

3-2013

Object-oriented Programming: Some History, and Challenges for the Next Fifty Years

Andrew P. Black

Portland State University, black@cs.pdx.edu

Let us know how access to this document benefits you.

Follow this and additional works at: https://pdxscholar.library.pdx.edu/compsci_fac

 Part of the [Programming Languages and Compilers Commons](#), and the [Software Engineering Commons](#)

Citation Details

Black, Andrew P., "Object-oriented Programming: Some History, and Challenges for the Next Fifty Years" (2013). *Computer Science Faculty Publications and Presentations*. 94.

https://pdxscholar.library.pdx.edu/compsci_fac/94

This Pre-Print is brought to you for free and open access. It has been accepted for inclusion in Computer Science Faculty Publications and Presentations by an authorized administrator of PDXScholar. For more information, please contact pdxscholar@pdx.edu.

Object-oriented programming: some history, and challenges for the next fifty years

Andrew P. Black

Portland State University, Portland, Oregon, USA

Abstract

Object-oriented programming is inextricably linked to the pioneering work of Ole-Johan Dahl and Kristen Nygaard on the design of the SIMULA language, which started at the Norwegian Computing Centre in the Spring of 1961. However, object-orientation, as we think of it today — fifty years later — is the result of a complex interplay of ideas, constraints and people. Dahl and Nygaard would certainly recognise it as their progeny, but might also be amazed at how much it has grown up.

This article is based on a lecture given on 22nd August 2011, on the occasion of the scientific opening of the Ole-Johan Dahl hus at the University of Oslo. It looks at the foundational ideas from SIMULA that stand behind object-orientation, how those ideas have evolved to become the dominant programming paradigm, and what they have to offer as we approach the challenges of the next fifty years of informatics.

Keywords: Ole-Johan Dahl, class, object, prefix, Simula, process, cost-model,
2010 MSC: 01A60, 6803, 68N19, 68N19

Email address: black@cs.pdx.edu (Andrew P. Black)

URL: <http://www.cs.pdx.edu/~black> (Andrew P. Black)

1. Introduction

On 22nd August 2011, a public event was scheduled to open both the 18th International Symposium on Fundamentals of Computation Theory and the Ole-Johan Dahl hus [1], the new building that is home to the University of Oslo’s Department of Informatics, and which is shown in Figure 1. The morning session opened with an Introduction by Morten Dæhlen, followed by two invited talks, one by myself, Andrew Black, and one by Jose Meseguer. These talks were followed by a panel discussion on the future of object-orientation and programming languages, chaired by Arne Maus, and comprising Andrew Black, Yuri Gurevich, Eric Jul, Jose Meseguer, and Olaf Owe.

As it happened, none of these events took place in the Dahl hus, because the beautiful lecture room that had been scheduled for the conference was put out of commission, less than half an hour before the start of the session, by an electrical fault: the scientific opening of the Dahl hus was actually conducted in the neighbouring Kristen Nygaard building. Thus, nine years after their deaths, Dahl and Nygaard were still able to form a symbolic partnership to solve a pressing problem.



Figure 1: The Ole-Johan Dahl hus. (Photograph © the author)

This article is based on the invited talk that I delivered at this event. It is not a transcript; I have taken the opportunity to elaborate on some themes and to précis others, to add references, and to tidy up some arguments that seemed, in hindsight, a bit too ragged to set down in print.

2. The Birth of Simula

In American usage, the word “drafted” has many related meanings. It can mean that you have been conscripted into military service, and it can mean

that you have been given a job that is necessary, but that no one else wants to take on.

In 1948, Kristen Nygaard was drafted, in both of these senses. He started his conscript service at the Norwegian Defence Research Establishment, where his assignment was to carry out calculations related to the construction of Norway's first nuclear reactor [2]. Years later, Nygaard recalled that he had no wish to be responsible for the first nuclear accident on the continent of Europe¹.

After extensive work on a traditional numerical approach, Nygaard turned to Monte Carlo simulation methods. He was made head of the “computing office” at the Defence Establishment, and in 1952 turned full-time to operational research. He earned a Master of Science degree from the University of Oslo in 1956, with a thesis on probability theory entitled “Theoretical Aspects of Monte Carlo Methods” [3]. In 1960, Nygaard moved to the Norwegian Computing Centre (NCC), a semi-governmental research institute that had been established in 1958. His brief was to expand the NCC's research capabilities in computer science and operational research. He wrote “Many of the civilian tasks turned out to present the same kind of methodological problems [as his earlier military work]: the necessity of using simulation, the need of concepts and a language for system description, lack of tools for generating simulation programs” [2]. In 1961, he started designing a simulation language as a way of attacking those problems.

In January 1962, Nygaard wrote what has become a famous letter. It was addressed to the French operational research specialist Charles Salzmann. Nygaard wrote: “The status of the Simulation Language (Monte Carlo Compiler) is that I have rather clear ideas on how to describe queueing systems, and have developed concepts which I feel allow a reasonably easy description of large classes of situations. I believe that these results have some interest even isolated from the compiler, since the presently used ways of describing such systems are not very satisfactory. . . . The work on the compiler could not start before the language was fairly well developed, but this stage seems now to have been reached. The expert programmer who is interested in this part of the job will meet me tomorrow. He has been rather optimistic during our previous meetings.”

¹David Ungar, private communication.

The “expert programmer” was of course Ole-Johan Dahl, shown in Figure 2, and now widely recognised as Norway’s foremost computer scientist. Along with Nygaard, Dahl produced the initial ideas for object-oriented programming, which is now the dominant style of programming for commercial and industrial applications. Dahl was made Commander of the Order of Saint Olav by the King of Norway in 2000, and in 2001 Dahl and Nygaard received the ACM Turing Award “for ideas fundamental to the emergence of object-oriented programming, through their design of the programming languages SIMULA I and SIMULA 67.” In 2002, Dahl and Nygaard were awarded the IEEE von Neumann medal. Dahl died on 29th June 2002.

There has been some confusion about the naming of the various SIMULA languages. The first version of SIMULA, whose preliminary design was presented in May 1963 and whose compiler was completed in January 1965, was a process description and simulation language, and was at the time named just SIMULA. This language had neither classes nor inheritance. In their paper presented at the first History of Programming Languages Conference [2], Nygaard and Dahl refer to this language as SIMULA I, and I shall do the same here. Combining their experience with this language with some important new ideas, Dahl and Nygaard went on to design a general purpose language, which was finished in 1967. This language was at the time named SIMULA 67, although there is now a Simula ’87 standard, and this language is now usually referred to simply as Simula; it was SIMULA 67 that explored the ideas that were to become “fundamental to the emergence of object-oriented programming.” I will use the name SIMULA when referring to the common thread of ideas that pervade both languages, and the forms with post-scripted numerals when referring to a specific language.

In this article, I will try to identify the core concepts embodied in Dahl and Nygaard’s early languages, and see how these concepts have evolved in the fifty years that have passed since their invention. I will also hazard some guesses as to how they will adapt to the future.

3. History in Context

In seeking guidance for the perilous process of predicting the future, I looked at another event that also took place in 1961: the Centennial Celebration of the Massachusetts Institute of Technology. Richard Feynman joined Sir John Cockcroft (known for splitting the atom), Rudolf Peierls (who first conceptu-

alised a bomb based on U-235) and Chen Ning Yang (who received the 1957 Nobel Prize for his work on parity laws for elementary particles) to speak on “The Future in the Physical Sciences.” While the other speakers contented themselves with predicting 10, or perhaps 25, years ahead, Feynman, who was not to receive his own Nobel Prize for another four years, decided to be really safe by predicting 1000 years ahead. He said “I do not think that you can read history without wondering what is the future of your own field, in a wider sense. I do not think that you can predict the future of physics alone [without] the context of the political and social world in which it lies.” Feynman argued that in 1000 years the discovery of fundamental physical laws would have ended, but his argument was based on the social and political context in which physics must operate, rather than on any intrinsic property of physics itself [4].

If social and political context is ultimately what determines the future of science, we must start our exploration of the SIMULA languages by looking at the context that gave them birth. Nygaard’s concern with modelling the phenomena associated with nuclear fission meant that SIMULA was designed as a *process description* language as well as a *programming* language. “When SIMULA I was put to practical work it turned out that to a large extent it was used as a system description language. A common attitude among its simulation users seemed to be: sometimes actual simulation runs on the computer provided useful information. The writing of the SIMULA program was almost always useful, since . . . it resulted in a better understanding of the system” [2]. Notice that modelling means that the actions and interactions of *the objects created by the program* model the actions and interactions of the real-world objects that they are designed to simulate. It is not the SIMULA code that models the real-world system, but the objects created by that



Figure 2: Ole-Johan Dahl (photograph courtesy of Department of Informatics, University of Oslo).

code². The ability to see past the code to the objects that it creates seems to have been key to the success of SIMULA’s designers.

Because of Nygaard’s concern with modelling nuclear phenomena, both of the SIMULA languages followed Algol 60 in providing what was then called “security”: the languages were designed to reduce the possibility of programming errors, and so that those errors that remained could be cheaply detected at run time [5]. A combination of compile-time and run-time checks ensured that a SIMULA program could never give rise to machine- or implementation-dependent effects, so the behaviour of a program could be explained entirely in terms of the semantics of the programming language in which it was written [2]. We would call this property “type-safety”.

This was quite different from assembly languages, and even some contemporary high-level languages such as NPL [6], where references were untyped, and the consequences of mistaking a reference to a string for a reference to a floating-point number could not be explained in terms of the language itself, but could be understood only once one knew the representations of strings and floating-point numbers chosen by the language implementer.

SIMULA addressed “security” by a combination of static and dynamic checks. References were *qualified* to refer to objects of a particular class, so the attributes of those objects — their variables and methods — were known to the compiler. These attributes could be accessed by using a *connection statement* (more commonly called an **inspect** statement), and in SIMULA 67 by means of the now-ubiquitous “dot” notation. If an attribute were present in some subclasses of the statically-known class, but not in others, then accessing it required a dynamic check, which could be requested by means of one or more **when** clauses in the **inspect** statement [7]. This form of the **inspect** statement, containing multiple **when** clauses, is thus very similar to a type-case expression in a language like Haskell.

Dahl’s view of the contributions of SIMULA 67, as he explained in his rôle as discussant at the first HOPL conference [8], was the generalisation and liberalisation of the Algol 60 block to create the class. This gave SIMULA 67:

²More precisely, the classes of a SIMULA 67 program abstract from its objects, and those objects abstract from the real world. So the program can be regarded as a *meta*-model of the real world. I return to this topic in Section 10.

1. record structures (class blocks with variable declarations but no statements);
2. procedural data abstraction (class blocks with variable and procedure declarations, but liberated from the stack discipline so that instances of a class could outlast their callers);
3. processes (class blocks that continue to execute after their callers have resumed execution);
4. incremental abstraction, through *prefixing* one class block with the name of another class block; and
5. modules, obtained by prefixing an ordinary, Algol-60–style, in-line block with the name of a class block; this let the prefix block play the rôle of what Dahl called a context object.

Nowadays, item 4 is called inheritance, or subclassing; I will use the terms interchangeably. It remains one of the signature features of object-oriented programming; I’ll return to inheritance in Section 5. In contrast, item 5—using classes as a packaging or modularity construct—has been forgotten by mainstream languages. Dahl wrote: “This [item 5] turned out to be a very interesting application, rather like collecting together sets of interrelated concepts into a larger class, and then using that larger class as a prefix to a block. The class would function as a kind of predefined context that would enable the programmer to program in a certain style directed toward a certain application area. Simulation, of course, was one of the ideas that we thought of, and we included a standard class in the language for that particular purpose” [8]. Bracha’s Newspeak language has revived this idea [9], and the Grace language [10] does something very similar by using *objects* as modules.

This list of what could be done with a SIMULA class might lead one to think that the real contribution of SIMULA 67 was the class construct, and that classes, rather than the programming style that they enabled, were the real contribution of SIMULA 67. However, I don’t believe that this is correct, either technically or historically. Dahl himself wrote “I know that SIMULA has been criticized for perhaps having put too many things into that single basket of class. Maybe that is correct; I’m not sure myself. . . . It was great fun to see how easily the block concept could be remodeled and used for all these purposes. It is quite possible, however, that it would have been wiser to introduce a few more specialized concepts, for instance, a “context” concept

for the contextlike classes” [8]. I interpret these remarks as Dahl saying that the important contributions of SIMULA 67 were the five individual features, and not the fact that they were all realised by a single piece of syntax. Using identical syntax for features that users perceive as different may be a mistake.

4. The Origin of Simula’s Core Ideas

Dahl was inspired to create the SIMULA 67 class by visualising the *runtime representation* of an Algol 60 program. To those with sufficient vision, objects were already in existence inside every executing Algol program — they just needed to be freed from the “stack discipline”. In 1972 Dahl and Hoare wrote in *Hierarchical Program Structures* [11]:

In ALGOL 60, the rules of the language have been carefully designed to ensure that the lifetimes of block instances are nested, in the sense that those instances that are latest activated are the first to go out of existence. It is this feature that permits an ALGOL 60 implementation to take advantage of a stack as a method of dynamic storage allocation and relinquishment. But it has the disadvantage that a program which creates a new block instance can never interact with it as an object which exists and has attributes, since it has disappeared by the time the calling program regains control. Thus the calling program can observe only the results of the actions of the procedures it calls. Consequently, the operational aspects of a block are overemphasised; and algorithms (for example, matrix multiplication) are the only concepts that can be modelled.

In SIMULA I, Dahl made two changes to the Algol 60 block: small changes, but changes with far-reaching consequences. First, “a block instance is permitted to outlive its calling statement, and to remain in existence for as long as the program needs to refer to it” [11]. Second, references to those block instances are treated as data, which gives the program a way to refer to them as independent objects. As a consequence of these changes, a more general storage allocation mechanism than the stack is needed: a garbage collector is required to reclaim those areas of storage occupied by objects that can no longer be referenced by the running program. Dahl and Nygaard may not have been the first to see the possibilities of generalising the Algol block,

but they were the first to realise that the extra complexity of a garbage collector was a small price to pay for the wide range of concepts that could be expressed using blocks that outlive their calling creators. It was Dahl’s creation of a storage management package “based on a two-dimensional free area list” that made possible the SIMULA process [2]. This package seems to have been completed by May 1963, the date of a preliminary presentation of the SIMULA I Language. The SIMULA I compiler was completed in January 1965; a year later it was accepted by Univac, who had partially financed its development — see Figure 3.

SIMULA 67’s class concept was developed as a solution to a number of shortcomings of SIMULA I, but particularly the cumbersome nature of the `inspect` mechanism for accessing the local variables of a process, and the need to share properties between different simulation models. Hoare had previously introduced the notion of a class of record instances; one of Hoare’s classes could be specialised into several subclasses, but the class and its subclasses had to be declared together, with the subclass definitions nested inside

the class definition [12]. Dahl probably became familiar with Hoare’s proposals when they both lectured (Hoare on Record Handling and Dahl on Simulation Languages) at the NATO Summer school at Villard-de-Lans in 1966 [13]. Dahl and Nygaard realised that it would be useful if a class could be declared separately from its subclasses, so that it could be specialised separately for diverse purposes. However, Hoare’s nesting of subclass declarations within class declarations precluded such a separation. “The solution came with the idea of class prefixing: using `C` as a prefix to another class, the latter would be taken to be a subclass of `C` inheriting all properties of



Figure 3: Dahl and Nygaard during the development of SIMULA, around 1962. The Univac 1107 was obtained by Nygaard for the NCC at a favourable price in exchange for a contract to implement SIMULA I. (Photograph from the Virtual Exhibition *People Behind Informatics*, <http://cs-exhibitions.uni-klu.ac.at/>)

C” [14]. The semantics of prefixing was defined as concatenation, but the syntactic separation of class and subclass turned the class of SIMULA 67 into a unit of reuse. This was so successful that the special-purpose simulation facilities of SIMULA I were replaced in SIMULA 67 by a SIMULATION class, which was defined in SIMULA 67 itself. To make use of simulation facilities in a class, all that a programmer then needed to do was to prefix the definition of their class by the name of the class SIMULATION.

An important part of prefixing was the notion of *virtual*, although Dahl writes that it was “a last minute extension” [14]. In SIMULA 67, the binding of attribute identifiers is static. This means that the syntax of the reference, rather than the identity of the target object, determines which variable to reference or which procedure to call. The binding can be made dynamic by marking the attribute as *virtual*; this means that the choice of variable or procedure depends only on the target object. Making static binding the default was, I think, a concession to efficiency; with the benefit of hindsight, it seems clear to me that it was a mistake. Smalltalk made all methods virtual, and raised dynamic binding to the status of a fundamental part of object-orientation. Even Java made dynamic binding the default, although in Java the `final` keyword can be used to specify static binding and to prohibit overriding.

5. Why Inheritance Matters

Inheritance, combined with dynamic binding, has become the *sine qua non* of object-orientation; indeed, an influential 1987 paper by Wegner defined “object-oriented” as “objects + classes + inheritance” [15]. Is this inevitable? Can one have a viable object-oriented programming language without inheritance? I was one who objected to Wegner’s definition at the time: it seemed more reasonable for the term “object-oriented” to imply only the use of objects. We had just published the first papers on Emerald [16, 17], a language that we certainly thought of as object-oriented, even though it had no inheritance, so perhaps we were not unbiased observers. Wegner’s definition excluded not only Emerald, but also any delegation-based language from being object-oriented, even though such languages were of growing interest. Wegner did admit that delegation, which he defined as “a mechanism that allows objects to delegate responsibility for performing an operation or finding a value to one or more designated ‘ancestors’,” was actually the purer feature.

Inheritance has stood the test of time as a useful feature, and I now believe that some code-reuse mechanism *like* inheritance or delegation is a very important part of any object-oriented programming system. Since 1989, thanks to Cook, we have known that inheritance can be explained using fix-points of generators of higher-order functions [18]. What this means is that, in theory, functions parametrized by functions are “as good as” inheritance, in the sense that they can enable equivalent reuse. Nevertheless, in practice they are not as good, because to enable reuse, the programmer has to plan ahead and make every part that could possibly change an explicit parameter.

Functional programmers call this process of explicit parametrization “abstraction”. It has two problems: life is uncertain, and most people think better about the concrete than the abstract. Let’s examine these problems briefly. We have all experienced trying to plan ahead for change. Inevitably, when change comes, it is not the change we anticipated. We find that we have paid the price of generality and forward planning, but still have to modify the software to adapt to the change in requirements. Agile methodologies use this: they tell us to eschew generality, and instead make the software ready to embrace change by being changeable. That is, they tell us that rather than trying to write software that is general enough to cope with all possible situations without need of modification, we should write software that is specific to the current requirements, but easily modifiable to other possible requirements. Inheritance aids in this process by allowing us to write software in terms of the concrete, specific actions that the application needs now, and to abstract over them only when we know—because the change request is in our hands—that the abstraction is needed.

Let’s explain how this works using an example. Suppose that in the original program the concrete actions are represented as methods in a class.³ An example might be a `Text` class with a method `emphasise` that italicises a segment of the text. If the requirements change so that the visual appearance of emphasis should be controlled by a preference, then a new class can be created that inherits from `Text` and overrides the `emphasise` method with a

³The *methods* of an object are the set of named operations that it can perform. They execute in response to *messages* from itself or another object; each message contains the name of a method, and the necessary arguments. A *class* can be thought of as a template that describes the structure of a group of similarly-structured objects; in particular, it defines their methods.

new method; this new method would use a preference object to determine the visual appearance of the emphasised text. In doing so, the `emphasise` method, previously a constant, has become a parameter. Notice that we didn't have to anticipate that this particular change would be requested, and the original `Text` class didn't have to make every possible locus of change a parameter. Even if the process of emphasising text had not been broken out into a separate method in the original code, it could have been extracted into an `emphasise` method using a simple, behaviour-preserving refactoring before any change to the functionality of the program was attempted.

As with so much of object-oriented programming, Dahl foresaw the value of turning methods—which he called virtual procedures—into parameters. Indeed, virtual procedures were originally conceived, in part, as a replacement for what he felt to be the excessive power of Algol 60's call-by-name parameters. In the discussion at the Lysebu conference following the presentation of the paper on Classes and Subclasses, Dahl is reported to have remarked [7] “virtual quantities are in many ways similar to Algol's name parameters, but not quite as powerful. It turns out that there is no analogy to Jensen's device. This, I feel, is a good thing, because I hate to implement Jensen's device. It is awful.”

The second problem with abstraction derives from the way that we think. Most of us can grasp new ideas most easily if we are first introduced to one or two concrete instances, and are then shown the generalisation. To put this another way: people learn best from examples. So we might first solve a problem for $n = 4$, and *then* make the changes necessary for 4 to approach infinity. There *are* people who are able to first grasp abstractions directly, and then apply them: such people tend to become mathematicians, and sometimes programming language designers. Languages designed by and for such people tend to provide good support for creating abstractions. Languages for the rest of us should support inheritance, which lets the programmer provide a concrete instance first, and then generalise from it later.

To illustrate the power of inheritance to make complex abstractions easier to understand, let's look at a case study from the functional programming literature. In *Programming Erlang* [19, Chapter 16], Armstrong introduces the OTP (Open Telecom Platform) generic server. He comments: “This is the most important section of the entire book, so read it once, read it twice, read it 100 times—just make sure the message sinks in”. To make sure that the message about the way that the OTP server works *does* sink in,

Armstrong presents us with

four little servers . . . each slightly different from the last. `server1` runs some supplied code in a server, where it responds to remote requests; `server2` makes each remote request an atomic transaction; `server3` adds hot code swapping, and `server4` provides both transactions and hot code swapping.

Each of these “four little servers” is self-contained: `server4`, for example, makes no reference to any of the preceding three servers.

Why does Armstrong describe the OTP sever in this way, rather than just presenting `server4`, which is his destination? Because he views `server4` as too complicated for the reader to understand in one go. Something as complex as `server4` needs to be introduced step-by-step. However, his language, lacking inheritance (and higher-order functions) does not provide a way of capturing this stepwise development.

In an effort to understand the OTP server, I coded it up in Smalltalk. As George Forsythe is reputed to have said: “People have said you don’t understand something until you’ve taught it in a class. The truth is you don’t understand something until you’ve taught it to a computer — until you’ve been able to program it” [20]. First I translated `server1` into Smalltalk; I called it `BasicServer`, and it had three methods and 21 lines of code. Then I needed to test my code, so I wrote a name server plug-in for `BasicServer`, set up unit tests, and made them pass. In the process, as Forsyth had predicted, I gained a *much* clearer understanding of how Armstrong’s `server1` worked. Thus equipped, I was able to implement `TransactionServer` by subclassing `BasicServer`, and `HotSwapServer` by subclassing `TransactionServer`, bringing me to something that was equivalent to Armstrong’s `server4` in two steps, each of which added just one new concern. Then I refactored the code to increase the commonality between `TransactionServer` and `BasicServer`, a commonality that had been obscured by following Armstrong in re-writing the whole of the main server loop to implement transactions. This refactoring added one method and one line of code.

Once I was done, I discussed what I had learned with Phil Wadler. Wadler has been thinking deeply, and writing, about functional programming since the 1980s; amongst other influential articles he has authored “The Essence of Functional Programming” [21] and “Comprehending Monads” [22]. His first reaction was that the Erlang version was simpler because it could be described

in straight-line code with no need for inheritance. I pointed out that I could refactor the Smalltalk version to remove the inheritance, simply by copying down all of the methods from the superclasses into `HotSwapServer`, but that doing so would be a *bad idea*. Why? Because the series of three classes, each building on its superclass, *explained* how `HotSwapServer` worked in much the same way that Armstrong explained it in Chapter 16 of *Programming Erlang*. This was an “ah-ha moment” for Phil.

To summarise: most people understand complex ideas incrementally, by starting with a simple concrete example, and then taking a series of generalisation steps. A program that uses inheritance can *explain* complex behaviour incrementally, by starting with a simple class or object, and then generalising it in a series of inheritance steps. The power of inheritance is that it enables us to organise our programs incrementally, that is, in a fashion that corresponds to the way that most people think.

6. Object-Oriented Frameworks

In my view, one of the most significant contributions of SIMULA 67 was the idea of the object-oriented framework. By the 1960s, subroutine libraries were well-known. A subroutine, as its name suggests, is always subordinate to the main program. The relationship is always “don’t call us, we’ll call you”. The only exception to this is when a user-written subroutine p is passed as an argument to a library subroutine q ; in this case, q can indeed call p . However, in the languages of the time, any such “callback” to p could occur only during the lifetime of the original call from the main program to q . An object-oriented framework allows subroutines, but also enables the reverse relationship: the user’s program can provide one or more components that are invoked by the framework, and the framework takes on the rôle of “main program”.

This reversal of rôles enabled the special-purpose simulation constructs of SIMULA I to be expressed as a framework within SIMULA 67. The programmer of a simulation wrote code that described the behaviour of the individual objects in the system: the nuclear fuel rods, or the air masses, or the products and customers, depending on the domain being simulated. The simulation framework then *called* those user-defined objects. The framework was in control, but users were able to populate it with objects that achieved their diverse goals.

Of course, for this scheme to work, the user-provided objects called by the framework must have methods for all of the requests made of them. Providing these methods from scratch could be a lot of work, but the work can be made much more manageable by having the user-provided objects *inherit* from a framework object that already has methods for the majority of these requests; all the user need do is override the inherited behaviour selectively.

One of the most common instantiations of this idea is the user-interface framework. In most object-oriented languages, objects that are to appear on the computer's screen are sent requests by the display framework. They must respond to enquires such as those that ask for their bounds on the screen, and to draw themselves within these bounds. This is normally achieved by having displayable objects subclass a framework-provided class, such as Squeak's `Morph` or `java.awt.Component`.

Dahl and Nygaard realised that these ideas could be used to provide the simulation features of SIMULA I. SIMULA 67 is a general-purpose language; its simulation features are provided by an object-oriented framework called SIMULATION. SIMULATION is itself implemented using a more primitive framework called SIMSET.

7. From Simula to Smalltalk

Smalltalk-72 was an early version of Smalltalk, designed by Alan Kay and used only within Xerox PARC. It is described in Kay's paper on the Early History of Smalltalk [23]. Kay was inspired by the simplicity of LISP and the classes and objects of SIMULA 67, and took from SIMULA the ideas of classes, objects and object references.

Smalltalk-72 refined and explored the idea of objects as little computers, or, as Kay puts it: "a recursion on the notion of computer itself". Objects combined data with the operations on that data, in the same way that a computer combines a memory to store data with an arithmetic and logic unit to operate on the data.

Smalltalk-72 did not include inheritance, not because it was thought to be unimportant, but because SIMULA's single, static inheritance seemed too limiting; Kay believed that inheritance could be simulated within Smalltalk using its LISP-like flexibility. Smalltalk-72 also dropped some other ideas from SIMULA, notably the idea that objects were also processes, and that classes,

used as prefixes to inline blocks, were also modules. Smalltalk-76 included a Simula-like inheritance scheme that allowed superclasses to be changed incrementally; this was kept in Smalltalk-80. However, neither Smalltalk-76 nor Smalltalk-80 supported objects as processes or classes as modules; as a result, what we may call the “North-American School” of object-orientation views these features as relatively unimportant, and perhaps not really part of what it means to support objects.

By the late 1980s, objects had attracted significant industrial interest. Many different object-oriented languages had been developed, differing in the weight that they gave to classes, objects, inheritance, and other features. In 1991, Alan Snyder of Hewlett Packard wrote an influential survey paper “The Essence of Objects” [24], which was later published in *IEEE Software* [25].

Unlike SIMULA and Smalltalk, Snyder’s is a descriptive work, not a prescriptive one. He describes and classifies the features of contemporary object-oriented languages, including Smalltalk, C++, the Common Lisp Object System, and the Xerox Star, as well as a handful of other systems. In Snyder’s view, the essential concepts were as follows:

- An object embodies an abstraction.
- Objects provide services.
- Clients issue requests for those services.
- Objects are encapsulated.
- Requests identify operations.
- Requests can identify objects.
- New Objects can be created.
- The same operation on distinct objects can have different implementations and observably different behaviour.
- Objects can be classified in terms of their services (interface hierarchy).
- Objects can share implementations.
 - Objects can share a common implementation (multiple instances).
 - Objects can share partial implementations (implementation inheritance or delegation).

Table 1: Evolution of the features of an “object”. SIMULA’s idea of procedural encapsulation became more complete, both over time and as objects migrated across the Atlantic. However, other features of the SIMULA Class, such as Active Objects and Classes as Modules, have not persisted in mainstream object-oriented languages.

Feature	Simula 67	Smalltalk-80	Snyder (1991)
Procedural Abstraction	Attributes exposed	Attributes encapsulated	Objects characterised by services
Active Objects	Yes	No	“Associated Concept”
Dynamic object creation	Yes	Yes	Yes
Classes	Yes	Yes	Shared implementations
Inheritance	Class prefixing	Subclassing	“Shared partial implementations”
Classes as Modules	Yes	No	No

We see that objects as processes and classes as modules are not on this list. Snyder does mention “Active Objects” as an “associated concept”, that is, an idea “associated with the notion of objects, but not essential to it”. The idea that classes can serve as modules does not appear at all.

The evolution of the features that make up object-orientation is summarised in Table 1. Encapsulation has become more important, and the mechanisms for sharing implementation have become more diverse and refined. The idea that multiple objects can share an implementation, and that new kinds of objects can be defined that share parts of existing objects, have become recognised as important, although both kinds of sharing can be implemented either using SIMULA-like classes, or through other mechanisms such as delegation. In contrast, active objects, and classes as modules, are no longer regarded as key to object-orientation.

8. Objects as Abstractions

The word “abstraction” is used in Informatics in many different senses. As I mentioned previously, it is used by functional programmers to mean “param-

eterisation”. In the context of objects, I am going to use the word abstraction to capture the idea that what matters about an object is its protocol: the set of messages that it understands, and the way that it behaves in response to those messages. Nowadays, this is sometimes also referred to as the object’s interface. The key idea is that when we use an object, we focus on how it appears from the outside, and “abstract away from” its internal structure: more simply, that the internal structure of an object is hidden from all other objects. That’s why I said “the set of messages that it understands”, and not “the set of methods that it implements”. In many languages they are the same, but I wanted to emphasise the external rather than the internal view.

Abstraction, in this sense, is one of the key ideas behind objects, but it does not appear explicitly until Snyder’s 1991 survey. SIMULA doesn’t mention abstraction specifically; it speaks instead of modelling, which captures the importance of interacting with the object, but not the information hiding aspect. Dahl remarks [14] “According to a comment in [the SIMULA user’s manual, 1965] it was a pity that the variable attributes of a *Car* process could not be hidden away in a subblock”, but this was not possible in SIMULA without also hiding the procedures that defined the *Car*’s behaviour — and exposing those procedures was the whole reason for defining the *Car*. (Around 1972, SIMULA addressed this problem by adding the **protected** and **hidden** keywords to control the visibility of a name declared in a class [26].)

Smalltalk-80 approached abstraction by protecting the instance variables that contain the data representing an object: these variables are accessible only by the object itself. The Smalltalk-80 programming environment also provided a way for the programmer to say that a method should be private, but that was merely a convention, and not enforced by the system. Nevertheless, Dan Ingalls’ sweeping 1981 *Byte* article “Design Principles behind Smalltalk” [27] refers to abstraction only indirectly. Ingalls singles out *Classification*, not abstraction, as the key idea embodied in the Smalltalk class. Ingalls writes:

Classification is the objectification of *nessness*. In other words, when a human sees a chair, the experience is taken both literally as “that very thing” and abstractly as “that chair-like thing”. Such abstraction results from the marvellous ability of the mind to merge “similar” experience, and this abstraction manifests it-

self as another object in the mind, the Platonic chair or chairness.

The fact remains that Smalltalk classes do not classify objects on the basis of their behaviour, but according to their representation. It seems that the idea of separating the internal (concrete) and external (abstract) view of an object was yet to mature, although it can be glimpsed elsewhere in Ingalls' article. After Classification, his next principle is “*Polymorphism*: A program should specify only the behaviour of objects [that it uses], not their representation”. Also, in the “Future Work” section, Ingalls remarks that

message protocols have not been formalised. The organisation provides for protocols, but it is currently only a matter of style for protocols to be consistent from one class to another. This can be remedied easily by providing proper protocol objects that can be consistently shared. This will then allow formal typing of variables by protocol without losing the advantages of polymorphism.

Borning and Ingalls did indeed attempt just that [28], but it proved not to be as easy to get the details right as Ingalls had predicted. Still, Smalltalk as defined in 1980 did succeed in abstracting away from representation details: it never assumes that two objects that exhibit the same protocol must also have the same representation.

The related concept of *data abstraction* can be traced back to the 1970s. Landmarks in its evolution are of course Hoare's “Proof of Correctness of Data Representations” [29], Parnas' “On the Criteria to be used in Decomposing Systems into Modules” [30], both published in 1972, and the CLU programming language, which provided explicit conversion between the concrete data representation seen by the implementation and the abstract view seen by the client [31]. However, none of these papers pursues the idea that multiple representations of a single abstraction could co-exist. They assume that abstraction is either a personal discipline backed by training the programmer, or a linguistic discipline backed by types. The latter view — that abstraction comes from types — has been heavily marketed by the proponents of ML and Haskell. It is based on the idea that types capture representation as well as interface, but that by hiding the actual type from code that lies outside the abstraction boundary, the programmer can be prohibited from taking advantage of the representation. The mechanism used for this hiding is the notion of an *existentially quantified type*.

Objects follow a different route to abstraction, which can be traced back to Alonzo Church’s work on representing numbers in the lambda-calculus [32]. Church represented the number 2 in the pure lambda-calculus by the function that *did something twice*; the number 3 by the function that did something three times, and so on. In 1975 John Reynolds named this route “procedural abstraction” and showed how it could be used to support data abstraction in a programming language, as well as contrasting it with the approach based on existential types [33]. Procedural abstraction is characterised by using a *computational* mechanism — a representation of a function as executable code — rather than type discipline or self-discipline — to enforce abstraction boundaries. Reynolds writes:

Procedural data structures provide a decentralised form of data abstraction. Each part of the program which creates procedural data will specify its own form of representation, independently of the representations used elsewhere for the same kind of data, and will provide versions of the primitive operations (the components of the procedural data item) suitable for this representation. There need be no part of the program, corresponding to a type definition, in which all forms of representation for the same kind of data are known.

Unfortunately, Reynolds’ paper is not widely known, even though it has been reprinted in two collections of notable papers [34, 35]; many practitioners do not understand the fundamental distinction between type abstraction and procedural abstraction. This distinction is important because there is a trade-off between the costs and benefits of the two approaches. The benefit of procedural abstraction — the approach used by objects — is that it provides secure data abstraction without types, and allows multiple implementations of the same abstraction to co-exist; this supports change, and helps to keep software soft. The cost is that it is harder to program “binary operations” — those that work on two or more objects — with maximal efficiency.

William Cook made another attempt to explain this distinction in a 2009 OOPSLA Essay [36]. Cook coined the term “autognostic”, meaning “self-knowing”, for what Reynolds called “decentralisation”. Autognosis means that an object can have detailed knowledge only of itself: all other objects are encapsulated. Cook remarks: “The converse is quite useful: any programming model that allows inspection of the representation of more than

one abstraction at a time is not object-oriented.”

To see the benefits of procedural abstraction, consider a selection of objects implementing numbers. The objects 1 and 2 might include a representation as a machine integer, and methods $+$, $-$, *etc.*, that eventually use hardware arithmetic. However, the object representing 2^{67} might use a representation comprising three radix 2^{32} digits, and the methods for $+$ and $-$ would then perform multi-digit arithmetic. This scheme, augmented with operations to convert machine integers to the multi-digit representation, is essentially how Smalltalk implements numbers. Notice that new representations of number objects can be added at any time, without having to change existing code, so long as the new objects provide the necessary methods. Because the representation of an object can never be observed by any other object⁴, there is no need to protect it with an existential type.

The idea of procedural data abstraction was certainly in SIMULA I—indeed, I believe that it was the genesis of SIMULA I. As Dahl and Hoare observe in “Hierarchical Program Structures” [11], the key concept is already present in Algol 60: an Algol 60 block is a single mechanism that can contain both procedures and data. The limitation of Algol is that blocks are constrained not to outlive their callers. Dahl saw this, and “set blocks free” by devising a dynamic storage allocation scheme to replace Algol 60’s stack [2].

However, procedural data structures were not the only kind of data aggregate supported by SIMULA 67. As I mentioned in Section 3, Dahl realised that SIMULA’s class construct could be used to generate both records (unprotected, or protected by type abstraction) and objects (protected by procedural abstraction), and saw that as a strength of the mechanism; some later languages, notably C++, also support this dualism. Records were made possible by SIMULA I’s `inspect` statement, and later by SIMULA 67’s more flexible “dot” notation; these mechanisms were introduced specifically to expose the variables that would otherwise have been local to an object. In contrast, Smalltalk made the instance variables of an object inaccessible to any other object. If a Smalltalk object wanted to expose an instance variable, it could provide *methods* that permitted the reading and writing of the variable, but the client would nevertheless see only methods, never variables.

⁴To take advantage of the hardware arithmetic unit, that piece of hardware must naturally be allowed to “see” the representation of small integers. Nevertheless, the representation of a object remains hidden from all other objects in the program.

The modern view of an object is that only the *behaviour* of the object matters: whether that behaviour is ultimately supported by methods or variables is one of the hidden implementation details of the object. The text-book illustration of this consists of two `Point` objects, one encapsulating a representation using cartesian coordinates, and the other encapsulating a representation using polar coordinates. Both objects can offer methods to retrieve, and to change, the x , y , ρ and θ of the points. Some of these methods will involve direct access to single instance variables, while other will involve trigonometric calculations and access to multiple instance variables. Nevertheless, clients of the two objects will be able to treat them identically.

It seems to me that Dahl himself eventually came to believe that procedural abstraction was the fundamental contribution of the Simula languages. In 2002, Dahl wrote [14]:

The most important new concept of Simula 67 is surely the idea of data structures with associated operators ... called objects. There is an important difference, except in trivial cases, between

- *the inside view* of an object, understood in terms of local variables, possibly initialising operations establishing an invariant, and implemented procedures operating on the variables maintaining the invariant, and
- *the outside view*, as presented by the remotely accessible procedures, including some generating mechanism, dealing with more “abstract” entities.

There is one more consequence of the difference between type abstraction and procedural abstraction. Type abstraction makes a fundamental distinction between the collection of representation variables—the value that is protected by the existential type—and the functions that operate on it. Procedural abstraction makes no such distinction: it exposes *only* methods, with the consequence that there may *be no representation variables at all!*

Perhaps the clearest example of this is the implementation of the Booleans in Smalltalk, which is outlined in Figure 4. The objects `true` and `false` implement methods `&`, `|`, `not`, `ifTrue:ifFalse:`, *etc.* (The colon in a Smalltalk method name introduces an argument, in the same way that parentheses are used to indicate arguments in other languages, and `↑` means “return from this method with the following answer”.) For example, the `&` method of `true`

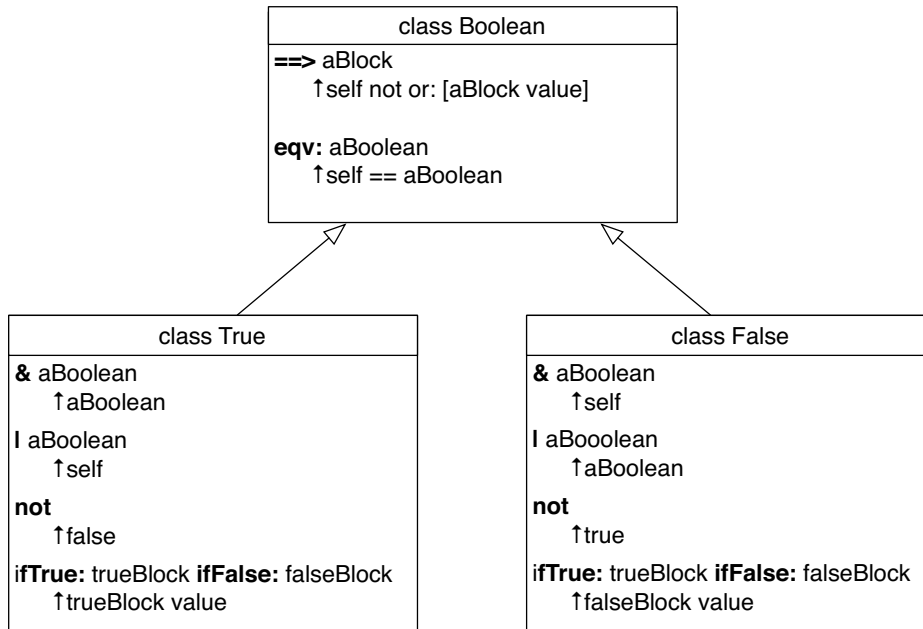


Figure 4: An extract from the Boolean classes in Smalltalk. Boolean is an abstract class with no instances; true and false are the only instances of the classes True and False, which inherit from Boolean. There is no “hidden state”; the methods themselves are sufficient to define the expected behaviour of the Booleans.

answers its argument, while the | method always answers self, *i.e.*, true. The ifTrue:ifFalse: methods are a bit more complicated, because the arguments must be evaluated conditionally. The ifTrue:ifFalse: method of false ignores its first argument, evaluates its second (by sending it the message value), and returns the result, while the ifTrue:ifFalse: method of true does the reverse.

9. Active Objects

Active objects seem to be one idea from SIMULA that has been lost to the object-oriented community. Activity was an important part of SIMULA; after all, the original purpose of the language was to simulate activities from the real world. SIMULA’s “quasi-parallelism” was a sweet-spot in 1961: it allowed programmers to think about concurrency while ignoring synchronisation. Because another task could execute only when explicitly resumed, programmers could be confident that their data would not change “out from

under them” at unexpected times.

Hewitt’s Actor model [37] built on this idea, as did Emerald [16], in which every object could potentially contain a process. Other languages, like Erlang, have made the process the main focus of the language rather than the object [38]. However, active objects have disappeared from “mainstream” object-oriented languages. Since Smalltalk, processes and objects have been independent concepts.

Why are Smalltalk objects passive? I don’t know. Perhaps Kay and Ingalls had a philosophical objection to combining what they saw as separate ideas. Perhaps the realities of programming on the Alto set limits as to what was possible at that time. Perhaps they wanted real processes, not coroutines; this requires explicit synchronisation, and thus a significantly different language design. Smalltalk does indeed have processes, rather than coroutines. However, these processes are realised as separate objects, rather than being a part of every object.

10. Classes and Objects

It’s interesting to revisit Wegner’s idea, discussed in Section 5, that object-orientation is characterised by the existence of classes. As we have seen, most of the concepts that we recognise as making up object-orientation came from the class concept of SIMULA. Nevertheless, as both Dahl and Nygaard would be quick to point out, it’s the dynamic objects, not the classes, that form the system model, and modelling was the *raison d’être* of SIMULA. Classes were interesting only as a way of creating the dynamic system model, which comprised interacting objects. Classes are indeed a most convenient tool if you want hundreds of similar objects. But what if you want just one object? In that case they impose a conceptual overhead. The fact is that classes are meta, and meta isn’t always better!

Classes are meta, relative to objects, because a class describes the implementation and behaviour of those objects that we call instances of the class, as well as providing a factory (procedure or method) that realises that description by actually making an instance. If you need just one or two objects, it’s simpler and clearer to describe those objects in the program directly, rather than to describe a factory (that is, a class) to make them, and then use it once or twice. This is the idea behind Self, Emerald, NewtonScript and JavaScript, which take the object, rather than the class, as

their fundamental concept. Which is not to say that classes are not useful: just that, at least in their rôle of object factories, they can be represented as objects.

In class-based languages, classes typically play many rôles. Borning, in an early paper contrasting class-based and object-based languages [39], lists eight rôles for Smalltalk classes. The primary rôle of a class is acting as a factory that creates new objects; some of the other rôles served by Smalltalk classes include describing the representation of those objects, defining their message protocol, serving as repositories for the methods that implement that protocol, providing a taxonomy for classifying objects, and serving as the unit of inheritance. Nevertheless, as we saw when discussing Ingalls' view of classification (Section 8 above), the taxonomy provided by classes, being based on implementation rather than behaviour, is fundamentally at odds with the autognosis that other researchers, such as Cook, believe to be fundamental to object-orientation.

If the only thing that we can know about an object is its behaviour, the most useful way of classifying objects must be according to their behaviour. This brings us back to types. Recall that types are not needed to provide abstraction in object-oriented systems. Nor are they needed to provide what Hoare called “security”: that comes from procedural encapsulation. The true rôle of types is to provide a *behavioural* taxonomy: types provide a way of classifying objects so that programmers can better understand what rôles a parameter is expected to play, or what can be requested of an object returned as an answer from some “foreign” code.

Why is such a taxonomy useful? To add redundancy. A type annotation is an assertion about the value of a variable, no different in concept from asserting that a collection is empty or that an object is not nil. Redundancy is in general a good thing: it provides additional information for readers, and it means that more errors — at least those errors that manifest themselves as an inconsistency between the type annotation and the code — can be detected sooner. Nevertheless, I am about to argue that types can be harmful.

On the surface, this is a ridiculous argument. If types add redundancy, and redundancy is good, how can types be harmful? The problem is not that adding type annotations will mess up your program, but that adding types to a language can, unless one is very careful, mess up the *language design*.

I have been an advocate of types for many years. The Emerald type system [17, 40] was one of the ways in which the Emerald project attempted

to improve on Smalltalk. Our original perception was that types would enable us to generate more efficient code. However, as we implemented the language, we realised that behavioural types, that is, types that constrain the interface of an object, but not its implementation, did not help us to generate more efficient code. Emerald was efficient, but that was because we inferred implementation information, for example, that no reference to a particular object ever escaped an enclosing object.⁵ I spent many years working on type-checking and type systems, but have only recently really understood the sage advice that I was given by Butler Lampson in the late 1980s: stay away from types — they are just too difficult.

The difficulty springs from two sources. One is Gödel’s second incompleteness theorem, which tells us that there *will be* true facts about any (sufficiently rich) formal system that cannot be proved, no matter what system of proof we adopt. In this case, the formal system is our programming language, and the system of proof is that language’s type system. The other source is our very natural desire to *know* that our program won’t go wrong, or at least that it won’t go wrong in certain circumscribed ways. This has led most statically-typed languages to adopt an interpretation of type-checking that I am going to refer to as “The Wilson interpretation”. The name honours Harold Wilson, prime minister of the UK 1964–70 and 1974–76. Wilson fostered what has been called “the Nanny State”, a social and political apparatus that said the government will look after you: if it is even remotely possible that something will go wrong, we won’t even let you try. The Wilson interpretation of type-checking embraces two tenets: first, that the type system must be complete, that is, every type assertion that can be made in it must be provably true or false, and second, that every part of the program must be type-checked before it can be executed. The consequence is that if it is even remotely possible that something may go wrong when executing your program, then the language implementation won’t even let you try to run it — assuming that the “something” that might go wrong is a thing over which the type system believes it has control.

⁵In one place we did get efficiency from types, but that was because we cheated. Emerald made it impossible to re-implement integers: a user-defined object would *never* be recognised as having type `Integer`, even though it provided all of the necessary operations. This was cheating because Emerald types are supposed to constrain interface, not implementation. The consequence was that we could indeed infer representation information from the `Integer` type, and directly invoke machine operations to perform arithmetic.

The Wilson interpretation is not the only interpretation of type-checking. At the other extreme is what I will call the “Bush interpretation”, honouring George W. Bush, president of the USA 2001–2009, who abolished many government regulations and weakened the enforcement of those that remained. The Bush interpretation of type checking is that type information is advisory. The program should be allowed to do what you, the programmer, said that it should do. The type-checker won’t stop it; if you mess up, the PDIC — the Program Debugger and Interactive Checker — will bail you out. The Bush interpretation amounts to not having static type-checking at all: the programmer is allowed to write type annotations in the program, but they won’t be checked until runtime.

There is a third interpretation of type-checking that lies between these extremes. This might be called the “Proceed with caution” approach. If the type-checker has been unable to prove that there are no type errors in your program, you are given a warning. “It may work; it may give you a run-time error. Good night, and good luck.” I’m going to name this interpretation in honour of Edward R. Murrow, 1908–1965, an American broadcaster, whose signature line that was. Like Murrow, I intend that phrase to be taken positively.

I’m referring to Wilson, Murrow, and Bush as *interpretations* of type-checking, because I’m taking the view that the semantics of the language, in the absence of errors, is independent of its type system. In other words, the language has an untyped semantics, and the rôle of the type system is to control how and when errors are reported, not to determine the meaning of the code. There is an alternative view in which the meaning of a program depends on the types of its terms; programs that don’t type-check *have no semantics*. In this view, Wilson, Murrow, and Bush define different languages, connected by a superset relation. Since, as we have seen, types are not necessary to define object-based abstraction, I feel that the untyped view is simpler and more appropriate, but the comparison of these views is a topic worthy of its own essay, and I’m not going to discuss it further here.

So, under all three interpretations, an error-free program has the same meaning. Under Wilson, we have conventional static typing: an erroneous program will result in a static error, and won’t be permitted to run. Because of Gödel’s theorem, this means that *some* error-free programs won’t be permitted to run either, no matter how complex the type system becomes.

Under Bush, we have conventional dynamic typing: all checks will be

performed at runtime—even those that are guaranteed to fail. Enthusiasts for dynamic typing point out that a counter-example is often more useful than a type-error message. The disadvantage is that the programmer will not receive any compile-time warnings. Under the Murrow interpretation, the programmer will get a mix of compile-time warnings and run-time checks. A Murrow warning may say that a certain construct is provably a type error, or it may merely say that the type-checker has been unable to prove that it is type safe. The difference between these two warnings is dramatic, but Wilson treats them as if they were the same, conflating a deficiency in the type system with a programming error.

Let me say here that I'm for Murrow! I believe that the Murrow interpretation of types is the most useful, not only for programmers, but also for language designers. Wilson's "Nanny Statism" is an invitation to mess up your language design. The temptation for the language designer is to exclude any construct that can't be statically checked, regardless of how useful it may be to the practicing programmer.

I believe that SIMULA was for Murrow too. Recall that, according to Nygaard, the core ideas of SIMULA were first modelling, and second, security. Modelling meant that the actions and interactions of the objects that the program created represented the actions and interactions of the corresponding real-world objects. Security meant that the behaviour of a program could be understood and explained entirely in terms of the semantics of the programming language in which it is written. Modelling came first! SIMULA did not compromise its modelling ability to achieve security; it compromised its run-time performance, by incorporating explicit checks when a construct necessary for modelling was not statically provable to be safe.

I feel that many current language designers are suffering from a Wilson obsession. This obsession has resulted in type systems of overwhelming complexity, and languages that are larger, less regular, and less expressive than they need to be. The fallacy is the assumption that, because type checking is good, *more* type-checking is necessarily better.

Let's consider an example of how the insistence on static type-checking can mess up a language design. I'm currently engaged, with Kim Bruce, James Noble and their students, in the design of a language for teaching programming. The language is called Grace, both to honour Rear Admiral Grace Hopper, and because we hope that it will enable the *Graceful* expression of programs. Grace is an object-oriented language, and contains an

inheritance facility, inspired by early proposals by Antero Taivalsaari [41]. Here is an example:

```
class aDictionary.ofSize( initialSize ) {  
  inherits aHashtable.ofSize( initialSize )  
  method findIndex (predicate) is override { ... }  
  method at (key) put (value) is public { ... }  
  ...  
}
```

In Grace, classes are nothing more than objects that have a single method that, when executed, creates a new object. This example declares a class object named `aDictionary`, which has an object-creation method with the name `ofSize`. Executing `aDictionary.ofSize(10)` will create and return a new object that inherits (that is, contains copies of) all of the methods and instance variables of the object `aHashtable.ofSize(10)`, except that the given method for `findIndex()` will override that inherited from the `Hashtable` object, and the given method for `at()put()` will be added.

Checks that one might expect to take place on such a definition would include that `aHashtable` actually does have a method `ofSize`, and that the object answered by that method does have a `findIndex()` method and does *not* have an `at()put()` method. There is no difficulty constructing a static type system to perform such checks so long as `aHashtable` is a globally known class. However, suppose that I want to let the client choose the actual object that is extended, perhaps because prefix trees are better than hash tables for some applications. In other words, suppose that I want to make the super-object a parameter:

```
class aDictionary.basedOn(superObj) {  
  inherits superObj  
  method findIndex (predicate) is override { ... }  
  method at (key) put (value) is public { ... }  
  ...  
}
```

Although the same implementation mechanisms will still work, constructing a static type system to perform the checking is no longer simple. What arguments may be substituted for `superObj`? The requirement that `superObj` have a (possibly protected) method `findIndex` (because of the `override` annotation) and that it *not* have a method `at()put()` (because of the absence of an `override` annotation) cannot be captured by the usual notion of type, which

holds only the information necessary to *use* an object, not the information necessary to inherit from it. One solution to the problem of type-checking this definition in a modular way is to invent a new notion of “heir types”, that is, types that capture the information needed to check inheritance, such as the presence of a method that is either public or protected, or the absence of a method.

Another possible solution is to make classes a new sort of entity, different from objects, and to give them their own, non-Turing-complete sublanguage, including class-specific function, parameter and choice mechanisms. A third possibility is to ban parametric super-objects. How then does one use inheritance, that is, how does the programmer gain access to a superclass to extend? The classical solution is to refer to classes by global variables, and to require that the superclass be globally known. This is a premature commitment that runs in opposition to the late-binding that otherwise characterises object-orientation; its effect is to reduce reusability. Some languages alleviate this problem by introducing yet another feature, known as open classes, which allows the programmer to add methods to, or override methods in, an existing globally-known class.

Virtual classes, as found in BETA [42], are yet another approach to the problem of parametrizing inheritance. A virtual class can be overridden in a subclass by a “more specialised” class. Virtual classes don’t by themselves *solve* our problem: they need to be accompanied by a set of type rules that specify when it is legal to override a virtual class. The difficulty is that one set of rules — conventional subtyping — is appropriate for clients of the class, and another set of rules — heir typing — is appropriate for inheritors. BETA chooses to use co-variant method subtyping for virtual classes, that is, to allow the arguments of methods to be specialised along with their results. This means that virtual classes cannot be guaranteed to be safe by a modular static analysis. Nevertheless, virtual classes are useful for modelling real systems; BETA follows SIMULA in ranking modelling above static checking.

Regardless of the path chosen, the resulting language is both larger and less-expressive than the simple parametric scheme sketched above. Indeed, in the scheme suggested by my original example, a new object could extend *any* existing object, and classes did not need to be “special” in any way; they could be ordinary objects that did not require any dedicated mechanisms for their creation, categorisation or use.

Another example of this sort of problem is collections that are parametrized

by types. If types are desirable, it certainly seems reasonable to give programmers the opportunity to parameterize their collection objects, so that a library implementor can offer **Bags of Employees** as well as **Bags of Numbers**. The obvious solution is to represent types as objects, and use the normal method and parameter mechanisms to parameterize collections—which is exactly what we did in Emerald [43]. Unfortunately, the consequence of this is that type checking is no longer decidable [44]. The reaction of the Wilsonian faction to this is to recoil in shock and horror. After all, if the mission of types is to prevent the execution of any program that might possibly go wrong, and the very act of deciding whether the program will go wrong might not terminate, what is the language implementor to do?

A language that takes the Murrow interpretation of type can react to the possibility that type-checking is statically undecidable more pragmatically. The language implementation will inevitably use run-time type checks, to deal with situations in which the program cannot be guaranteed to be free of type errors statically. If a type assertion cannot be proved true or false after a reasonable amount of effort, it suffices to insert a run-time check.

An alternative way of parameterising objects by types is to invent a new parameter passing mechanism for types, with new syntax and semantics, but with restricted expressivity to ensure decidability. Nevertheless, because of Gödel’s second incompleteness theorem, some programs will still be untypeable. The consequence is that the language becomes larger, while at the same time less expressive.

11. The Future of Objects

Now I’m going to turn to some speculations about the programming languages of the future. I could follow Feynman, and predict, with a certain confidence, that in 1000 years object-oriented programming will no longer exist as we know it. This is either because our civilisation will have collapsed, or because it has not; in the latter case humans will no longer be engaged in any activity resembling programming, which will instead be something that computers do for themselves.

I think that it is probably more useful, and certainly more within my capabilities, to look 10 or 20 years ahead. The changes in computing that will challenge our ability to write programs over that time span are already clear: the trend from multicore towards manycore, the need to control energy

consumption, the growth of mobility and “computing in the cloud”, the demand for increased reliability and failure resilience, and the emergence of distributed software development teams.

11.1. Multicore and Manycore

Although the exact number of “cores”, that is, processing units, that will be present on the computer chips of 2021–31 is unknown, we can predict that with current technology trends we will be able to provide at least thousands, and perhaps hundreds of thousands. It also seems clear that manycore chips will be much more heterogeneous than the two- and four-processor multicore chips of today, exactly because not all algorithms can take advantage of thousands of small, slow but power-efficient cores, and will need larger, more power-hungry cores to perform acceptably. The degree of parallelism that will be available on commodity chips will thus depend both on the imperatives of electronic design and on whether we can solve the problem of writing programs that use many small cores effectively.

What do objects have to offer us in the era of manycore? The answer seems obvious. Processes that interact through messages, which are the essence of SIMULA’s view of computation, are exactly what populate a many-core chip. The similarity becomes even clearer if we recall Kay’s vision of objects as little computers. Objects also offer us heterogeneity. There are many different classes of objects interacting in a typical program, and there is no reason that these different classes of objects need be coded in the same programming language, or compiled to the same instruction set. This is because the “encapsulation boundary” around each object means that no other object need know about its implementation. For example, objects that perform matrix arithmetic to calculate an image might be compiled to run on a general-purpose graphical processing unit, while other parts of the same application might be compiled to more conventional architectures.

11.2. Controlling Energy Consumption

How does one write an energy-efficient program? How, indeed, can one compare two programs for efficiency? A common approach is to use a *cost model*: a way of thinking about computation that is closely-enough aligned with the real costs of execution that it can help us to reason about performance. Objects can provide us with a cost model for manycore computing.

Most current computing models date from the 1950s and treat computation as the expensive resource — which was the case when those models were developed, and computation made use of expensive and power-hungry valve or discrete-transistor logic. In contrast, data movement is regarded as free, since it was implemented by copper wires, which were inexpensive compared to valves or transistors, and consumed little energy. As a consequence, these models lead us to think, for example, of moving an operation out of a loop and caching the result in memory as an “optimisation”. The reality today is very different: computation is essentially free, because it happens “in the cracks” between data fetch and data store; on a clocked processor, there is little difference in energy consumption between performing an arithmetic operation and performing a no-op. In contrast, data movement is expensive: it takes both time and energy to propagate signals along wires. So controlling the cost of computation on a manycore chip requires one to exercise control of data movement. Today’s programming languages are unable to even express the thing that needs to be carefully controlled: data movement.

I would like to propose an object-oriented cost model for manycore computing. To SIMULA’s concurrent object model we need add only three things. The first addition is to treat objects as enjoying true parallelism, rather than SIMULA’s quasi-parallelism. The second and third additions are to attribute to each object a size and a spatial location. Size is important because an object needs to fit into the cache memory available at its location. “Size” includes not just the data inside the object, but also the code that makes up its method suite, since both must be cached.

With this model, local operations on an object can be regarded as free, since they require only local computations on local data. Optimising an object means reducing its size until it fits into the available cache, or, if this is not possible, partitioning the object into two or more smaller objects. In contrast to local operations, requesting that a method be evaluated in another object may be costly, since this may require communicating with an object at a remote location. The cost of such a message is proportional to the product of the amount of data in it and the distance to the receiver; we can imagine measuring cost in byte-nanometers. The cost of a whole computation can be estimated as proportional to the number of messages sent and the cost of each message and reply. We can reduce this cost by relocating objects so that frequent communication partners are close together, and by recomputing answers locally rather than requesting them from a remote location. Such a

cost model does not, of course, capture all of the costs of a real computation, but it seems to me that it will provide a better guide for optimisation than models from the 1950s.

11.3. Mobility and the Cloud

The world of mobile computing is a world in which communication is transient, failure is common, and replication and caching are the main techniques used to improve performance and availability. The best models for accessing objects in such an environment seem to me to be those used for distributed version control, as realised in systems like subversion and git [45, 46]. In such systems, all objects are immutable; updating an object creates a new version. Objects can be referred to in two ways: by version identifier and by name. A version identifier is typically a long numeric string and refers to (a copy of) a particular immutable object. In contrast, a name is designed for human consumption, but the binding of a name to a version object changes over time: when a name is resolved, it is by default bound to a recent version of the object that it describes.

Erlang uses a similar naming mechanism to realise its fail-over facility. Erlang messages can be sent either to a process identifier or to a process name. A particular process identifier will always refer to the same process; sending a message to a process identifier will fail if the process is no longer running. In contrast, a process name indirects through the name server, so messages sent to a process name will be delivered to the most recent process to register that name. When a process crashes and its monitor creates a replacement process, the replacement will usually register itself with the name sever under the same name that was used by the crashed process. This ensures that messages sent to the process name will still be delivered, although successive messages may be delivered to different processes [19]. Perhaps a similar mechanism should be part of every object-oriented system? This would mean that we would be able to reference an object either by a descriptor, such as “Most recent version of the Oslo talk”, or by a unique identifier, such as Object16x45d023f. The former would resolve to a replica of a recent version; the latter would resolve to a specific object.

11.4. Reliability

Total failures are easy for the programmer to handle, because there is nothing that the programmer can do! However, they can be disastrous for the users of

a computer system, who are prevented from doing their work. As distributed systems have become ubiquitous, total failures have become rare, and partial failures much more common. This trend will accelerate as the size of the features on a chip decreases, and massively manycore chips proliferate. The reliability of individual transistors on such chips will be much lower than we are accustomed to today: this is a consequence of the physics of doping. By 2014, if there are 100 billion transistors on a chip, 20 percent of them will not work when it is fabricated, and another 10 percent will fail over the life of the chip⁶. This means two things: that processor design will have to change quite radically, and that single chips will have reliability characteristics similar to today's Internet, where computers may fail at any time.

What techniques are available for dealing with partial failures? It is often possible to mask a failure using replication in time or space, but this is not always wise. Resending lost messages is not a good idea if they contain out-of date information; it may be better to send a new message containing fresh information. Similarly, it may be more cost-effective to recompute lost data than to retrieve it from a remote replica. So automatic masking of failure by the programming model is not a good idea. Instead, it seems to me that a variety of facilities for dealing with failures must be visible in the programming model, so that programmers can choose whether or not to use them. What, then, should be the unit of failure in an object-oriented model of computation? Is it the object, or is there some other, larger, unit? Whatever unit we choose, it must "leak" failure, in the sense that its clients must be able to see it fail, and then take some action appropriate to the context.

11.5. Distributed Development Teams

With the ubiquity of computing equipment and the globalisation of industry, distributed software development teams have become common. You may wonder what this trend has to do with objects. The answer is packaging: collaborating in loosely-knit teams demands better tools for packaging code and sharing it between the parts of a distributed team. Modules, by which I mean the unit of code sharing, are typically parametrized by other modules.

⁶These predictions are from Shekhar Borkar, an Intel fellow and Director of the Microprocessor Technology lab. They were presented by Bob Colwell at FCRC, San Diego, 2007.

To take full advantage of the power of objects, module parameters should be bound at the time that a module is used, not when it is written. In other words, module parameters should be “late bound”. This means that there is no need for a global namespace in the programming language itself; URLs or versioned objects provide a perfectly adequate mechanism for referring to the arguments of modules.

12. The Value of Dynamism

Before I conclude, I’m going to offer some speculations on the value of dynamism. These speculations are motivated by the fact that in designing SIMULA, Dahl recognised that the *runtime structures* of Algol 60 already contained the mechanisms that were necessary for simulation, and by Nygaard and Dahl’s understanding that it is the *runtime behaviour* of a simulation program that models the real world, not the program’s text. I see a similar recognition of the value of dynamism in Agile software development, a methodology in which a program is developed in small increments, in close consultation with the customer. The idea is that a primitive version of the code runs at the end of the first week, and new functionality is added every week, under the guidance of the customer, who determines which functions have the greatest value. Extensive test suites make sure that the old functionality continues to work while new functionality is added.

If you have never tried such a methodology, you may wonder how it could possibly work! After all, isn’t it important to *design* the program? The answer is yes: design is important. In fact, design is *so* important, that Agile practitioners don’t do it only at the start of the software creation process, when they know nothing about the program. Instead, they design every day: the program is continuously re-designed as the programmers learn from the code, and from the behaviour of the running system.

What I’ve learned from watching this process is that the program’s runtime behaviour is a powerful teaching tool. This is obvious if you view programs as existing to control computers, but perhaps less obvious if you view programs as static descriptions of a system, like a set of mathematical equations. As Dahl and Nygaard discovered, it is a program’s behaviour, not the program itself, that models the real-world system. The program’s text is a meta-description of the program behaviour, and it is not always easy to infer the behaviour from the meta-description.

As a teacher of object-oriented programming, I know that I have succeeded when students anthropomorphise their objects, that is, when they turn to their partners and speak of one object asking another object to do something. I have found that this happens more often, and more quickly, when I teach with Smalltalk than when I teach with Java: Smalltalk programmers tend to talk about objects, while Java programmers tend to talk about classes. I suspect that this is because Smalltalk is the more dynamic language: the language and the programming environment are designed to help programmers interact with objects, as well as with code. Indeed, I am tempted to define a “Dynamic Programming Language” as one designed to help the programmer learn from the run-time behaviour of the program.

13. Summary

Fifty years ago, Dahl and Nygaard had some profound insights about the nature of both computation and human understanding. “Modelling the world” is not only a powerful technique for designing and re-designing a computer program: it is also one of the most effective ways yet found of communicating that design to other humans.

What are the major concepts of object-orientation? The answer depends on the social and political context. Dahl listed five, including the use of objects as record structures and as procedurally-encapsulated data abstractions, and the use of inheritance for incremental definition. Some of what he thought were key ideas, such as active objects and objects as modules, have been neglected over the last 30 years, but may yet be revived as the context in which we program becomes increasingly distributed and heterogeneous, in terms of both execution platform and programming team. It certainly seems to me that, after 50 years, there are still ideas in SIMULA that can be mined to solve twenty-first century problems. There may not be any programming a thousand years from now, but I’m willing to wager that some form of Dahl’s ideas will still be familiar to programmers in fifty years, when SIMULA celebrates its centenary.

Acknowledgements

I thank Stein Krogdahl, Olaf Owe and the committee of FCT11 for honouring me with the invitation to speak at the scientific opening of the Ole-Johan Dahl hus, as well as supplying me with information about Dahl and

SIMULA. I also thank Olaf for convincing me to turn my lecture into this article. William Cook was an encouraging and through critic; Ivan Sutherland helped me understand where the costs lie in modern processors, and provided useful feedback on a draft of my manuscript. David Ungar shared many insights with me and helped to improve both my lecture and this article. The anonymous referees were especially helpful in pointing out errors of fact and passages in need of clarification.

References

- [1] FCT 2011, Scientific opening of the Ole-Johan Dahl building, 2011. Web page last visited 2 March 2013. <http://fct11.ifi.uio.no/index.php?n=General.SocialEvents>.
- [2] K. Nygaard, O.-J. Dahl, The development of the SIMULA languages, in: R. L. Wexelblat (Ed.), *History of programming languages I*, ACM, New York, NY, USA, 1981, pp. 439–480.
- [3] University of Oslo, Department of Informatics, Kristen Nygaard—Education and Career, 2011. Web page last visited 25 December 2011. <http://www.mn.uio.no/ifi/english/about/kristen-nygaard/career/>.
- [4] R. Feynman, Speech at Centennial Celebration of Massachusetts Institute of Technology (version C), 1961. The Feynman Archives at California Institute of Technology.
- [5] C. Hoare, Hints on Programming Language Design, Memo AIM-224, Stanford Artificial Intelligence Laboratory, 1973. Invited address, 1st POPL conference.
- [6] G. Radin, The early history and characteristics of PL/I, *SIGPLAN Notices* 13 (1978) 227–241.
- [7] O.-J. Dahl, K. Nygaard, Class and subclass declarations, in: J. N. Buxton (Ed.), *Simulation Programming Languages*, Proceedings from the IFIP working conference in Oslo, North Holland, Amsterdam, 1967, pp. 158–174.
- [8] O.-J. Dahl, Transcript of discussant’s remarks, in: R. L. Wexelblat (Ed.), *History of programming languages I*, ACM, New York, NY, USA, 1981, pp. 488–490.

- [9] G. Bracha, P. von der Ahé, V. Bykov, Y. Kashi, W. Maddox, E. Miranda, Modules as objects in newspeak, in: T. D'Hondt (Ed.), ECOOP, volume 6183 of *Lecture Notes in Computer Science*, Springer, 2010, pp. 405–428.
- [10] A. P. Black, K. B. Bruce, M. Homer, J. Noble, Grace: the absence of (inessential) difficulty, in: *Onward! '12: Proceedings 12th SIGPLAN Symp. on New Ideas in Programming and Reflections on Software*, ACM, New York, NY, 2012, pp. 85–98.
- [11] O.-J. Dahl, C. Hoare, Hierarchical program structures, in: *Structured Programming*, Academic Press, 1972, pp. 175–220.
- [12] C. A. R. Hoare, Record handling, *ALGOL Bull.* 21 (1965) 39–69.
- [13] F. Genuys (Ed.), *Programming Languages: NATO Advanced Study Institute*, Academic Press, London, New York, 1968.
- [14] O.-J. Dahl, The roots of object-oriented programming: the Simula language, in: M. Broy, E. Denert (Eds.), *Software Pioneers: Contributions to Software Engineering*, Springer-Verlag, Berlin, Heidelberg, 2002, pp. 79–90.
- [15] P. Wegner, Dimensions of object-based language design, in: N. Meyrowitz (Ed.), *Proceedings Second ACM Conference on Object-Oriented Programming Systems, Languages and Applications*, ACM Press, Orlando, Florida, 1987, pp. 168–182.
- [16] A. P. Black, N. Hutchinson, E. Jul, H. Levy, Object structure in the Emerald system, in: *Proceedings First ACM Conference on Object-Oriented Programming Systems, Languages and Applications*, ACM Press, Portland, Oregon, 1986, pp. 78–86.
- [17] A. P. Black, N. Hutchinson, E. Jul, H. M. Levy, L. Carter, Distribution and abstract types in Emerald, *IEEE Trans. Software Engineering SE-13* (1987) 65–76.
- [18] W. R. Cook, *A Denotational Semantics of Inheritance*, Ph.D. thesis, Brown University, Department of Computer Science, 1989.

- [19] J. Armstrong, *Programming Erlang: Software for a concurrent world*, Pragmatic Bookshelf, 2007.
- [20] B. Archer, *Programming quotations*, 2011. Web page last visited 22 January 2012. http://www.bobarcher.org/software/programming_quotes.html.
- [21] P. Wadler, *The essence of functional programming*, in: *Conference Record of the Nineteenth ACM Symposium on Principles of Programming Languages*, ACM Press, Albuquerque, NM, 1992, pp. 1–14.
- [22] P. Wadler, *Comprehending monads*, *Mathematical Structures in Computer Science* 2 (1992) 461–493. Originally published in *ACM Conference on Lisp and Functional Programming*, June 1990.
- [23] A. C. Kay, *The early history of Smalltalk*, in: *The second ACM SIGPLAN conference on History of programming languages, HOPL-II*, ACM Press, New York, NY, USA, 1993, pp. 511–598.
- [24] A. Snyder, *The Essence of Objects: Common Concepts and Terminology*, Technical Report HPL-91-50, Hewlett Packard Laboratories, 1991.
- [25] A. Snyder, *The essence of objects: Concepts and terms*, *IEEE Software* 10 (1993) 31–42.
- [26] S. Krogdahl, *Concepts and terminology in the Simula programming language—an introduction for new readers of Simula literature*, 2010. <http://folk.uio.no/simula67/Archive/concepts.pdf>.
- [27] D. H. Ingalls, *Design principles behind Smalltalk*, *Byte* 6 (1981) 286–298.
- [28] A. H. Borning, D. H. H. Ingalls, *A type declaration and inference system for Smalltalk*, in: *Conference Record of the Ninth ACM Symposium on Principles of Programming Languages*, ACM Press, Albuquerque, NM, USA, 1982, pp. 133–141.
- [29] C. Hoare, *Proof of correctness of data representations*, *Acta Informatica* 1 (1972) 271–281.
- [30] D. L. Parnas, *On the criteria to be used in decomposing systems into modules*, *Comm. ACM* 15 (1972) 1053–1058.

- [31] B. Liskov, A. Snyder, R. Atkinson, C. Schaffert, Abstraction mechanisms in CLU, *Comm. ACM* 20 (1977) 564–576.
- [32] A. Church, The Calculi of Lambda-Conversion, volume 6 of *Annals of Mathematical Studies*, Princeton University Press, 1941.
- [33] J. C. Reynolds, User-defined types and procedural data structures as complementary approaches to data abstraction, in: S. A. Schuman (Ed.), *Conference on New Directions in Algorithmic Languages*, IFIP Working Group 2.1 on Algol, INRIA, Munich, Germany, 1975, pp. 157–168.
- [34] J. C. Reynolds, User-defined types and procedural data structures as complementary approaches to data abstraction, in: D. Gries (Ed.), *Programming Methodology: A Collection of Articles by members of IFIP WG2.3*, Texts and Monographs in Computer Science, Springer Verlag, New York, Heidelberg, Berlin, 1978, pp. 309–317.
- [35] J. C. Reynolds, User-defined types and procedural data structures as complementary approaches to data abstraction, in: C. A. Gunter, J. C. Mitchell (Eds.), *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, Cambridge, MA, 1994, pp. 13–23.
- [36] W. R. Cook, On understanding data abstraction, revisited, in: S. Arora, G. T. Leavens (Eds.), *Proceedings 24th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, ACM, 2009, pp. 557–572.
- [37] C. Hewitt, P. Bishop, R. Steiger, A universal modular actor formalism for artificial intelligence, in: N. J. Nilsson (Ed.), *IJCAI*, William Kaufmann, Stanford University, California, 1973, pp. 235–245.
- [38] J. Armstrong, Concurrency oriented programming in Erlang, in: *Frühjahrsfachgespräch 2003 (FG 2003)*, German Unix Users Group, Köln, Germany, 2003.
- [39] A. Borning, Classes versus prototypes in object-oriented languages, in: *FJCC*, IEEE Computer Society, 1986, pp. 36–40.

- [40] A. P. Black, N. Hutchinson, Typechecking Polymorphism in Emerald, Technical Report Technical Report CRL 91/1 (Revised), Digital Equipment Corporation Cambridge Research Laboratory, 1991.
- [41] A. Taivalsaari, Delegation versus concatenation, or cloning is inheritance too, SIGPLAN OOPS Mess. 6 (1995) 20–49.
- [42] O. L. Madsen, An overview of BETA, in: J. Knudsen, O. Madsen, B. Magnusson (Eds.), Object-Oriented Environments, Prentice Hall, Englewood Cliffs, NJ, 1993, pp. 99–118.
- [43] A. P. Black, N. C. Hutchinson, E. Jul, H. M. Levy, The development of the Emerald programming language, in: B. G. Ryder, B. Hailpern (Eds.), HOPL III: Proceedings Third ACM SIGPLAN conference on History of Programming Languages, ACM Press, San Diego, CA, 2007, pp. 11–1–11–51.
- [44] A. R. Meyer, M. B. Reinhold, ‘Type’ is not a Type: Preliminary report, in: Conference Record of the Thirteenth ACM Symposium on Principles of Programming Languages, ACM Press, St Petersburg Beach, FL, USA, 1986, pp. 287–295.
- [45] M. Mason, Pragmatic Version Control Using Subversion, Pragmatic Bookshelf, Pragmatic Programmers, 2005.
- [46] C. Duan, Understanding Git: Repositories, 2012. <http://www.eecs.harvard.edu/~cdan/technical/git/git-1.shtml>. Accessed on 3 Aug 2012.