

Understanding and Designing Recursive Functions via Syntactic Rewriting

Tom VERHOEFF

*Mathematics and Computer Science, Eindhoven University of Technology
Groene Loper 5, 5612 AE, Eindhoven, Netherlands
e-mail: t.verhoeff@tue.nl*

Abstract. Recursion is considered a challenging programming technique by many students. There are two common approaches intended to help students understand recursion. One of them is based on the operational semantics of function execution involving a stack, where students trace the execution of a recursively defined function for some concrete arguments. The other approach is based on the axiomatic semantics involving inductive reasoning with the contract of the recursively defined function. The former approach is not so helpful when designing recursive functions, whereas the latter can be helpful (being a special case of divide and conquer) but contracts can be hard to discover.

In this article, I will show a third approach. It is based neither on an operational nor an axiomatic semantics. Rather, it involves a rewriting semantics using program transformation by substitution, thereby inlining function calls. We show that this approach not only may help understanding, but can also be used to design recursive functions.

Keywords: computer science, education, programming, recursion.

1. Introduction

I have written about recursion before (Verhoeff, 2018, 2021) and until recently I would not have thought to have something significant to add. But while preparing a programming Q&A session for first-year mathematics students, I was struck by a (for me) new idea, which is the topic of this article.

Verhoeff (2018) presents the two common approaches to understand recursion. As example, consider the following Python code for the recursively defined function `print_bit_strings` (we left it undocumented for now on purpose):

```
1 def print_bit_strings(n: int, s: str = "") -> None:
2     if n == 0:
3         print(s)
4     else:
```

```

5     for b in "01":
6         print_bit_strings(n - 1, s + b)

```

The call `print_bit_strings(2)` prints

```

00
01
10
11

```

One could wonder why there is a need for this extra parameter `s`, and why the recursive call on line 6 has argument `s + b` rather than `b + s`.

To understand this, students could use the *operational* approach, where they trace the execution of a specific call of the recursive function in question through multiple levels of subsequent recursive calls. Beginning programmers perceive this execution as magical and confusing, because there is a single definition of the given recursive function, but multiple calls are simultaneously executing that same piece of code independent of each other. That is, each invocation can be at a different location in that code, and the parameters and other local variables can have different values. The execution of a recursive function traverses an imaginary *dynamic call tree* (Verhoeff, 2018, §3.1), where each node corresponds to the execution of a call, which then gives rise to zero or more subsequent recursive calls. The active invocations form a *root path* in this tree, and the ‘instruction pointer’ and values of local variables are stored on a *stack*, that grows and shrinks as the root path.

For example, the call `print_bit_strings(2, "s")` gives rise to the call tree in Fig. 1. This approach helps in understanding how a stack machine can correctly execute a recursively defined function in an imperative programming language. But in my experience it is neither helpful for reasoning about (the correctness) of recursive function definitions nor for designing them.

Alternatively, there is the *axiomatic* approach (Verhoeff, 2018, §4), which requires a specification of the recursive function in terms of a *contract* consisting of a *precondition* and a *postcondition*, such that

- if the precondition is satisfied *before* the function call,
- then the postcondition is satisfied *after* the function call.




Fig. 1. Call tree for `print_bit_strings(2, "s")`, only showing parameters.

The docstring for `print_bit_strings(n, s)` could read:

```

7     """Print s + t for all strings t over "01" of length n,
8     in lexicographic order.
9
10    Assumption: n >= 0
11    """

```

That is, its precondition is $n \geq 0$ and its postcondition is: ‘for all strings t over “01” of length n , strings $s + t$ have been printed’. To reason about the call with parameters (n, s) , we assume as *induction hypothesis* that calls with parameters $(n_, s_)$ where $n_ < n$ work as specified by the contract. The design of function `print_bit_strings` can be argued as follows. The goal is to prove that lines 2–6 satisfy the contract, under the assumption that the recursive call on line 6 satisfies its contract; that is, it satisfies ‘if $n > 0$, then for all strings u of length $n - 1$ over “01”, it prints strings $s + b + u$ ’.

- If $n == 0$, then there is only one string t of length n , viz. the empty string. Observe that s extended with the empty string equals s (line 3). Thus, the contract is fulfilled by printing just s .
- If $n > 0$ then extensions of length n over “01” start with a single bit, say b , in “01”, followed by $n - 1$ more bits over “01”. Thus, by the induction hypothesis, the loop on lines 5–6 prints all required strings.

In this reasoning style, the induction hypothesis is sometimes referred to as the *recursive leap of faith* (Roberts, 1986; Rubio, 2018). I find this terminology unfortunate, because faith has nothing to do with it. There is already good terminology, viz. *induction hypothesis*. Trusting the compiler, runtime system, and hardware to execute recursive definitions faithfully, could be called a leap of faith. Students need to understand this implementation only once, e.g., by tracing an execution through the call tree and observing the role of the stack. But this is not needed (nor helpful) to understand specific recursive definitions, and certainly not for designing them.

2. Syntactic Rewriting

There is a third kind of semantics for programming languages, viz. based on *rewriting*. It is typically used to describe the semantics of more advanced language constructs in terms of simpler language constructs. For example, $i += 1$ can be rewritten into $i = i + 1$. That way, the meaning of $+=$ is defined, without the need to speak of how $+=$ works operationally, nor how to reason about it axiomatically. This rewriting is a purely syntactic operation and is also called a *semantics-preserving program transformation*. One sometimes calls the notation $+=$ *syntactic sugar*, because it does not make the language more expressive. It is a mere abbreviation that can be eliminated by rewriting, also known as *syntactic desugaring*.

2.1. Rewriting Function Calls, by Inlining

In the absence of recursion, the function mechanism is syntactic sugar for abbreviations that can be eliminated by a program transformation. Consider a void¹ function definition without `return` statements of the form

```
13 def func(param_1, param_2, ...):
14     body # containing param1, param2, ..., but no return
```

The call `func(expr_1, expr_2, ...)` can be eliminated as follows.

1. Replace the call `func(expr_1, expr_2, ...)` by

```
15     param_1, param_2, ... = expr_1, expr_2, ...
16     body
```

where `param_1, param_2, ...` are local variables.

2. Systematically rename any local variables in `body` whose name clashes with a name occurring in the context of the call.

When the expressions `expr_i` are free of *side effects*, also known as *referential transparency* (which will be the case in our examples), Step 1 can be replaced by a *double substitution*:

- 1.a. Replace the call `func(expr_1, expr_2, ...)` by `body`,
- 1.b. in which all occurrences of parameters `param_i` are simultaneously replaced by the corresponding argument expressions `(expr_i)`. The parentheses are needed to guarantee the proper evaluation order.

The result of this transformation is a program that is semantically equivalent to the original program. As a *compiler optimization* technique and as a *code refactoring* technique, it is also known as *inlining*. In Lambda Calculus, the double substitution corresponds to *β -reduction*.

For example, consider the Python function definition

```
17 def f(x, y):
18     z = x * y
19     print(z)
```

Then we can rewrite as follows

```
20     f(z - 1, z + 1)
21     print(z)
22
23 # inline call: x, y = z - 1, z + 1
24
```

¹ In some programming languages void functions are known as procedures. In Python, their body does not contain `return expr`. For non-void functions and `return`, see Appendix A.

```

25     z1 = (z - 1) * (z + 1) # z1: fresh local variable
26     print(z1)
27     print(z)

```

Note that in this case the parentheses are really needed to preserve the evaluation order of operators. When they are not needed, we silently omit them.

Non-recursive function definitions can be completely eliminated by inlining all their calls. This may result in faster execution, at the expensive of a larger code foot print. Recursive function definitions cannot be completely eliminated this way. Well, they can, provided we allow *infinite* program texts. Conceptually, there is nothing wrong with an infinite program text. Termination of the recursion corresponds to guaranteeing that only a finite (though unbounded) part of that infinite program gets executed.

2.2. Rewriting *if*-statements, by Deleting Dead Branches

It turns out that for the examples in §3 we also need rewriting rules for *if*-statements with constant conditions. These are the two relevant rewrite rules:

```

28     if True:
29         statement_suite_1
30     else: # unreachable
31         statement_suite_2
32
33 # delete dead else-branch
34
35     statement_suite_1

```

and

```

36     if False: # unreachable
37         statement_suite_1
38     else:
39         statement_suite_2
40
41 # delete dead if-branch
42
43     statement_suite_2

```

2.3. Rewriting Expressions, by Constant Folding

The final rewrite step that we use in the examples below is that of evaluating an expression involving only constants. As a compiler optimization technique this is known as *constant folding*. We label such rewrites by `simplify`.

3. Example for Understanding Via Rewriting

Let's rewrite the call `print_bit_strings(2)`:

```

44     print_bit_strings(2)
45
46 # inline call: n, s = 2, ""
47
48     if 2 == 0:
49         print("")
50     else:
51         for b in "01":
52             print_bit_strings(2 - 1, "" + b)
53
54 # delete dead if-branch: 2 != 0 ; simplify: 2 - 1 == 1 ; "" + b == b
55
56     for b in "01":
57         print_bit_strings(1, b)
58
59 # inline call: n, s = 1, b; rename local variables
60
61     for b1 in "01":
62         # print_bit_strings(1, b1)
63         if 1 == 0:
64             print(b1)
65         else:
66             for b2 in "01":
67                 print_bit_strings(1 - 1, b1 + b2)
68
69 # delete dead if-branch: 1 != 0 ; simplify: 1 - 1 == 0
70
71     for b1 in "01":
72         for b2 in "01":
73             print_bit_strings(0, b1 + b2)
74
75 # inline call: n, s = 10, b1 + b2; rename local variable
76
77     for b1 in "01":
78         for b2 in "01":
79             # print_bit_strings(0, b1 + b2)
80             if 0 == 0:
81                 print(b1 + b2)
82             else:
83                 for b3 in "01":

```

```

84         print_bit_strings(0 - 1, b1 + b2 + b3)
85
86 # delete dead else-branch: 0 == 0
87
88     for b1 in "01":
89         for b2 in "01":
90             print(b1 + b2)

```

So, by *equational reasoning*, the recursive function applied to argument 2 is equivalent to 2 nested `for`-loops. It is now believable that for argument `n`, the call is equivalent to `n` nested `for`-loops, since each recursive call adds a level of nesting:

```

91 # print_bit_strings(n)
92 for b1 in "01":
93     for b2 in "01":
94         ...
95         for bn in "01":
96             print(b1 + b2 + ... + bn)

```

The role of (accumulation) parameter `s` can apparently be viewed as collecting the state of all enclosing `for`-loops. We see that recursion enables one to write programs with a variable number of nested `for`-loops. Without recursion this is often not possible in an imperative programming language. However, see §4 for a way of doing it in Python without recursion.

To my delight, the IntelliJ IDE can do the refactoring steps of inlining calls of recursive functions and deleting dead branches in `if`-statements, but only for Java. It can't do them for Python (neither can VS Code). It can inline nonrecursive Python function calls, but not for direct recursive functions. Surprisingly, if the recursive function is duplicated and turned into a pair of mutually recursive functions, the IDE will refactor their function calls properly.

Finally, it is important to note the difference between tracing the execution and rewriting the program as ways of understanding recursion. Execution tracing suffers from the exponential blow up in branching recursion, whereas syntactic program rewriting does not. The latter reasons about the (unexecuted) program as a whole.

4. Designing Recursive Function Definitions

The rewriting approach described in the previous section to help understand recursive function definitions can also be used to help design such functions. The steps are as follows.

1. Write down a non-recursive program that solves a particular instance of the problem.

2. Decide which part will be done in a single layer of the recursion, which part will be handled by recursion, and which part is handled by preceding layers. Focus on the perspective of that single layer.
3. Introduce appropriate parameters to feed in data that comes from the preceding layers, and modify and pass them on to the lower layers. In particular, there will also be a parameter for the problem size.
4. Decide on the base case(s).
5. Define the recursive function; in particular, introduce an `if`-statement to distinguish the base case(s) and the ‘general’ case that adds a single layer.

4.1. First Design Example

As an example, consider the problem of printing all bit strings of length n .

1. A straightforward non-recursive program for $n = 3$ consisting of 3 nested `for`-loops:

```

97     for b1 in "01":
98         for b2 in "01":
99             for bn in "01":
100                 print(b1 + b2 + b3)

```

2. A single layer of the recursion does one `for`-loop, say with control variable `b2`, the loops nested inside will be handled by the recursive call, and the outer loops were done in the preceding layers:

```

101     # print_bit_strings(3)
102     #####
103     for b1 in "01": # preceding layers
104         #####
105         for b2 in "01": # < this will be done in one layer
106             #####
107             for b3 in "01": # handled by
108                 print(b1 + b2 + b3) # recursion
109             #####

```

Note that the part handled by recursion concerns a *generalization*, because those `for`-loops do not just contain `print(b3)`. Rather, they contain `print(b1 + b2 + b3)`. So, the generalization is that the recursive call must print `s = b1 + b2` extended with `b3`. That first part must come in via an extra parameter, say `s`. This also means that the call itself will receive that parameter, but then its value will be `b1`, received from the preceding layers. For the top-level call, we can take `s = ""`, because it is the unit of string concatenation.

3. Thus, we obtain the following structure:

```

110     #####
111     for b1 in "01": # preceding layers
112     #####
113         # print_bit_strings(2, b1) # call being designed
114         for b2 in "01":
115             print_bit_strings(1, b1 + b2) # recursive call
116             # which should inline as
117             # for b3 in "01":
118                 #     print(b1 + b2 + b3)

```

4. In case $n == 0$, there are no `for`-loops, and only a `print` statement, which can be fed with the parameter `s`.

5. This leads to the following definition, which we have seen before:

```

119 def print_bit_strings(n: int, s: str = "") -> None:
120     """Print s + t for all strings t over "01" of length n,
121     in lexicographic order.
122
123     Assumption: n >= 0
124     """
125
126     if n == 0:
127         # t == ""
128         print(s)
129     else:
130         # n > 0, write t == b + u for b in "01"
131         for b in "01":
132             print_bit_strings(n - 1, s + b)

```

Note the default value `s = ""`, corresponding to an empty context of preceding recursive layers.

In this approach, one can use the IDE refactoring technique known as *extract function*. It will introduce appropriate parameters to feed in values that are must be supplied.

4.2. Second Design Example

Consider a binary tree of depth 2 (Fig. 2, left). How can it be used to create a binary tree of depth 3? One way is to copy the binary tree of depth 2 and combine the two instances with a fork on top (Fig. 2, middle). This is the view we followed in the preceding example. But one can also grow a binary fork on each of the leaves of the binary tree of depth 2 to get a binary tree of depth 3 (Fig. 2, right).




Fig. 2. Binary tree of depth 2 (left); two ways (middle, right) of constructing a binary tree of depth 3 from trees of depth 2.

Let's see how that alternative choice can be worked out in case of function `print_bit_strings`.

1. The non-recursive program is the same as above in §4.1.
2. A single layer of the recursion does one `for`-loop, but now the recursion will do the *outer* `for`-loops and the preceding layers did the *inner* loops:

```

133     # print_bit_strings(3)
134     #####
135     for b1 in "01": # handled by recursion
136         #####
137         for b2 in "01": # < this will be done in one layer
138             #####
139             for b3 in "01": # preceding
140                 print(b1 + b2 + b3) # layers
141             #####

```

This may look strange, but bear with me. Apparently, in the recursion, bit strings of length one shorter are produced, and these still need to be extended and printed. So, again, we see a *generalization*: rather than just print all bit strings, we need to apply a function to them, and this function is going to be an extra parameter, say `f`. The call we are designing receives this parameter, which represents the work to be done in the preceding layers (inner loops) For the top-level call, we can take `f = print`. The recursive call will receive a function (as parameter) that adds one `for`-loop around the given `f`.

3. Thus, we obtain the following structure:

```

142     # print_bit_strings(2, f)
143     def g(s: str) -> None:
144         for b2 in "01":
145             f(s + b2)
146

```

```

147     print_bit_strings(1, g)
148     # which should eventually inline as
149     # for b1 in "01":
150     #     g(s + b1)
151
152     # where def f(s: str) -> None:
153     #####
154         for b3 in "01": # preceding
155             print(s + b3) # layers
156     #####

```

4. In case $n == 0$, there are no `for`-loops and `f` can just be applied to `""`.

5. This leads to the following recursive definition:

```

157 def print_bit_strings(n: int,
158                       f: Callable[[str], None] = print
159                       ) -> None:
160     """Apply f to each string over "01" of length n,
161     in lexicographic order.
162
163     Assumption: n >= 0
164     """
165     if n == 0: # only the empty string has length 0
166         f("")
167     else: # n > 0
168         def g(s: str) -> None:
169             for b in "01":
170                 f(s + b)
171
172         print_bit_strings(n - 1, g)

```

Let's now see if we can understand this recursive definition by rewriting the call `print_bit_strings(2)`:

```

173     print_bit_strings(2)
174
175 # inline call: n, f = 2, print
176
177     if 2 == 0:
178         print("")
179     else:
180         def g(s: str) -> None:
181             for b in "01":
182                 print(s + b)

```

```

183
184     print_bit_strings(2 - 1, g)
185
186 # delete dead if-branch: 2 != 0 ; simplify: 2 - 1 == 1
187
188     def g(s: str) -> None:
189         for b in "01":
190             print(s + b)
191
192     print_bit_strings(1, g)
193
194 # inline call: n, f = 1, g; rename new local function g -> g2
195
196     def g(s: str) -> None:
197         for b in "01":
198             print(s + b)
199
200     if 1 == 0:
201         g("")
202     else:
203         def g2(s: str) -> None: # renamed local function
204             for b in "01":
205                 g(s + b)
206
207         print_bit_strings(1 - 1, g2)
208
209 # delete dead if-branch: 1 != 0 ; simplify: 1 - 1 == 0
210
211     def g(s: str) -> None:
212         for b in "01":
213             print(s + b)
214
215     def g2(s: str) -> None:
216         for b in "01":
217             g(s + b)
218
219     print_bit_strings(0, g2)
220
221 # inline call g(s + b); rename control variables
222
223     def g2(s: str) -> None:
224         for b1 in "01": # renamed control variable
225             for b2 in "01": # renamed control variable
226                 print((s + b1) + b2)

```

```

227
228     print_bit_strings(0, g2)
229
230 # inline call: n, f = 0, g2; rename new local function g -> g3
231
232     def g2(s: str) -> None:
233         for b1 in "01":
234             for b2 in "01":
235                 print((s + b1) + b2)
236
237     if 0 == 0:
238         g2("")
239     else:
240         def g3(s: str) -> None:
241             for b in "01":
242                 g2(s + b)
243
244         print_bit_strings(0 - 1, g3)
245
246 # delete dead else-branch: 0 == 0
247
248     def g2(s: str) -> None:
249         for b1 in "01":
250             for b2 in "01":
251                 print((s + b1) + b2)
252
253     g2("")
254
255 # inline call g2("")
256
257     for b1 in "01":
258         for b2 in "01":
259             print("" + b1 + b2)
260
261 # simplify: "" + b1 = b1
262
263     for b1 in "01":
264         for b2 in "01":
265             print(b1 + b2)

```

Now, the extra parameter (*f*) accumulates the work to be done, and in the base case it is applied to the empty string. Such a parameter is called a *continuation*. Observe that this definition of `print_bit_strings` is *tail recursive* and thus can easily be transformed into a **while**-loop:

```

266 def print_bit_strings(n: int) -> None:
267     """Print all strings over "01" of length n,
268     i   n lexicographic order.
269
270     Assumption: n >= 0
271     """
272     f = print
273
274     while n > 0:
275         def g(s: str, f=f) -> None:
276             for b in "01":
277                 f(s + b)
278
279             n, f = n - 1, g
280
281     f("")

```

Note that here `g` needs an extra parameter `f` with default value `f`, to ensure that the definition of `g` is a *closure* that properly captures the function object currently bound to the name `f` at the moment of definition. Without that `f` parameter, the definition of `g` would contain an ‘open’ (un-dereferenced) name `f`, which will be looked up during execution of `g` to find the value bound to `f` at the moment of execution, rather than at the moment of definition.

By redefining `f` directly, `print_bit_strings` can be simplified to

```

282     f = print
283
284     for _ in range(n):
285         def f(s: str, f=f) -> None:
286             for b in "01":
287                 f(s + b)
288
289     f("")

```

Note that `f` here is not defined recursively, since the `f` in its body is the parameter, which is bound to the earlier value of `f`.

Apparently, in Python one can write a non-recursive program that behaves like a variable number of nested loops. Actually, the program constructs a function that behaves like those nested loops. Thus, this is a form of *metaprogramming*.

5. Conclusion

I have described and illustrated how a rewriting semantics can help understand and design recursive function definitions. This approach complements the traditional approaches based on operational semantics (execution tracing) and axiomatic semantics (contracts). I am not claiming that the approach via rewriting is the best, but I do find it better than execution tracing, because (i) it is purely syntactic, (ii) does not need a stack to distinguish the states of concurrently active recursive calls, (iii) nor does it suffer from any exponential blow up. Furthermore, the rewriting approach may help in discovering and formulating generalized contracts, which are needed for the approach via axiomatic semantics.

There is a clear relationship to *functional programming*, whose semantics can be based on Lambda Calculus, which has a rewriting semantics via β -reduction (inlining). Note that a rewriting semantics does not need a stack. In this article, I have shown that this approach also can work for imperative programs. It thus allows *equational reasoning* on the level of whole programs. However, some care is needed, in particular when **return** statements are used and when expressions can have side effects, causing a lack of *referential transparency*.

The approach via rewriting would benefit from IDE support for the relevant *refactoring techniques*, because manual rewriting is tedious and error-prone. Unfortunately, such support is currently rather limited (JetBrains IntelliJ can do it for Java).

A word of warning is in place concerning the examples. The rewriting approach can help to come up with recursive definitions. But these definitions may still have performance issues. There are other techniques to improve the performance of recursively defined functions, e.g., see Verhoeff (2018). Appendix B offers better ways of printing bit strings in Python. To show the power of a purely functional programming language, I have included some Haskell programs for generating bit strings in Appendix C.

I hope also that I have shown some nifty uses of Python. Maybe the recent Python compiler named Codon (Shajii, 2023) can make Python interesting for use in programming contests such as the IOI.

Acknowledgment

I would like to thank Berry Schoenmakers and Sten Wessel (TU Eindhoven, Netherlands) and Radu Negulescu (Ontario, Canada) for helping me improve this article.

References

- Roberts, E. (1986). *Thinking Recursively* (1st Ed.). Wiley.
- Rubio-Sánchez, M. (2018). *Introduction to Recursive Programming*. Taylor & Francis.

- DOI: 10.1201/9781315120850
- Shajii, A. *et al.* (2023). Codon: A Compiler for High-Performance Pythonic Applications and DSLs. In: *Proceedings of the 32nd ACM SIGPLAN International Conference on Compiler Construction*. ACM, pp. 191–202. DOI: 10.1145/3578360.3580275
- Verhoeff, T. (2018). A Master Class on Recursion. In: *Adventures Between Lower Bounds and Higher Altitudes*. Lecture Notes in Computer Science Vol. 11011. Springer, pp. 610–633. DOI: 10.1007/978-3-319-98355-4_35
- Verhoeff, T. (2021). Look Ma, Backtracking without Recursion. (IOI Conference 2021). *Olympiads in Informatics*, 15, 119–132. DOI: 10.15388/ioi.2021.10
- Verhoeff, T. (2023). Git repository with source code for “Understanding and Designing Recursive Functions via Syntactic Rewriting”. (Accessed 29 April 2023)
<https://gitlab.tue.nl/t-verhoeff-software/code-for-understanding-recursion>



T. Verhoeff is Assistant Professor in Computer Science at Eindhoven University of Technology, where he works in the group Software Engineering & Technology. His research interests are support tools for verified software development and model driven engineering. He received the IOI Distinguished Service Award at IOI 2007 in Zagreb, Croatia, in particular for his role in setting up and maintaining a web archive of IOI-related material and facilities for communication in the IOI community, and in establishing, developing, chairing, and contributing to the IOI Scientific Committee from 1999 until 2007.

Appendix A. Python Example with Non-void Function

Eliminating the call of a non-void² function is a bit more involved than for void functions. To keep things simple, we will assume that the return statements in the function body occur at *tail positions*, that is, if the `return` statement would have been a function call, it would be the last thing done in the body (a so-called tail call). Now consider a non-void function definition of the form

```
290 def func(param_1, param_2, ...):
291     body # containing 'return expr' in tail positions only
```

A call of this function occurs as a (sub)expression, which is part of some statement, such as for example

```
292     print(func(expr_1, expr_2, ...) + 1)
```

In general, such a call takes the form

```
293     stmt(func(expr_1, expr_2, ...))
```

where `stmt` is a void function. Assuming, for simplicity, that the argument expressions have no side effects, it can be eliminated as follows.

1. Replace `stmt(func(expr_1, expr_2, ...))` by `body`.
2. Replace every occurrence of `return expr` in `body` by `stmt(expr)`. N.B. The expression may need to be parenthesized as `(expr)`.
3. Simultaneously replace all occurrences of parameters `param_i` by their corresponding argument expressions `(expr_i)`.
4. Systematically rename any local variables in `body` whose name clashes with a name occurring in the context of the call.

If a return statement would not occur in a tail position, then a ‘jump’ to the end of the body would also be needed. Note, however, that Python does not support `goto` statements (except on April’s Fool Day 2004).

Here is an example involving the famous recursive factorial function:

```
294 def fac(n: int) -> int:
295     """For n >= 0, return n factorial.
296     """
297     if n == 0:
298         return 1
299     else:
300         return n * fac(n - 1)
```

² Non-void functions are sometimes known as fruitful functions. In Python, they contain `return expr`.

Note that the two return statements occur in tail positions. Let's rewrite the statement `print(fac(2))`:

```

301     print(fac(2))
302
303 # inline call: n = 2
304
305     if 2 == 0:
296         print(1)
297     else:
298         print(2 * fac(2 - 1))
299
300 # delete dead if-branch: 2 != 0 ; simplify: 2 - 1 == 1
301
302     print(2 * fac(1))
303
304 # inline call: n = 1
305
306     if 1 == 0:
307         print(2 * 1)
308     else:
309         print(2 * 1 * fac(1 - 1))
310
311 # delete dead if-branch: 1 != 0 ; simplify: 1 - 1 == 0
312
313     print(2 * 1 * fac(0))
314
315 # inline call: n = 0
316
317     if 0 == 0:
318         print(2 * 1 * 1)
319     else:
320         print(0 * 2 * * 1 * fac(0 - 1))
321
322 # delete dead else-branch: 0 == 0 ; simplify: 2 * 1 == 2
323
324     print(2 * 1 * 1)

```

In general, `print(fac(n))` rewrites to

```

335     print(n * (n - 1) * ... * 2 * 1 * 1)

```

where there are $n + 1$ factors in the expression.

If the body of a non-void function consists of single return statement, then one can do equational reasoning directly on the level of expressions rather than statements. For instance, consider the following definition of function `fac`:

```
336 def fac(n: int) -> int:
337     return 1 if n == 0 else n * fac(n - 1)
```

Now, rewriting call `fac(2)` is less complicated:

```
338     fac(2)
339
340 # inline call: n = 2
341
342     1 if 2 == 0 else 2 * fac(2 - 1)
343
344 # delete dead if-branch: 2 != 0 ; simplify: 2 - 1 == 1
345
346     2 * fac(1)
347
348 # inline call: n = 1
349
350     2 * (1 if 1 == 0 else 2 * 1 * fac(1 - 1))
351
352 # delete dead if-branch: 1 != 0 ; simplify: 1 - 1 == 0
353
354     2 * 1 * fac(0)
355
356 # inline call: n = 0
357
358     2 * 1 * (1 if 0 == 0 else 0 * 2 * 1 * fac(0 - 1))
359
360 # delete dead else-branch: 0 == 0 ; simplify: 2 * 1 == 2
361
362     2 * 1 * 1
```

Appendix B. Better Python Solutions for Bit Strings

For completeness sake, let me note that the problem of printing all bit strings of a given length can be solved in a better way. Decompose the problem into:

1. Constructing all bit strings of a given length.
2. Printing them all.

In Python, one might be tempted to construct all those bit strings, by putting them in a list. But then they are all stored in memory before printing them (or doing whatever one wants, for instance, count them). For that reason, generator expressions were introduced in Python. They allow on-demand construction. Here is a Python solution based on Fig. 2 (middle):

```

363 from typing import Iterator
364
365 def generate_bit_strings(n: int) -> Iterator[str]:
366     """Yield all strings over "01" of length n,
367     in lexicographic order.
368
369     Assumption: n >= 0
370     """
371     if n == 0:
372         yield ""
373     else: # n > 0
374         yield from (b + u
375                     for b in "01"
376                     for u in generate_bit_strings(n - 1)
377                    )

```

It can be invoked to print the bit strings like this:

```

378 print(*generate_bit_strings(3), sep='\n')

```

And here is a solution based on Fig. 2 (right):

```

379 def generate_bit_strings(n: int) -> Iterator[str]:
380     if n == 0:
381         yield ""
382     else: # n > 0
383         yield from (u + b
384                     for u in generate_bit_strings(n - 1)
385                     for b in "01"
386                    )

```

Appendix C. Haskell Solutions for Bit Strings

It is illustrative to see the same definitions in a pure non-strict functional programming language like Haskell. Using the recursive decomposition in Fig. 2 (middle), combining two recursively grown trees of size one smaller:

```

387 bitStrings :: Int -> [String]
388 -- bitStrings n = list of strings over "01" of length n (n >= 0),
389 -- in lexicographic order
390 bitStrings 0 = [""]
391 bitStrings n = [b : u | b <- "01", u <- bitStrings (n - 1)]

```

and when growing one tree of size one smaller and splitting all its leaves (Fig. 2, right):

```

392 bitStrings 0 = [""]
393 bitStrings n = [u ++ [b] | u <- bitStrings (n - 1), b <- "01"]

```

The latter is not efficient, because appending at the end of a list is not efficient in Haskell. But this can be improved by introducing an accumulation parameter:

```

394 gbitStrings :: [String] -> Int -> [String]
395 -- gbitStrings s n = [t ++ u | t <- bitStrings n, u <- s]
396 -- hence, gbitStrings [""] = bitStrings
397 gbitStrings s 0 = s
398 gbitStrings s n = gbitStrings [b : t | b <- "01", t <- s] (n - 1)

```

Observe that this definition is more efficient, because it now prepends to a list. Moreover, it is tail recursive, and thus the recursion can be compiled into a loop. Also note that in Haskell, lists are lazy, that is, they are only constructed in so far as needed (like the generator expressions in Python used in Appendix B).

