i_{old}, i_{adjusted})$ designates the symmetric set-difference between the set of modules in i_{old} and the set of modules in $i_{adjusted}$.

The generalized metric is parameterised by the variables *Modules* and μ :

- ◆ *Modules* represents any notion of module that is appropriate for the circumstances, such as class, procedure, method, and package. For example, $\mathcal{C}_{Packages}^1$ is a coarse metric counting the steps in the evolution process in terms of number of *packages* in the implementation.
- ◆ μ represents any *software complexity metric* that is meaningful with relation to a particular module m . For example, fixing $\mu = LoC$ (Lines of Code) yields the metric $\mathcal{C}_{Modules}^{LoC}$ which incorporates information about the size of each module, whereas fixing $\mu = CC$ (Cyclomatic Complexity) incorporates information about the control flow.

Naturally, Definition 4 gives rise to any number of evolution cost metrics. Various such metrics are demonstrated in the remainder of this paper. More generally, an evolution cost metric \mathcal{C} is a function which associates an evolution step ε with a non-negative number which attempts to quantify the cost of the process of executing ε . Formally:

$$\mathcal{C} : \mathbb{P} \times \mathbb{P} \times \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{R}^+$$

Discussion

In proposing the metric $\mathcal{C}_{Classes}^I$ we assumed that the costs of (1) adding, (2) removing, and (3) changing a module commensurate. How can this assumption be justified? Experience teaches that the cost of evolving a module often exceeds that of developing it in the first place ⁽¹³⁾. [1] It is therefore not unreasonable to assume that removing and changing a module costs about as much as adding a new module.

Alternatively, consider a situation where only adding a module is significant. For example, let us assume that—

- ♦ the cost of evolution effort is proportional to the module’s size, which is determined by the simplistic count of Lines of Code (LoC), and that
- ♦ removing and changing a module are inconsequential.

Under such circumstances, the appropriate evolution cost metric is $\mathcal{C}_{Modules}^{Add/LoC}$, which can be defined as follows:

Definition 5. The **evolution cost metric** $\mathcal{C}_{Modules}^{Add/LoC}$ measures the cost of executing evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$ as the total sum of the lines of code in each one of the new modules:

$$\mathcal{C}_{Modules}^{Add/LoC}(\varepsilon) \triangleq \sum_{m \in Modules(i_{adjusted}) \setminus Modules(i_{old})} LoC(m)$$

where $Modules(i_{adjusted}) \setminus Modules(i_{old})$ returns the set of modules *added* to the implementation and $LoC(m)$ designates the number of ‘lines of code’ of module m .

Let us demonstrate that the metric $\mathcal{C}_{Modules}^{Add/LoC}$ can be obtained by refining the generalized evolution cost metric. We fix the software complexity metric μ to be the metric Add/LoC , defined as follows:

⁽¹³⁾ The reason is usually because of the inherent complexity of large systems, where every change in an existing module can potentially have a “domino effect”, which is precisely the reason for quantifying flexibility.

$$Add/LoC(m) \triangleq \begin{cases} LoC(m) & \text{When } m \text{ is a new module} \\ 0 & \text{Otherwise} \end{cases}$$

This demonstrates that $C_{Modules}^{Add/LoC}$ as formulated in Definition 5 can be obtained by fixing $\mu(m) = Add/LoC(m)$ in the generalized evolution cost metric.

The evolution cost metrics $C_{Modules}^{Add/LoC}$ and $C_{Classes}^1$ demonstrate that the generalized evolution cost metric $C_{Classes}^1$ can yield any number of specific metrics. In other words, the generalized metric can be refined in any number of ways of costing of the evolution process. Additional applications of the generalized metric are demonstrated in the following section.

3 Case studies in evolution complexity

In this section, we demonstrate the use of evolution cost metrics to corroborate informal claims on the flexibility of various programming paradigms, architectural styles, and design patterns.

3.1 Architectural styles

In 1972, Parnas [24] presented the problem of Key Word in Context (KWIC) for the purpose of demonstrating the flexibility of two ‘modular decomposition policies’. The problem has become a classic in software design literature, and in 1993 Garlan and Shaw [27] encoded each modular decomposition policy as an architectural style. In this section, we use evolution complexity to establish and to quantify the informal claims made in [27] on these architectural styles.

Problem

The KWIC problem is defined as follows:

The KWIC [Key Word in Context] index system accepts an ordered set of lines, each line is an ordered set of words, and each word is an ordered set of characters. Any line may be “circularly shifted” by repeatedly removing the first word and appending it at the end of the line. The KWIC index system outputs a listing of all circular shifts of all lines in alphabetical order. [24]

The description can be encoded as follows:

- ◆ $KWIC \triangleq$ Represent an indexing system with algorithms

$$Alg = \{Input, Shift, Alphabetize, Output\}$$

Implementations

Parnas discusses several design policies that can possibly guide the solution to the KWIC problem, encoded by Garlan and Shaw [27] as the following three architectural styles: Shared Data, Abstract Data Structure and Pipes and Filters ⁽¹⁴⁾. Below we briefly summarize each architectural style. Note that they differ not in the number of modules, which is $|Alg|$ in all implementations, but in the way data and functionality are distributed:

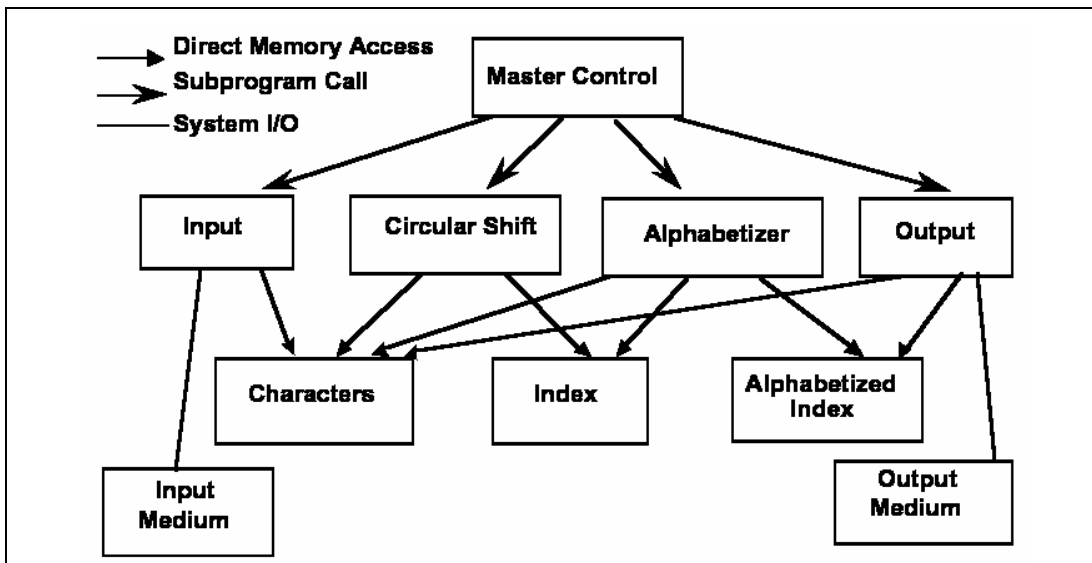


Figure 3. A Shared Data implementation to the KWIC problem (adapted from [27]). Horizontal rectangles represent the elements of the implementation (modules, as chunks of procedures).

⁽¹⁴⁾ In addition, Garlan and Show suggest other architectural styles that may also be used to solve the KWIC problem such as *blackboard architecture*, omitted from discussion for lack of space.

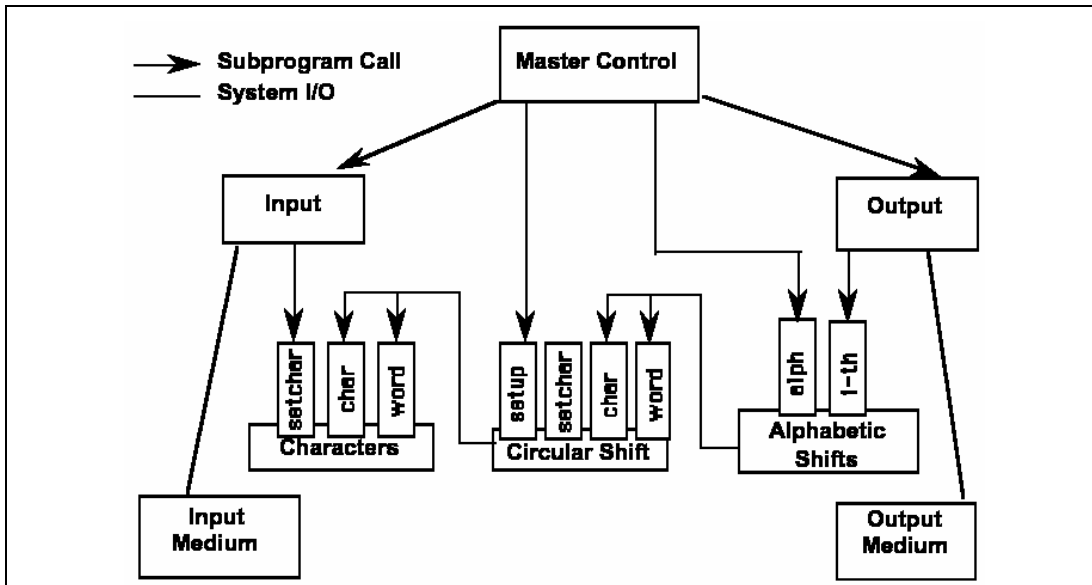


Figure 4. An Abstract Data Type implementation to the KWIC problem (adapted from [27]). Horizontal rectangles represent the elements of the implementation (abstract data types) and vertical rectangles represent elements in the ADT's interface.

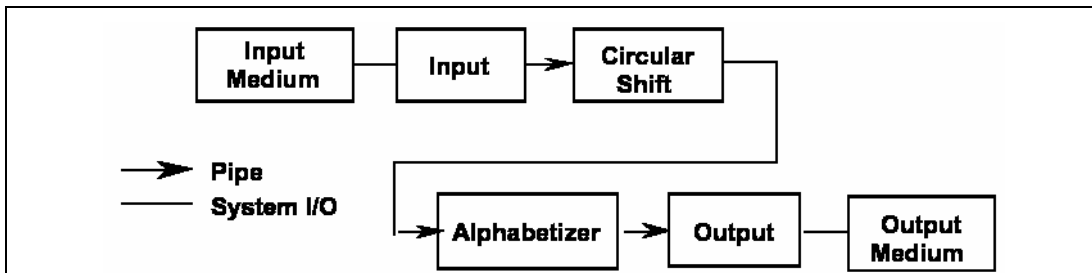


Figure 5. A Pipes and Filters implementation to the KWIC problem (adapted from [27]). Horizontal rectangles represent the elements of the implementation (filters).

- ◆ Shared Data (Figure 3): A functional decomposition policy yields one module-per-functionality, all of which operate on some shared representation of the data.
- ◆ Abstract Data Type (Figure 4): A policy that conforms to the principles of data abstraction (a.k.a. information hiding), where operations over data are only allowed via an abstract interface.
- ◆ Pipes and Filters (Figure 5): A modular decomposition policy of encapsulating each algorithm in an independent module, or a 'filter', which is a stateless process that accepts an input stream and produces an output stream.

Shifts

Consider the following shifts to the *KWIC* problem:

- ◆ Compact the data representation to an efficient format, e.g., by packing four letters to a byte.
- ◆ Parallelize the processes, thereby allowing simultaneous, distributed processing of multiple documents.

Analysis

Informal claims about the flexibility of each design policy were made by Parnas [24], later refined by Shaw and Garlan [27], summarized in Table 3.

Regarding the Shared Data policy they claim:

- ◆ “A change in data storage format will affect almost all of the modules.”
- ◆ “Changes in the overall processing algorithm and enhancements to system function are not easily accommodated.”

Regarding the Abstract Data Type policy they claim:

- ◆ “Both algorithms and data representations can be changed in individual modules without affecting others.”
- ◆ “the solution is not particularly well-suited to [functional] enhancements.”

Regarding Pipes & Filters policy they claim:

- ◆ “it supports ease of modification [of the algorithm]”
- ◆ “it is virtually impossible to modify the design to support an interactive system [because] decisions about data representation will be wired into the assumptions about the kind of data that is transmitted along the pipes.”

These claims were summarized in [27] using a comparative matrix, depicted in Table 3.

We may use the metric $C_{Modules}^1$ to corroborate and quantify these claims. The results of this analysis are summarized in Table 4.

Table 3. Informal claims about the flexibility of three design policies towards shifts in the KWIC problem (adapted from [27]). A minus — symbol stands for ‘easy to evolve’ and a plus symbol + for ‘difficult to evolve’.

| Design policy \ Evolution step | Change Data Representation | Enhance functionality |
|--|----------------------------|-----------------------|
| Shared Data | — | — |
| Abstract Data Type | + | — |
| Pipes and Filters | — | + |

Table 4. $\mathcal{C}_{Modules}^1$ -complexity of evolving the design policies in Table 3, corroborating the claims made in [27].

| Design policy \ Shift | Compact Data Representation | Enhance functionality |
|-------------------------------------|-----------------------------|-----------------------|
| Shared Data | $\mathcal{O}(Alg)$ | $\mathcal{O}(Alg)$ |
| Abstract Data Type | $\mathcal{O}(1)$ | $\mathcal{O}(Alg)$ |
| Pipes and Filters | $\mathcal{O}(Alg)$ | $\mathcal{O}(1)$ |

3.2 Programming paradigms

Object-oriented programming (OOP) is hailed, among other reasons, for promoting flexibility. Experienced programmers, however, observe that object-oriented mechanisms such as inheritance and dynamic binding make programs more flexible only towards the particular shifts they specifically were tailored to accommodate. For example, it has been established that gratuitous use of inheritance may lead to the problem of ‘fragile base class’ [29] and yield highly inflexible programs.

In this subsection we establish the intuition regarding the flexibility of object-oriented programs expressed above. We formulate the problem of representing a deterministic finite-state automaton (DFSA) and analyze the complexity of evolving object-oriented vs. procedural implementations towards shifts in this problem.

Problem

Consider the problem of representing the behaviour of a digital clock with three display states: *DisplayHour*, *DisplaySeconds*, *DisplayDate* and two setting states *SetHour*, *SetDate*. The clock accepts input from two buttons b_1 and b_2 , which are used perform a specific action depending on the current state. The clock's behaviour can be modelled as a deterministic finite state automaton (DFSA), illustrated in Figure 6.

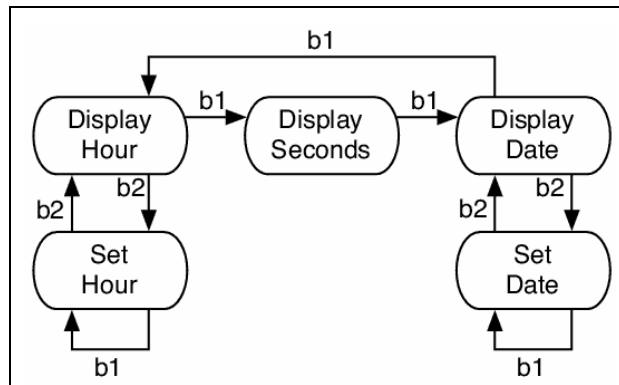


Figure 6. DFSA representation of a digital clock

We may encode this problem as follows:

- ◆ $Clock \triangleq$ Represent a clock with a set of states

$States = \{ DisplayHour, DisplaySeconds, DisplayDate, SetHour, SetDate \}$

and buttons

$Buttons = \{ b_1, b_2 \}$

Shifts

Consider two revisions to problem p_{Clock} :

- ◆ Add button b to $Buttons$
- ◆ Add state s to $States$

Implementations

Consider the following two implementations of p_{Clock} :

- ♦ $OO \triangleq$ Implement *Clock* in OOP policy. The State pattern [8] manifests a classical object-oriented solution to the representation of a clock. Overall, there are $|States|$ classes and $|States| \times |Buttons|$ methods in this implementation, sketched in Table 5.
- ♦ $PROC \triangleq$ Implement *Clock* in procedural policy, described as follows:

An alternative is to use data values to define internal states and have context operations check the data explicitly. But then we'd have look-alike conditional or case statement scattered throughout the context's implementation. ([8], p. 307)

Overall, there are $|Buttons|$ procedures (ANSI C functions) in such an implementation, each consisting of a 'switch' state with $|States|$ 'cases', sketched in Table 6.

Table 5. An object-oriented implementation to the clock problem (in Java™)

```
interface ClockState {
    void b1(); // button 1 pressed
    void b2(); // button 2 pressed
}
class DisplayHour implements ClockState {
    public void b1() { /* b1 pressed */ }
    public void b2() { /* b2 pressed */ }
}
class DisplaySecond implements ClockState {
    public void b1() { /* b1 pressed */ }
    public void b2() { /* b2 pressed */ }
}
class DisplayDate implements ClockState {
    public void b1() { /* b1 pressed */ }
    public void b2() { /* b2 pressed */ }
}
class SetHour implements ClockState {
    public void b1() { /* b1 pressed */ }
    public void b2() { /* b2 pressed */ }
}
class SetDate implements ClockState {
    public void b1() { /* b1 pressed */ }
    public void b2() { /* b2 pressed */ }
}
```

Table 6. A procedural implementation of the Clock problem (in ANSI C)

```

enum states = {DisplayHour, DisplaySecond, DisplayDate, SetHour, SetDate};

struct clock_r {
    states state;           // Current state
} aClock;

void b1(aClock) {          // button 1 pressed
    switch (aClock.state) {
        case DisplayHour: /*...*/;
        case DisplaySecond: /*...*/;
        case DisplayDate: /*...*/;
        case SetHour: /*...*/;
        case SetDate: /*...*/; }
    }

void b2(aClock) {         // button 2 pressed
    switch (aClock.state) {
        case DisplayHour: /*...*/;
        case DisplaySecond: /*...*/;
        case DisplayDate: /*...*/;
        case SetHour: /*...*/;
        case SetDate: /*...*/; }
    }
}

```

Analysis

Informal claims about the cost of each evolution step are made in [8].

Regarding the object-oriented implementation they claim:

- ◆ “new states ... can be added easily”
- ◆ “decentralizing the transition logic in this way makes it easy to modify or extend the logic” (*ibid.*, pp. 307–308)

Regarding the procedural implementation they claim:

- ◆ “Adding a new state ... complicates maintenance.”

The difficulty at this point lies in comparing adjustments to an implementation in Java™ with adjustments to an implementation in ANSI C. The question is, Which commensurable units of modularity can be used to compare such adjustments? As the simplest answer, we define the function $Class/Func(i)$, which counts the number of modules (ANSI C functions or Java classes) in the implementation i :

$$Class/Func(i) \triangleq \begin{cases} \#Functions(i) & \text{When } i \text{ is written in Ansi C} \\ \#Classes(i) & \text{When } i \text{ is written in Java} \end{cases}$$

Fixing the software complexity metric $\mu(m)=1$ and $Modules(i) = Class/Func(i)$ in the generalized evolution cost metric yields the metric $\mathcal{C}_{Class/Func}^1$. The results of calculating the complexity of each evolution step using $\mathcal{C}_{Class/Func}^1$ are summarized in Table 7.

Table 7. $\mathcal{C}_{Class/Func}^1$ -complexity of evolving object-oriented vs. procedural implementations towards shifts in the clock problem.

| Evolution step Design policy | Add letter | Add state |
|---|-------------------------|--------------------------|
| O-O programming | $\mathcal{O}(States)$ | $\mathcal{O}(1)$ |
| Procedural programming | $\mathcal{O}(1)$ | $\mathcal{O}(Buttons)$ |

However, it may be (justly) argued that adjusting ANSI C functions and Java™ classes do not commensurate, and that a more refined approach may be in place. An alternative (and possibly more sophisticated) software complexity metric, such as the cyclomatic complexity metric [18], may provide a more faithful representation of the cost of evolving of different modules. Fixing $\mu(m)=CC(m)$ and $Modules(m)=Class/Func(m)$ in the generalized evolution cost metric yields the metric $\mathcal{C}_{Class/Func}^{CC}$, formulated as follows:

Definition 6. The **evolution cost metric** $\mathcal{C}_{Class/Func}^{CC}$ measures the cost of evolution step $\varepsilon = \langle \langle p_{old}, p_{shifted}, i_{old} \rangle, i_{adjusted} \rangle$ in terms of the total sum of the cyclomatic complexity of the modules affected by the adjustment:

$$\mathcal{C}_{Class/Func}^{CC}(\varepsilon) \triangleq \sum_{m \in \Delta Class/Func(i_{adjusted}, i_{old})} CC(m)$$

where $CC(m)$ of module m is its Cyclomatic Complexity [18], which is defined as follows:

$$CC(m) \triangleq \begin{cases} \text{\#independent paths in the flow graph of } m & \text{When } m \text{ is an Ansi C function} \\ \text{Total sum of \#independent paths in the} \\ \text{flow graph of all methods of } m & \text{When } m \text{ is a Java class} \end{cases}$$

Analyzing the complexity of the same evolution steps using the metric $\mathcal{C}_{Class/Func}^{CC}$ yields slightly different results, summarized in Table 8.

Table 8. $\mathcal{C}_{ClassFunc}^{CC}$ -complexity of evolving an object-oriented vs. procedural implementation of clock problem.

| Evolution step Design policy | Add letter | Add state |
|---------------------------------|--|--|
| O-O programming | $\mathcal{O}(States \times Buttons)$ | $\mathcal{O}(Buttons)$ |
| Procedural programming | $\mathcal{O}(States)$ | $\mathcal{O}(Buttons \times States)$ |

Conclusions

We may conclude the following from the results obtained from using different evolution cost metrics:

- (a) Neither programming paradigm is flexible in absolute terms, regardless of the metric chosen.
- (b) The answer to the question *Which programming paradigm is more flexible?* can be reduced depends on the answer to the question *Which shifts to the problem are most likely to occur?*
- (c) The accuracy of an evolution cost metric varies with the amount of information available about the implementation.

More generally, we also conclude the following:

- (d) Metrics of different granularity levels are useful during both early and late phases in the software lifecycle: A coarse metric (such as $\mathcal{C}_{ClassFunc}^I$) is useful during the design process, namely before the implementation is complete, whereas a refined metric (such as $\mathcal{C}_{ClassFunc}^{CC}$) is useful in later stage, namely once the implementation phase has been completed.

A third solution to the clock problem has not been discussed: States and buttons can be represented in a data structure rather than being hard-coded as classes or functions. This solution is more flexible because the complexity of either evolution step is constant. This solution however does not affect the conclusions drawn above.

3.3 Design patterns

In [19], we analyzed the complexity of evolving two design patterns. In this section, we summarize the conclusions drawn from our analysis.

Visitor

The patterns catalogue [8] discusses the problem of representing abstract syntax trees and operations thereon. It is argued that, ideally, the most flexible implementation is one written in a programming language that supports *double dispatch* [4]. However the programming languages at the focus of the authors (C++ and Smalltalk in [8], but also neither in most other O-O programming languages such as Java and C#) do not support such a mechanism, which motivates the *Visitor* pattern as a design policy. The pattern consists of two class hierarchies, representing (a) the set of elements in abstract syntax (e.g., *constant*, *variable*, *addition expression*) and (b) the set of operations thereof (e.g., ‘print this tree’). Regarding the Visitor design policy they claim:

- ◆ “[it] makes adding new operations easy”
- ◆ “Adding new concrete element classes is hard”

We may use $\mathcal{C}_{Classes}^1$ to corroborate these claims and to make them precise. The results of this analysis are summarized in Table 9.

Table 9. $\mathcal{C}_{Classes}^1$ -complexity of evolving the Visitor pattern. *Operations* stands for the number of ‘visit’ operations.

| Evolution step Design policy | Add operation | Add element |
|---------------------------------|------------------|-----------------------------|
| Visitor pattern | $\mathcal{O}(1)$ | $\mathcal{O}(Operations)$ |

Abstract Factory

The same catalogue [8] also discusses the problem of providing “an interface for creating families of related or dependent objects without specifying their concrete classes.” The actual object that need be created depends on the global context (“current configuration”). For example, when a GUI (Graphical User Interface) client seeks to create a new ‘dialogue box’, considerations of flexibility determine that it must remain independent from the question which windowing systems have been implemented and from the particulars of how dialogue boxes are generated in windowing system.

The authors discuss two design policies solving this problem: The first (the ‘anti-pattern’) uses a ‘switch’ statement (multiple conditional branching) to determine which dialogue box to create. As an alternative which conforms to the general spirit of the OOP paradigm, the Abstract Factory pattern dictates that subtyping and dynamic binding should be used to hide the concrete class of the object created behind a uniform interface. In [19], we use the evolution cost metric $\mathcal{C}_{Classes}^1$ to analyze the flexibility of each design policy, the results of which are summarized in Table 10.

Table 10. $\mathcal{C}_{Classes}^1$ -complexity of evolving the Abstract Factory vs. Switch implementations towards shifts in the Object Creation problem.

| Evolution step Design policy | Add configuration | Add state |
|---|---------------------------|--|
| Abstract Factory | $\mathcal{O}(Products)$ | $\mathcal{O}(Configurations + Products)$ |
| 'Switch' | $\mathcal{O}(Clients)$ | $\mathcal{O}(Clients)$ |

These results also corroborate the more general claims we made: that flexibility depends on which class of shifts that are most likely to occur. This suggests that a software architect must weigh carefully the question which shifts exactly are most likely to occur before s/he may choose the appropriate design policy.

4 Empirical support

A metric is validated when it can be shown to measure what it is supposed to measure. The most obvious test to the evolution cost metrics is to investigate how the actual costs of a particular evolution step grow in a controlled environment. Below, we describe the consolidated results obtained from conducting several small-scale experiments to this extent at the University of Essex and at the Université de Mons-Hainaut. These experiments, described in [6] and [20], corroborate the predictions made in Section 3.2 regarding the complexity of evolving object-oriented and procedural implementations towards shifts in the clock problem.

In the first part of all experiments, subjects were presented with an implementation of a state machine (p_{clock}) with three states ($|States|=3$) and one button ($|Buttons|=1$). Next, the subjects were asked to evolve the implementation in a series of tasks, and to measure the time required to complete each task:

- ◆ In task 2, subjects added a second button to *Buttons*
- ◆ In task 3, subjects added three more states to *States*
- ◆ In task 4, subjects added a third button to *Buttons*

Specifically, the experiment was designed to test the predictions made (Section 3.2) using the evolution cost metrics $\mathcal{C}_{ClassFunc}^1$ and $\mathcal{C}_{ClassFunc}^{CC}$ regarding the cost of evolving an object-oriented (i.e., along the list of the program in Table 5) vs. a procedural (Table 6) implementation. Table 11 contrasts the theoretical predictions with the empirical results obtained in the two experiment types.

Table 11. Theory vs. practice: Summary comparing the predictions made vs. the empirical results obtained in experiments regarding the cost of evolving object-oriented and procedural implementations.

| | Prediction by $\mathcal{C}_{ClassFunc}^1$ | Prediction by $\mathcal{C}_{ClassFunc}^{CC}$ | Median measured | Standard deviation |
|-----------------------------------|--|---|--------------------|-----------------------|
| Object-oriented implementation | 2 | 2 | 2.50 | 1.40 |
| Procedural implementation | 1 | 2 | 1.50 | 0.47 |

Table 11 should be interpreted as follows, depending on the implementation to be evolved:

- ◆ Both metrics $\mathcal{C}_{Class/Func}^1$ and $\mathcal{C}_{Class/Func}^{CC}$ predict that the complexity of adding a letter to the alphabet in an object-oriented implementation grows proportionally with the number of states $|States|$, namely that task 4 will cost about twice as much as task 2. In practice, our results show that task 4 took 2.5 times as much as task 2 (calculated as the median of the ratios between the tasks). We believe that these initial results, collected from seven subjects (standard deviation: 1.4), strongly corroborate both metrics as predictors of flexibility.
- ◆ The metric $\mathcal{C}_{Class/Func}^1$ predicts that the complexity of adding a letter to the procedural implementation is fixed, namely that task 4 will cost about as much as task 2. The metric $\mathcal{C}_{Class/Func}^{CC}$ predicts that the complexity of the task of adding a letter to the alphabet in a procedural implementation grows proportionally with the number of states $|States|$, namely that task 4 will cost about twice as much as task 2. In practice, our results show that task 4 took 1.5 times as much as task 2 (calculated as the median of the ratios between the tasks). These initial results, collected from seven subjects (standard deviation: 0.47) indicate that the cost grew less than predicted by $\mathcal{C}_{Class/Func}^{CC}$ and more than predicted by $\mathcal{C}_{Class/Func}^1$.

Statistically sounder results should be obtained to explore the possibility that a different evolution cost metric may be in place.

5 Related work

Switching between design policies is an expensive and time-consuming activity, but tool support for automating it has the potential of reducing its complexity. For example Kerievsky [14] offers a catalogue of well-defined ‘refactorings’, such as switching for example from a procedural ‘switch’ policy to an object-oriented (e.g., the State pattern) design policy. Such catalogues of well-defined transformations constitute a first step towards automating changes in the design policy, a process which may greatly reduce the cost of executing the corresponding class of evolution steps.

Software complexity

Curtis [5] suggests that “In the maintenance phase [software] complexity determines ... how much effort will be required to modify program modules to incorporate specific changes.” Zuse [31] counts over two hundred metrics for software complexity in literature. Three prominent examples are the following:

- ◆ Lines of Code (LoC) counts the number of lines in the source code of the implementation.
- ◆ McCabe’s Cyclomatic Complexity (CC) [18] measures the number of nodes in the flow graph of the program (demonstrated in Section 3.2).
- ◆ Halstead’s Volume [11] metric is given by the equation $(N_1 + N_2) \times \lg_2(n_1 + n_2)$, where n_1 is the number of distinct operators, n_2 is the number of distinct operands, N_1 is the total number of operators and N_2 is the total number of operands in the respective module.

Each one of these software complexity measures can replace the parameter μ in the general metric for evolution complexity (Definition 4):

- ◆ In Sections 2.5 we demonstrated the result of fixing $\mu=\text{LoC}$ in the generalized metric.
- ◆ In Section 3.2 we demonstrated the result of fixing $\mu=\text{CC}$ in the generalized metric.
- ◆ The result of fixing $\mu=\text{Volume}$ (Halstead’s Volume metric) in the generalized metric yields the evolution cost metric $\mathcal{C}_{\text{Modules}}^{\text{Volume}}$, defined as follows:

$$\mathcal{C}_{\text{Modules}}^{\text{Volume}}(\varepsilon) = \sum_{m \in \Delta \text{Modules}(i_{\text{old}}, i_{\text{adjusted}})} (N_1(m) + N_2(m)) \times \lg_2(n_1(m) + n_2(m))$$

where $N_1(m)$ is the total number of operators in module m , $N_2(m)$ is the total number of operands in m , $n_1(m)$ is the number of distinct operators in m , and $n_2(m)$ is the number of distinct operands in m .

Unfortunately, neither simple nor sophisticated software complexity metrics have been proven accurate as indicators of productivity, comprehensibility or maintainability [31]. This suggests that the accuracy of an evolution cost metric is only limited by the accuracy of the software complexity. We hope that the future will bring more accurate software complexity metrics.

Metrics for software evolution

Quantifying the actual cost of the software evolution remains a relatively unexplored problem. Jørgensen [13] used several models to predict the effort that randomly selected software maintenance tasks require. The size of individual maintenance tasks was measured in LOC. Sneed [28] proposed a number of ways to extend existing cost estimation methods to the estimation of maintenance costs. Ramil et al. [26] provided and validated six different models that predict software evolution effort as a function of software evolution metrics. None of these approaches however suggests an obvious ways in which it is tied to the notion of software flexibility.

Metrics for software flexibility

We are unaware of alternative approaches for quantifying software flexibility, nor of any formal criteria for establishing this quality. It has been suggested that a more accurate way to measure flexibility relies on algorithms or measures that compute the impact of changes [17]. For example, Chaumon et al. [3] report on experimental results with a change-impact model for object-oriented systems. Because the cost and complexity of software evolution may depend on the type of evolution activity, we also require a finer granularity of recognition of types of software evolution activities. Such an attempt to make an objective classification of evolution activities was carried out in [2].

6 Conclusions

We proposed evolution cost metrics and suggested that flexibility can be measured as the complexity of executing particular evolution steps. We studied the complexity of evolving implementations of five recognized programming paradigms, architectural styles, and design patterns, and demonstrated that evolution complexity corroborates intuitions and established observations on the flexibility of these design policies. As stipulated, we also demonstrated the *flexibility* of a particular implementation is relative to the change, and that while a particular implementation is flexible towards one class of changes, it is also inflexible towards another. In particular, the benefits from the measurements proposed are the following:

- (a) evolution complexity can be used to corroborate and quantify informal claims on the flexibility of particular programming paradigms, architectural styles, and design patterns.
- (b) evolution complexity can be used to measure flexibility with varying degrees of accuracy.
- (c) evolution complexity can be used to choose the most flexible design policy, given the class of the most likely shifts to the problem.

6.1 Future directions

The small-scale experiment described in Section 4 should be expanded in all aspects, e.g., testing predictions made w.r.t. other problems, as well as for the purpose of establishing statistically sounder results, using larger sample groups and larger implementations. Of particular interest is to examine the validity of coarse (such as $\mathcal{C}_{Classes}^I$) vs. refined (such as $\mathcal{C}_{Class/Func}^{CC}$) evolution cost metrics.

Evolution complexity can be used to analyze the flexibility of design policies beyond the examples given here. For example, it can be used to throw light on the claims made on the recent introduction of generics to Java and in comparing the flexibility of particular technologies (e.g., CORBA vs. .NET). In particular, evolution complexity can be used in supporting the decision whether to apply a particular refactoring [7], possibly by incorporating a range of evolution cost metrics into integrated development environments which support refactoring, such as IBM Eclipse and Borland JBuilder.

An investigation of the relation between evolution complexity and actual cost of the evolution process is also of interest, albeit more of socio-economic nature than from the software engineering perspective.

Note, however, that given the similarity between the concepts, evolution complexity is no more dependent on empirical validation than computational complexity. So defined, it remains to be examined whether polynomial, exponential, and logarithmic complexity functions are meaningful in the context of software evolution.

Evolution complexity tradeoffs

Studying the flexibility of different programming paradigms towards shifts in the DFSA problem (Section 3.2) suggests that a trade-off may exist between evolution complexity and computational complexity, a relation that is analogous to the trade-off between space and time (computational) complexities. While evidence at this stage is anecdotal, it remains to be examined whether decreased evolution complexity (increased flexibility) leads to time and/or space penalties. Trade-off may also exist between the development effort (early design) and the evolution effort, in the spirit of the adage “weeks of programming can save you hours of planning”. But metrics for quantifying the complexity of the software development process are yet to be proposed.

Acknowledgements

This research has been carried out in the context of the scientific network RELEASE financed by the European Science Foundation (ESF). We wish to thank Nguyen Long, Jon Nicholson, Peter Ebraert, Khalid Allem, Valérie Fiolet, Arista Benoit, Chevalier Jérémie, Michot Arnaudm, Pistone Bruno, Bastien Grégory, Hoste Michael, Jonckers Jérôme, Lavry Aline, Salame Maroun, and Scolas Laurent for taking part in our experiment. We also wish to thank Jeff Reynolds and Mehdi Jazayeri for their comments and suggestions, Naomi Draaijer for her support, and Mary J. Anna for her inspiration.

References

- [1] B. Boehm. *Software Engineering Economics*. Englewood Cliffs: Prentice-Hall, 1981.
- [2] N. Chapin, J. Hale, K. Khan, J. Ramil, W.G. Than. “Types of software evolution and software maintenance”. *J. Software Maintenance and Evolution*, Vol. 13 (2001), pp. 3–30. New York: John Wiley & Sons.
- [3] M.A. Chaumon, H. Kabaili, R.K. Keller, F. Lustman. “A change impact model for changeability assessment in Object-Oriented Software Systems.” *Science of Computer Programming*, Vol. 45, Nos. 2–3 (2002), pp. 155–174. Amsterdam: Elsevier.
- [4] I. Craig. *The Interpretation of Object-Oriented Programming Languages*. New York: Springer-Verlag, 2000.
- [5] B. Curtis. “In search of software complexity.” *Proc. Workshop on Qualitative Software Models for Reliability, Complexity and Cost* (Oct. 1979), pp. 95–106.
- [6] A.H. Eden. “An experiment in evolution complexity: instructions to subjects.” Technical report CSM-431 (Jun. 2005), Department of Computer Science, University of Essex, ISSN 1744-8050.
- [7] M. Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, 2003.
- [8] E. Gamma, R. Helm, R. Johnson, J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading: Addison-Wesley, 1995.
- [9] D. Garlan, G.E. Kaiser, D. Notkin. “Using tool abstraction to compose systems.” *IEEE Computer*, Vol. 25, No. 6 (June 1992), pp. 30–38. Los Alamitos: IEEE Computer Society.
- [10] W.W. Gibbs. “Software’s chronic crisis”. *Scientific American* (Sep. 1994), p. 86–95.
- [11] M.H. Halstead. *Elements of Software Science*. New York: Elsevier, 1977.
- [12] IEEE. *Standard Glossary of Software Engineering Terminology* 610.12-1990, Vol. 1. Los Alamitos: IEEE Press, 1999.
- [13] M. Jørgensen. “Experience with the accuracy of software maintenance task effort prediction models.” *IEEE Trans. Software Engineering*, Vol. 21, No. 8 (1995), pp. 674–681. Los Alamitos: IEEE Computer Society Press.
- [14] J. Kerievsky. “Refactoring to Patterns.” Reading: Addison Wesley Professional, 2004.
- [15] M. Lehman. “Laws of Software Evolution Revisited.” *Lecture Notes in Computer Science* 1149 (Proc. 5th European Workshop on Software Process Technology), pp. 108–124. Berlin: Springer, 1996.
- [16] M.M. Lehman, J.F. Ramil, P. Wernick, D.E. Perry, W.M. Turski. “Metrics and laws of software evolution—the nineties view.” *Proc. Int’l Symposium Software Metrics* (5–7 Nov., 1997), Albuquerque, NM, pp. 20–32. Los Alamitos: IEEE Computer Society Press.
- [17] L. Li, A. Offutt. “Algorithmic analysis of the impact of changes to object-oriented software.” *Proc. Int’l Conf. Software Maintenance—ICSM* (1996), pp. 171–184. Los Alamitos: IEEE Computer Society Press.
- [18] T. McCabe. “A software complexity measure.” *Trans. Software Engineering* Vol. 2 (1976), pp. 308–320. Los Alamitos: IEEE Computer Society Press.

- [19] T. Mens, A.H. Eden. “On the evolution complexity of design patterns”. *Electronic Lecture Notes in Computer Science*, Vol. 127, No. 3, pp. 147–163. Amsterdam: Elsevier, 2004.
- [20] T. Mens, A.H. Eden. “Revised Experiment in Evolution Complexity: Instructions to Subjects.” Technical report CSM-439 (Nov. 2005), Department of Computer Science, University of Essex, ISSN 1744-8050.
- [21] P. Naur, B. Randell, (eds.) *Software Engineering: Report of a conference sponsored by the NATO Science Committee* (7–11 Oct. 1968), Garmisch, Germany. Brussels, Scientific Affairs Division, NATO, 1969.
- [22] L. Osterweil. “Software processes are software too” *Proc. 9th Int’l Conference Software Engineering—ICSE* (Mar. 1987), pp. 2–13. Los Alamitos: IEEE Computer Society.
- [23] L. Osterweil. “Software processes are software too revisited.” *Proc. 19th Int’l Conference Software Engineering—ICSE* (May 1997), pp. 540–548. Los Alamitos: IEEE Computer Society.
- [24] D.L. Parnas. “On the criteria to be used in decomposing systems into modules.” *Communications of the ACM*, Vol. 15, No. 12 (Dec. 1972), pp. 1053–1058. New York: ACM Computing Society.
- [25] D.L. Parnas. “Software Aging.” *Proc. Int’l Conf. Software Engineering—ICSE* (May 1994), pp. 279–287. Los Alamitos: IEEE Computer Society Press.
- [26] J.F. Ramil, M.M. Lehman. “Metrics of software evolution as effort predictors—a case study”. *Proc. Int’l Conf. Software Maintenance* (Oct. 2000), pp. 163–172. Los Alamitos: IEEE Computer Society Press.
- [27] M. Shaw, D. Garlan. *Software Architecture—Perspectives on an Emerging Discipline*. Upper Saddle River: Prentice Hall, 1996.
- [28] H. Sneed. “Estimating the costs of software maintenance tasks.” *Proc. Int’l Conf. Software Maintenance* (1995), pp. 168–181. Los Alamitos: IEEE Computer Society Press.
- [29] A. Taivalsaari. “On the notion of inheritance”. *ACM Computing Surveys*, Vol. 28, No. 3 (Sep. 1996), pp. 438–479. New York: ACM Computing Society.
- [30] A. Urquhart. “Complexity”. Ch. in L. Floridi (ed.), *The Blackwell Guide to Philosophy of Computing and Information*. Oxford: Blackwell, 2004.
- [31] H. Zuse. *Software Complexity*. Berlin: Walter de Gruyter, 1998.