

# Tail call elimination

Michel Schinz  
Advanced Compiler Construction – 2009-04-03

# Tail calls and their elimination

## Loops in functional languages

Several functional programming languages do not have an explicit looping statement. Instead, programmers resort to *recursion* to loop.

For example, the central loop of a Web server written in Scheme might look like this:

```
(define web-server-loop
  (lambda ()
    (wait-for-connection)
    (fork handle-connection)
    (web-server-loop)))
```

3

## The problem

Unfortunately, recursion is not equivalent to the looping statements usually found in imperative languages: recursive function calls, like all calls, consume stack space while loops do not...

In our example, this means that the Web sever will eventually crash because of a stack overflow – this is clearly unacceptable!

A solution to this problem must be found...

4

## The solution

In our example, it is obvious that the recursive call to `web-server-loop` could be replaced by a jump to the beginning of the function. If the compiler could detect this case and replace the call by a jump, our problem would be solved! This is the idea behind **tail call elimination**.

5

## Tail calls

The reason why the recursive call of `web-server-loop` could be replaced by a jump is that it is the last action taken by the function :

```
(define web-server-loop
  (lambda ()
    (wait-for-connection)
    (fork handle-connection)
    (web-server-loop)))
```

Calls in terminal position – like this one – are called **tail calls**. This particular tail call also happens to target the function in which it is defined. It is therefore said to be a **recursive tail call**.

6

## Tail calls examples

In the functions below, all calls are underlined. Which ones are tail calls?

```
(define map
  (lambda (f l)
    (if (null? l)
        1
        (cons (f (car l))
                 (map f (cdr l))))))

(define fold
  (lambda (f z l)
    (if (null? l)
        z
        (fold f (f z (car l)) (cdr l)))))
```

tail call

recursive tail call

7

## Tail call elimination

When a function performs a tail call, its own activation frame is dead, as by definition nothing follows the tail call.

Therefore, it is possible to first free the activation frame of a function about to perform such a call, then load the parameters for the call, and finally *jump* to the function's code.

This technique is called **tail call elimination** (or **optimization**), here abbreviated **TCE**.

8

## TCE example

Consider the following function definition and call:


```
(define sum
  (lambda (z l)
    (if (null? l)
        z
        (sum (+ z (car l)) (cdr l)))))
(sum 0 (list3 1 2 3))
```

How does the stack evolve, with and without tail call elimination?

9

## TCE example


Without tail call elimination, each recursive call to `sum` makes the stack grow, to accommodate activation frames.



10

## TCE example

With tail call elimination, the dead activation frames are freed before the tail call, resulting in a stack of constant size.



time →

11

## Tail call optimization?

Tail call elimination is more than just an optimization! Without it, writing a program that loops endlessly using recursion and does not produce a stack overflow is simply impossible.

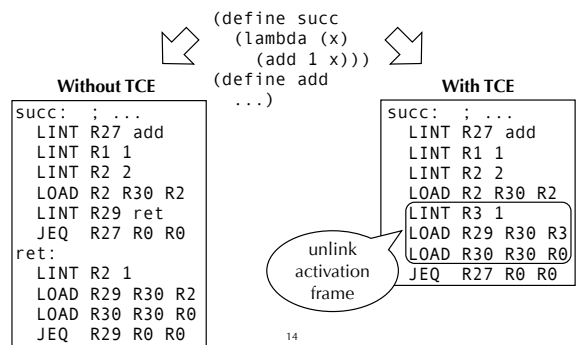
For that reason, full tail call elimination is actually *required* in some languages, e.g. Scheme.

In other languages, like C, it is simply an optimization performed by some compilers in some or all cases.

12

## Tail call elimination for minischeme

## TCE example



## Implementing TCE

Tail call elimination is implemented by:

1. identifying tail calls in the program,
2. compiling those tail calls specially, by deallocating the activation frame of the caller before jumping to the called function.

We already know how to compile tail calls, but we did not explain yet how to identify them.

15

## Identifying tail calls

To identify tail calls, we first assume that all calls are marked with a unique number. We then define a function  $T$  that returns the marks corresponding to the tail calls.

For example, given the following expression:

```
(lambda (x)
  (if (even? x) (g (h x)) (h (g x))))
```

$T$  produces the set  $\{2, 4\}$ .

16

## Identifying tail calls

$T[(\text{lambda } (args) \ body)] = T[body]$

where the auxiliary function  $T'$  is defined as follows:

$T'[(\text{let } ((n \ v)) \ body)] = T[body]$

$T'[(\text{begin } body_1 \ \dots \ body_n)] = T[body_n]$

$T'[(\text{if } e_1 \ e_2 \ e_3)] = T[e_2] \cup T[e_3]$

$T'[(\text{and } e_1 \ e_2)] = T[e_2]$

$T'[(\text{or } e_1 \ e_2)] = T[e_2]$

$T'[^m(e_1 \ \dots \ e_n)]$  when  $e_1$  is not a primitive =  $\{m\}$

$T'$ [anything else] =  $\emptyset$

17

## Removing the epilogue

Usually, functions end with an epilogue that de-allocates the stack frame. But this epilogue is not necessary for functions that end only with tail calls.

The functions  $E$  below determines whether a function requires an epilogue or not, using an auxiliary function  $E'$ .

$E[(\text{lambda } (args) \ body)] = E'[body]$

$E'[(\text{let } ((n \ v)) \ body)] = E'[body]$

$E'[(\text{begin } body_1 \ \dots \ body_n)] = E'[body_n]$

$E'[(\text{if } e_1 \ e_2 \ e_3)] = E'[e_2] \vee E'[e_3]$

$E'[(e_1 \ \dots \ e_n)]$  when  $e_1$  is not a primitive = False

$E'$ [anything else] = True

18

## Tail call elimination in uncooperative environments

## TCE in various environments

When generating assembly language, it is easy to perform TCE, as the target language is sufficiently low-level to express the deallocation of the activation frame and the following jump. When targeting higher-level languages, like C or the JVM, this becomes difficult – although recent VMs like .NET's support tail calls. Let's explore several techniques that have been developed to perform TCE in such contexts.

20

## Benchmark program

To illustrate how the various techniques work, we will use a benchmark program in C that tests whether a number is even, using two mutually tail-recursive functions.

When no technique is used to manually eliminate tail calls, it looks as follows. And unless the C compiler performs tail call elimination – like GCC does with full optimization – it crashes with a stack overflow at run time.

```
int even(int x) {
    return x == 0 ? 1 : odd(x - 1);
}
int odd(int x) {
    return x == 0 ? 0 : even(x - 1);
}
int main(int argc, char* argv[]) {
    printf("%d\n", even(300000000));
}
```

21

## Single function approach

The “single function” approach consists in compiling the whole program to a single function of the target language. This makes it possible to compile tail calls to simple jumps within that function, and other calls to recursive calls to it. This technique is rarely applicable in practice, due to limitations in the size of functions of the target language.

22

## Single function approach in C

```
typedef enum { fun_even, fun_odd } fun_id;

int wholeprog(fun_id fun, int x) {
    switch (fun) {
        case fun_even: goto even;
        case fun_odd:  goto odd;
    }

    even:
    if (x == 0) return 1;
    x = x - 1;
    goto odd;
    odd:
    if (x == 0) return 0;
    x = x - 1;
    goto even;
}

int main(int argc, char* argv[]) {
    printf("%d\n", wholeprog(fun_even, 300000000));
}
```

23

## Trampolines

With trampolines, functions never perform tail calls directly. Rather, they return a special value to their caller, informing it that a tail call should be performed. The caller performs the call itself.

For this scheme to work, it is necessary to check the return value of all functions, to see whether a tail call must be performed. The code which performs this check is called a **trampoline**.

24

## Trampolines in C

```
typedef void* (*fun_ptr)(int);
struct { fun_ptr fun; int arg; } resume;
void* even(int x) {
    if (x == 0) return (void*)1;
    resume.fun = odd;
    resume.arg = x - 1;
    return &resume;
}
void* odd(int x) {
    if (x == 0) return (void*)0;
    resume.fun = even;
    resume.arg = x - 1;
    return &resume;
}
int main(int argc, char* argv[]) {
    void* res = even(300000000);
    while (res == &resume)
        res = (resume.fun)(resume.arg);
    printf("%d\n", (int)res);
}
```

25

## Extended trampolines

Extended trampolines trade some of the space savings of standard trampolines for speed.

Instead of returning to the trampoline on every tail call, the number of successive tail calls is counted at run time, using a tail call counter (tcc) passed to every function. When that number reaches a predefined limit *l*, a non-local return is performed to transfer control to a trampoline "waiting" at the bottom of the chain, thereby reclaiming *l* activation frames in one go.

26

## C's setjmp / longjmp

Extended trampolines are more efficient when a non-local return is used to free dead stack frames.

In C, non-local returns can be performed using the standard functions `setjmp` and `longjmp`, which can be seen as a form of `goto` that works across functions:

- `setjmp(b)` saves its calling environment in `b`, and returns `0`,
- `longjmp(b, v)` restores the environment stored in `b`, and proceeds like if the call to `setjmp` had returned `v` instead of `0`.

In the following slides, we use `_setjmp` and `_longjmp`, which do not save and restore the signal mask and are therefore much more efficient.

27

## Extended trampolines in C

```
typedef int (*fun_ptr)(int, int);
struct { fun_ptr fun; int arg; } resume;
jmp_buf jmp_env;

int even(int tcc, int x) {
    if (tcc > TC_LIMIT) {
        resume.fun = even;
        resume.arg = x;
        _longjmp(jmp_env, -1);
    }
    return (x == 0) ? 1 : odd(tcc + 1, x - 1);
}
int odd(int tcc, int x) { /* similar to even */ }

int main(int argc, char* argv[]) {
    int res = (_setjmp(jmp_env) == 0)
        ? even(0, 300000000)
        : (resume.fun)(0, resume.arg);
    printf("%d\n", res);
}
```

28

## Baker's technique

(Henry) **Baker's technique** consists in first transforming the whole program to continuation-passing style.

One important property of CPS is that all calls are tail calls. Consequently, it is possible to periodically shrink the whole stack using a non-local return.

29

## Baker's technique in C


```
typedef void (*cont)(int);
typedef void (*fun_ptr)(int, cont);
int tcc = 0;
struct { fun_ptr fun; int arg; cont k; } resume;
jmp_buf jmp_env;
void even_cps(int x, cont k) {
    if (++tcc > TC_LIMIT) {
        tcc = 0;
        resume.fun = even_cps;
        resume.arg = x;
        resume.k = k;
        _longjmp(jmp_env, -1);
    }
    if (x == 0) (*k)(1); else odd_cps(x - 1, k);
}
void odd_cps(int x, cont k) { /* similar to even_cps */ }
int main(int argc, char* argv[]) {
    if (_setjmp(jmp_env) == 0) even_cps(300000000, main_1);
    else (resume.fun)(resume.arg, resume.k);
}
void main_1(int val) { printf("%d\n", val); exit(0); }
```

30

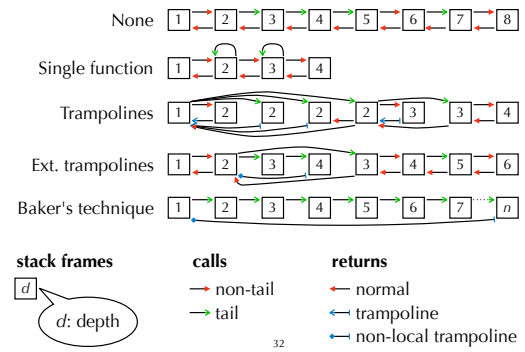
## Benchmark results

The programs presented earlier were compiled with GCC v4.01 and two different optimization settings (-O0 and -O3). The normalized running times observed on an Intel Core 2 Duo are presented below.

Notice that the initial version compiled without optimization produces a stack overflow, hence the absence of timing.



## Techniques summary



## Summary

Tail call elimination consists in compiling tail calls specially, so that the activation frame of the caller is freed before the called function is invoked.

This technique reduces memory usage and makes it possible to write loops using recursion without overflowing the stack.

Tail call elimination can be hard to implement efficiently when the target platform is uncooperative.