

Why Denotational?

Remarks on applied denotational semantics

By Andrzej Blikle
February 1990

This is a very early preprint of my paper published under the same title in *Fundamenta Informaticae* 28, 1996, p. 55-85.

Andrzej Blikle

Why denotational?

Remarks on applied denotational semantics

679

February 1990

WARSZAWA

INSTYTUT PODSTAW INFORMATYKI POLSKIEJ AKADEMII NAUK
INSTITUTE OF COMPUTER SCIENCE POLISH ACADEMY OF SCIENCES
00-901 WARSAW, P.O. Box 22, POLAND

Andrzej Blikle

WHY DENOTATIONAL ?
REMARKS ON APPLIED DENOTATIONAL SEMANTICS

679

Warsaw, February 1990

R a d a R e d a k c y j n a

A.Blikle (przewodniczący), L.Bolc, J.Borowiec,
S.Bylka, M.Dabrowski (zastępca przewodniczącego),
J.Lipski (sekretarz), L.Lukaszewicz, R.Marczynski,
A.Mazurkiewicz, T.Nowicki, J.Seidler, M.Stolarski,
Z.Szoda, J.Winkowski

Prace zgłosił Antoni Mazurkiewicz

Mailing address: Andrzej Blikle
Institute of Computer Science
Polish Academy of Sciences
PKiN P.O.Box 22
00-901 Warszawa
Poland

ISSN 0138-0648

Printed as a manuscript
Na prawach rękopisu

Nakład 400 egz. Ark. wyd. 3.45; Ark. druk. 3.50.
Papier offset. kl. III, 70 g, 70 x 100. Oddano do
druku w lutym 1990 r. W. D. N. Zam. nr 81/90

ABSTRACT * STRESZCZENIE

This is an essay where the author expresses his views on applied denotational semantics. In the author's opinion, whether a software system has or does not have a sufficiently abstract denotational semantics should be regarded as a pragmatic attribute of the system rather than merely as a mathematical attribute of its description. In a software system with denotational semantics structured programming is feasible and for such systems there is a routine method of developing program-correctness logic. All that may not be the case if denotationality is not ensured. On the other hand, a non-denotational semantics can be always artificially made denotational on the expense of lowering its level of abstraction. This leads to an important pragmatic question: to what extent and in which situations can we sacrifice denotationality and/or abstraction of a semantics? All discussions are carried on an algebraic ground but the paper is not very technical and contains a short introduction into the algebraic approach to denotational semantics.

Dlaczego Denotacyjna? Uwagi o Stosowanej Semantyce Denotacyjnej

Niniejsza praca stanowi esej, w którym autor przedstawia swoje poglądy na stosowaną semantykę denotacyjną. To czy system oprogramowania ma lub nie ma dostatecznie abstrakcyjnej semantyki denotacyjnej powinno być, zdaniem autora, traktowane jako właściwość samego systemu, a nie jedynie jako matematyczna własność jego opisu. W systemie oprogramowania mającym semantykę denotacyjną można programować strukturalnie, a także istnieje pewna standardowa metoda budowania logiki dla dowodzenia poprawności programów. Oba te fakty mogą okazać się nieprawdziwe jeżeli semantyka systemu nie jest denotacyjna. Z drugiej jednak strony, semantykę niedenotacyjną można zawsze sztucznie uczynić denotacyjną kosztem obniżenia jej poziomu abstrakcji. Prowadzi to do ważnego pragmatycznie problemu: do jakiego stopnia i w jakich sytuacjach możemy poświęcić denotacyjności semantyki na korzyść innych własności? Dyskusja przedstawionych problemów prowadzona jest na gruncie algebraicznym, praca nie jest jednak bardzo techniczna i zawiera krótkie wprowadzenie w algebraiczny model do semantyki denotacyjnej.

CONTENTS

1) Introduction	5
2) About the concept of semantics	6
3) Basic notation	11
4) Denotational semantics from an algebraic perspective	13
5) Why denotational?	19
6) Why not non-denotational?	25
7) Denotationality versus abstraction	29
8) Repairing semantics	34
9) On the borderline of denotationality	41
10) Copy-rule semantics	44
11) An operational definition of a denotational semantics	48
12) Final remarks	50
References	52

Everything which is evident should be given a special attention at the beginning in order to avoid all possible misunderstandings in the future.

[a popular wisdom]

1. INTRODUCTION

This essay has been addressed to readers interested in the applications of denotational semantics in software engineering and is devoted to the discussion of the attribute of denotationality. In the author's opinion, whether a software system has or does not have a sufficiently abstract denotational semantics should be regarded as a pragmatic attribute of the system rather than merely as a mathematical attribute of its description. Denotationality in system design is like structurality in programming: it makes the final product easy to understand and to prove correct and therefore constitutes a prerequisite of a sound engineering style.

Our main discussion is preceded by some clarification of concepts. In Sec.2 we point out that the word *semantics* may have at least three different meanings and we try to convince the reader that such attributes of a semantics as *denotational*, *operational* and *algebraic* should be regarded as orthogonal rather than as alternative or contrasted. In Sec.3 we define our notation and in Sec.4 we briefly introduce the reader into an algebraic framework of denotational semantics. That framework is then used throughout the paper.

The main part of our discussion starts in Sec.5 from an argument that a software system with denotational semantics provides an adequate ground for structured programming and for a systematic development of a program-correctness logic. Then we show that this need not to be the case if

A short conference version of the paper was published earlier as [Blikle 89a]. The reported research has been sponsored by the Polish Academy of Sciences under grant CPBP 02.17.

denotationality is not ensured (Sec.6). This is followed by the analysis of such properties of a semantics that can make the semantics non-denotational. We discuss the trade-off between denotationality and abstraction (Sec.7) and we show that a non-denotational semantics can be always artificially "made denotational" on the expense of lowering its level of abstraction (Sec.8). This leads us to an important pragmatic question: to what extent and in which situations can we sacrifice denotationality and/or abstraction of a semantics? A few typical cases are discussed where the denotationality of a semantics can be spoiled (Sec.9 and Sec.10). One of them is a copy-rule mechanism. We complete the discussion with an example of a structured operational definition (in the style of G.Plotkin) of a denotational semantics of a simple programming language (Sec.11). The last Sec.12 contains some concluding remarks.

2. ABOUT THE CONCEPT OF SEMANTICS

This section is devoted to the clarification of the concept of semantics and to a discussion on the relationship between three common attributes of a semantics: denotational, operational and algebraic. We shall argue that these attributes are orthogonal rather than alternative or contrasted.

Let us start from a remark that the issue of semantics is not restricted - as is frequently thought - to programming languages, but applies to all kinds of software including system software, tools and applications. In fact each software system contains some programming language which allows the user to communicate his requests to the system. Although in all examples which follow in the paper we analyze a toy programming language, this choice has been dictated only by the merit of simplicity. A toy programming language can be made much simpler than a toy operating system or a toy data-base management system.

In any software system we can always identify some syntax, which is used to formulate our requests to the system, some denotations, which are the meanings of these requests and some semantics, which assigns denotations to syntax. In other words, in the mathematical model of a software system we can always identify three following components:

- a set of syntactic objects Syn,
- a set of denotations Den,
- a function of semantics $S: \text{Syn} \rightarrow \text{Den}$

Unfortunately in the current literature the word *semantics* is used ambiguously to mean four different things:

- (1) the functions S ,
- (2) the definition of that function,
- (3) the underlying theory.
- (4) or even the denotations (e.g. "the semantics of commands are state-to-state transformations")

Case (4) is, of course, only a linguistic sloppiness. Case (3) can easily be recognized from a context and therefore does not cause problems. Cases (1) and (2), however, if not distinguished properly may lead to a confusion. For instance, if an author says that: "*we call a semantics denotational if it is compositionally defined and tackles recursion with the help of fixed points*", it is not clear if he is talking about a function (which is "compositionally defined..."), about its definition (which "tackles recursion...") or, maybe, about both at the same time?

In this paper by a *semantics* we always mean a function and when we want to talk about its definition, then we say that explicitly, unless the context indicates clearly enough what we mean. Below we briefly explain our understanding of the attributes *denotational*, *operational* and *algebraic*. We start from the attribute of denotationality.

Syn usually represents a context-free language described either by a context-free grammar, or by BNF equations, or by a set of syntactic domain equations or by a signature of an algebra. In each of these cases one may construct a unique many-sorted algebra over Syn. Now we say that S is *denotational* if it has the property of compositionality, i.e. if one can construct such a many-sorted algebra over Den that S becomes a homomorphism. As we shall see in the sequel (Sec.7) that understanding of denotationality is a little too weak for applications. In fact, a denotational semantics should be also sufficiently abstract in order to be of some use. That is, however, a rather pragmatic issue and cannot be easily formalized.

In the literature the term **denotational semantics** is usually associated with two specific techniques of constructing the algebra of denotations: **reflexive domains** and **continuations** (see e.g. [Stoy 77], [Gordon 79] or [Schmidt 86]). These techniques have been introduced by the pioneers of denotational semantic ([Scott 71] and [Scott,Strachey 71]) in order to construct a denotational model of an untyped lambda-calculus mixed with unrestricted goto's (Algol-60). Reflexive domains and continuations have so much attracted the attention of researchers that for many years it became customary to assume that every denotational semantics must involve these concepts. In fact, however, their authors have always emphasized that the most relevant attribute of a denotational semantics is **compositionality**:

In this approach the semantical functions give mathematical values to expressions - values related to some given model. The values of expressions are determined in such a way that the value of a whole expression depends functionally on the values of its parts - the exact connection being found through the clauses of the syntactical definition of the language.

[Scott,Strachey 71]

The idea of compositionality has been later formalized on an algebraic ground by a group of American authors known as ADJ, see e.g. [Thacher,Wagner,Wright 78]. On the other hand it has been pointed out in [Blikle,Tarlecki 83] that reflexive domains are needed only if we wish to describe self-applicable functions, i.e. functions that can assume themselves as arguments. Such functions appear e.g. in Algol'60 - due to typeless procedural parameters - or in Lisp - due to dynamic recursion. Since, however, self-applicability has turned out to be an unsafe programming mechanism, it has been abandoned in modern software systems including the majority of contemporary programming languages. This has led to a conclusion that the denotational models of software may be conveniently constructed in a framework where the domains of denotations are usual sets. It has also been known from the applications of VDM [Bjørner,Jones 78] that jumps in programming languages - even such anarchic jumps as in Algol-60 - can be described without continuations (cf. also [Blikle,Tarlecki 83]).

In the majority of projects devoted to the development of software-specification systems - and in particular in the project dedicated to denotational semantics such as e.g. MetaSoft [Blikle 88b], RAISE [Nielsen et al. 88] or BSI/VDM [Larsen et al. 89] - a set-theoretic continuation-free style has been assumed. Although the discussion which follows

applies essentially to all styles of denotational semantics, it refers mainly to - and it has been largely stimulated by - the mentioned above recent developments of that theory and its applications.

The attribute of denotationality may be associated not only to a function of semantics but also to its definition. When we say that a programming language Pascal has been given a denotational semantics we mean that the semantics of Pascal has been given a denotational definition, i.e. a definition which expresses the compositionality of S explicitly by the equations of the form:

$$S[op(syn_1, \dots, syn_n)] = [op](S[syn_1], \dots, S[syn_n]) \quad (2.1)$$

where op is an operation in the algebra Syn and [op] is the corresponding operation in Den. Usually the operations op and [op] are not explicit in these equations but implicit in the (meta)expressions which appear on both sides of (2.1). For instance, we write:

$$S[com_1; com_2](sta) = S[com_2]((S[com_1](sta)))$$

where com stands for a command and sta for a state.

Of course, a non-denotational semantics cannot be given a denotational definition, although a converse situation is possible (Sec.11). If, therefore, we are designing a new software system - rather than formalizing an existing one - and if we choose a denotational form of its definition, then we may be sure that the semantics of the system will become denotational, and therefore that the system will enjoy some important properties. If we choose a non-denotational style of the definition, then we still may be able to construct a denotational semantics, but in that case we have to prove that our semantics has indeed this property. Of course, in such a case there is always a certain risk that our semantics may "come out of control" and become non-compositional.

Now let us discuss the concept of an operational semantics. In contrast to denotationality, the attribute of operationality is not very sharp and applies in the first place to the definition of S rather than to S itself. Moreover, that attribute has never been formalized. We can only point to some techniques and/or metalanguages which are regarded as related to an operational style. For instance the Vienna Definitional Language (VDL) [Lucas, Walk 69], the

structured operational semantics (SOS) [Plotkin 81] or the natural semantics [Kahn 87] belong to that group.

In the opinion of the author we can say that a function of semantics has an operational definition if that definition describes some algorithm of "executing" the underlying syntax. A code of an interpreter is an example of a very operational definition. Another such example, much more abstract in fact, may be a definition written in Plotkin's SOS. Also definitions written in VDM [Bjorner, Jones 82] are to a large extent operational in that sense (although they are denotational at the same time!), since they usually describe an abstract interpreter of a syntax in question. In contrast to the attribute of denotationality, where we can always formally decide whether a given definition is or is not denotational, when we talk about operationality we can only argue about the degree to which our definition is operational in a given context.

In this paper we do not discuss much of the idea of operational semantics. We only wish to express an opinion that operational should not be contrasted to denotational. The author believes that a definition of a software system must always be operational to some degree, that degree depending on several factors such as e.g. the stage of the system development, the target programming language where the system is to be implemented, the expected reader of the definition, etc.

The attribute algebraic, similarly to operational, refers to the definition of S , rather than to S itself, and is not very sharp. Usually an algebraic definition consists of a set of axioms which define a class of the algebras of denotations Den . The corresponding syntax is common for all these algebras and is implicit in their common signature. For any Den the function of semantics is the unique homomorphism from the algebra of ground terms over that signature into Den . An algebraic (axiomatic) definition of a semantics always guarantees that the defined semantics is denotational.

The dichotomy between algebraic and non-algebraic - or better between axiomatic and non-axiomatic - although relevant for applications, is not quite mathematical. From a mathematical viewpoint each definition is axiomatic, since each definition is an axiom which we add to some already existing axioms of an underlying theory. For instance, a VDM-style definition is nothing but a set of axioms added to the (implicit in the definition of VDM) axioms of an appropriate set-theory or domain-theory. In a VDM-style definition we can also

identify an algebra of syntax and an algebra of denotations, although they are frequently not explicit in the definition. In an OBJ-style semantics [Goguen, Meseguer, Plaisted 83], the algebra Den is defined explicitly - although in an axiomatic way - whereas Syn and S are implicit.

In the opinion of the author a semantics should be always sufficiently denotational and should be described at an appropriate level of operationality, the latter depending on the current application. At the same time the algebraic framework seems most appropriate for the description of software whether or not we are using axiomatic techniques and independently of the degree of operationality of the used semantics.

3. BASIC NOTATION

Most definitions of semantics which we discuss in this paper are written in a so called model-oriented style typical of the Danish dialect of VDM and of MetaSoft. In that style the algebras of syntax and of denotations are constructed within some set-theory, rather than described axiomatically as e.g. OBJ [Goguen et al. 83] or ACT-ONE [Ehrig, Mahr 85]. Below we briefly introduce a notation which is used in our definitions. For more details we refer the reader to [Blikle 87b].

For any sets A and B:

- A|B denotes the union of A and B,
- A→B denotes the set of all total functions from A into B,
- A↗B denotes the set of all partial functions from A into B,
- $A \xrightarrow[m]{B}$ denotes the set of all mappings from A into B, i.e. partial functions defined over a finite subset of A
- A^{C*} denotes the set of all finite strings (including the empty string) of the elements of A

Domain equations are written in the form e.g.:

sta : State = Identifier → Integer

by which we mean that each state is a total function from identifiers to integers and that a typical element of the set State is denoted by sta possibly with indices. By $f:A \rightarrow B$ and $f:A \rightarrow B$ we denote the fact that f is a total resp. partial function from A to B . In our paper the formula $a:A$ is used synonymously with $a \in A$. It can be read as "a is of type A", which in some software specification languages means more than to be an element (cf. Sec.9). By $\text{dom} f$ we denote the domain of the function f . For curried functions like $f:A \rightarrow (B \rightarrow (C \rightarrow D))$ we write $f:A \rightarrow B \rightarrow C \rightarrow D$. We also write $f.a$ for $f(a)$ and $f.a.b.c$ for $((f.a).b).c$. For uniformity reasons each many-argument non-curried function is regarded as a one-argument function on tuples. Consequently we write $f.\langle a_1, \dots, a_n \rangle$ for $f(a_1, \dots, a_n)$. Formally this should have led us to writing $f.\langle a \rangle$ rather than $f.a$, but we keep the latter notation as more natural and simpler. We also assume - for a better readability of semantic clauses - that the syntactic argument of a function of semantics is always closed in square brackets. E.g. we write $C.[x:=x+1]$ rather than $C.x:=x+1$. If $f:A \rightarrow B$ and $g:B \rightarrow C$, then $f \circ g:A \rightarrow C$ where $f \circ g = \{(a,c) \mid (\exists b)(f.a=b \ \& \ g.b=c)\}$. In the definitions of functions we frequently use conditional expressions of the form $b \rightarrow c, d$ which stand for:

if b then c else d .

This construction may be nested in which case the expression $b_1 \rightarrow (a_1, (b_2 \rightarrow \dots (b_n \rightarrow a_n, a_{n+1}) \dots))$ is written in a column:

$b_1 \rightarrow a_1,$
 \dots
 $b_n \rightarrow a_n,$
 $\text{TRUE} \rightarrow a_{n+1}$

Sometimes in conditional expressions we are nesting "local constant declarations" of the form $\text{let } x = \text{exp}_1 \text{ in } \text{exp}_2$. The scope of such a declaration is the expression exp_2 .

For any partial function $f:A \rightarrow B$, by $f[b/a]$, where $a \in A$, $b \in B$, we denote a function $f[b/a]:A \rightarrow B$ such that:

$f[b/a].x = x=a \rightarrow b, f.x$

By $[b_1/a_1, \dots, b_n/a_n]$ we denote a total function on $\{a_1, \dots, a_n\}$

which assigns b_i to a_i for $i=1, \dots, n$.

4. DENOTATIONAL SEMANTICS FROM AN ALGEBRAIC PERSPECTIVE

As we have already mentioned in Sec.2, a denotational semantics may be understood as a homomorphism between two many-sorted algebras. This section is devoted to a short introduction of a mathematical theory of such semantics. More on an algebraic framework of denotational semantics may be found in [Blikle 89b] and a general introduction to an axiomatic setting of algebraic semantics is [Ehrig, Mahr 85]).

By a signature we mean a quadruple:

$$\text{Sig} = (\text{Sn}, \text{Fn}, \text{sort}, \text{arity})$$

where Sn is a nonempty set of sort names, Fn is a nonempty set of function names and:

$$\begin{aligned} \text{sort} &: \text{Fn} \rightarrow \text{Sn} \\ \text{arity} &: \text{Fn} \rightarrow \text{Sn}^{c^*} \end{aligned}$$

are functions which associate the sorts of the results and of the arguments respectively to any function name. By an algebra over the signature Sig , or shortly by a Sig -algebra, we mean a triple $\text{Alg} = (\text{Sig}, \text{car}, \text{fun})$ where car and fun are functions interpreting sort names as nonempty sets and function names as total functions on these sets respectively. More precisely, for any $\text{sn} \in \text{Sn}$, car.sn is a set called the carrier of sort sn , and for any $\text{fn} \in \text{Fn}$ with $\text{sort.fn} = \text{sn}$ and $\text{arity.fn} = \langle \text{sn}_1, \dots, \text{sn}_n \rangle$, fun.fn is a total function between corresponding carriers, i.e.

$$\text{fun.fn} : \text{car.sn}_1 \times \dots \times \text{car.sn}_n \rightarrow \text{car.sn}$$

If $\text{arity.fn} = \langle \rangle$, then fun.fn is a nullary function, i.e. it accepts only the empty tuple " $\langle \rangle$ " as an argument. The fact that f is a nullary function with values in A is denoted by $f: \rightarrow A$ and the unique value of f is denoted by $f.\langle \rangle$. Nullary functions are also called constants

In applications, and also in the examples, algebras are treated a little less formally and are defined as collections of carriers and operations between them. In that case it is understood that the signatures are implicit in the notation used.

Two algebras with the same signature are called similar. Given two similar algebras $Alg_1 = (Sig, car_1, fun_1)$ for $i=1,2$, we say that Alg_1 is a subalgebra of Alg_2 if for any $sn \in Sn$,

$$car_1.sn \subseteq car_2.sn$$

and for any $fn \in Fn$, $fun_1.fn$ is the restriction of $fun_2.fn$ to the carriers of Alg_1 . By a homomorphism from Alg_1 (a source algebra) into Alg_2 (a similar target algebra) we mean a many-sorted function H that assigns to any sort $sn \in Sn$ a function:

$$H.sn : car_1.sn \rightarrow car_2.sn \tag{4.1}$$

called the sn -component of H , such that for any $fn \in Fn$ with $sort.fn=sn$:

$$\begin{aligned} &\text{if } arity.fn = \langle sn_1, \dots, sn_n \rangle \text{ with } n \geq 0, \text{ then for any tuple} && (4.2) \\ &\text{of arguments } \langle a_1, \dots, a_n \rangle \in car.sn_1 \times \dots \times car.sn_n \text{ we have} \\ &H.sn.(fun_1.fn.\langle a_1, \dots, a_n \rangle) = \\ &fun_2.fn.\langle H.sn_1.a_1, \dots, H.sn_n.a_n \rangle \end{aligned}$$

By $H : Alg_1 \rightarrow Alg_2$ we denote the fact that H is a homomorphism from Alg_1 into Alg_2 .

With every many-sorted function which satisfies (4.1) we may associate a Sn -sorted relation \equiv_H in Alg_1

$$\equiv_H.sn \subseteq car_1.sn \times car_1.sn$$

called the kernel of H and defined as follows:

$$a_1 \equiv_H.sn a_2 \text{ iff (def) } H.sn.a_1 = H.sn.a_2.$$

Formally the kernel is a function \equiv_H which to every sort $sn \in Sn_1$ assigns a binary relation $\equiv_H.sn$ in $car_1.sn$. It is a well known fact that each

$\equiv_H.sn$ is an equivalence relation. The many-sorted relation \equiv_H is said to be a congruence in Alg_1 , if for any $fn \in Fn_1$ with $arity_1.fn = \langle sn_1, \dots, sn_n \rangle$ and $sort_1.fn = sn$ and for any $a_i, b_i \in car_1.sn_i, i=1, \dots, n$:

if $a_i \equiv_{\langle H, sn_i \rangle} b_i$ for $i=1, \dots, n$
 then $fun_1.fn.\langle a_1, \dots, a_n \rangle \equiv_{H.sn} fun_1.fn.\langle b_1, \dots, b_n \rangle$

A congruence relation is also said to have the extensionality property. In the sequel we shall frequently refer to the following well-known fact:

Fact 4.1 If H is an arbitrary many-sorted function which satisfies (4.1) then \equiv_H is a congruence in Alg_1 iff one can construct such an algebra Alg_2 over the many-sorted family of carriers $\{car_2.sn \mid sn \in Sn\}$ that H becomes a homomorphism from Alg_1 into Alg_2 .

An element of an algebra is said to be *reachable* if it can be constructed from the constants of the algebra by means of the operations of the algebra. An element which is not reachable is called a *junk*. For instance, in an algebra of integers $\langle Int, 1, + \rangle$ all positive integers are reachable and all non-positive are junks. A subset of $car.sn$ which consists of all reachable elements is called a *reachable carrier* of the sort sn . All reachable carriers are closed under the operations of the algebra and therefore, if, all of them are not empty, then they constitute a subalgebra. We call it the *reachable subalgebra*. It is the (unique) least subalgebra of a given algebra. If all the elements of an algebra are reachable, i.e. if the algebra has no junk, then it is called a *reachable algebra*. The following well-known fact is important for the theory of denotational semantics:

Fact 4.2 If $H : Alg_1 \rightarrow Alg_2$ is a homomorphism and Alg_1 is reachable, then:

- (1) the image of Alg_1 in Alg_2 is the reachable subalgebra of Alg_2 .
- (2) H is a unique homomorphism between Alg_1 and Alg_2 .

If for a given algebra Alg there is exactly one homomorphism into any algebra with the same signature, then we say that Alg is *initial*, or - precisely speaking - that it is initial in the class of all algebras similar with Alg . An algebra is called *unambiguous* if each of its reachable elements may be

constructed from the constants of that algebra in exactly one way.

Fact 4.3 An algebra is initial iff it is reachable and unambiguous.

For instance, the algebra $\langle \text{Int}^+, 1, + \rangle$ of positive integers is reachable but not unambiguous, hence not initial, since e.g. 4 may be constructed as $((1+1)+1)+1$ or as $((1+1)+(1+1))$. If, however, we replace "+" by "+1" (successor), then the new algebra becomes initial.

On an algebraic ground a denotational model of a software system is represented by a triple $\langle \text{Syn}, \text{Den}, S \rangle$ where Syn is a reachable algebra of syntax, Den is an algebra of denotations and

$$S : \text{Syn} \rightarrow \text{Den} \tag{4.3}$$

is a unique corresponding denotational semantics. The algebra of syntax is always reachable - a junk syntax makes no practical sense - but does not need to be unambiguous. The admission of ambiguous syntax contrasts our approach from some other algebraic approaches to denotational semantics. It allows us to regard our algebra of syntax as a close approximation of so called concrete syntax rather than merely as an abstract syntax. For a discussion of that problem see [Blikle 89b].

Now, let us explain the introduced concepts in the context of a toy programming language. That example constitutes also a starting point for many other examples which we discuss in the paper. Let the syntax of our language be defined by the following CF-grammar written in the form of a fixed-point set of equations [Blikle 72]:

ide : Ide = {x,y,z}	(identifiers)	(4.4)
exp : Exp = {1 Ide ((Exp {+} Exp {})	(expressions)	
com : Com = Ide {:=} Exp Com {;} Com	(commands)	

The carriers of the corresponding algebra Syn are the three sets (formal languages) defined above and the operations, which are implicit in this grammar, are the following:

set-x	: → Ide	(and the same for y and z)
set-1	: → Exp	

make-exp : Ide → Exp
 plus : Exp × Exp → Exp
 asg : Ide × Exp → Com
 follow : Com × Com → Com

where

set-x.< > = x (4.5)
 set-1.< > = 1
 make-exp.ide = ide (make-exp transforms an identifier into an expr.)
 plus.<exp₁,exp₂> = (exp₁+exp₂)
 asg.<ide,exp> = ide:=exp
 follow.<com₁,com₂> = com₁;com₂

In order to define the corresponding algebra of denotations we first define a domain of states:

sta : State = Ide → Real

and then we define three carriers of Den:

ide : Ide = {x,y,z} (the denotations of identifiers)
 eva : Evaluator = State → Real (the denotations of expressions)
 exe : Executor = State → State (the denotations of commands)

Observe that the denotations of identifiers are the identifiers themselves. Now we define the operations of Den. Since Den is to be a homomorphic image of Syn, for each operation op from Syn we define an operation [op] in Den, such that when [op] is applied to the denotations of the arguments of op, it gives the denotation of the corresponding value of op:

[set-x] : → Ide (and the same for y and z)
 [set-1] : → Evaluator
 [make-exp] : Ide → Evaluator
 [plus] : Evaluator × Evaluator → Evaluator
 [asg] : Ide × Evaluator → Executor
 [follow] : Executor × Executor → Executor

where

```
[set-x].<>           = x
[set-1].<>           = 1.0
[make-exp].ide.sta   = sta.ide
[plus].<eva1,eva2>.sta = (eva1.sta) [+] (eva2.sta)
[asg].<ide,eva>.sta  = sta[(eva.sta)/ide]
[follow].<exe1,exe2>.sta = exe2.(exe1.sta)
```

In these equations 1.0 denotes the real number which corresponds to the numeral (name) "1" and [+] denotes the arithmetical operation of addition. The equations are written in an implicit lambda-notation. For instance, the third equation can be read as follows: [make-exp] is a function that given an identifier ide returns an evaluator [make-exp].ide which given a state sta returns a real number stored under ide in sta.

As is easy to prove, the algebra Syn is reachable and therefore a homomorphism $S : \text{Syn} \rightarrow \text{Den}$ is unique if it exists. The proof of the existence of S is, however, not quite trivial (cf. [Blikle 89b]) since the grammar which underlies (4.4) is ambiguous and therefore so is Syn.

Our homomorphism is a many-sorted function and therefore it may be regarded as a triple of functions $\langle I, E, C \rangle$ where:

```
I : Ide → Ide
E : Exp → Evaluator
C : Com → Executor
```

The definitions of these functions are of the form (4.2) and therefore are implicit in the correspondence between [op] and op. Hence we do not need to write them explicitly. In a more traditional approach, however, these functions are usually defined explicitly whereas the definitions of [op]'s are implicit. In such a case we write:

```
I.[ide]           = ide
E.[1].sta         = 1.0
E.[ide].sta       = sta.ide
E.[(exp1+exp2)].sta = E.[exp1].sta [+] E.[exp2].sta
C.[ide:=exp].sta  = sta[(E.[exp].sta)/ide]
C.[com1;com2].sta = C.[com2].(C.[com1].sta)
```

where the syntactic arguments of the function of semantics are traditionally closed in square brackets (cf. Sec. 3).

If we are developing a denotational model of a software system in a traditional order - i.e. first Syn, then Den and finally S - then we have to prove that S is indeed a homomorphism. In real-life situations this need not be a simple task. Besides, we can easily make a mistake and construct an S which is not compositional. In the sequel we shall see some typical sources of such mistakes. If, therefore, we want to be sure that the mathematical model of a software system becomes denotational, we should start the development of that model from the algebra Den. In that case we can always develop a custom-made syntax Syn such that the existence of the corresponding unique homomorphism (denotational semantics) is guaranteed by the way in which Syn has been developed [Blikle 89b].

5. WHY DENOTATIONAL?

The role of denotationality in software engineering is similar to the role of structurality in programming. Both improve the readability, the comprehensibility and the maintainability of a final product and both allow for the decomposition of a large task into a number of independent subtasks. Also in both cases the non-convinced may give thousands of "clever examples" where the principle of denotationality, respectively structurality, has been violated in the derivation of a "smart" program. It has been known for years, however, that in large-scale applications an anarchic cleverness brings always more disasters than benefits.

Readability is, of course, a property of a definition of semantics. The advantages of writing such definitions in a denotational form are widely known - even if not widely appreciated - and therefore we shall not discuss them here. This section is devoted to a claim that the decision whether the (function of) semantics of a system is to be compositional should be regarded as a design decision since it is relevant for some properties of the future system.

Claim 5.1 The denotationality of semantics allows structured top-down

programming in the corresponding syntax.

Consider an arbitrary software system represented by a denotational model

$$S : \text{Syn} \rightarrow \text{Den}$$

where for the sake of simplicity we assume that the algebras Syn (of syntax) and Den (of denotations) are one-sorted and that Syn and Den denote the corresponding unique carriers. Each programming task in our system consists of the construction of a syntactic object $\text{syn} \in \text{Syn}$ which for a given prespecified denotation $\text{den} \in \text{Den}$ satisfies the equation:

$$S.[\text{syn}] = \text{den}$$

Of course, if den is to be programmable at all, it must belong to the reachable part of Den , since the image of Syn in Den is always reachable (Fact 4.2). This means that there must be $\text{den}_1, \dots, \text{den}_n$ in Den , $\text{syn}_1, \dots, \text{syn}_n$ in Syn and an operation $\text{op} : \text{Syn} \times \dots \times \text{Syn} \rightarrow \text{Syn}$ such that:

- (1) $\text{den} = [\text{op}].\langle \text{den}_1, \dots, \text{den}_n \rangle$ and
- (2) $S.[\text{syn}_i] = \text{den}_i$ for $i=1, \dots, n$.

This means in turn that the task of finding a program for den may be split into the subtasks of finding programs for each of den_i . These subtasks may be assigned to independently working group of programmers, and when programming in the groups is completed, the denotationality of S guarantees that

$$\text{op}.\langle \text{syn}_1, \dots, \text{syn}_n \rangle$$

realizes the global task, i.e. that

$$S.[\text{op}.\langle \text{syn}_1, \dots, \text{syn}_n \rangle] = [\text{op}].\langle \text{den}_1, \dots, \text{den}_n \rangle = \text{den}$$

Of course, the problem of splitting den into the "right" den_i 's does not need to be easy since in general there exist such splits where den_i 's are not reachable. A correct split, however, always exists. As we shall see in an example which comes in Sec.6, if S is not denotational, then it may be impossible to split some programming tasks into independently programmable subtasks.

Claim 5.2 For each software systems with denotational semantics there is a systematic (although not algorithmic) method of deriving a sound and a relatively complete set of program-correctness proof-rules.

Below we give a rough, half-formal, justification of that claim. A full discussion, which requires the introduction of many technical arguments, would go outside the scope of the present paper.

To say that a program is correct is to say that its denotation possess a certain property. Mathematically the properties of denotations are represented by total functions on denotations called predicates:

$\text{pred} : \text{Pred} = \text{Den} \rightarrow \text{Bool}$

where $\text{Bool} = \{\text{true}, \text{false}\}$. Of course, in general predicates may be many-argument functions from several - possibly different - sorts into Bool. We may have, therefore, more than one domain of predicates. When established all chosen domains of predicates are added to the algebra of denotations and then we define some constructors on them. These constructors identify a class of properties that we want to express by means of predicates, such as e.g. partial or total correctness of sequential programs, liveness and deadlock freeness of concurrent systems etc. As we shall see in an example this may also require the introduction of some auxiliary carries to the algebra. The new algebra of denotations corresponds to a language in which besides writing programs we can also express their properties.

When we are done with the extended algebra we proceed to establishing the proof rules which in our case are lemmas of, roughly, the following form:

for any reachable pred and any reachable $\text{den}_1, \dots, \text{den}_n$: (5.1)

$\text{pred}.\langle [\text{op}].\langle \text{den}_1, \dots, \text{den}_n \rangle \rangle = \text{true}$

iff

there exist predicates $\text{pred}_1, \dots, \text{pred}_n$ such that,

(1) $\text{pred}_i.\text{den}_i = \text{true}, i=1, \dots, n$

(2) $\Phi.\langle \text{pred}, \text{pred}_1, \dots, \text{pred}_n \rangle$

where $\Phi.\langle \text{pred}, \text{pred}_1, \dots, \text{pred}_n \rangle$ expresses a certain relationship between all introduced predicates. We establish such a lemma for each operation $\text{op}:\text{Syn} \times \dots \times \text{Syn} \rightarrow \text{Syn}$ of the algebra of syntax (notice that [op] denotes the

counterpart of op in Den). Of course, if $op \rightarrow Syn$, then $i=0$ and "iff" is followed only by $\$pred$.

The lemmas constructed in that way can be used in structured-inductive proofs of programs' correctness. Each of them facilitates the reduction of a global correctness problem - for a compound object - into a number of local correctness problems - for the components of that object. Of course, we can also develop a formalized proof system (a logic) in which our lemmas become inference rules. The "if" part of each lemma guarantees then the soundness of such a rule and the "only if" - a relative completeness in the sense close to that of [Cook 78].

In general, by a completeness of a logic we mean the fact that every true statement which can be expressed in the language of that logic can be proved. In the case of a logic of programs we can only expect a relative completeness, by which we mean that in every concrete situation the applicability of each of our lemmas depends on the following three factors:

- (i) that the required $pred_i$'s are reachable, i.e. expressible in our logic,
- (ii) that we are able to express condition (2) of (5.1) in our logic,
- (iii) that we are able to prove that condition in our logic, i.e. that we can prove it on the ground of the corresponding theory of data (such as e.g. integers, reals, records, etc.)

Now, let us discuss an example. Consider the programming language defined in Sec.4: We shall construct a corresponding Hoare-like proof-system for the partial correctness of commands. First we expand the algebra Den of our language by two new carriers: Condition - which constitutes an auxiliary carrier - and Predicate. The new algebra has five carriers:

```
ide : Ide      = {x,y,z}
eva : Evaluator = State  $\rightarrow$  Real
exe : Executor  = State  $\rightarrow$  State
con : Condition = State  $\rightarrow$  Bool
pred : Predicate = Executor  $\rightarrow$  Bool
```

Then we define the constructors of conditions and predicates. As the constructors of conditions we choose e.g.:

[less] : Evaluator x Evaluator → Condition.
 [less].⟨eva₁,eva₂⟩.sta = eva₁.sta < eva₂.sta

[and] : Condition x Condition → Condition
 [and].⟨con₁,con₂⟩.sta = con₁.sta & con₂.sta

etc.

Of course, in both definitions "=" denotes the equality in Bool. For predicates we define just one constructor, which corresponds to the property of partial correctness of commands. Since the commands in our language represent total functions, this constructor is defined as follows:

[parcor] : Condition x Condition → Predicate
 [parcor].⟨con₁,con₂⟩.exe =
 (∀sta∈State)[con₁.sta=true → con₂.exe.sta=true]

Now we can formulate our lemmas. For a better readability we assume that:

PRE con₁ : exe POST con₂

stands for [parcor].⟨con₁,con₂⟩.exe = true. We formulate one lemma for each of our two constructors of commands - "follow" and "asg" - introduced in Sec.4:

for any con₁, con₂ and any exe₁, exe₂: (5.2)

PRE con₁: follow.⟨exe₁,exe₂⟩ POST con₂

iff

there exist conditions con₁₁ and con₁₂ such that,

(1) PRE con₁: exe₁ POST con₁₁

PRE con₁₂: exe₂ POST con₂

(2) (∀sta)(con₁₁.sta=true → con₁₂.sta = true)

for any con₁, con₂ and for any ide, eva: (5.3)

PRE con₁: asg.⟨ide,eva⟩ POST con₂

iff

(∀sta∈State)(con₁.sta=true → con₂.sta[(eva.sta)/ide]=true)

Observe that the form of (5.3) is such as if the syntactic operation asg were nullary, which is, of course, not the case. In our example we have not

introduced predicates for identifiers and evaluators, and therefore we do not introduce any statement that corresponds to (1) of (5.1), and in the clause corresponding to (2) of (5.1) we refer directly to *ide* and *eva* rather than to the predicates on them.

It should be also noticed that in (5.2) and (5.3) we have omitted an explicit assumption about the reachability of predicates. This makes our lemmas a little stronger than in the general case (5.1). In order to make them exact analogues of (5.1) we should have assumed that our con_1 's are reachable, which in this particular case is not necessary. On the other hand all the predicates that appear in our lemmas are of the form:

[parcor].<con₁,con₂>

and hence they are not quite arbitrary.

Our lemmas may be now used in the construction of program-correctness logic, in that case a Hoare's partial-correctness logic. Assume that the expanded syntax of our programming language covers the following syntax of formulas and of correctness statements:

for : For = Exp {less} Exp | For {and} For (formulas)
cst : Cst = {pre} For {;} Com {post} For (correctness statements)

This syntax has an obvious semantics and the proof rules corresponding to (5.2) and (5.3) are the following:

```
pre for1 : com1 post for11
pre for12 : com2 post for2
for11 implies for12
-----
pre for1 : com1;com2 post for2

for1 implies for2[exp/ide]
-----
pre for1 : ide := exp post for2
```

Of course, $for_2[exp/ide]$ denotes a formula which results from for_2 after substituting *exp* for all free occurrences of *ide*. Each "enumerator" expresses a

conjunction of metaformulas from which one may infer the "denominator". Observe that since in these rules we are talking about formulas rather than about conditions we have now implicitly introduced the assumption about the reachability of the involved predicates.

The readers interested in the mathematical problems related to the construction of logics for the denotational models of software may find more technical material in [Blikle 87b] and [Blikle 88a].

6. WHY NOT NON-DENOTATIONAL?

In the former section we have discussed some advantages of denotational semantics. Here we show that these advantages may disappear if a semantics is not denotational. Consider as an example our little programming language of Sec.4 which we now modify by setting

$$\begin{aligned}
\text{com} : \text{Com} &= \text{Ide} \{:=\} \text{Exp} \mid \{(\} \text{Com} \{;\} \text{Com} \{)\} && \text{and} \\
C.[(\text{com}_1; \text{com}_2)] &= && (6.1) \\
& \text{fai.com}_1 + \text{fai.com}_2 + C.[\text{com}_1] \cdot C.[\text{com}_2], \\
& \text{TRUE} + \text{nullsta}
\end{aligned}$$

where

fai.com = first assignable identifier in com, i.e. the left-hand side identifier in the first assignment of com,
nullsta = a function that transforms any state to [0.0/x, 0.0/y, 0.0/z]

Similarly to our notation for 1.0 vs. 1 (Sec.4) 0.0 stands for the real number "zero", whereas "0" denotes the corresponding syntax (symbol).

In the new version of the language the syntax has been modified by introducing parentheses into compound commands (for otherwise (6.1) would be ambiguous) and the semantics has been modified by making the effect of commands dependent on fai's. The latter modification makes our semantics not denotational since now C.[(com₁; com₂)] depends on more than just C.[com₁] and C.[com₂]. We

prove that fact by showing that the equivalence relation \equiv_C is not a congruence (cf. Fact 4.1). Indeed:

$$\begin{aligned} C.[(x:=1; y:=2)] &= C.[(y:=2; x:=1)] \text{ but} \\ C.[(x:=1; y:=2); y=(x+y)] &\neq C.[(y:=2; x:=1); y=(x+y)] \end{aligned}$$

Although our semantics is not denotational, its definition is still in a structural inductive form, hence it is easy to understand and implement. Why then should we bother about the non-denotationality? As we are going to see, structured programming and structured-inductive proofs are not possible in the new language.

Consider first the problem with structured programming and take as an example a task of writing a program com that loads 2.0 to x and to y , i.e. a program which satisfies the equation:

$$C.[com].sta = sta[2.0/x, 2.0/y] \tag{6.2}$$

Now assume that we want to split this task in two new ones, described by two following specifications:

$$\begin{aligned} C.[com_1].sta &= sta[1.0/x, 1.0/y] \\ C.[com_2].sta &= sta[(E.[(x+y)].sta)/x, E.[(x+y)].sta)/y] \end{aligned}$$

If we assign these subtasks to two programmers, then the first has (at least) a choice of writing (we use a simplified intuitive notation for commands):

either $(x:=1 ; y:=1)$ or $(y:=1 ; x:=1)$

and the other has (at least) a choice of writing:

either $(y:=(x+y) ; x:=y)$ or $(x:=(x+y) ; y:=x)$

As is easy to see unless our programmers communicate about the syntax of their target programs, they cannot guarantee that $(com_1; com_2)$ will satisfy (6.2). Hence in our language top-down structured programming is not feasible.

Now consider the problem of structured-inductive proofs. First observe that in the new language none of the two implications in (5.2) is true. Indeed,

although

PRE true : C.[(x:=1;x:=1)] POST x=0

is certainly satisfied, there are no intermediate assertions which can be used in a proof of that fact based on (5.2), since

PRE con : C.[x:=1] POST x=0

is false for any con:Condition. Similarly, although both

PRE true : C.[x:=1] POST x=1

PRE x=1 : C.[x:=1] POST x=1

are true, the statement

PRE true : C.[(x:=1;x:=1)] POST x=1

is not true.

It is also not very easy to modify the logic of Sec.5 to the new language. The major problem consists in the fact that in the present case we do not have - and we cannot have - an algebra of denotations. We cannot apply, therefore, the routine way for the construction of logic described in Sec.5. In fact, we cannot use here any of the known techniques of formal logic, since all these techniques are inherently based on the assumption that the underlying language of terms and formulas has a denotational semantics. In the usual formal logic - whether classical, algorithmic, temporal or any other - we never talk about the properties of (the syntax of) formulas. And here we should have to do that in order to formulate a proof rule for ";".

In the opinion of the author the only rational way of solving the problem consists of "repairing" a non-denotational semantics by making it denotational and then applying the method of Sec.5. We shall discuss this method in Sec.8.

Our example of a non-denotational programming language has been made a little artificial in order to be sufficiently simple. Similar examples may be shown, however, on a more natural ground. Take, for instance, Pascal and its concept of a type. As we can read in [Jensen,Wirth 75]: "A data type determines a set

of values which variables of that type may assume...". This suggests that types in Pascal are just the sets of values. If we assume that, then the type definitions of Fig.6.1 have the same denotation. At the same time, however, according to the standard of Pascal the blocks of Fig.6.2 have different denotations since the first block generates a type error whereas the second does not.

<pre>TYPE DEFINITION ONE: type item = record no,size : integer; end; object = record no,size : integer; end</pre>	<pre>TYPE DEFINITION TWO: type item = record no,size : integer; end; object = item;</pre>
---	---

Fig.6.1

<pre>"TYPE DEFINITION ONE" var x : item; y : object; begin x := y; end</pre>	<pre>"TYPE DEFINITION TWO" var x : item; y : object; begin x := y; end</pre>
--	--

Fig.6.2

As the type-definition parts of the blocks of Fig.6.2 are semantically equivalent and the remaining parts of the blocks are identical, the equivalence relation which corresponds to our semantics is not a congruence. This means that the semantics of Pascal, in which types are regarded as sets of values, is not denotational.

Our example should not be understood as an argument that Pascal cannot be given a denotational semantics. That example only indicates that in the context of Pascal the interpretation of types as sets is too abstract to be denotational. Next two sections are devoted to a general discussion of the relationship between the denotationality and the abstraction of a semantics.

7. DENOTATIONALITY VERSUS ABSTRACTION

In the context of our applications each semantics describes the effect of the execution of syntax. Of course, such an effect may be described in a more or less detailed way. For instance, the effect of the execution of an imperative program may be described by the set of sequences of memory states generated by all executions of that program, or - more abstractly - by a corresponding input-output function on states, or - even more abstractly - by a function on states truncated to the global variables of the program. The more abstract is a semantics, the less information is carried by denotations. The class of all semantics of a given syntax may be viewed as a spectrum preordered by a reflexive and a transitive relation of abstraction. On one end of that spectrum we have the least abstract semantics, which maps each syntactic object identically into itself, and on the other end - the most abstract semantics, which maps all syntactic object of the same sort into a common denotation.

Of course, both extremes of our spectrum are trivial. A good semantics should provide all the relevant information about the effect of the execution of a piece of syntax, but at the same time it should hide all the irrelevant details of executions. Of course, what is relevant and what is not depends on the current application. We should also be aware of the fact that if a semantics is to be denotational, then the abstraction levels of its components (e.g. of I, E, C, in the example of Sec.4) must be mutually balanced. For instance, if the denotations of expressions do not carry enough information in order to compute the denotation of commands in which they appear, then the semantics is not denotational.

In Sec.6 we have seen two examples of non-denotational semantics. In the first example the denotations of commands do not include an information about $fai's$, although that information is relevant for the behavior of compound commands. It seems rather clear intuitively that since we have made the aforesaid behavior dependent on $fai's$, we should have put an appropriate information on $fai's$ into the denotations, i.e. we should have made our semantics less abstract. In the example with Pascal, the set-theoretical meaning of a type is not sufficiently informative since each compiler of Pascal discriminates between two

compound types with different names, unless they have been explicitly declared to be equal. The compilers of Pascal compare the definitions of types rather than their set-theoretical interpretations. We need, therefore, more information about a type - more than just the corresponding set of values - in order to predict the behavior of a program where that type has been used. Again, in order to make our semantics denotational we have to make it less abstract.

This section is devoted to a formal discussion of a trade-off between denotationality and abstraction. Let us start from introducing a few basic concepts. Consider two, not necessarily denotational, semantics of the same syntax:

$$\begin{aligned} S &: \text{Syn} \rightarrow \text{Den} \\ S^* &: \text{Syn} \rightarrow \text{Den}^* \end{aligned}$$

where $\text{Syn} = \langle \text{Sig}, \text{car-s}, \text{fun-s} \rangle$ and $\text{Sig} = \langle \text{Sn}, \text{Fn}, \text{sort}, \text{arity} \rangle$. We assume that if S or S^* are not denotational, then Den and/or Den^* are just families of sets rather than algebras. If for any $sn \in \text{Sn}$ and any $\text{syn}_1, \text{syn}_2 \in \text{car-s.sn}$:

$$S^*.sn.syn_1 = S^*.sn.syn_2 \rightarrow S.sn.syn_1 = S.sn.syn_2 \quad (7.1)$$

i.e. if $\equiv_{S^*} \subseteq \equiv_S$ in a componentwise way, than we say that

S^* is less abstract than S

We also say that S^* is adequate for S , since S^* bears - in a certain sense - at least as much information as S . Whenever we change a semantics by enriching denotations, the new semantics becomes less abstract (i.e. more informative) than the former.

We say that S and S^* are equally abstract if $\equiv_{S^*} = \equiv_S$. Of course, equally abstract semantics need not be equal. However, if S and S^* are equally abstract and one of them is denotational, then so must be the other one. In other words: we cannot adequately repair non-denotationality without a loss of abstraction.

Proposition 7.1 For any S there exists a maximally abstract S^* which is both denotational and adequate for S .

Proof. Let \equiv^* be the transitive closure of the union of all congruences included in \equiv_S . This relation is the largest congruence included in \equiv_S . The corresponding homomorphism $S^* : \text{Syn} \rightarrow \text{Den}/\equiv^*$ is the maximally abstract denotational semantics which is adequate for S. []

Of course, the semantics S^* constructed in the proof above is not the unique semantics that satisfies Proposition 7.1. There are many such semantics, but all of them are equally abstract with S^* .

When we construct a semantics we should care not only about its global level of abstraction, but also about a balance between the abstraction levels of the semantics assigned to different sorts of **Syn**. If syntactic objects of sort sn_1 are used in the construction of syntactic objects of sort sn_2 , then on one hand the denotations of sort sn_1 should be sufficiently informative in order to enable the calculation of the denotations of sort sn_2 , but on the other hand they should not carry more information than necessary, i.e. they should be as abstract as possible. In order to express that claim in a more formal way we introduce a few technical concepts.

Let **Syn** be a many-sorted algebra of syntax. By a context of arity $\langle sn_1 \rangle$ and sort sn_2 we mean a function of the type

$$ct : \text{car-s.}sn_1 \rightarrow \text{car-s.}sn_2$$

which represents a "term with a hole" like e.g. (cf. Sec.4):

$$\begin{aligned} (\lambda com)(com;x:=f) & : Com \rightarrow Com \\ (\lambda exp)(y:=exp;x:=f) & : Exp \rightarrow Com \\ (\lambda ide)(y:=ide+x;z:=f) & : Ide \rightarrow Com \\ (\lambda ide)(ide+x) & : Ide \rightarrow Exp \end{aligned}$$

For a more formal definition of that concept, based on the notion of derived operators, see e.g. [Stoughton 88]. Now, let

$$S : \text{Syn} \rightarrow \text{Den}$$

be an arbitrary - not necessarily denotational - semantics. Two syntactic objects syn_1 and syn_2 of the same sort sn_1 are said to be context-equivalent w.r.t. a sort sn_2 , in symbols

$\text{syn}_1 \text{ CE.} \langle \text{sn}_1, \text{sn}_2 \rangle \text{ syn}_2$,

if one of them may be replaced by the other in any context of sort sn_2 without changing the meaning of the whole phrase, i.e. if for any context $\text{ct:car-s.sn}_1 \rightarrow \text{car-s.sn}_2$

$$S.\text{sn}_2.(\text{ct}.\text{syn}_1) = S.\text{sn}_2.(\text{ct}.\text{syn}_2)$$

Given two sorts sn_1 and sn_2 we say that the semantics $S.\text{sn}_1$ of car-s.sn_1 is sufficiently abstract w.r.t. the semantics $S.\text{sn}_2$ of car-s.sn_2 , if any two syntactic objects of the sort sn_1 which are context-equivalent w.r.t. sn_2 , have the same meaning, i.e. if for any $\text{syn}_1, \text{syn}_2 \in \text{car-s.sn}_1$:

$$\text{syn}_1 \text{ CE.} \langle \text{sn}_1, \text{sn}_2 \rangle \text{ syn}_2 \rightarrow S.\text{sn}_1.\text{syn}_1 = S.\text{sn}_1.\text{syn}_2 \quad (7.2)$$

For instance, we say that the semantics of expressions of a programming language is sufficiently abstract w.r.t. the semantics of commands, if any two expressions which are context-equivalent in commands have the same denotation.

If in (7.2) the opposite implication holds, then we say that $S.\text{sn}_1$ is sufficiently informative w.r.t. $S.\text{sn}_2$.

Proposition 7.2 A semantics S is denotational iff any two of its components $S.\text{sn}_1, S.\text{sn}_2$ are sufficiently informative w.r.t. each other.

The proof of that proposition is quite routine, but since it requires the formalization of the concept of a context we omit it.

In many software systems there exists a sort - usually called "programs" - such that the syntactic objects of that sort may be executed without any context, whereas the objects of all other sorts - e.g. expressions, declarations, commands - can be executed only in the context of programs. The user of such a system is, of course, mainly interested in the behavior of programs and he wants to see these behaviors at an abstraction level adequate for the intended applications. For instance, in a general-purpose programming language an adequate abstraction level for programs may be represented by I/O functions, whereas in a programming language for controlling robots - by sets of sequences of states. The designer of the system should, therefore, choose an

adequate abstraction level for programs in the first place, and then he should "tune" to it the abstraction level of all other sorts in such a way that the whole semantics becomes denotational and maximally abstract.

If in a semantics S all its components $S.sn_i$, including the component for programs, are sufficiently abstract w.r.t. the semantics of programs, then R.Milner calls S fully abstract [Milner 77]. For instance, the semantics of our toy programming language of Sec.4 is fully abstract, if we assume that commands play the role of programs, but a semantics of a programming language with blocks, where the denotations of blocks include an information about local variables, is in general not fully abstract. In the latter case we can construct two blocks which have different denotations although they are interchangeable in any program.

The term "full abstraction" has not been chosen very adequately with respect to what the word "full" means in colloquial English. When we say "fully abstract" we could have expected that the semantics in question is "as abstract as it can be". In fact, however, it is only sufficiently abstract with respect to the semantics of programs. If, therefore, the semantics of programs is little abstract, then the other semantics may be also little abstract. For instance, a trivial semantics which maps syntax identically into itself is fully abstract. Since it is also denotational and adequate with respect to any other semantics of the same syntax, we can formulate the following trivial proposition about the reparation of non-denotational semantics:

Proposition 7.3 For any semantics $S:Syn \rightarrow Den$ there exists a semantics $S^*:Syn \rightarrow Den^*$ such that:

- (1) S^* is denotational,
- (2) S^* is adequate w.r.t. S ,
- (3) S^* is fully abstract.

Of course, (3) above makes only sense if the common signature of Syn and Den contains a sort of programs. For the sake of further investigations (in Sec.8) we slightly generalize the concept introduced by R.Milner and say that a given semantics S is fully abstract w.r.t. a given sort sn if all $S.sn_i$'s are sufficiently abstract w.r.t. $S.sn$. We also assume that the concept of full abstraction is applicable to any semantics rather than only to denotational semantics, as it was the case with Milner's definition.

8. REPAIRING SEMANTICS

In this section we discuss the problems of repairing non-denotational or not fully abstract semantics. We start from non-denotationality.

As Proposition 7.3 indicates, the repairability problem for non-denotational semantics should be formulated with a certain care, since otherwise it may become trivial. Even if we request that besides satisfying (1)-(3) of that proposition the new semantics preserves all the information that was carried by S , we can still find a trivial solution, namely

$$S^+ : \text{Syn} \rightarrow \text{Den} \times \text{Syn}$$

where $S^+.\text{syn} = \langle S.\text{syn}, \text{syn} \rangle$. The semantics S^+ is as little abstract as S^* , i.e. $\varepsilon_{S^+} = \varepsilon_{S^*}$. The latter is denotational on the expense that it says nothing about the behavior of programs. The former says everything that S does, but it adds the whole syntax to that information. If we want to use S^+ in order to tell a programmer what his program is supposed to do, we have to give to him that program explicitly. The semantic S^+ is therefore as useless as S^* .

Although of no value for applications, S^+ indicates a certain way of searching for the most abstract denotational semantics which is adequate for S (cf. Proposition 7.1). In fact, we can always try to repair S by a semantics S^\wedge , where $S^\wedge.\text{syn} = \langle S.\text{syn}, A.\text{syn} \rangle$ and where $A.\text{syn}$ provides the additional information which is necessary to make S^\wedge denotational. In the worst case $A.\text{syn} = \text{syn}$, but frequently we can do much better.

Consider as an example the non-denotational language described in Sec.6, and let - a little informally - $S = \langle I, E, C \rangle$. In this semantics C is not sufficiently informative w.r.t. itself since the denotations of commands lack the information about the first assignable identifier and therefore the denotation of a compound command cannot be "computed" from the denotations of its subcommands. In order to repair S we have to add the missing information to the denotations of commands. This leads to a semantics $S^\wedge = \langle I, E, C^\wedge \rangle$, where I and E

are the same as in S and where:

$$C^{\wedge} : Com \rightarrow Executor \times Ide$$

$$C^{\wedge}.[ide:=exp] = \langle (\lambda sta) sta[E.[exp].sta/ide], ide \rangle \tag{8.1}$$

$$C^{\wedge}.[(com_1;com_2)] = \begin{aligned} & \text{let } \langle exe_1, ide_1 \rangle = C^{\wedge}.[com_1] \text{ in for } i=1,2 \\ & ide_1 \# ide_2 \rightarrow \langle exe_1 \circ exe_2, ide_1 \rangle \\ & \text{TRUE} \quad \rightarrow \langle (\lambda sta) nullsta, ide_1 \rangle \end{aligned} \tag{8.2}$$

The semantics S^{\wedge} is, of course, both adequate for the former one and denotational. We can also show that it is a maximally abstract such semantics. Indeed, assume that $S'=(I',E',C')$ is denotational, inherently more abstract than S^{\wedge} and adequate for S. By these assumptions I' and E' must be equally abstract with I and E respectively. Therefore C' must be inherently more abstract than C^{\wedge} , i.e. there must be two commands com_1 and com_2 such that:

- (1) $C^{\wedge}.[com_1] \neq C^{\wedge}.[com_2]$ and
- (2) $C'.[com_1] = C'.[com_2]$

Let $C^{\wedge}.[com_i]=\langle exe_i, ide_i \rangle$ for $i=1,2$. From the adequacy of C' for C and by (2) we may conclude that $C.[com_1]=C.[com_2]$, i.e. $exe_1=exe_2$. Therefore, by (1), $ide_1 \# ide_2$. In that case:

$$C.[(com_1;ide_2:=1)] \neq C.[(com_2;ide_2:=1)]$$

and hence, again by the adequacy of C' for C,

$$C'.[(com_1;ide_2:=1)] \neq C'.[(com_2;ide_2:=1)]$$

which in the virtue of (2) contradicts the denotationality of S' .

It is also not difficult to show that our semantics S^{\wedge} is fully abstract w.r.t. commands. One should only prove that any two identifiers, any two expressions and any two commands which have different denotations can be discriminated by a certain command-context. An easy proof is left to the reader.

It should be stressed that although in our example the most abstract

denotational semantics that adequately repairs S turned out to be fully abstract, it does not need to be so in general. In fact, S^* is fully abstract because the original semantics S was so. Assume, however, that we repair a semantics which is similar to that of Sec.6, but where the denotations of expressions include also their syntax. In that semantics E is not sufficiently abstract w.r.t. C and that property will be preserved in the corresponding most abstract S^* , since on the way from S to S^* we are always lowering the level of abstraction.

The denotationality and the full abstraction of a semantics may be both repaired (or spoiled) either by lowering or by raising the abstraction level. It is so since both these properties require a certain balance between the abstraction levels of the components of a semantics. Let us illustrate this statement by an example where we compare four different semantics.

Assume first that we extend the syntax of the language defined in Sec.4 by a new sort called programs:

prog : Program = {begin} Com {end}

For the extended syntax we define four different semantics:

(1) Semantics S , denotational and fully abstract, which results from the (natural) semantics defined in Sec.4 by setting

$P : \text{Program} \rightarrow \text{Executor}$
 $P.[\text{begin com end}] = C.[\text{com}]$

(2) Semantics S^{nd} , not denotational but fully abstract, which results from the semantics of (1) by assuming that C^{nd} has been spoiled as in (6.1), by making the denotation of a compound command dependent on first assignable identifiers. This semantics is not denotational since C^{nd} is not sufficiently informative with respect to itself; C^{nd} does not give enough information for predicting the behavior of a command in the context of another command.

(3) Semantics S^{nfa} , denotational but not fully abstract, which results from the semantics of (1) by assuming that:

$C^{\text{nfa}} : \text{Com} \rightarrow \text{Executor} \times \text{Ide}$

```

cnfa.[ide:=exp] =
  <(\sta)sta[(E.[exp].sta)/ide], ide>
cnfa.[(com1;com2)] =
  let <exe1,ide1> = cnfa.[com1] in for i=1,2
  <exe1·exe2, ide1>

```

```

pnfa : Program → Executor
pnfa.[begin com end] = first.(cnfa.[com])

```

Here the denotations of commands include the information about fai's, but unlike in (8.2) this information is irrelevant for the "executorial effect" of commands. The semantics S^{nfa} is not fully abstract since C^{nfa} is not sufficiently abstract with respect to p^{nfa} . Indeed, C^{nfa} gives more information than one needs to calculate the denotations of programs.

(4) Semantics S^{\wedge} , denotational and fully abstract, which results from the semantics defined at the beginning of this section (i.e. C^{\wedge} is defined by (8.1) and (8.2)) by setting the semantics of programs as follows:

```

P^ : Program → Executor × Ide
P^.[begin com end] = C^.[com]

```

If by $S_1 \longrightarrow S_2$ we denote the fact that S_1 is inherently less abstract than S_2 , then our four semantics fit into the diagram of Fig.8.1, where in parentheses we indicate the current domains of the denotations of commands and of programs, respectively.

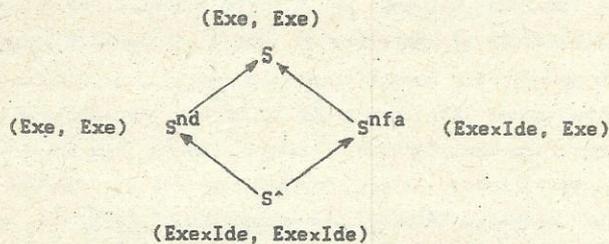


Fig.8.1

Let us analyze this diagram more carefully, since it puts some additional light into the nature of denotationality and full abstraction.

In S^{nd} the denotationality of S has been spoiled as a result of lowering the abstraction of only one component of S , namely C , by changing the behavior of commands without changing their denotations. The new behavior requires more information than carried by the old denotations and therefore C^{nd} is not sufficiently informative w.r.t. itself. In turn, the denotationality of S^{\wedge} has been spoiled as a consequence of raising the abstraction of C^{\wedge} - by changing the denotations of commands from $Executor_{Ide}$ to $Executor$ - without changing their behavior. The new denotations carry less information than required by the the old behavior.

In S^{nfa} the full abstraction of S has been spoiled as a result of lowering the abstraction of C - by changing the denotations of commands from $Executor$ to $Executor_{Ide}$ - without lowering the abstraction of P ; C^{nfa} is not sufficiently abstract w.r.t. P^{nfa} . In turn, the full abstraction of S^{\wedge} has been spoiled by raising the abstraction of P^{\wedge} - as a result of changing the denotations of programs - without raising the abstraction of C^{\wedge} .

Our analysis indicates two different strategies of repairing a semantics. One consists in changing the "behavior" of software, the other consists in changing the denotations. The choice of strategy is a pragmatic issue. We change denotations whenever we feel that the original semantics describes the intended behavior of our software and that all we want to do is to describe that behavior in a compositional and/or fully abstract way. We change the behavior of software if we decide that the assumed denotations adequately express the effect of the execution of software.

In our example the designer of the language should probably conclude that the denotations of commands assumed in S^{nd} are adequate for applications and that the non-denotationality of semantics is due to an awkward behavior of commands. He should then correct his semantics by passing to S rather than to S^{\wedge} . The case where the other strategy seems to be more appropriate is that of Pascal (Sec.6). In that case the designer should change the denotations rather than the behavior of software. Pascal types should not be regarded as sets since they have to be compared during the execution of programs, and the comparison of sets may be computationally too expensive, if computable at all. Readers interested in a denotational model of Pascal types may refer to [Blikle 87b]).

Of course, a method of repairing non-denotationality is only needed if at the beginning we have unintentionally defined a non-denotational semantics. As was

already pointed out in Sec.4, when we design a software system from a scratch, a safe systematic way of giving it a denotational semantics consists in starting the designing process from the algebra of denotations. In that case we have no syntax around and therefore we have no chance of making the constructors of denotations dependent on syntax. Of course, it may happen that we will not be able to express adequately some of our constructors on the ground of assumed denotations. For instance, we will not be able to define a parallel composition of processes if we assume that the denotations of processes are I/O functions or even sets of computations. In such a case we have to modify appropriately our model. e.g. by introducing new sorts (see [Blikle 87a] for an example of a denotational semantics of a language with concurrency). In any case, a design mistake like that described in Sec.6 cannot happen if we start the design of a software system from the algebra of denotations.

To complete the discussion of the repairability of semantics let us briefly comment on the construction of a logic for the non-denotational language defined in Sec.6. As we have mentioned there, the only rational way of tackling the problem of logic in such a case seemed to consist in first repairing the non-denotationality of the semantics and then constructing the logic in the usual way, as described in Sec.5. For the sake of our discussion let us assume that we repair the semantics S^{nd} by changing it into S^{\wedge} , since otherwise - i.e. in the case of S - the logic has been described already in Sec.5. For simplicity we omit programs in our investigations.

Let us start from defining explicitly two constructors of command denotations for which we intend to develop our proof rules:

$$\begin{aligned} [assg]^{\wedge} &: \text{Ide} \times \text{Evaluator} \rightarrow \text{Executor} \times \text{Ide} \\ [assg]^{\wedge} \cdot \langle \text{ide}, \text{eva} \rangle &= \langle [assg] \cdot \langle \text{ide}, \text{eva} \rangle, \text{ide} \rangle \end{aligned}$$
$$\begin{aligned} [follow]^{\wedge} \cdot \langle \langle \text{exe}_1, \text{ide}_1 \rangle, \langle \text{exe}_2, \text{ide}_2 \rangle \rangle &= \\ \text{ide}_1 \# \text{ide}_2 &\rightarrow \text{exe}_1 \circ \text{exe}_2, \\ \text{ide}_1 = \text{ide}_2 &\rightarrow \text{nullsta} \end{aligned}$$

In order to construct a Hoare-like logic of partial correctness for the new semantics we have to modify the concept of partial correctness in such a way that it captures also an information about $\text{fal}'s$. Let:

[parcorfai] : Condition x Condition x Ide \rightarrow Predicate
[parcorfai].<con₁,con₂,ide>.<exe,ide_e> =
(\forall sta \in State)[con₁.sta=true \rightarrow con₂.(exe.sta)=true] & ide=ide_e

and let

PRE con₁ : exe POST con₂ FAI ide

stand for

[parcorfai].<con₁,con₂,ide>.<exe,ide_e>=true.

Let allzero:Condition where:

allzero.sta=true \leftrightarrow (\forall ide)(sta.ide=0.0)

The modified rules are as follows:

for any con_i and <exe_i,ide_i> for i=1,2 and for any ide: (8.3)

PRE con₁:

[follow][^].<<exe₁,ide₁>,<exe₂,ide₂>

POST con₂ FAI ide

iff

there exist con₁₁, con₁₂ and ide' such that:

(1) PRE con₁: exe₁ POST con₁₁ FAI ide

PRE con₁₂: exe₂ POST con₂ FAI ide'

(2) [ide \neq ide' \rightarrow

(\forall sta)(con₁₁.sta=true \rightarrow con₁₂.sta = true)] &

[ide=ide' \rightarrow

(\forall sta)(con₂.sta=true \rightarrow zero.sta=true)]

for any con₁, con₂ and for any ide, ide' and eva: (8.4)

PRE con₁: [asg][^].<ide,eva> POST con₂ FAI ide'

iff

(\forall sta \in State)(con₁.sta=true \rightarrow con₂.sta[(eva.sta)/ide]=true) &

ide=ide'

For the sake of brevity we shall not transform these rules into formalized inference rules as in Sec.5. Let us only mention that although the rules

developed here correspond in some sense to the language of Sec.6, they are not formally applicable to that language. The reason is that

PRE con₁ : exe POST con₂ FAI ide

describes a property of $\langle \text{exe}, \text{ide} \rangle$ rather than of exe , whereas exe , rather than $\langle \text{exe}, \text{ide} \rangle$, are the denotations of commands in the language of Sec.6. We can also see quite clearly now that neither for S^{\wedge} , hence also nor for S^{nd} , there exist predicates on executors that could be used in the proofs of their partial correctness. This means that there is no logic of structured-inductive proofs of the partial correctness of commands for the language of Sec.6.

9. ON THE BORDERLINE OF DENOTATIONALITY

No theory can capture all the reality. In some applications a very orthodox attitude to the principle of denotationality may lead to overcomplicated or to non-implementable models. In such cases a pragmatic solution may consist of adding a non-denotational supplement to a denotational core of the model. As long as the bulk of the system is described in a compositional way, such a style may be still acceptable. Below we discuss two typical examples. A third, rather singular, is discussed in Sec.10.

Our first example is related to types in programming languages. As we have mentioned already in Sec.5 and Sec.6, types that we want to think about may be sets with a non-computable equality relation, whereas types that we implement are usually less abstract objects, e.g. some equivalence classes of expressions with a computable equality relation. The former are called **domain types** and the latter **symbolic types**. Each domain type may be, in general, represented by many symbolic types, but each symbolic type represents exactly one domain type. The reader is referred to [Blikle 87b] for an example of the use of symbolic- versus domain types in the semantics of Pascal, and to [Bednarczyk et al. 90] for a more general treatment of that problem.

There are two different strategies of constructing a mathematical semantics of a programming language with types. One - which is probably most common today - consists of assuming that the domain types appear explicitly in the semantics,

i.e. are assigned to variables, whereas the symbolic types are implicit. In that case, whenever we compare the types of two variables we check if a certain equivalence relation between the definitions of these types is satisfied. That strategy is most frequently associated with the technique of partitioning a semantics into a static semantics and a dynamic semantics, which is typical for VDM (see e.g. [Bjorner, Jones 82]). The corresponding semantics is then in general not denotational (cf. Sec. 6).

The second strategy is dual to the former and consists of assuming that symbolic types are explicit in the semantics, whereas domain types are implicit. In that case variables are typed by symbolic types and whenever we compare the types of two variables we compare the corresponding symbolic types rather than their definitions. This makes our semantics denotational. In addition to the definition of semantics we define in that case a function

$D : \text{SymbolicType} \rightarrow \text{DomainType}$

which tells the user what the symbolic types stand for. Of course, our language should have the following adequacy property:

"if in the execution of a program a variable v has been declared to be of a type syntype , then everywhere in the scope of that declaration the value of v belongs to $D.\text{syntype}$ ".

In general SymbolicType and DomainType constitute two algebras over the same signature. That signature usually contains several sorts which correspond to different classes of types, such as e.g. scalar types, record types, file types, etc., plus a sort possessing the same carrier Bool in both algebras. The latter sort is needed in order to define an equivalence relation, a subtype relation, etc. between types. Now, if we forget about the boolean sort, then D is usually a homomorphism, hence it may be regarded as a denotational semantics of symbolic types. However, in the full algebra of types D is not a homomorphism since two non-equal symbolic types may denote the same domain type. The (partial) non-compositionality of D is inherent to the difference between symbolic- and domain types and therefore it cannot be avoided.

The definition of D may be regarded as a non-denotational supplement of the main definition of semantics. This fact is not very harmful since it neither affects the feasibility of structured programming in the language nor the

construction of the corresponding logic.

Another typical situation where we may wish to slightly relax the principle of denotationality corresponds to the case where we modify an existing syntax by introducing some notational conventions. For instance, we may wish to allow for the optionality of parentheses in expressions while introducing some priority rules for operators. In that case the mathematical model of our system is usually described by the following diagram:

$$\begin{array}{ccc} \text{SynE} & \xrightarrow{\quad} & \text{Syn} & \xrightarrow{\quad} & \text{Den} \\ & P & & D & \end{array}$$

where SynE denotes the extended syntax and P is a preprocessing function. In general P is not a homomorphism in the strict sense of the word since SynE may have a different signature than Syn. However, P is usually a homomorphism in a generalized sense, namely a homomorphism over a morphism of signatures. We shall not go here into any technical details, but explain our remark on a simple example. Consider two syntaxes of arithmetic expressions described by the following equational grammars (cf.[Blikle 72]):

SynE	Syn
ExpE = Cpn Cpn+ExpE	Exp = {x}
Cpn = Fac Fac*Cpn	(Exp+Exp)
Fac = {x} (ExpE)	(Exp*Exp)

Each of these grammars defines unambiguously an algebra of syntax (cf. [Blikle 89b]). The former has three carriers: ExpE - extended expressions, Cpn - components and Fac - factors. The latter has only one carrier, Exp - expressions. Notice that SynE imposes a priority of * over +. Now, the preprocessing P is a many-sorted function that maps all the three carriers of SynE into the unique carrier of Syn. Formally, it is represented by three functions:

$$\begin{array}{l} E : \text{ExpE} \rightarrow \text{Exp} \\ C : \text{Cpn} \rightarrow \text{Exp} \\ F : \text{Fac} \rightarrow \text{Exp} \end{array}$$

defined by the following equations:

E.[cpn] = C.[cpn]
E.[cpn+expe] = (C.[cpn]+E.[expe])
C.[fac] = f.[fac]
C.[fac*cpn] = (F.[fac]*C.[cpn])
F.[x] = x
F.([expe]) = E.[expe]

Observe that although $P=\langle E,C,F \rangle$ is not a homomorphism in a strict sense, it certainly has a compositional character and in fact is not very far from a usual homomorphism.

The moral of our two stories is that if the major mechanisms of a system are described within a denotational model, then some "peripheral" information may be given in a not strictly compositional way. Of course, what is peripheral and what is not is an informal question and therefore in all such cases the designer has to rely on his/her own professional experience and common sense understanding. Little non-denotationalities are not too harmful, but it is clear that a too rich pre- or post-processing of the "main semantics" may completely destroy the compositional effect of the latter.

10. COPY-RULE SEMANTICS

A copy-rule semantics has been known for many years as a technique for providing a mathematical semantics of a typeless lambda-calculus in which a function may take itself as an argument. In applications this allows one to formalize such programming mechanisms as e.g. Algol-60 procedural parameters or Lisp dynamic recursion. The idea of copy rule is rather simple (cf. [Landin 64]), and corresponds closely to the way in which the original informal semantics of Algol-60 was described. A procedure declaration assigns the text of the procedure body - rather than the corresponding state-transition function - to the procedure name in the environment. At the call time this text is retrieved, and its denotation is applied to the current state.

Copy-rule semantics is not denotational. A denotational semantics of lambda-calculus may be defined by using Scott's model of reflexive domains [Scott 72] or one of its later versions, e.g. information systems [Scott 82].

Although mathematically very elegant, these models are not simple and therefore - at least in the opinion of some authors - not very convenient in applications. Consequently, the majority of software specification systems, such as e.g. BSI/VDM [Larsen et al. 89], MetaSoft [Blikle 87b] or RAISE [Nielsen et al. 87] are based on set-theoretic domains rather than on Scott's models. In all such systems one cannot define a denotational semantics of Algol-60 or of Lisp, but one can give them a copy-rule semantics.

Although the use of self-applicable functions in software systems is certainly not recommendable, and the problem of giving a mathematical semantics to Algol-60 or Lisp is today not very urgent, it may be of some interest to know what is the price of using a copy-rule semantics. As we shall try to argue below, that price - measured by the loss of denotationality and abstraction - does not seem very high, especially if the only alternative are reflexive domains.

Let us analyze an example of a programming language with self-applicable procedures, i.e. a language in which every procedure may take any other procedure - even itself - as an actual parameter. To simplify our example we assume that every procedure has exactly one parameter which is always a procedure, and that all variables in a procedure body are global. The syntax of our language is the following:

ide : Ide = a b ... z	identifiers
exp : Exp = Ide (Exp+Exp) ...	expressions
dec : Dec = proc Ide(Ide)=Com Dec;Dec	declarations
com : Com = Ide:=Exp call Ide(Ide) Com;Com ...	commands
pro : Pro = begin Dec;Com end Pro;Pro	programs

States in our language are triples consisting of an environment, used for storing procedures, a store, for storing values and a message which is either an OK message or an error message:

sta : State	= Environment x Store x Message
env : Environment	= Ide \xrightarrow{m} Procedure
sto : Store	= Ide \xrightarrow{m} Value
mes : Message	= {OK, ERROR}
prc : Procedure	= Ide x Com
val : Value	= Integer ...

Observe that procedures are pairs consisting of an identifier (formal parameter) and a command (body). The functions of semantics have the following signatures:

```
I : Ide → Ide
E : Exp → Store → Value
D : Dec → Environment → Environment
C : Com → State → Store
P : Pro → Store → Store
```

Below we analyze only the interesting semantic clauses i.e. the clauses for declarations, for calls and for single-block programs. All others have the usual form.

```
D.[proc idepn(idefp)=com].env = env[⟨idefp,com⟩/idepn]
D.[dec1;dec2].env = D.[dec2].(D.[dec1].env)
```

A procedure declaration assigns the corresponding procedure, i.e. both the formal parameter and the body, to the procedure name in the current environment. Semicolon is interpreted in the usual way.

```
C.[call idepn(ideap)].⟨env,sto,mes⟩ =
  mes=ERROR → ⟨env,sto,ERROR⟩,
  not idepn ∈ dom.env → ⟨env,sto,ERROR⟩
  not ideap ∈ dom.env → ⟨env,sto,ERROR⟩
  let ⟨idefp,com⟩=env.idepn in
  let prc=env.ideap in
  TRUE → C.[com].⟨env[prc/idefp],sto,mes⟩
```

After all necessary error checks a procedure call applies the denotation C.[com] of the corresponding procedure body to a state in which the current environment has been modified by assigning the value of the actual parameter to the formal parameter. We recall that all identifiers in the procedure body are global and therefore there is no need to rename them.

```
P.[begin dec:com end].sto =
  let env=D.[dec].[] in
  C.[com].⟨env,sto,OK⟩
```

The execution of a program consists of three steps. First an environment is created by applying the declaration part of the program to an empty environment. Then a state is created by combining that environment with the current store and the OK message. Finally, the denotation of the command part of the program is applied to that state.

As is not difficult to see, our semantics has the following two properties:

(1) E and C are not sufficiently informative w.r.t. D, since the denotation of a declaration depends on the syntax rather than on the denotations of its components.

(2) D is not sufficiently abstract w.r.t. P, because the denotation of a program depends on the denotation of a procedure body rather than on that procedure body itself.

By (1), our semantics is not denotational, and, by (2), it is not fully abstract. If we assume that the implicit part of the language follows the usual style of e.g. Algol-60, then (1) and (2) are the only violations of denotationality and full abstraction, respectively. What does that mean for the user of the language?

From a formal viewpoint, our language does not provide a fully adequate framework for structured programming (cf. Sec.5). For instance, it is impossible to define a programmer's task of writing a declaration of a procedure $\langle \text{ide}, \text{com} \rangle$ by giving only ide and $\text{C}[\text{com}]$ and it is not worthwhile to do the same with a program begin dec:com end by giving $\text{D}[\text{dec}]$ and $\text{C}[\text{com}]$. However a "local structured programming" is possible. A project coordinator may split the task of writing a compound command with a given denotation into the subtasks of writing subcommands (and/or subexpressions) with appropriate denotations or to split the task of writing a compound program into the tasks of writing a sequence of subprograms. Moreover, he/she may split the task of writing a program begin dec:com end into a task of writing dec with a given denotation $\text{C}[\text{com}_p]$ of the procedure body and a given denotation $\text{C}[\text{com}]$. Both these tasks may be further structurally decomposed.

As regards the logic for our language the related problems and their solutions are similar to the former. We can infer the properties of a compound command from the properties of its component expressions and commands, and the

properties of a compound program from the properties of its subprograms. We cannot infer the properties of a (denotation of a) declaration from the properties of its components, but we can infer the properties of a program from the properties of the components of its declaration. Readers interested in a detailed discussion of Hoare-like logics for different copy-rule schemes are referred to a very elegant paper [Olderog 81].

11. AN OPERATIONAL DEFINITION OF A DENOTATIONAL SEMANTICS

As it was mentioned in the introduction, a denotational semantics may have a non-denotational definition. In this section we briefly discuss a typical example of such a definition, written in the style of structured operational semantics (SOS) introduced in [Plotkin 81].

Consider our little programming language and its semantics as defined in Sec.4. We shall discuss a SOS-definition of the semantics of commands C . For that sake we introduce a few auxiliary concepts. Let:

```
pcom : Pseudocommand = Com | {nil}
conf : Configuration = Pseudocommand x State
      Terminal-conf = {nil} | State
```

A configuration may be interpreted as a global memory state of a von Neuman machine, in which we store both, data and programs. By a transition relation between configurations we mean the least transitive and reflexive relation

$$\longrightarrow \subseteq \text{Configuration} \times \text{Configuration}$$

with three following properties:

$$\langle \text{ide} := \text{exp}, \text{sta} \rangle \longrightarrow \langle \text{nil}, \text{sta}[(E.[\text{exp}].\text{sta})/\text{ide}] \rangle \quad (11.1)$$
$$\langle \text{com}_1, \text{sta} \rangle \longrightarrow \langle \text{com}_1', \text{sta}' \rangle$$

$$\langle \text{com}_1; \text{com}_2, \text{sta} \rangle \longrightarrow \langle \text{com}_1'; \text{com}_2, \text{sta}' \rangle \quad (11.2)$$

$$\langle com_1, sta \rangle \longrightarrow \langle nil, sta' \rangle$$

(11.3)

$$\langle com_1; com_2, sta \rangle \longrightarrow \langle com_2, sta' \rangle$$

The latter two formulas should be, of course, read as top-down implications. The transition relation describes a way in which our machine transforms a global state when executing a program. If $conf_1 \longrightarrow conf_2$ holds, then we say that $conf_1$ reduces to $conf_2$. A sequence of configurations:

$$conf_1 \longrightarrow conf_2 \longrightarrow \dots \longrightarrow conf_n$$

is called a **computation**. If $conf_n$ is a terminal configuration, then the above sequence is called a **terminating computation**. Below we show a simple example of a terminating computation:

pseudocommand	state
$\langle x:=1; y:=2; z:=x+y, sta \rangle$	
$\langle y:=2; z:=x+y, sta[1.0/x] \rangle$	
$\langle z:=x+y, sta[1.0/x, 2.0/y] \rangle$	
$\langle nil, sta[1.0/x, 2.0/y, 3.0/z] \rangle$	

Properties (11.1) and (11.3) imply - by structural induction - that for each non-terminal configuration $\langle com, sta \rangle$ there is exactly one terminal configuration $\langle nil, sta' \rangle$ such that:

$$\langle com, sta \rangle \longrightarrow \langle nil, sta' \rangle$$

This proves the existence of a function of semantics:

$$C : Com \rightarrow State \rightarrow State$$

$$C.[com].sta = sta' \text{ iff } \langle com, sta \rangle \longrightarrow \langle nil, sta' \rangle$$

Now, (11.1) and (11.3) imply (respectively) two following properties of C:

$$C.[ide:=exp].sta = sta[(E.[exp].sta)/ide]$$

$$C.[com_1; com_2].sta = C.[com_2].(C.[com_1].sta)$$

This means that our function is a component of a homomorphism (I, E, C) between Syn and Den as defined in Sec.4.

Some authors prefer the SOS style as more appealing to the intuition than denotational equations. When defining a denotational semantics in the SOS style one has to remember, however, about two proof obligations:

- (1) that the corresponding function of semantics exists and is total,
- (2) that the function of semantics has the compositionality property.

Of course, in practical situations both these proofs may be far from trivial.

12. FINAL REMARKS

In the Introduction we have formulated a claim that the semantics of a software system should be denotational, unless we agree to give up structured programming and structured-induction correctness-proofs. We have discussed some arguments that denotationality does guarantee these possibilities, and an example indicating that this may not be the case if denotationality is not insured. At the same time, however, we have shown that the principle of denotationality may be always trivially insured by adding syntax to denotations. Doesn't that mean that denotationality is merely a property of the definition of a system rather than - as we have claimed earlier - a property of the system itself?

Each software system is a tool for constructing of some applications (programs), and therefore an adequate mathematical model of a system should provide a ground for a convenient description and validation of these applications. Whether a model is sufficiently adequate depends, generally speaking, on two factors: on the choice of denotations and on the compositionality of semantics. Compositionality is important but it is not a goal in itself. It is worth of a care only if denotations adequately express the behavior of programs and of its components. And once we fix denotations, the compositionality of semantics becomes a property of a system rather than of its definition.

Our requirement of denotationality should be understood as a pragmatic rule. By choosing some "clever" denotations we may be able to construct an elegant algebraic model for a mechanism which is neither elegant nor algebraic. We have

seen a simple example of such a situation in Sec.8. Some more interesting examples are related to the well-known technique of continuations. Using that technique one may construct a denotational model of a very unstructured programming language with most anarchic goto's. Formally, such a language guarantees the feasibility of structured programming but this is on the price that the denotations of programs are even more difficult to read, decompose and analyze than unstructured flowcharts. By the use of continuations the principle of denotationality is "cheated" in a very subtle way by putting all the conceptual mess of a language deep into denotations.

One of the advantages of a formal mathematical semantics is the possibility of discovering awkward mechanisms of a software system at the stage of its design rather than at the stage of its use. It has been rather generally agreed that a complicated definition of a mechanism should be regarded as a warning that the use of such a mechanism may be not easy. However, the simplicity of a definition - especially if this is a local simplicity, as e.g. in the case of a continuation-style definition of goto's - does not guarantee a sufficient simplicity of applications. The latter may be adequately estimated by analyzing the corresponding proof rule. Therefore the construction of a program-correctness logic should be regarded as an inherent part of the process of system design. Whenever we cheat on the subject of denotationality too much, the price to be paid is the complexity of proof-rules.

ACKNOWLEDGMENTS

The author expresses his thanks to Marek Gondzio, Beata Konikowska, Jacek Leszczykowski, Peter Mosses and Wieslaw Pawlowski for many relevant comments about the early drafts of this paper. Special thanks are addressed to Andrzej Tarlecki who suggested the inclusion of Proposition 7.1.

STOUGHTON A. (1988) **Fully Abstract Models of Programming Languages**, Res.
Notes in Theoretical Comp. Sci., Pitman and John Wiley & Sons, Inc.

THATCHER J.W., WAGNER E.G., WRIGHT J.B. (1978) **Notes on algebraic fundamentals
of theoretical computer science**, Proc. of The 3rd Advanced Course on
Foundations of Computer Science, Amsterdam, 21 August - 1 September, 1978