Specifications Are Necessarily Informal or: The Ultimate Myths of Formal Methods *

Baudouin Le Charlier
Institut d'Informatique
Facultés Universitaires de Namur, B - 5000 Namur, Belgium
Email: ble@info.fundp.ac.be

Pierre Flener
Department of Computer Engineering and Information Science
Bilkent University, 06533 Bilkent, Ankara, Turkey
Email: pf@cs.bilkent.edu.tr

Abstract

We reconsider the concept of specification in order to bring new insights into the debate of formal versus non-formal methods in Computer Science. In our view, a specification is the link between the program (formality) and its purpose (informality). Since the purpose of a program must be something directly understandable, specifications are the essential tool for constructing, in practice, correct real-world programs through explicit, yet not completely automatable reasoning. This allows us to explain why formal specifications cannot meet the demands of specifications in our sense, since this would be a contradiction in terms. Our view of specifications does not imply a rejection of all ideas put forward in the literature on formal methods. On the contrary, we agree with the proponents of formal methods on most of their arguments, except on the fact that specifications had better be written in a formal specification language. And, inevitably, we also disagree with arguments following from this assumption. Finally, we examine why the role and nature of specifications are so often misunderstood.

1 Introduction

> *Gist specifications were nearly as hard to read as those in other formal specification languages. We soon realized that the problem was not particular to Gist, but extant across the entire class of formal specification languages. In their effort to be formal, all these languages have scrubbed out the mechanisms which make informal languages understandable, such as summaries and overviews, alternative points of view, diagrams, and examples.*
> — R. Balzer, in [1]

Recently, there has been a flurry of papers advocating the use of "formal methods" in the software industry (see [3, 4, 5, 11, 13, 14, 16, 19], some in [23], ...). Similar opinions were sporadically published before (some in [7], [10, 15, 18, 21, 24], ...). Academicians, with and without industrial experience, apologetically missionarize for formal methodism, under various degrees of radicalism. Sometimes, they even berate industrial software engineers for not using such supposedly formal methods, warning them of imminent disasters if these methods are not adopted, as they are perceived to be a key solution to the chronic software crisis (or plague, rather). In some cases, even practicing engineers are found preaching the gospel of formal methods to their fellow members of the industrial congregation ([16, 19], some in [23]).

*Contrary to the previous papers of the "myths series" [16, 4], this paper is not about industry-level myths on the uselessness of formal methods, but rather about the academic myths on their usefulness. True to the tradition, we discuss our own list of seven myths on formal methods, namely in Section 3.2.

Other industry members, with and without academic experience, have been fighting back, exposing fallacies in the academicians' assumptions (such as that there are essential differences between engineering and mathematics in general, and between computing and mathematics in particular), if not chastising them for their bad attitude (others in [23]). Indeed, why is there so much aggressive publishing in favor of formal methods as of late? If these methods are worth their money, they will sooner or later speak for themselves. The fact that they have not after over twenty years of research speaks volumes.

The same debate is raging inside academe, as far as the teaching [7] of computer science is concerned. Should the curriculum include formal methods or not? To what extent?

Myths circulating in industry on the inadequacy of formal methods for real-world applications have been "debunked" with great zeal and even greater lack of convincing arguments [16, 4].[1] But curiously, very few voices *inside academe* are heard arguing against formal methods, or at least begging for caution (but see DeMillo, Lipton, and Perlis [6], Fetzer [8], Karp in [7], Parnas [22] and in [23], Winograd in [7]).

Simultaneously, there is a debate on whether formal specification languages ought to be executable or not [17, 12]. But curiously again, almost nobody *inside academe* challenges the now nearly universal contention that specifications ought to be formal in the first place (but see Balzer [2, 1], Parnas [22] and in [23]).

Does this relative silence of academe result from mere conformism, or is Hilbert's ghost still alive and making ignorant converts by the legion?

Our (academicians') objective is to stir the academic community from its Cinderella sleep by (re-)raising a few itchy issues that are almost always (conveniently or accidentally) overlooked. As proponents of mathematical rigor based on real (i.e., non-formal specifications), we thus also provide industry with new arguments against formal methods. Admittedly, the software industry rarely takes the right decisions in its methodological choices, witness the current fad about abstract datatypes and object orientation. But as far as formal methods are concerned, the reluctance of industry seems justified, although we doubt this near-rejection is based on the right reasons.

We have chosen a pious vocabulary so far to illustrate that we are here touching on nearly religious values, if not tackling what very much looks like a Byzantine discussion. We are aware that very few people will thus be open-minded enough to lucidly examine our arguments. Since we have not seen our analysis elsewhere (at least not in print, nor recently), we have taken the pen to write a much-needed, contemporary, intra-academic rebuttal of the prevalent formalist current. We thus simply hope motivating formalists to reconsider their position, so that some crucial issues be reexamined and, hopefully, the lopsidedness of printed academic opinion in favor of formal methods become corrected.

In this paper, we reconsider the concept of specification in order to bring new insights into the debate of formal versus non-formal methods in computer science. In our view, a specification is the link between the program (formality) and its purpose (informality). Since, as we will argue, the purpose of a program must be something directly understandable, specifications are the essential tool for constructing, in practice, correct real world programs through explicit reasoning. Additionally, our discussion of specifications allows us to explain why formal specifications (i.e., specifications written in a formal specification language) are not really specifications, since this would be a contradiction in terms.

Our view of specifications does not imply a rejection of all ideas put forward in the literature on formal methods. On the contrary, we agree with the proponents of formal methods on most of their arguments, except on the fact that specifications had better be written in a formal, i.e., completely pre-defined and syntactically checkable, language. And, inevitably, we also disagree with other arguments that are a consequence of this assumption that formal specification languages are desirable.

Formal methods are in general introduced as being the use of mathematics in the process of constructing computer software (including the elaboration of specifications). We agree that mathematics are extremely useful in this context, but we disagree on reducing the concept of

[1] Indeed, overly general consequences are drawn from a mere handful of supposedly successful formal methods-based projects, completely and conveniently ignoring that said projects were led by formalists, probably even staffed by formalists, and their evaluation inevitably skewed because of the Hawthorne effect (the teams did extra well because they knew they were being observed and that they were about to design showcase examples).

mathematics for computer science to the restricted framework of any formal specification language.

Program verification is advocated by most distinguished computer scientists as the only way to improve the quality of software. We agree that program verification or, better, systematic program construction is the only way to build satisfactory computer software, but we disagree on the fact that program proofs must be automated, since, as we try to demonstrate, this would imply a vicious circle.

Requirements engineering is viewed by most authors as the most crucial stage in the development of a large software system. We agree on this viewpoint and especially on the importance of the elicitation process, but we disagree with the opinion that writing formal specifications is the best basis for the elicitation process: such a process is best achieved in a language as expressive as possible, i.e., a natural language enhanced with any desired notational conventions.

Finally, it is generally accepted that formal methods should be supported by corresponding software tools. We argue that formal descriptions of any kind (programs, finite-state automata, "declarative" descriptions, and the like) can be useful only because they can be the input of an automated process whose output provides directly understandable information that could not be realistically discovered by manual calculation. Nevertheless, the elaboration of any formal description (of whatever nature) requires a careful construction process that cannot be formalized in any way since this would entail a *regressum ad infinitum*.

This paper is based on the Ph.D. dissertation of the first author [20] (and includes translations of tracts of this thesis), as well as on numerous discussions between the two authors. The second author has used some of these ideas for debunking some of the myths on deduction-based and induction-based approaches to the (semi-)automatic synthesis of (logic) programs [9].

The remainder of this paper is organized as follows. In Section 2, we elaborate on our view of specifications. Based on the evidence that a program must be useful in some sense and must therefore have an understandable purpose, we argue that the specification of a program is precisely the link that allows the user to understand the results of the program in the most direct way. We show that such specifications can be the basis of explicit, yet not completely automatable reasoning allowing us to construct programs in a systematic way with optimal confidence (obviously, absolute confidence is never achievable). Finally, we explain that good specifications in our sense require the existence or the elaboration of an adequate theory and we relate this issue to the classical notion of requirements specifications. In Section 3, we use the previous discussion to demonstrate that formal specifications cannot meet the demands of specifications in our sense and we answer a number of frequently asked questions about formal methods. Section 4 contains the conclusion, which examines why the role and nature of specifications are so often misunderstood.

2 The Role and Nature of Specifications

In this section, we more closely examine specifications of programs. Experience shows that their role and nature are extremely poorly understood by most computer scientists. Program specifications are, in general, very abysmally written because there was no understanding of what to put into them, and what to omit from them. But specifications are the essential pivot of the whole programming activity: without good specifications, it is impossible to understand the concept of correctness of a program, and hence to reason *rigorously* while constructing it or constructing a program using it.

In the software engineering literature, the word 'specification' is used to designate many different kinds of things (such as requirements specifications —for an entire software— and detailed-design specifications —for its modules—), and yet there is something in common to all of them. For the moment, we deliberately do not make precise the kind of specification that we consider, but we will come back to this issue in Section 2.7.

2.1 Why and How can a Program be Useful?

Despite all the doubts one might have about the purpose of computers for the resolution of real problems such as the creation of a more just and harmonious society, if one writes and uses programs then it is because one believes they are useful. This fact is so evident that one never wonders why

and how a program can be useful. However, it is the answer to that question that leads to an understanding of what programming is and why specifications play a fundamental role in it.

If a program is useful, it is not because its execution results in displaying certain strings on the screen or in changing the contents of the computer memory in a certain way. It is because this execution yields useful information or provides substantial help in the realization of a task. But, to take advantage of the program, other things than its text and the format of its data need to be known. Even observing its behavior for some time does not suffice. It must be possible to interpret the produced results, but the knowledge necessary to this cannot be part of the program text nor of its results. It is relative to concepts totally alien to the objects manipulated by the program, and to the conventions according to which these objects represent these concepts.

Example: The Belgian National Lottery. Suppose all we know about a certain program is how to launch it on a certain computer and that its execution only results in displaying the string:

$$5, 11, 15, 22, 29, 46$$

No information can be drawn from this; our lives are unaffected by the knowledge that the execution of a certain program gives exactly this result. Now suppose, to the contrary, that we know from an informed source that the execution results in displaying the next draw of the Belgian national lottery. This changes everything: everybody now sees how such a program can be used advantageously ...

This single example shows why a program is "not useful" by itself, but only in conjunction with some knowledge that is totally outside of it, of which neither its text nor its results can give the slightest clue. Some will now object that it is easy to change that program so that it exhibits its own purpose, say by displaying the following string instead:

5, 11, 15, 22, 29, 46 is the next draw of the Belgian national lottery.

But this objection is flawed for two reasons. First, it is not the simple observation of the result that allows us to understand it. The act of "seeing" the string above cannot possibly give the necessary knowledge to the understanding of the sentence it represents. This knowledge must be available before or must be acquired by other means. Second, it is not enough to be able to interpret the result of a program by an assertion in order to deduce from it whether it is true. To do so, there should be other good reasons to believe that an execution of a program can only produce outputs that represent true assertions.

Finally, if a program can be useful, even though its manipulated objects have by themselves no meaning, it is because it is possible to use these objects to represent useful information so as to be able, first, to write the program so that it computes the representations in a correct way (according to chosen conventions), and, second, to "easily finish the job" by interpreting the results. Programming is a worthwhile activity because we are able to imagine a huge variety of representation conventions that are satisfactory from these two viewpoints (it suffices to think that everything that can be expressed in text form can be represented by a string) and because it is unnecessary to express these conventions inside the program. (It is impossible anyway.)

Example: A payroll program. Let us now consider the payroll program of a company. It is useful to the extent that it is easier to (correctly) solve the payroll problem with it than without it. In any case, it is not the program that solves the problem. The problem is solved if and only if the whole personnel gets their due salary at the deadline. This happens or does not happen independently of the existence of a payroll program and its results. The responsibility of the solving of the payroll problem belongs to the corresponding accountant. The program can only help her as an intermediary and is only really useful if it noticeably reduces the amount of work the accountant has to do to solve the problem. The accountant's task is, on the one hand, to prepare the inputs to the program, and on the other hand, to exploit its results so that everybody gets their salary. So she must know how to use the program. This also means that she must be able to make a *reasoning* by which, knowing the inputs, knowing the usage she made of the outputs, and knowing "sufficiently many things" about the program itself, she can conclude that everybody's exact salary is paid at the deadline. Nowadays, the accountant may have almost nothing to do

to complete her task, but some verification (of whether the program performs its task) has to be done nevertheless.

Example: A search sub-program. Let us finally consider a sub-program that locates a value in an array. It is useful because one can use it as a primitive for writing a larger program, and this without worrying about how the search is done. However, to use it properly, some supplementary information must be available: how to call the sub-program and how the results are represented. One might think this example is fundamentally different from the first one. In this case, some will say, to understand the purpose of the program it suffices to know the programming language and to have the text of the program. Indeed, the latter would be so simple that one will "immediately see" what the program does. The text would define the purpose of the program. This opinion is incorrect: to understand the purpose of the program, the concept of membership in an array must be known in advance, but it is not a concept of the programming language because otherwise it would not have been necessary to write a sub-program representing it. The opinion above stems from the fact that one might recognize quite easily an array search in the program text, provided one has already done some programming beforehand, hence already knows what an array search is, for what it can be used, and what form one generally gives to programs performing it. But this does not mean that this knowledge can be derived from the program text.

This example has been chosen on purpose among the most simple and "classical" ones. It is clear, however, that in general one does not write programs solving known problems. Therefore, the knowledge of some programming concepts and methods is totally insufficient for understanding not only the purpose of a "large" program but also the one of most of its components. To understand the use of a program computing $sin(x)$ according to given representation conventions and a given precision, trigonometry and analysis notions must be known. Pretending that the program defines the corresponding approximation is only a pleasant joke, because it is not the scrutiny of this text that can give the slightest idea about trigonometry to somebody who does not already have it.

Finally, it often happens that the concepts necessary to understanding the purpose of a (sub-) program cannot be found in our "preliminary knowledge" but must be invented *ad hoc*. It is well-known that the resolution of a simple problem may necessitate the introduction of completely new ideas. Such invention is done via definitions. But there would be a vicious circle to try and explain the purpose of a program by referring to concepts only known by their definitions: this would almost amount to saying that this purpose can be understood by examining another program. To leave this vicious circle, it is necessary to give these newly defined concepts an intuitive and objective "substance," by shaping them into a theory allowing their understanding without any definitions. These ideas will be further developed in Section 2.5.

Note that there is an important difference between our notion of specification and the notion of requirements specification, which mainly consists of defining new concepts. Again, we refer to Section 2.7 for more details on this issue.

2.2 What is a Specification?

"Definition." A program specification is a statement whose role is to say (1) what purpose the program serves *and* (2) how it can be correctly used.

This "definition" is not a mathematical one, but the previous discussion will help us to understand it in detail. The definition means that the specification of a program is the necessary link between what the program computes and the information that we can deduce from its results. This link is exactly what we need to use the program or to construct it.

A specification must be simple and directly understandable. The objective of a specification is to transmit information. So there is a parallel between the notions of specification and program output. The output is meaningless by itself: it must be interpreted in order to extract the information it carries. This does not mean the particular form of the outputs is irrelevant as long as the representation conventions are known. For instance, if the task of a teller machine in Belgium is to display the balance of a bank account, not all representations are equivalent: a decimal representation of the amount expressed in Belgian Franks is acceptable, but a binary representation of the square root of the amount expressed in Turkish Lira is not. The good representation is the

one that minimizes the work that remains to be done to transform the output into the desired information. In the example above, the first representation is the only acceptable one because the customer immediately knows how much money can be withdrawn from the account, whereas a long and tedious computation would be necessary from the second representation. Similarly, the "good" specification of a program is the text that can be transformed as directly as possible into a correct understanding of the purpose of the program and of the way of using it.

Besides this analogy, there also is a fundamental difference between a specification and the results of a program. The principal role of the specification precisely is to state how to interpret the results, but there is no need for a text explaining how to interpret the specification, as otherwise one would need a specification of the specification, and a specification of the specification of the specification, ad infinitum. *Therefore, unless one completely denies the pertinence of this notion, one has to admit that a specification is a text that must be comprehensible by itself. Hence it must be written in the only language adapted to this end: natural language. We do not say that specifications ought to be written in pure natural language. It can be a technical language including problem-specific concepts and notations. But it cannot be a formal language, in the strict sense of the word (i.e., whose syntax and semantics are defined a priori).*

A specification need *not* be correct, but only correctly understandable. Since the role of a specification is to communicate the purpose of a program, the only correct means of judging the quality of a specification is to ask whether it allows every potential reader to understand conveniently and in the most direct possible way the purpose of the program.

The notion of "correctness" of a specification is thus less important than the one of "being correctly understandable." A specification can perfectly play its role, even if it lacks style, or has unorthodox phrases, if not even mistakes and contradictions. A reader may well have understood it even though she estimates it to be "incorrect" or poorly written, because it does not follow her own stylistic criteria or contains some obvious mistakes. But how is it possible to correctly understand a specification while judging it incorrect? The answer lies in the observation that the role of a specification is not to define everything that ought to be known to understand the purpose of the program, but only to *state* this purpose. Where is the difference? According to the first viewpoint, one would suppose that the knowledge necessary to use the program is entirely inside the specification (i.e., would be derivable from the specification). It would, then, be evident that an incorrect specification cannot be satisfactorily understood by itself because it would be the only reference. According to the second viewpoint, one supposes that the reader already knows almost everything on what makes the program interesting, the role of the specification being somehow to say "this is the program that you needed." In this case, the presence of some errors or quirks in the specification would not really be an insurmountable obstacle to its understanding, because the enormous quantity of things already known allows the reader to fill the gaps.

All this does not imply that specifications can be written carelessly, but only that the quality of specifications cannot be judged according to hypothetical correctness criteria. The key issue is that they communicate "the message" in the most direct way.

2.3 Why are Adequate Specifications Necessary?

The specification of a program is an indispensable aid for remembering details. After close consideration, it is even *only* such an aid, as it only has to state the purpose of the program but not all the knowledge necessary to understand its meaning. The customer must thus already know, before reading the specification of a program for the first time, everything that makes the program useful to her. She will then know that a program with this purpose exists and how it can be used. Later, she can occasionally re-read the specification, not because she has forgotten its purpose, but because she does not recall with certainty some representation details that are too arbitrary to be possible (or useful) to remember.

Specifications are not only absolutely necessary for documentation of already existing programs, but also before and during the construction of programs, for three reasons.

First, one can only construct small programs at a time. The difficulty observed in the rigorous construction (à la Dijkstra, Gries, etc.) of small programs is inherent to programming (and there is no way such techniques can ever be scaled up to constructing "real" programs), so small programs

exactly represent the limit that should not be crossed if the programming activity is ever to be mastered. The only realistic approach is thus to build "large" programs from "small" ones that are constructed independently of each other, and recursively so on (no matter whether one proceeds top-down or bottom-up). This is possible only because the specifications attached to programs allow us to consider them as new primitives of the programming language, no matter how large these programs are. *All specifications should be of the same level of complexity, namely of the utmost simplicity.*

Second, intermediate specifications are necessary as a basis for the discussion between the computer scientist and the customer, because they are, in general, of too different backgrounds for coming up with the good specification at the first time. From the specification, the computer scientist must be able to make a reasoning to convince herself that she can construct the required program, whereas the customer must be able to make a reasoning to make sure the program will provide the expected service. The specification thus takes the role of a contract.

Third, intermediate specifications are necessary during the design of an architecture for the program. Strictly linear top-down design is difficult, and the implementation of certain sub-problems may reveal inadequacies in earlier choices, forcing backtracking in the design, if not the deletion of already written code. Since programming is costly, there is a risk of trying to preserve at all cost what has already been done, even if this means going into blind alleys. A more reasonable approach is thus to write all specifications of all sub-programs before writing the first line of code. This requires mental persuasion that the program can be written using all and only the specified sub-problems. Designing such an architecture may still require backtracking, but it is less tedious to rewrite specifications than programs, and easier to persuade oneself that a program can be written than actually writing it.

2.4 Can there be Adequate Specifications (for Real-World Problems)?

We think that adequate specifications, according to our criteria, *can* be written, even for real-world problems. However, the quality of specifications depends much more on the competence of those who write them than on the usage of "methodological" guidelines.

A specification is not meant for just anybody. Only a program with a purpose should have a specification. Saying that a program has a purpose amounts to saying that somebody is able to exactly understand this purpose. So the specification of a "useful" program will always exist because somebody must be able to *say* what its purpose is. But this does not mean that just anybody can understand this specification. It is only comprehensible by somebody having the "same background" as its author, at least as far as the application domain is concerned. The existence of satisfactory specifications is thus only possible because they are only meant to be read and understood by people already knowing almost *everything* of the application domain in which the program has its purpose. This does not imply that only the specifier will be able to understand it or that this privilege is reserved for a select few. It simply means that every user of the program must first make a careful and sufficiently long study of its application domain.

In practice, it is unfortunately rare that a person understanding the purpose of a program can express it simply. Programmers, for instance, tend to give incomprehensible technical gibberish about the implementation technique and run-time behavior when prompted to explain what their programs do, instead of talking about the essentials. The absence of specifications for many actually used programs stems from an inability of many people to express themselves clearly. (As already said by others before:) *Instead of including specification rules or formalisms in computer science curricula, it would be much better to teach students how to correctly use their native language (or natural language, in general).*

Another reason for the absence of convenient specifications is that programs are often constructed by successive approximations, by trial and error, so that there cannot possibly be a convenient specification, because nobody is able to understand how to use it. But it is precisely because the programmer was unable, or thought useless, to write a specification that she, not knowing what to do and hoping to find it out progressively, constructed a mysterious program to which no specification can be attached.

A specification should have an objective meaning. Some will object to our notion of specification by saying that two different people never understand things in exactly the same way, so that we can never be sure whether a specification is correctly understood by all concerned people. However, it is not necessary that the programmer and all users of a program understand its specification in the same way. Note that such a condition is insufficient anyway, because it does not matter whether all people have understood exactly the same thing, but rather whether everybody has understood what is needed to do their job. And this new condition can be fulfilled because the specification of a program must express a property that has an objective meaning. It is true that nobody understands this meaning completely and in the same way as their neighbor, but everybody should understand that the question of correctness of the program with respect to its specification corresponds to a *fact*, and not to personal interpretation. The programmer must be able to construct the program by making a reasoning to persuade herself that it has the desired property; whereas the users must be able to derive other facts from it, such as the possibility of doing their job using the program.

For instance, consider a program computing the sine function under certain precise conventions. The programmer need not completely know the "essence" of this function, but only sufficient properties for constructing a correct program. The users need not understand the function in the same way as the programmer, but only other properties allowing them to solve their problems. So it is because of its objective nature that the specification of this program will be satisfactory: it expresses a fact, the same for everybody, even though they may understand it differently, and hence can derive other facts from it. A not completely unfounded objection to this example is that it is not realistic because the sine concept has been studied for such a long time that it would be foolish to deny its objective nature, but that not all specifications can be expressed in terms of such well-established concepts. Indeed, this objection pinpoints one of the fundamental difficulties of programming compared to, say, mathematics: one never has the time to polish all the needed concepts for a specification, because the program is needed urgently.

Nevertheless, the objectivity condition for specifications seems absolutely necessary for the correct communication of the purpose of programs, and, hence, for mastering the programming activity. According to us, without this condition, one would have to admit that the usage of programs for achieving a certain activity amounts to redefining that activity as being the exploitation of the results of the program without giving a satisfactory link between this redefinition and the initial concrete problem. Moreover, to us, this condition seems largely achievable, if one admits that the objectivity of the concepts necessary to the writing of good specifications can be founded on the creation of a "theory" of these concepts, with more or less detail according to the imperatives of the problem, a theory that can be studied by all concerned people until each of them has convinced themselves *personally* that it really corresponds to the intended object.

This perception of course has the "disadvantage" of founding the mastery of programming and its usage on the competence and responsibility of people, whereas some would prefer to found them on rules that are easy to apply and verify. The issue thus is whether one had better try and do one's best or believe in miracles.

2.5 Role and Content of the "Theory" of a Problem

Good specifications are useful and understandable because of some theory of the problem at hand. The objective of this "theory" is to establish the conceptual framework necessary for allowing everybody involved in the construction and usage of the program to know enough about it to do their job (or at least to know what it amounts to). We will not, and cannot, give precise methodological rules, but will rather draw some attention to certain points whose misunderstanding might complicate things uselessly.

There are two categories of concepts and objects of such a theory: those whose identity was determined before and independently of the problem, and those that are defined (or, better, identified) especially for the problem at hand. They should all have the same final status, namely to be known not by their definitions but by a sufficiently rich set of properties linking them to numerous other concepts of the problem. They thus have their own individuality, equivalent to an objective status. The theoretical development necessary for achieving this status is different and more or less long and difficult according to the category of concept.

2.5.1 On the Study of "Long-Established" Concepts

Defining once again preexisting concepts is common practice in formal methods of program design. It is however unwise to start the study of a predetermined concept by defining it. Indeed, what is necessary is to study the concept as it is, but not another concept given the same name through a definition. Even if a "predetermined" concept can be considered completely determined by a certain property (i.e., all other properties useful to the problem at hand can be derived from that property), one cannot consider it a definition of the concept. On the contrary, one would have to ensure that the concept really has that property. The objective of the theory to be built is to ensure that things are sufficiently well-understood by all involved people. If one started redefining all the fundamental concepts of the problem, nothing would be known about the relationship between the problem and what has been done. In any case, all involved people have a preliminary understanding of the problem. The role of the theory is to make things precise, if not to correct them, but not to reconstruct everything from nothing. It is thus more important to stress the difficult or delicate issues than to try and found everything already known.

The case of mathematical concepts. Suppose the concept of "greatest common divisor" is needed in the resolution of a programming problem. It is not the following redefinition of this concept that makes its role in the problem more precise:

> **Definition 2.1** The *greatest common divisor* of two natural numbers m and n is a natural number p, denoted $gcd(m, n)$, such that p divides m and n, and, for every natural number i, if i divides m and n, then i divides p.

Indeed, if one does not already know the concept of greatest common divisor (gcd) and its applications, this definition will not, by itself, help one understand its purpose. But let us consider a person who already has a good idea about it. The only information she can draw from this definition is that it probably is the definition of the notion of gcd that she already knows. Therefore, *the only immediately useful part of this definition is the only word that is theoretically arbitrary*! Indeed, one could define the same concept by naming it "foo" or "Nabuchodonosor." Two things are possible from here. Either this person is satisfied with her conclusion, and then the definition has not brought any new information, or she wants to verify this first impression by examining whether the definition is compatible with her existing knowledge of the concept of gcd. In this case, she might not be able to do so immediately, because her definition rather says that $gcd(m, n)$ is the greatest of the divisors of m and n, according to the usual ordering relation. To show that the two concepts coincide, she actually has to make a long reasoning, which should by the way conclude negatively, because they do not coincide when $m = n = 0$ (where the greatest common divisor is usually considered undefined, but the definition above gives $gcd(0, 0) = 0$). Anyway, at the end of this superb intellectual effort, she will still not know whether this definition was introduced for the fun of scrambling the message or for some better reason. To conclude, it would have been better to admit that the concept of gcd is predetermined beyond all definitions and to show why the very close concept of greatest common divisor according to the "divides" ordering relation was substituted for it. For instance, it could have been because one wanted to be able to apply, in all cases, the formula $gcd(m, gcd(n, p)) = gcd(gcd(m, n), p)$. (For $m = n = 0$ and $p \neq 0$, only the left-hand side of this equality is defined according to the usual definition.)

The case of "non-mathematical" concepts. The preceding precept applies unchanged to any kind of problem. It is not because the program to be written has its purpose in, say, an accounting setting, that one has to start by defining all involved concepts in order to understand its purpose.

For instance, in the payroll program, the "theory" of the problem should not start with definitions of employees, salaries, companies, etc. What is necessary is to arrive at a sufficient understanding of these concepts (which are perfectly determined, even if they might be poorly understood at the beginning) in order to solve the problem. It would not be acceptable either to define the effect of the program by the rules of computing the salaries in terms of the employee database. One should study the legislation, the structure of the company, etc, in sufficient detail so as to be able to *deduce* (i.e., to justify, by a rigorous reasoning) an adequate structure for said database as well as valid computation rules. The user of the program (i.e., the accountant) need not have

studied all the details of the "theory" that the programmers have had to elaborate, but she should understood it sufficiently for correctly using the program. It would be hard to say where the limit is: it is her responsibility to decide herself how far to go in order to reach a sufficient understanding.

2.5.2 On the Study of "Problem-Specific" Concepts

The writing and understanding of "good" specifications of programs nearly always requires the definition and study of problem-specific concepts, discovered or created especially for constructing the program. Such concepts can only be introduced by definitions, and it is crucial to understand their role correctly.

The case of simple concepts that are close to known ones. One often has to deal with concepts that can be considered already implicitly known and understood by all people who have to use them, but whose relevancy is insufficient for having been given a name that is universally admitted. It is then necessary to have recourse to a definition for naming the concept and making everybody agree on some important details whose identification is necessary for correctly using it. When reading such definitions, it should be possible to "immediately see what they are about." The concept-specific theory then reduces to only a few things, because the concept "naturally takes its place" among already known ones.

Let us illustrate this with a specification of the classical plateau problem.

> **Definition 2.2** Let $S = (s_1, s_2, \ldots, s_n)$ be a finite non-empty sequence of integers. A *plateau* of S is an interval[2] $\langle i : j \rangle$ such that:
>
> 1. $1 \leq i \leq j \leq n$
> 2. $s_i = s_{i+1} = \ldots = s_j$
> 3. $\langle i : j \rangle$ is not strictly included in any other interval with properties (1) and (2).
>
> **Problem:** Given a non-empty initialized array $a[1..n]$ of integers, construct a program that assigns to integer variable *np* the number of plateaus of the sequence $(a[1], a[2], \ldots, a[n])$, and to integer variable *maxlp* the maximum of their lengths (the length of a plateau is the number of its elements).

The definition in the specification above is sufficient for a satisfactory problem statement, for two reasons. First, the "technical" concept of plateau is not brand-new, but rather a particular and precise occurrence of a more general concept that we already know (the choice of the name "plateau" is thus *not* arbitrary). Second, this definition is sufficiently simple for linking this particular concept to the general one, that is for verifying whether the chosen terminology really corresponds to something intuitive. Moreover, the definition is necessary, because the intuitive notion of plateau is too vague for being able to rule out, in its absence, a misunderstanding of the notions of number and length of the plateaus of a sequence.

On the usefulness of examples. Specifications may be accompanied by carefully chosen examples, so as to facilitate their understanding. However, many people despise the usage of examples and consider that, in the best case, they are only noise, and, in the worst case, there is a risk of introducing contradictions. This attitude corresponds to a confusion of the end and the means. The role of a definition, as considered here, is not to be formally irreproachable (i.e., non-contradictory, for instance), but to help understand something. From this perspective, there is no reason to reject other means of communication that might have other qualities. Some well-chosen examples often provide an intuitive understanding that no definition could achieve. The latter then only makes more precise the exact contours of the concept. Other examples could help eliminate certain risks of ambiguity in the definition by illustrating delicate issues that are likely to be misunderstood for whatever reason.

As far as the risk of contradiction between definition and examples is concerned, note that this kind of contradiction would only be a real disaster if it were the non-contradiction of a

[2] We assume already known the concept of interval: $\langle i : j \rangle = \{x | x \text{ is an integer and } i \leq x \leq j\}$, where i, j are integers.

definition that would lend value to a concept. This is the usual confusion between truth and non-contradiction. What is important is to make known what one wants to say, not to escape contradiction. One could even argue that the discovery of a contradiction between an example and a definition is the best thing that can happen in some cases, because it carries an undeniable message: something is wrong somewhere!

Personal experience shows that a "poorly defined" concept can be perfectly understood thanks to examples, especially when the concept can be considered already implicitly known. Definition and examples are thus complementary means of designating the concept. And one may well conclude that there is only one concept corresponding to both the definition and the examples, even if one has spotted an apparent contradiction between them. What one already knows helps understand the error. Finally, note that the error risk is much higher in a definition than in an example, because it has to cover all cases. Examples are more reliable, because more "local," and are thus an ideal means of getting things straight.

Let us illustrate this on the plateau problem. Assume condition (3) was omitted from the definition above, but that the following example was added:

> **Example 2.1** If $S = (1, 1, 3, 3, 3, 2, 3, 5, 5)$, then there are 5 plateaus of S, namely $\langle 1 : 2\rangle$, $\langle 3 : 5\rangle$, $\langle 6 : 6\rangle$, $\langle 7 : 7\rangle$, and $\langle 8 : 9\rangle$. Also, its longest plateau is $\langle 3 : 5\rangle$, its length being 3.

From the definition and the example, one easily understands that plateaus are the *longest* non-empty intervals $\langle i : j\rangle$ included in $\langle 1 : n\rangle$ such that (2) holds. One could even have understood this without noticing that the definition is incomplete.

On the usefulness of remarks, or, better, of a "reasoned" presentation of definitions. All this shows that it is difficult to correctly define a concept in order to explain it to somebody else. In a sense, writing the definition is already a programming act. A definition is thus always the product of a more or less explicit reasoning process. So if one wants to facilitate the correct understanding of a definition, one could point out delicate issues in remarks, or, better, make explicit this reasoning process.

For the plateau problem, one might want to point out that the notion of plateau only makes sense with respect to a sequence S, that a plateau of S is always a non-empty interval, and that the set of plateaus of S partitions the interval $\langle 1 : n\rangle$, where n is the number of elements of S.

To do even better, one might show how the given definition was reached from a "reasonable" intuition of the concept of plateau of a sequence. This could go as follows.

> Let $S = (s_1, s_2, \ldots, s_n)$ be a finite non-empty sequence of integers. Let us draw a coordinate system, and mark the points at coordinates (i, s_i), for $1 \leq i \leq n$. Let us now draw, from each of these points, a horizontal segment of unit length. Some of these segments can be merged, giving rise to disjoint segments of integer length, called *plateaus* of the sequence:

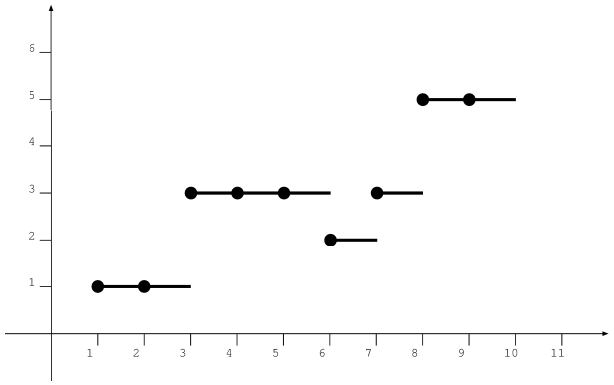

The plateaus of the sequence $(1, 1, 3, 3, 3, 2, 3, 5, 5)$.

> The objective is to write a program for computing the number of plateaus of a sequence and the maximum of their lengths. It is obvious that to each plateau corresponds an interval

verifying conditions (1) to (3) [included here as above], and vice-versa. Moreover, the length of a plateau is the number of elements of such an interval. This leads us to the following redefinition of the plateau concept, for our problem: [here follows the definition above].

This presentation clearly allows one not only to understand the definition more quickly, but also to "verify" it according to one's criteria. We even claim that such an intuitive presentation is self-sufficient and even preferable to the definition, as the latter is only a property of plateaus that one might discover by oneself.

A final remark: in practice, it is not always useful, nor possible (for "financial" reasons), to discuss all introduced "simple" concepts in this much detail. An acceptable compromise between the quality of the presentation and the time invested to its tuning must be found. Only experience shows where to situate this compromise. The most important thing is to understand that the objective is always the same: capture the posed problem as well as possible, as it is.

The case of more complicated or "new" concepts. Sometimes, the solving of a programming problem requires the invention of relatively original concepts that one cannot pretend having known before tackling the problem. They can thus not be imagined from nothing, but only constructed in small steps after comparing the problem to what is already known. The role of a definition is then to anchor some ideas for further investigation: one must be able to deduce many other properties from it, establishing thus the usefulness of the concept for solving the problem at hand. The concept thus offers economies of thought and reveals ways of solving the problem. The choice of a definition is guided by an intuition, i.e. the impression of having perceived an analogy with something already known. This definition is, in general, not the good one, because it may later turn out that not all the "desirable" properties can be derived from it, so that it does not play an efficient role in the problem solving process. The definition then has to be modified, in the light of these first conclusions, and so on, until the "good" concept has been obtained, namely the one that holds the key to the solution, or part of it. At the end of this process, whose essential steps must be reconstituted by all "clients" of the concept, the latter is known beyond the finally adopted definition. It is known by numerous properties linking it to other concepts. It has become "intuitive and well-known" and its definition is only one of its properties, among many others.

Remark. The distinction made here between "simple, implicitly known concepts" and "complex, new concepts" is of course too crude. They are only the extremes between which intermediates can be found, corresponding to a gradation of the effort to be done in order to construct a theory.

2.5.3 Only Representations of Concepts Can Be Defined, not the Concepts Themselves

To justify even more that the definitions introduced during the analysis of a problem are not the ultimate reference point for judging the value of a "solution," but only (imperfect) means of communication or "transient" starting points that can be (or actually should be) forgotten once the concept is well-understood, it is interesting to remark that a definition never really defines a concept, but only a certain representation thereof. This remark ruins, by itself, the "absolute" character of definitions by showing why they can be "wrong": whereas a concept cannot be something else than itself, its representations can be incorrect, i.e., fail to respect the (implicit or explicit) rules according to which they are supposed to represent the concept.

In our opinion, a concept worthy of this name must have a real and original identity that makes it indivisible, distinct from every more or less complex combination of "simpler" concepts. A concept is an atom of thought. Therefore, an interesting concept will always escape any particular definition, because one can define, from given concepts, only combinations thereof, i.e., nothing really new.

All this is particularly clear for "old," universally known concepts: for instance, whatever effort is undertaken to define natural numbers must be arbitrary. A natural number is what it is and cannot be reduced to anything else. Any definition thereof rests on representation conventions that had better be fixed very explicitly if one wants to wind up with a satisfactory definition.

But all this is still true for "new," problem-specific concepts. For instance, the concept of plateau of a sequence introduced above by a definition corresponds, in fact, to an intuitive concept that is very precise, but impossible to communicate as it is. This is why that definition of a plateau only defines a representation of this concept, namely as an interval of integers. Other representation choices would have led to different definitions. For example, one could have decided to represent the intuitive notion of plateau by couples of integers instead, as follows:

Definition 2.3 Let $S = (s_1, s_2, \ldots, s_n)$ be a finite non-empty sequence of integers. A *plateau* of S is a couple (i, j) such that:

1. $1 \leq i \leq j \leq n$
2. $s_i = s_{i+1} = \ldots = s_j$
3. $i = 1$ or ($i > 1$ and $s_{i-1} \neq s_i$)
4. $j = n$ or ($j < n$ and $s_j \neq s_{j+1}$).

This definition seems (to us) less good than the previous one, because it handles differently the plateaus at the extremities of the sequence. This is due to the fact that one cannot talk about the inclusion of a couple in another one. The reasoning to be made for constructing and understanding this second definition is thus slightly more tedious and error-prone. The plateau concept is thus more easily assimilated from the first definition. In any case, in both approaches, one has only defined a representation of the intuitive concept of plateau, which is the only really important thing to understand. One should however not believe that an axiomatic definition (e.g., an abstract datatype definition) would be immune from this. Consider, for example, the following definition:

Definition 2.4 Let $S = (s_1, s_2, \ldots, s_n)$ be a finite non-empty sequence of integers. A *plateau structure* on S is defined by choosing a set P and two functions lb, $rb : P \rightarrow \mathbb{N}$ such that the following conditions hold:

1. $\forall p \in P : 1 \leq lb(p) \leq rb(p) \leq n$
2. $\forall p \in P, \forall i : lb(p) \leq i \leq rb(p) : s_i = s_{lb(p)}$
3. $\forall i, j : (1 \leq i \leq j \leq n$ and $s_i = s_{i+1} = \ldots = s_j), \exists! p \in P : lb(p) \leq i$ and $j \leq rb(p)$

But this definition certainly says no more than the previous two about the essentials of the plateau concept. Refusing to say what plateaus are "made of" (be it intervals, couples, or beer bottles) is not sufficient for guaranteeing that the reader immediately understands the concept.

A concept is abstract not because it was introduced in a certain way, but because it has acquired an importance and identity in our thoughts. Therefore, the important issue is not to try and discover the good way of defining things, but to choose the adequate concepts, namely those that help us because we understand them the way they are.

2.5.4 On the Usage of Executable Definitions

Of course, not all ways of defining a concept are equally good. The good choice is always problem-specific. Rather than giving rules for writing definitions, we will criticize a commonly given one. According to some, a good definition ought to be declarative, i.e., written in a non-executable language [17]. This rule is, in general, absurd. To illustrate our point, we choose the very text formatting problem that was selected to show the virtues of declarative (and formal) specifications and was already discussed so much in the literature (see [21] for an overview).

Most people involved with this problem sought to specify it well, because, according to them, the correctness of a program can only be judged against its specification. But, then, against what should the specification be judged? Against arbitrary subjective criteria, of course, which entails that the value of specifications will be the object of endless discussions. According to us, on the contrary, the correctness of a program corresponds to an objective fact, independently of the way the problem is posed. Indeed, posing a problem means first of all admitting that there *is* a problem, next, trying to understand it sufficiently, and finally, writing a text allowing somebody else to understand it.

Posing the problem requires first of all the definition of the input and output texts. This can only be done here after making some hypotheses on the "environment" of the user. If we had to solve this problem for a real environment rather than for the sake of this article, then we could not make any such hypotheses but should learn about the environment of the user so as to replace these hypotheses by facts, which would be substantially more complicated than those used here. We thus suppose the user "sees" texts as sequences of lines (corresponding, in general, to lines on the screen or on paper), each line being a sequence of characters. This leads to the following definition of the input text:

> **Definition 2.5** A *word* is a finite, non-empty sequence of non-blank characters.[3] A *line* is a finite, possibly empty sequence of characters and blanks. Every line l can thus be uniquely decomposed as follows: $b_0\ w_1\ b_1\ w_2\ b_2\ \ldots\ w_n\ b_n$, where $n \geq 0$, the w_i are words, and the b_i are sequences of blanks that are non-empty except possibly for b_0 and b_n. The sequence of words $(w_1, w_2, \ldots, w_n)$ is the *sequence of words represented by* l, which we denote l *repr* $(w_1, w_2, \ldots, w_n)$. A *text* is a finite, possibly empty sequence of lines. Let $t = (l_1, l_2, \ldots, l_p)$ be the input text. The *sequence of words represented by* t is the sequence of words S such that $S = S_1 \bullet S_2 \bullet \ldots \bullet S_p$, where the S_i are the sequences of words represented by the l_i (and $\bullet$ denotes sequence concatenation). We denote this by t *repr* S. Two texts are *equivalent* if they represent the same sequence of words.

One can easily see that these definitions correctly capture the notion of input text suggested by the previous discussion. Now we must define the output text corresponding to a given input text. Therefore, we first have to capture this concept from an intuitive point of view. So, to what does it correspond? The answer is: to the result of applying an algorithm! The best way to understand this concept is to imagine a human typist having a listing of the input text and a terminal where every line has a length of *maxpos* characters. The job of the typist is to type the input text into the terminal by filling every line as much as possible, without trespassing the limit of the screen nor breaking words. This clearly amounts to the application of an algorithm whose execution uniquely determines the output text. Therefore, if one absolutely wants to "mathematically" define the output text in terms of the input text (i.e., if the previous explanations are deemed insufficient), the best one can do is to give a definition paraphrasing as closely as possible the typist's algorithm, because such a definition has the best chances of being correct and comprehensible. We thus propose the following definition:

> **Definition 2.6** Let $S = (w_1, w_2, \ldots, w_n)$ be a finite, possibly empty sequence of words that are each at most *maxpos* characters long. The *compact representation* of S, denoted *compact*(S), is the text defined as follows:
>
> 1. if $S = ()$, then *compact*$(S) = ()$;
> 2. else (i.e., if $n \geq 1$), let i be the largest integer such that $1 \leq i \leq n$ and the line $w_1 \sqcup w_2 \ldots \sqcup w_i$ has no more than *maxpos* characters, and let $l = w_1 \sqcup w_2 \ldots \sqcup w_i$ and $S' = (w_{i+1}, \ldots, w_n)$, so that *compact*$(S) = (l) \bullet$ *compact*(S').
>
> Given an input text t and the sequence of words S represented by t, the *output text corresponding to* t is defined if and only if no word in S is longer than *maxpos* characters. It is then equal to *compact*(S).

The definition of *compact*(S) is clearly executable. The only precise information it gives about *compact*(S) is how to compute it. But it is precisely this information that is the most precious one, for this example, for understanding the problem, because the output text essentially is, as argued above, the result of applying an algorithm. Note that the definition of *compact*(S) contains an over-specification according to [21], because we constrain the lines to be filled as much as possible in top-down order, rather than in non-determinate order. We do not see the utility of preferring a non-deterministic specification in this case. Of course, one should not conclude that executable definitions are always the best or always "as good as" others. Since the objective of a definition

[3] We consider an alphabet with a single blank character, denoted $\sqcup$, and no layout characters, such as for tabulation and end-of-line.

is to communicate a concept that is already implicitly known (which is clearly the case in the problem above), the best definition is the one that best captures the intuitive understanding of the concept while making precise its details. It just happens that the best way to explain a concept is, *sometimes*, to indicate how to compute it or how to generate its instances.

2.6 The "General Form" of Specifications

We now try to capture the "general form" of specifications, without however giving systematic rules for writing "good" specifications, as such cannot be an objective. All one can hope for is a sufficient understanding of the concept of specification for being able to use it satisfactorily in many cases.

The specification of a program should always have two parts that play very distinct roles:

1. a statement indicating the purpose of the program, i.e., the information that can be drawn from the results of its execution;

2. a list of representation conventions that are to be satisfied for using the program correctly and for interpreting its results correctly.

Statement (1) must always be very simple because the information produced by a program (after interpretation of its results) must have a simple meaning to the user. Without it, she would be unable to use the program to her advantage. The role of the "theory" of the problem is to make sure that this meaning exists and that it can be clearly and simply formulated. The list (2) must also be sufficiently simple to understand for the purpose of the program not to be completely annulled by its difficulty of usage and of interpretation of its results.

This is not always easy to achieve due to the formal character of programming languages. It is thus sometimes necessary to construct another theory before being able to simply state the representation conventions.

We now state what the specifications of the three problems in Section 2.1 should contain.

Example: The Belgian National Lottery. The specification reduces to the indication of how to start the program and to the statement that it results in displaying the next draw of the Belgian national lottery. (It is practically useless to state the exact format of the produced character string and the rules for decoding this information, because everybody immediately understands how to interpret the message when it appears.)

Example: A payroll program. The accountant user of the payroll program must know the necessary information as well as the rules of its representation by the input data. She must be able to verify the correctness of these data. She also must know enough about the rules of representation of the results in order to be able to finish the payroll task (this is actually the responsibility of a bank, nowadays). The specification thus reduces to the indication of how to start the program and to the statement that, from correct input data, the program produces correct results according to the used representation rules.

Example: A search sub-program. Depending on the desired generality, the used programming language, and the general context of the problem at hand, there is a tremendous variety of possible specifications for a program performing a search in an array.[4] A satisfactory specification, in some cases, could be the following:

> **Specification 2.1** The procedure *search* is a Pascal procedure declared as follows:
>
> $$function\ search(x : integer) : boolean$$
>
> Its declaration must figure within the scopes of the declarations of an integer constant n (such that $n \geq 1$) and an array a of type $array[1..n]\ of\ integer$, which also is in the scope of

[4] This shows why it is unreasonable to try and give precise general rules of writing specifications, even if one restricts oneself to a class of problems as particular as this one.

the former. When calling the procedure, the elements of array a must be in non-decreasing order. Let v be the actual value of the formal parameter x. If at least one of the elements of a is equal to v, then the call returns the value $true$, otherwise it returns $false$. (The contents of a will be unchanged.)

The bulk of this specification is dedicated to the statement of the representation conventions and to technical details. These details are tedious but unavoidable because the used programming language is a formalism. They do not, however, render the specification unusable because the problem of knowing where to put the various declarations and how to write them can be solved separately as well as once and for all. When reasoning about it in the future, it suffices to remember how to call the procedure, that it answers the question "does v belong to a?," and that the answer is given as a boolean value.

However, it is important to note that the introduction of general representation conventions that are specific to a particular problem (i.e., they are chosen for an application and used for the specifications of all the sub-programs of this application) can contribute to making much more manageable the amount of representation details specific to each specification.

2.7 Requirements Specifications and the Theory of the Problem (Are the Same Thing)

The process of elaborating requirements specifications is nowadays considered by many computer scientists as the most crucial stage of software development. Requirements engineering is thus emerging as a new and major branch of the software engineering discipline. It is primarily concerned with the identification of the user's needs, i.e., the so-called requirements elicitation process. As soon as the user's requirements are explicitly stated, they can (and must) be checked with respect to consistency and completeness. In fact, this is what we call "elaboration of the theory of the problem." Thus, requirements specifications are *not* specifications (in our sense), but rather an exposition of the very theory making it possible to specify the software system.

Formal specification languages are advocated by many researchers as the distinguished methodological tool for requirements engineers because they allow them to make the user's informal statements precise, to check the requirements specification for consistency and completeness, and to ease the discussion with the user by means of prototyping, to name only a few supposed advantages. In our opinion, the mechanical treatment of (formal translations of) the user's requirements can indeed possibly provide information that could not be easily inferred by hand. However, the formal translation process is completely similar to the writing of a program in that it necessitates giving precise specifications (in our sense) to most symbols and constructs of the formal text, in order to ensure that the formalization captures exactly what the user meant.[5] Thus, the writing of (so-called) formal requirements specifications presupposes the existence of an already fully understood theory of the problem, in our sense. Finally, as seen in Section 2.5, even the elaboration of the theory of the problem may benefit from the use of specifications in our sense, in order to make explicit the rationales underlying the concepts introduced by means of definitions.

3 On the Nonsense of Formal Specifications

3.1 Why Can't There Be Any Formal Specifications?

A "formal specification" is a statement in a formal specification language. Such a statement is unintelligible "by itself," primarily because the concepts of the problem are almost never primitive concepts of the used formal language. Therefore, a formula can only be "understood" as a representation of an intuitive statement, according to explicitly given conventions. These conventions are in general that the formula is true, in the chosen interpretation of the language, if and only if the intuitive statement is true. The enunciation of such conventions is precisely what we call a specification, in the sense that we discussed in Section 2, although not the specification of a program but rather of a formula. Its role is to give a meaning and thus a purpose to something

[5] Automated processing of formal texts could also be used only for psychological reasons, e.g., the computer is God and always tells the truth.

(the formula in this case) that would otherwise not have one. Whether an inextricable formula is true or false is of no interest whatsoever if this is the only thing we know about it. In general thus, a specification is necessary each time that, for good reasons or bad ones, one wants to represent a *known* property or concept by a text written in an artificial language (be it formal, mathematical, ...). This also shows that any "formal specification" of a (formal) program is always much closer to the program itself than to a specification in our sense. A noticeable difference may be that it is not "executable" because it is written in a "non-executable" language. In our opinion, it is not important whether the chosen language is executable or not, but whether it allows us to say in the most direct way what the purpose of the program is. Such a condition cannot be fulfilled by any formal language, given the extremely low expressiveness of such languages. A formal language is always almost as bad as a programming language for communicating the purpose of a program. In other words: providing a formal specification of a program amounts almost to considering that the text of the program (or of another program) allows us to understand its purpose.

3.2 Seven Frequently Asked Questions about Formal Specifications

Are informal specifications and formal ones complementary? Some people readily admit that it is necessary to add an "informal comment" to a program that helps understand the purpose of the program and that corresponds to our notion of specification. But, for these people, such a comment is insufficient to ensure that the effect of the program has been precisely defined. This corresponds to the frequent opposition of *intuition* and *rigor*, which considers that a fruitful intellectual activity should be driven by intuition (which is comprehensible because incorrect) so as to produce rigorous results (which are formal but incomprehensible). In our opinion, the correct usage of a program necessitates having understood intuitively *and* rigorously its purpose. There is no need to distinguish two notions of specification, one comprehensible and vague, the other precise and unintelligible. If a specification features delicate issues that are likely to be misunderstood, it is only necessary to give more details about them. There is no reason to believe such difficulties are best resolved, in all cases, by using a formal language chosen once and for all.

If one thinks it is not safe to directly and simply explain the purpose of a program, i.e., in the way one understands it oneself, and that one had better define with absolute precision the "effect" of the program, even under the risk of incomprehensibility, by giving the readers "indications" on how to reconstruct a comprehensible specification for themselves, then one is confronted with the following difficulties. *It is almost as difficult to write without errors a formal specification as the program itself, and it is barely easier to "decipher the message," in the opposite direction.* To write a correct formal specification, one has to make an explicit detailed reasoning that is very different from a vague informal comment. In order to convince oneself of having understood the formal specification, another reasoning has to be done, which is extremely tedious if the formal specification is not accompanied by such comments. So, for a couple ⟨ formal specification, informal specification ⟩ to suitably play its intended role, it would have to be accompanied by a detailed reasoning fixing their representation relationships. However, this is only meaningful if the informal specification has been explicitly and precisely stated. The role of the formal specification and the reasoning is then reduced to lifting the last doubts and ambiguities. But this can be achieved at lower cost by other means, such as the inclusion of significant examples, the provision of the reasoning process leading to the definitions in the specification, etc.

Are formal specifications a means of dividing the difficulty of programming? Other people would rather say that the recourse to formal specifications is, if not a panacea, at least a means of division of the difficulty. Indeed, it would allow, on the one hand, the formal and mechanical proof of correctness of programs, and, on the other hand, the intuitive justification that the formal specifications correctly represent the problem to be solved.[6] One could thus give much more confidence to programs, since everything reduces to the problem of validity of the formal specifications, formal correctness being established beyond all doubt.

This viewpoint rests on two forms of exaggerated optimism on formal methods. First, it is in general not significantly easier or safer to prove intuitively the correctness of formal specifications

[6] This second part of the division is completely ignored by some people who believe that only formal justifications are valuable, and who assume thus that formal specifications are satisfactory "by miracle."

than that of programs. Second, formal proofs of program correctness are almost often infeasible in practice, whatever the available mechanical aid (proof verifier or theorem prover). For example, note that a formal proof of program correctness amounts to proving a formula whose length is at least the sum of the lengths of the formal specification and the program. So what will be the length of the proof?! This also assumes a *complete* formalization of the semantics of the programming language, which is already by itself an almost unrealizable task. If one considers that the time and budget allocated to the verification of program correctness is necessarily limited, it can be easily seen that one had better spend a bit more time justifying intuitively the correctness of the program and carefully choosing test cases than make use of such formal methods.

However, many people (especially those who worry about improving the productivity of industrial scale programming) think that automatic verification of program correctness will soon be feasible and common. In their minds, there is no essential difference between the automatic proof of the correctness of a program and, say, its syntactic verification: both amount to entering a text into the computer and waiting for it to reply 'yes' or 'no'. If we object to this that the two problems are of completely different orders of difficulty, some will reply that it was also once "proven" that things heavier than air could not fly or that "he succeeded because he did not know it was impossible." But suppose such program verifiers were really used one day (no matter the exact value of the verifications they perform), there is a high risk that they will be used as sole criterion of the quality of a program. A "good" program will be the one that gets the blessings of the verifier, and the primary objective of programmers will be to write such programs. By this token, once again, personal judgment will have been replaced by arbitrary answers of a program!

Is it necessary to formalize specifications to prove their consistency and completeness? Some people say that formal specifications allow systematic verification of their consistency and completeness. This deserves several remarks.

First, if it is desirable that a statement be consistent and complete, the precise meaning of these notions always strongly depends on the context of the statement, that is on a lot of things that are known about the subject of the statement before even examining it. If a statement defines a problem that has no solution, it is sometimes judged inconsistent, but, at other times, it is considered a perfectly consistent statement of a problem that just happens to have no solution. Similarly for completeness, when the problem has many solutions. Since a formal statement only is, in general, a representation of a non-formal statement, which is the only one to be comprehensible, the consistency and completeness of a formal statement can only receive a precise meaning through this representation relation. As this relation is always chosen *ad hoc*, it is impossible to satisfactorily define (i.e., in a manner always corresponding to the intuitive concepts) consistency and completeness of formal specifications. Since this relation is thus totally exterior to the used formalism, consistency and completeness cannot be verified mechanically.

Second, why are so many people interested in the notion of consistent and complete formal specification? The only explanation to this phenomenon that we could find is that they have completely taken up Hilbert's belief in the foundation of mathematics through formalization (in the hope, of course, that they could do better than him). This belief can be summarized by the following assertions:

1. There must be (and somebody will define it sooner or later) a formal language in which every problem (or, to be more modest, every problem of a certain "class") has a natural expression. In other words, the "intrinsic structure" of every problem can be exhibited in this language, stripped of all its veils.

2. Better, this language will be so extraordinary that whoever uses it to state a problem will be incapable, ipso facto, of writing anything else than the "good" definition of the problem. Such a language will be said to be "thought-structuring."

3. Finally, if nevertheless an error slipped into the specification, it could only result from a distraction and would inevitably provoke an inconsistency or incompleteness in the specification, which will be easily detected, if not corrected, by a good verification program.

This belief is maybe not proclaimed in public, nor even completely conscious, but nevertheless underlies all the efforts of these people.

Are formal specifications more concise than informal ones? A common argument is that formal specifications are more concise than informal ones. However, some people argue to the contrary. Strictly speaking, the raised question is meaningless for specifications in our sense, since they are a link between formality and informality. So the question in fact only applies to requirements specifications, or, in other words, to the theory of the problem. During the elaboration of this theory, the usage and introduction of mathematical notations is certainly permitted and thus necessarily makes this first part of the specification more concise than without such notations. But an explanation of the link between these mathematical concepts and the concepts of the problem is also needed, and this second part cannot be made concise by means of mathematical notations. That a formalization (in the strict sense) of the first part of a requirements specification makes it more concise is obvious (as this is the very reason why notations have been introduced into mathematics in the first place), but one cannot possibly pretend that this formal description is a specification (in our sense), because the equally essential second part of a requirements specification would be missing.

Are formal specifications more pragmatic than informal ones? Some advocates of formal methods readily agree on the inevitability of informal specifications and informal verification, but also point out that formal and informal specifications have different purposes and qualities. Indeed, formal specifications, whether executable or not, would offer a means of early feedback from the customer —through execution of the specification (early prototyping) or through demonstration of desired properties—, and hence could allow significant cost saving. Otherwise, discrepancies between the specification and the customer's intentions might only be detected when the customer runs (an increment of) the final software. However, and again, in our opinion the question of pragmatism of formal specifications does not even arise. Indeed, one may certainly construct intermediate formal descriptions before constructing the final software, as they can help during the process of elaborating the theory of the problem. But one cannot call such a description a "formal specification" (and writing it is more of a programming activity than a specification activity), as it is not a specification at all (in our sense) and as it is incomprehensible by itself and must thus be explained to the customer (which explanation process provides the very part that is missing in the formalization), be it as a document or as an executable or demonstrable prototype.

Can formal specifications be automatically generated from informal ones? Some researchers advocate writing informal specifications in some supposedly "semi-formal" notation (such as SA/SD) or in some form of "controlled natural language" (in the sense that the vocabulary and grammar are restricted so as to give sentences a "clear" semantics), expecting that they can be automatically translated into (executable) formal specifications. The problem with these approaches is that these languages are formal ones, no matter what they are called. There is no such thing as "semi-formal languages" or "informal controlled natural languages." Since the descriptions are thus actually formal, it is only obvious that they can be automatically translated into some other formal languages. And, as formal statements, they cannot possibly be specifications, in our sense. For such specifications (in our sense), there is of course no way that they can be automatically formalized, as the link between the formal concepts and the real-life ones is not formalizable and as one would have to prove that the translation process is equivalent to the mechanisms of human knowledge acquisition.

Are formal specifications necessary for safety-critical systems? It is often argued that formal methods are necessary for the design of safety-critical systems, and some standards organizations even start imposing/recommending their usage for such projects. The rationale is that systems satisfying "specifications" in the form of, say, finite-state machines (that are deemed trivially correct after inspection) can be shown, say, to be free of deadlock and lifelock risks. Our objection to this formalist viewpoint is essentially the same as to the pragmatism issue above, because, once again, it is a delusion to believe that there can be "obviously correct formal specifications." But note that we do support the idea that extra care and rigor are needed in the design of safety-critical systems; however, we do not agree that designing intermediate formal descriptions in order to verify them is necessarily the best way to improve our confidence in the system. Explicit

reasoning based on specifications is often a better way; it is needed anyway to correctly build the intermediate formal descriptions when they really help.

4 Conclusion: Why are the Role and Nature of Specifications so Often Misunderstood?

We now explain why our notion of specifications is difficult to understand and admit by many practitioners and theoreticians of computer science. But let us first summarize our viewpoint:

1. A program is useful because its results can help to solve a problem. There is no limit to the class of problems that we can imagine in the "real" world. Therefore, the understanding of the purpose of a program may necessitate the knowledge of notions as distant as desired from programming concepts (or from concepts used in formal specification languages).

2. The specification of a program essentially is the statement of its purpose.

3. A specification should not, nor can it provide all the knowledge necessary to the understanding of the purpose of the program. It must just try to state it in the most satisfactory possible way, that is in the most simple and direct way. That is why a specification is not meant for just anybody, but only for those who can understand it.

4. For the specification to be comprehensible by sufficiently many people, it is, in general, necessary to "construct" a theory that can be studied and understood by all. Such a theory cannot be constructed from nothing, but assumes a considerable preliminary knowledge that is partly shared by all the considered people.

Now, there are at least two reasons for the reluctance of so many people to admit the pertinence of the assertions above.

First, there is the influence of the currently dominating ideas on the nature of mathematics. Mathematical theories are supposed to be founded on formal axiomatized theories. This means that every intuitive statement of the theory is supposed to be only an "abbreviation" of a formal statement that is itself mechanically deducible from the axioms. From there to infer that every interesting result of a theory can be discovered relatively quickly as soon as the axioms of its theory are known is only a small step. And this is the "step" made, consciously or not, when asserting that the specification of a program should, above all, define with absolute precision the effect of the executions of this program. Indeed, it is clear that from the input/output relation determined by the executions of a program, one can theoretically deduce all other interesting properties of this program. Therefrom, some conclude that a specification reduces to such a definition, assuming that every reader is sufficiently intelligent to derive from it all other "interesting" properties of the program. (This means the reader is assumed to be omniscient, because if a program outputs the string "380,000", she would, for instance, have to derive from this observation that one of the properties of the program is to give the distance between the Earth and the Sun, expressed in kilometers in the decimal system.)

Therefore, the idea that the specification of a program must be and can only be the definition of an input/output relation is a simple transposition of the idea that there is nothing more in a mathematical theory than in its axioms. But, in order to understand the exact role of specifications, one should realize that, to the contrary, there is infinitely more than that in an intuitive theory: every new concept, notation, or result adds value to it that is not at all contained in the statement of its axioms. The intuitive statement of an important theorem certainly is not a mechanical consequence of the axioms of a formal system, no more than the assertion of the "truth" equivalence between this statement and a formula. And this even holds for statements of the form "that formula is a theorem," because the meaning of the notions of formula and theorem is not derivable from the mechanical rules of the formal system.

In conclusion, *a correct understanding of the notion of specification necessitates, in our opinion, a return to a more intuitive and "transcendent" perception of mathematics.*

Second, there is the opinion according to which the mastery of the programming problem can only be achieved by recourse to effective and automatable methods. It seems (sadly) evident that

few people are ready to admit that the mastery of programming will always depend, above all, on the competence of the involved people. The manager wants effective criteria evaluating the quality of the work done by the programmers. The programmer expects the "theoreticians" to provide rules that can be followed blindly. Nobody wants to admit that the best way to realize whatever task is to do one's best by trying to stick to utmost intellectual honesty.

If, regarding specifications, we say that the best thing to do is to understand the exact role of this notion so as to be able to "see," in most cases, how to state them best, it will be considered that we have not brought anything interesting, because we have not given any rule or criterion for writing good specifications or for evaluating them. However, some people say that, as it is better to do something rather than nothing at all, it is better, all things considered, to give rules that are arbitrary but measurable.

For us, it is certain that little progress can be expected in programming as long as the opinion is so widespread that the value of a criterion is determined by its being measurable and computer readable. We think so because this idea can only prolong the illusions and avoid the real problems: thanks to such criteria, the manager can take decisions without having to get involved in the project, and this changes nothing to the quality of the programmers' work, except that they have to adjust themselves so as to respect these rules even if their application leads to absurdities.

Acknowledgments

The authors are indebted to Prof. Henri Leroy for his spiritual patronage. The central ideas of this paper have been deeply influenced by his teaching and the numerous nightly discussions with the first author.

References

[1] R. Balzer. A 15 year perspective on automatic programming. *IEEE Trans. on Software Engineering*, 11(11):1257–1268, Nov. 1985.

[2] R. Balzer, N. Goldman, and D. Wile. Informality in program specifications. *IEEE Trans. on Software Engineering* 4(2):94–102, March 1978. Also in C. Rich and R.C. Waters (eds), *Readings in Artificial Intelligence and Software Engineering*, pp. 223–232. Morgan Kaufmann, 1986.

[3] J.P. Bowen and M.G. Hinchey. Ten commandments of formal methods. *IEEE Computer* 28(4):56–63, April 1995.

[4] J.B. Bowen and M.G. Hinchey. Seven more myths of formal methods. *IEEE Software* 12(3):34–41, July 1995.

[5] D. Craigen, S.L. Gerhart, and T. Ralston. Formal methods reality check: Industrial usage. *IEEE Trans. on Software Engineering* 21(2):90–98, Feb. 1995.

[6] R.A. De Millo, R.J. Lipton, and A.J. Perlis. Social processes and proofs of theorems and programs. *Comm. of the ACM* 22(5):271–280, May 1979. Comments in *Comm. of the ACM* 22(11):621–630, Nov. 1979.

[7] P.J. Denning (ed). A debate on teaching computing science. *Comm. of the ACM* 32(12):1397–1414, Dec. 1989.

[8] J.H. Fetzer. Program verification: The very idea. *Comm. of the ACM* 31(9):1048–1063, Sept. 1988. Comments in *Comm. of the ACM* 32(3):374–381, March 1989.

[9] P. Flener and L. Popelínský. On the use of inductive reasoning in program synthesis: Prejudice and prospects. In L. Fribourg and F. Turini (eds), *Proc. of META'94 and LOPSTR'94*, pp. 69–87. LNCS 883, Springer-Verlag, 1994.

[10] M.D. Fraser, K. Kumar, and V.K. Vaishnavi. Informal and formal requirements specification languages: Bridging the gap. *IEEE Trans. on Software Engineering* 17(5):454–466, May 1991.

[11] M.D. Fraser, K. Kumar, and V.K. Vaishnavi. Strategies for incorporating formal specifications in software development. *Comm. of the ACM* 37(10):74–86, Oct. 1994.

[12] N.E. Fuchs. Specifications are (preferably) executable. *Software Engineering Journal* 7:323–334, Sept. 1992.

[13] S.L. Gerhart, D. Craigen, and T. Ralston. Experience with formal methods in critical systems. *IEEE Software* 11(1):21–28, Jan. 1994.

[14] W.W. Gibbs. Software's chronic crisis. *Scientific American* 271(3):86–95, Sept. 1994.

[15] J. Guttag, J. Horning, and J. Wing. Some notes on putting formal specifications to productive use. *Science of Computer Programming* 2(1):53–68, Oct. 1982.

[16] A. Hall. Seven myths of formal methods. *IEEE Software* 7(5):11–19, Sept. 1990.

[17] I.J. Hayes and C.B. Jones. Specifications are not (necessarily) executable. *Software Engineering Journal* 4(6):330–338, Nov. 1989.

[18] C.A.R. Hoare. An overview of some formal methods for program design. *IEEE Computer* 20(9):85–91, Sept. 1987.

[19] P.G. Larsen, J. Fitzgerald, and T. Brookes, Applying formal specification in industry. *IEEE Software* 13(7):48–56, May 1996.

[20] B. Le Charlier. *Réflexions sur le problème de la correction des programmes*. Ph.D. Thesis (in French), Facultés Universitaires Notre-Dame de la Paix, Namur (Belgium), 1985.

[21] B. Meyer. On formalism in specifications. *IEEE Software* 2(1):6–26, Jan. 1985.

[22] D.L. Parnas. Mathematical description and specification of software. In B. Pehrson and I. Simon (eds), *Proc. of IFIP'94*, pp. 354–359. Elsevier Science, 1994.

[23] H. Saiedian (ed). An invitation to formal methods. *IEEE Computer* 29(4):16–30, April 1996.

[24] J.M. Wing. A specifier's introduction to formal methods. *IEEE Computer* 7(5):8–24, Sept. 1990.