

Chapter 1

Fundamental Concepts

This book is about the mathematical foundations of programming, with a special attention on computing with infinite objects. How can mathematics help in programming? There are several important roles that mathematical modelling and reasoning plays in computer science.

First of all, mathematics helps us *understand the meaning of programs*. We can use the power of abstraction to interpret data structures and computational processes as mathematical entities. This view allows us to forget about some of the complicated practical details of implementation and to concentrate on the high-level concepts.

A second role of mathematics is in *specification*: stating clearly and precisely what a program must do. Usually the purpose of a program is expressed in common language with all its ambiguities and imprecisions. This opens the risk of the final software being inadequate to the requirements. But if these were not expressed in a clear and precise way, it may be very difficult to point out where the code went wrong. Logical mathematical specification allows us to state with absolute precision what a program must do. We can use the specification as a guide in development and as a test to verify the success of the final software.

A third role for mathematics is in *reasoning* about programs. Using symbolic logic we can prove properties of them. First of all, we can verify, with formal methods, that the software satisfies the specifications developed according to the second point. Further, we can also investigate high-level characteristics of the software and its behaviour and even its computational complexity. This can be done with the help of computers: *proof assistants* are systems that allow a user to write both specifications and programs and verify correctness properties interactively with the systems. If the process is successful it guarantees that the result are certain, eliminating the possibility of any human error.

Finally, the fourth role that mathematics plays in computer science is that of *abstract interpretation*. We see data structures as mathematical objects, to clarify their nature. Data is implemented in specific ways in various computer architecture. By abstracting from the specific details of the implementation, we can concentrate on their essential nature, which doesn't depend on the concrete

realization. This nature can be described as part of logical/mathematical reality.

Syntax and Semantics

Two important sides of the use of mathematical modelling in computer science are *syntax* and *semantics*.

Syntax describes the way we write expressions, terms, formulas, programs. It tells us how a programming language is defined precisely: what symbols we use; what are the reserved words and the identifiers; how we combine symbols to form correct expressions and instructions; how we combine instructions to create programs.

Semantics consists in giving the meaning of expressions, terms, formulas, programs. Without it a programming language would consist only of a set of meaningless strings of symbols. Semantics specifies what the symbols stand for and what happens when we combine them according to the rules of the syntax. There are two fundamental ways of giving meaning to a language.

Denotational Semantics interprets expressions and programs as objects in some mathematical space. The meaning of a string of symbols is some abstract entity. The advantage of this approach is that we can use powerful mathematical methods to derive properties of the programs.

Operational Semantics interprets expressions and programs as computations. The meaning of a program consists in what happens when we run it. Operational semantics doesn't move away from syntax to a mathematical reality. Instead it adds a dynamic element to the syntax, stating how syntactic entities are computed.

Arithmetic Expressions

To illustrate the basic functions of syntax and semantics, we use a very simple language of arithmetic expressions. It is extremely limited and not much interesting can be done with it. But it is sufficient to clarify the main notions. Later we'll move to much richer languages.

This language contains natural numbers, 0, 1, 2, et cetera, and the Boolean expressions **true** and **false**. We can combine them with simple operations: successor and predecessor for numbers, a test of whether an number is 0, a conditional (if-then-else) constructor.

We give the syntax in two different but equivalent ways: first in Backus-Naur Form (BNF), then by derivation rules. The BNF is simple and compact, while derivation rules are a bit lengthier and involved. However, there are richer languages that cannot be specified by a BNF, so derivation rules will be necessary. They are more flexible and powerful than BNF. For this reasons, we start introducing them immediately in this simple example to make their use clear in view of future more complex languages.

Backus-Naur Form for Arithmetic Expressions:

$$\begin{aligned} \text{Expr} ::= & \text{ true } \mid \text{ false } \mid \text{ zero } \mid \text{ succ Expr } \mid \text{ pred Expr} \\ & \mid \text{ isZero Expr } \mid \text{ if Expr then Expr else Expr} \end{aligned}$$

The meaning of this definition is that *Expr* represents the set of all arithmetic expressions. The left-hand side of the definition gives all the correct ways of defining an expression. There are three constants, **true**, **false** and **zero**, and some recursive constructors that allow us to make new expressions by inserting previously build expressions in the places where *Expr* occurs.

Here are some examples of correct expressions:

```

zero
true
succ zero
succ true  (correct expressions, though meaning unclear)
isZero (succ zero)
if (isZero (pred false)) then zero else (succ (isZero true))
    
```

Not all expressions seem to make sense. Some clearly denote a natural number, some other denote a Boolean value, but if we mix them in the wrong order (for example (**pred true**)) we have a meaningless expression. We'll see at the moment of giving the semantics what we can do with them.

Let us now describe the same language using derivation rules rather than the BNF. A derivation rule is a horizontal line with some expressions written above it and an expression written under it. For example:

$$\frac{e_1 \in \text{Expr} \quad e_2 \in \text{Expr} \quad e_3 \in \text{Expr}}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \in \text{Expr}}$$

The rule has three assumptions and a conclusion. It states that if we have already constructed expressions e_1 , e_2 and e_3 , then we can construct the new expression (**if e_1 then e_2 else e_3**).

There are rules for the base elements as well. They don't need any assumption, so the space above the line is empty:

$$\frac{}{\text{zero} \in \text{Expr}} \quad \frac{}{\text{true} \in \text{Expr}} \quad \frac{}{\text{false} \in \text{Expr}}$$

The rules for the unitary operators have just one assumption:

$$\frac{e \in \text{Expr}}{\text{succ } e \in \text{Expr}} \quad \frac{e \in \text{Expr}}{\text{pred } e \in \text{Expr}} \quad \frac{e \in \text{Expr}}{\text{isZero } e \in \text{Expr}}$$

To construct any element of *Expr* we must derive it completely, starting with his subterms. We must build a *derivation tree* in which all assumptions are

resolved by using further derivation rules. Here is an example:

$$\frac{\frac{\frac{\text{zero} \in Expr}{\text{succ zero} \in Expr}}{\text{isZero}(\text{succ zero}) \in Expr} \quad \frac{\text{false} \in Expr}{\text{succ false} \in Expr} \quad \frac{\text{zero} \in Expr}{\text{pred zero} \in Expr}}{\text{if}(\text{isZero}(\text{succ zero})) \text{ then } (\text{succ false}) \text{ else } (\text{pred zero}) \in Expr}$$

The derivation has the form of a tree with subterms as nodes and base cases (rules with no assumptions) as leaves. To avoid confusion we use parentheses around subexpressions.

Defining a language using derivation rules is long and boring. BNF is much simpler and more compact. However, derivation rules are useful in more complex languages and for other purposes than just defining syntax. When we go beyond a toy language like *Expr* or a simple imperative programming language, BNF is not flexible enough. We will soon see examples: several system of (dependently) typed λ -calculus can only be defined using derivation rules.

Here are the main areas in which use of derivation rules is essential.

Syntax: Definition of a language. BNF only supports languages with a finite number of *syntactic categories*. We had only one category in our example, the set *Expr*. We could change it a bit by having a category for numerical expressions and a category for Boolean expressions. This would also solve the problem of meaningless terms. However, when we introduce *types*, each of them needs to be a separate category. We cannot give a BNF for all of them. We will see instead how derivation rules make it quite easy to define a typed language.

Manipulation of Expressions: We may want to specify how expressions can be modified in various ways. One example is the definition of substitution. In general defining functions on terms, when they are not obviously recursive, may require giving them as relations generated by derivation rules.

Operational Semantics: Defining the computation rules of a language requires specifying some reduction relation. This is in general defined by derivation rules. Mostly the simple one-step derivation can be given directly, without need of assumption. But there are also *structural rules* that allow us to do a reduction in context: these usually require assumptions.

Logical Systems: Derivation rules were originally introduced in mathematical logic as a way of expressing symbolic reasoning. They are important in defining the specific logical systems that we use to reason about programs.

The general form of a derivation rule is

$$\frac{A_1 \quad A_2 \quad \dots \quad A_n}{B}$$

where A_1, A_2, A_n are the *assumptions* and B is the *conclusion*. Assumptions and conclusions are *judgments*, some sentences stating a fact. In our example a judgment is something in the form $e \in Expr$ states that e is a correct arithmetic expression.

The rule says that if we have already derived A_1, A_2, A_n , then we can derive B . All these judgments can contain *meta-variables*, identifiers that can be replaced by concrete expressions. In our examples, e, e_1, e_2, e_3 are meta-variables. In concrete derivations, the meta-variables are instantiated by expressions.

A correct judgment is a statement that has been derived by giving a full derivation tree whose leaves are instances of rules without assumptions.

Semantics

We have defined the syntax of arithmetic expressions. That is, we specified what strings of symbols are correct terms of the language $Expr$. Now we want to define the semantics. That is, we describe precisely what the expression means. There are two basic kinds of semantics. According to *denotational* semantics, the meaning of an expression is an object in some mathematical space. According to *operational* semantics, the meaning of an expression is the value that it computes when we execute it.

Denotational Semantics

Arithmetic expressions describe both integer numbers from zero and Boolean values. Therefore we choose two mathematical domains to interpret them, the set $\mathbb{N} = \{0, 1, 3, \dots\}$ of natural numbers and the set $\mathbb{B} = \{\text{True}, \text{False}\}$ of truth values.

We interpret every term of $Expr$ as an element of either of these sets. For example the term $(\text{succ}(\text{succ}(\text{succ zero})))$ denotes the number 3 and the term $(\text{isZero}(\text{succ zero}))$ denotes the truth value `False`. However, some terms have a meaningless combination of the operations for both domains, for example (succ true) . We leave the denotation of these meaningless terms undefined.

Formally, we use *semantic brackets* $\llbracket - \rrbracket$ around an expression to indicate its denotation, for example $\llbracket \text{succ}(\text{succ}(\text{succ zero})) \rrbracket = 3$ and $\llbracket \text{isZero}(\text{succ zero}) \rrbracket = \text{False}$. For meaningless terms we use the *bottom* symbol \perp , so $\llbracket \text{succ true} \rrbracket = \perp$ simply means that `succ true` is meaningless. The symbol \perp itself is not any mathematical object. (However, in a more advanced kind of denotational semantics, called *domain theory*, we define mathematical spaces in which \perp is an actual element.)

The interpretation is defined by *structural recursion* on the syntax: we define the meaning of an expression in terms of the meaning of its components. For every rule that generates terms, we assume that we already know the interpretation of the assumptions and we specify how to combine them to obtain the

interpretation of the conclusion.

$$\begin{aligned}
 \llbracket \text{zero} \rrbracket &= 0 \\
 \llbracket \text{false} \rrbracket &= \text{False} \\
 \llbracket \text{true} \rrbracket &= \text{True} \\
 \llbracket \text{succ } e \rrbracket &= \llbracket e \rrbracket + 1 \quad \text{if } \llbracket e \rrbracket \in \mathbb{N} \\
 \llbracket \text{pred } e \rrbracket &= \begin{cases} 0 & \text{if } \llbracket e \rrbracket = 0 \\ n - 1 & \text{if } \llbracket e \rrbracket = n > 0 \end{cases} \\
 \llbracket \text{isZero } e \rrbracket &= \begin{cases} \text{True} & \text{if } \llbracket e \rrbracket = 0 \\ \text{False} & \text{if } \llbracket e \rrbracket = n > 0 \end{cases} \\
 \llbracket \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rrbracket &= \begin{cases} \llbracket e_2 \rrbracket & \text{if } \llbracket e_1 \rrbracket = \text{True} \\ \llbracket e_3 \rrbracket & \text{if } \llbracket e_1 \rrbracket = \text{False} \end{cases}
 \end{aligned}$$

The expressions that don't fall into these cases are left undefined, that is $\llbracket e \rrbracket = \perp$. Also, the same side condition that we imposed for `succ`, that $\llbracket e \rrbracket \in \mathbb{N}$, must hold for `pred` and `isZero`; if not, the expression is meaningless. Similarly, in the last case, $\llbracket e \rrbracket$ must be in \mathbb{B} . We also impose that $\llbracket e_2 \rrbracket$ and $\llbracket e_3 \rrbracket$ are in the same set, either both in \mathbb{N} or both in \mathbb{B} .

Here is the calculation of the denotation of a complex expression:

$$\llbracket \text{if (if (isZero (pred zero)) then false else true) then (succ (pred zero)) else zero} \rrbracket = ?$$

Since it is a bit involved, we start calculating the denotation of its component parts first.

$$\begin{aligned}
 \llbracket \text{pred zero} \rrbracket &= 0 \quad \text{because } \llbracket \text{zero} \rrbracket = 0, \\
 \text{therefore } \llbracket \text{isZero (pred zero)} \rrbracket &= \text{True};
 \end{aligned}$$

$$\begin{aligned}
 &\llbracket \text{if (isZero (pred zero)) then false else true} \rrbracket \\
 &= \llbracket \text{false} \rrbracket \quad \text{because } \llbracket \text{isZero (pred zero)} \rrbracket = \text{True} \\
 &= \text{False};
 \end{aligned}$$

$$\begin{aligned}
 &\llbracket \text{if (if (isZero (pred zero)) then false else true) then (succ (pred zero)) else zero} \rrbracket \\
 &= \llbracket \text{zero} \rrbracket \quad \text{because } \llbracket \text{if (isZero (pred zero)) then false else true} \rrbracket = \text{False} \\
 &= 0.
 \end{aligned}$$

Operational Semantics

A different way to explain the meaning of expressions consists in specifying their computations. A complex term can be simplified by applying some *reduction* rules. Repeatedly reducing a term may eventually result in a *value*. This reduction of expressions to values is the *operational semantics* of a language.

We denote the reduction relation by a squiggly arrow \rightsquigarrow . When we write $e_1 \rightsquigarrow e_2$, we mean that we can do one step of simplification in e_1 and obtain e_2 . For example `pred (succ zero) \rightsquigarrow zero` or `isZero (succ zero) \rightsquigarrow true`. Sometimes several steps of reduction are possible. We use \rightsquigarrow^* to denote several steps of \rightsquigarrow .

There are two kinds of rules for \rightsquigarrow : the actual simplification rules, that specify how terms of a certain form reduce, and the *structural rules*, that say

that reduction can be done in context, that is, we can reduce a subterm inside a bigger term.

The simplification rules for *Expr* are:

if true then e_2 else $e_3 \rightsquigarrow e_2$
 if false then e_2 else $e_3 \rightsquigarrow e_3$
 pred (succ e) $\rightsquigarrow e$
 pred zero \rightsquigarrow zero
 isZero zero \rightsquigarrow true
 isZero (succ e) \rightsquigarrow false.

The structural rules are formulated as derivation rules. (The simplification rules can also be given as derivation rules with no assumptions.) For each of the expression constructors, we state that if one of its argument contains a reducible term, then we can perform the reduction inside the constructor. There are many structural rules, one for each argument of each constructor. The general shape is the same for all of them. Here are just a few of them to give the idea of the general pattern:

$$\frac{e \rightsquigarrow e'}{\text{succ } e \rightsquigarrow \text{succ } e'}$$

This rule states that if a term e reduces to e' , then we can perform this reduction step in the argument of the \succ constructor. For example $\text{succ (isZero zero)} \rightsquigarrow \text{succ true}$. Never mind that this term is itself meaningless. Reductions are purely syntactic operations, they don't worry whether the expressions are sensible.

Since the if – then – else – constructor has three argument, we have three structural rules for reduction:

$$\frac{e_1 \rightsquigarrow e'_1}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rightsquigarrow \text{if } e'_1 \text{ then } e_2 \text{ else } e_3}$$

$$\frac{e_2 \rightsquigarrow e'_2}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rightsquigarrow \text{if } e_1 \text{ then } e'_2 \text{ else } e_3}$$

$$\frac{e_3 \rightsquigarrow e'_3}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rightsquigarrow \text{if } e_1 \text{ then } e_2 \text{ else } e'_3}$$

The subterm of an expressions to which the reduction rule is applied are called the *redex*, the result of reducing it is called the *reductum*. For example, in the reduction

if (isZero zero) then (succ zero) else (pred zero)
 \rightsquigarrow if true then (succ zero) else (pred zero)

the redex is (isZero zero) and the reductum is true.

There may be several redexes in an expression. We have the choice of which one to reduce. For example, the previous expressions can be reduced with this alternative step (redex underlined):

if (isZero zero) then (succ zero) else (pred zero)
 \rightsquigarrow if (isZero zero) then (succ zero) else zero.

Now we defined the one step reduction relation by rules. The many-step reduction relation \rightsquigarrow^* consists in applying several one step reduction consecutively, in any of the allowed order. For example, we can reduce the example term that we used earlier to illustrate denotational semantics (at each step, the active redex is underlined):

```

if (if (isZero (pred zero)) then false else true) then (succ (pred zero)) else zero
 $\rightsquigarrow$  if (if (isZero zero) then false else true) then (succ (pred zero)) else zero
 $\rightsquigarrow$  if (if (isZero zero) then false else true) then (succ zero) else zero
 $\rightsquigarrow$  if (if true then false else true) then else (succ zero) zero
 $\rightsquigarrow$  if false then (succ zero) else zero
 $\rightsquigarrow$  zero
    
```

After performing reduction steps as far as possible, we obtain a term that doesn't contain any redexes any more.

Definition 1 *A term is in normal form if it contains no redexes.*

A value is a term belonging to one of the following languages, respectively of numerals and of Boolean values:

$$\begin{aligned}
 nv &::= \text{zero} \mid \text{succ } nv \\
 bv &::= \text{true} \mid \text{false}
 \end{aligned}$$

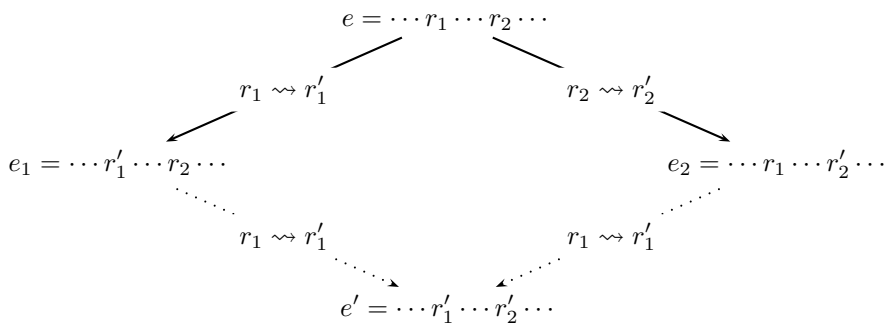
For a natural number n , we use the notation \bar{n} for the corresponding numeral: $\bar{0} = \text{zero}$, $\bar{1} = \text{succ zero}$, $\bar{2} = \text{succ (succ zero)}$ and so on.

All values are normal forms, but there are also normal forms that are not values, for example $\text{isZero (pred true)}$. This happens because these terms are meaningless. When we define languages with variables, we will also have normal forms containing free variables, which are not values because of the indeterminacy of the values of the variables.

We have seen that a term can be reduced in several ways. If it contains more than one redex, we can choose which one to reduce first. The final result of the computation, that is, the normal form that we obtain at the end, does not depend on the order of reduction.

Theorem 2 (Confluence) *Given an expression $e \in \text{Expr}$; if there exist two expressions $e_1, e_2 \in \text{Expr}$ such that $e \rightsquigarrow e_1$ and $e \rightsquigarrow e_2$, then there exist a common reduct $e' \in \text{Expr}$ such that $e_1 \rightsquigarrow^* e'$ and $e_2 \rightsquigarrow^* e'$.*

Proof. In the case that the redexes that were reduced in e_1 and e_2 are independent, that is, they occur in separated places inside e , then it is enough to perform the other reduction in each of the two:



But it can also happen that the redexes are contained inside each other. For example in the term

if true then (pred zero) else zero

we have two redexes inside each other. In this particular example it is still true that after reducing one, the other will still be present, so we can adopt the same strategy as for when they are separated.

However, it is also possible that the reduction of one redex makes the other disappear. This is what happens in the following redex:

e = if true then zero else (pred (succ zero)).

When we reduce each redex separately, we get:

$e_1 = \text{zero}, \quad e_2 = \text{if true then zero else zero.}$

We can still perform the reduction of the first redex in e_2 , obtaining zero, but the second redex has disappeared from e_1 , so there is nothing to do. Clearly, we still have a common reduct zero, and this will be the true in similar cases of overlapping redexes.

This is the reason we formulated the theorem using many-step reduction: $e_1 \rightsquigarrow^* e'$ and $e_2 \rightsquigarrow^* e'$. Specifically, one of the two many-step reductions could be empty. \square

While the proof of the confluence theorem is straightforward in this case, it becomes more complicated for richer languages. In particular, if the language has reduction rules that can duplicate some terms, then one reduction step may generate many copies of an existing redex. We will need more sophisticated proof techniques to show that confluence still holds.

If we keep reducing some redexes in a term, we will eventually get rid of them all and arrive at a normal form.

Theorem 3 (Normal Form) *For every expression $e \in Expr$ there exist a (unique) normal form $e \in Expr$ such that $e \rightsquigarrow^* t$.*

Proof. Notice that every reduction rule makes the term smaller. Therefore it is necessary that eventually the reduction process will terminate. This is a proof by induction on the length of the term. \square

In more complex systems, reduction steps do not necessarily decrease the size of the term. In some cases they can increase it in an explosive way. Therefore a simple proof by induction on the length of the term will not work any more and we will have to invent more sophisticated proof techniques.

Soundness and Completeness

We have defined two radically different ways to give meaning to expressions: denotational semantics interpret them as elements of a mathematical space, operational semantics interpret them as computation processes that lead to a normal form.

We want to show that these two interpretation of our language agree with each other in some sense. We have seen an example with the term

if (if (isZero (pred zero)) then false else true) then (succ (pred zero)) else zero

We calculated its denotation to be the number 0. We computed its normal form to be the numeral zero. This is good news: the two interpretations agree.

We want to show that this always happens. First of all, the reduction process doesn't change the denotation of a term.

Theorem 4 (Soundness) *For every expression $e \in Expr$, if $e \rightsquigarrow e'$ for some other $e' \in Expr$, then $\llbracket e \rrbracket = \llbracket e' \rrbracket$. (For meaningless terms this means that they are both \perp .)*

Proof. **TO DO** [By induction on the structure of the term.] \square

If we reduce an expression all the way to normal form, the denotation will be preserved, so if the normal form is a value, the denotation must be the corresponding object in the mathematical space.

Corollary 5 *If $e \rightsquigarrow^* v$ for some value v , then $\llbracket e \rrbracket = \llbracket v \rrbracket$.*

The vice versa is also true: the denotation of an expression (when it exists) will correspond exactly to a value to which the expression reduces.

Theorem 6 (Completeness) *For every expression $e \in Expr$, we have:*

- *If $\llbracket e \rrbracket = n$ for some $n \in \mathbb{N}$, then $e \rightsquigarrow^* \bar{n}$;*
- *If $\llbracket e \rrbracket = \text{True}$, then $e \rightsquigarrow^* \text{true}$;*
- *If $\llbracket e \rrbracket = \text{False}$, then $e \rightsquigarrow^* \text{false}$;*
- *If $\llbracket e \rrbracket = \perp$ then there is no value v such that $e \rightsquigarrow^* v$.*

Proof. TO DO [By induction on the structure of e and by definition of $\llbracket - \rrbracket$ and \rightsquigarrow .] \square

The combination of the soundness and completeness theorems tell us that there is complete agreement between the definitional and operational semantics. This also suggests a clever method to compute the normal form of a term (when it exists): First find the result of $\llbracket e \rrbracket$ and then take the syntactic value corresponding to it. We may succinctly write: $e \rightsquigarrow^* \llbracket e \rrbracket$ (with the additional definitions $\overline{\text{True}} = \text{true}$ and $\overline{\text{False}} = \text{false}$). This process is called *Normalization by Evaluation*.

Reduction Strategies

We have noticed that certain terms can contain many redexes and therefore the order of reduction steps is not unique. Although the confluence property ensures that any reduction sequence will produce the same normal form, there are issues of efficiency that make certain reduction strategies more convenient than others. Also, for other languages in which there is no guarantee that a normal form exists, the choice of reduction strategy may make the difference between a terminating computation and an infinite one.

Let's take as example the following term:

$$e = \text{if } (\underline{\text{isZero zero}}) \text{ then } (\underline{\text{pred (succ zero)}}) \text{ else } (\underline{\text{succ (pred zero)}}).$$

It contains three redexes, which we can choose to reduce in any possible order.

Two of the most important reduction strategies are called *Eager Evaluation* and *Lazy Evaluation*. They differ in the order in which they reduce the arguments of a function. In e there are three arguments, the *if* test condition, the *then* branch and the *else* branch.

In eager evaluation we reduce all arguments of a function to normal form before reducing the main expression:

$$e \rightsquigarrow^* \text{if true then zero else succ zero} \rightsquigarrow \text{zero}.$$

In lazy evaluation, we only reduce the arguments that are needed to make a computation step in the main expression:

$$e \rightsquigarrow \text{if true then (pred (succ zero)) else (succ (pred zero))} \rightsquigarrow \text{pred (succ zero)} \rightsquigarrow \text{zero}.$$

Notice that we saved one reduction step: we didn't have to reduce the *else* branch of the expression.