

Executing T_EX

Hans Hagen

Much of the Lua code in ConT_EXt originates from experiments. What survives in the source code is probably either used, waiting to be used, or kept for educational purposes. The functionality that we describe here has already been present for a while in ConT_EXt, but has been improved a little starting with LuaT_EX 1.08 due to an extra helper. The code shown here is generic and is not used in ConT_EXt as such.

Say that we have this code:

```
for i=1,10000 do
  tex.sprint("1")
  tex.sprint("2")
  for i=1,3 do
    tex.sprint("3")
    tex.sprint("4")
    tex.sprint("5")
  end
  tex.sprint("\\space")
end
```

When we call `\directlua` with this snippet we get some 30 pages of 12345345345. The printed text is saved until the end of the Lua call so basically we pipe some 170 000 characters to T_EX that get interpreted as one paragraph.

Now imagine this:

```
\setbox0\hbox{xxxxxxxxxxxx} \number\wd0
```

which gives 3652935 (the width of box0). If we check the box in Lua, with:

```
tex.sprint(tex.box[0].width)
tex.sprint("\\enspace")
tex.sprint("\\setbox0\hbox{!}")
tex.sprint(tex.box[0].width)
```

the result is 3652935 3652935 i.e. the same number repeated, which is not what you would expect at first sight. However, if you consider that we just pipe to a T_EX buffer that gets parsed *after* the Lua call, it will be clear that the reported width is each time the width that we started with. Our code will work all right if we use:

```
tex.sprint(tex.box[0].width)
tex.sprint("\\enspace")
```

```
tex.sprint("\\setbox0\\hbox{!}")
tex.sprint("\\directlua{tex.sprint(tex.box[0].width)}")
```

and now we get: 3652935 354000, but this use is a bit awkward.

It's not that complex to write some support code that is convenient and this can work out quite well but there is a drawback. If we add references to the status of the input pointer:

```
print(status.input_ptr)
tex.sprint(tex.box[0].width)
tex.sprint("\\enspace")
tex.sprint("\\setbox0\\hbox{!}")
tex.sprint("\\directlua{print(status.input_ptr)\
  tex.sprint(tex.box[0].width)}")
```

we then get 6 and 7 reported. You can imagine that when a lot of nested `\directlua` calls happen, this can lead to an overflow of the input level or (depending on what we do) the input stack size. Ideally we want to do a Lua call, temporarily go to TeX, return to Lua, etc. without needing to worry about nesting and possible crashes due to Lua itself running into problems. One charming solution is to use so-called coroutines: independent Lua threads that one can switch between — you jump out from the current routine to another and from there back to the current one. However, when we use `\directlua` for that, we still have this nesting issue and what is worse, we keep nesting function calls too. This can be compared to:

```
\def\whatever{\ifdone\whatever\fi}
```

where at some point `\ifdone` would be false so we quit, but we keep nesting when the condition is met and eventually we will end up with some nesting related overflow. The following:

```
\def\whatever{\ifdone\expandafter\whatever\fi}
```

is less likely to overflow because there we have tail recursion which basically boils down to not nesting but continuing. Do we have something similar in LuaTeX for Lua? Yes, we do. We can register a function, for instance:

```
lua.get_functions_table()[1] = function() print("Hi there!") end
```

and call that one with:

```
\luafunction 1
```

This is a bit faster than calling a function such as:

contextgroup > context meeting 2018

```
\directlua{HiThere()}
```

which can also be achieved by

```
\directlua{print("Hi there!")}
```

and is sometimes more convenient. Don't overestimate the gain in speed because `directlua` is quite efficient too (and on an average run a user doesn't call it that often, millions of times that is). Anyway, a function call is what we can use for our purpose as it doesn't involve interpretation and effectively behaves like a tail call. The following snippet shows what we have in mind:

```
tex.routine(function()  
  tex.sprint(tex.box[0].width)  
  tex.sprint("\\enspace")  
  tex.sprint("\\setbox0\\hbox{!}")  
  tex.yield()  
  tex.sprint(tex.box[0].width)  
end)
```

We start a routine, jump out to T_EX in the middle, come back when we're done and continue. This gives us: 3652935 186020, which is what we expect.

This mechanism permits efficient (nested) loops like:

```
tex.routine(function()  
  for i=1,10000 do  
    tex.sprint("1")  
    tex.yield()  
    tex.sprint("2")  
    tex.routine(function()  
      for i=1,3 do  
        tex.sprint("3")  
        tex.yield()  
        tex.sprint("4")  
        tex.yield()  
        tex.sprint("5")  
      end  
    end)  
    tex.sprint("\\space")  
    tex.yield()  
  end  
end)
```

We do create coroutines, go back and forwards between Lua and T_EX, but avoid memory being filled up with printed content. If we flush paragraphs (instead of

e.g. the space) then the main difference is that instead of a small delay due to the loop unfolding in a large set of prints and accumulated content, we now get a steady flushing and processing.

However, even using this scheme we can still have an overflow of input buffers because we still nest them: the limitation at the T_EX end has moved to a limitation at the Lua end. How come? Here is the code that we use defining the function `tex.yield()`:

```

local stepper = nil
local stack   = { }
local fid     = 2 -- make sure to take a free slot
local goback  = "\\luafunction" .. fid .. "\\relax"

function tex.resume()
    if coroutine.status(stepper) == "dead" then
        stepper = table.remove(stack)
    end
    if stepper then
        coroutine.resume(stepper)
    end
end

lua.get_functions_table()[fid] = tex.resume

function tex.yield()
    tex.sprint(goback)
    coroutine.yield()
    texio.closeinput()
end

function tex.routine(f)
    table.insert(stack, stepper)
    stepper = coroutine.create(f)
    tex.sprint(goback)
end

-- Because we protect against abuse and overload of functions, in
ConTeXt we
-- need to do the following:

if context then
    fid     = context.functions.register(tex.resume)
    goback = "\\luafunction" .. fid .. "\\relax"
end

```

contextgroup > context meeting 2018

The routine creates a coroutine, and `yield` gives control to \TeX . The resume is done at the \TeX end when we're finished there. In practice this works fine and when you permit enough nesting and levels in \TeX then you will not easily overflow.

When I picked up this side project and wondered how to get around it, it suddenly struck me that if we could just quit the current input level then nesting would not be a problem. Adding a simple helper to the engine made that possible (of course figuring this out took a while):

```
local stepper = nil
local stack   = { }
local fid     = 3 -- make sure to take a free slot
local goback  = "\\luafunction" .. fid .. "\\relax"

function tex.resume()
  if coroutine.status(stepper) == "dead" then
    stepper = table.remove(stack)
  end
  if stepper then
    coroutine.resume(stepper)
  end
end

lua.get_functions_table()[fid] = tex.resume

if texio.closeinput then
  function tex.yield()
    tex.sprint(goback)
    coroutine.yield()
    texio.closeinput()
  end
else
  function tex.yield()
    tex.sprint(goback)
    coroutine.yield()
  end
end

function tex.routine(f)
  table.insert(stack, stepper)
  stepper = coroutine.create(f)
  tex.sprint(goback)
end

-- Again we need to do it as follows in ConTeXt:
```

```

if context then
  fid      = context.functions.register(tex.resume)
  goback   = "\\luafunction" .. fid .. "\\relax"
end

```

The trick is in `texio.closeinput`, a recent helper to the engine and one that should be used with care. We assume that the user knows what she or he is doing. On an older laptop with a i7-3840 processor running Windows 10 the following snippet takes less than 0.35 seconds with Lua_T_EX and 0.26 seconds with Lua_{jit}_T_EX.

```

tex.routine(function()
  for i=1,10000 do
    tex.sprint("\\setbox0\\hpack{x}")
    tex.yield()
    tex.sprint(tex.box[0].width)
    tex.routine(function()
      for i=1,3 do
        tex.sprint("\\setbox0\\hpack{xx}")
        tex.yield()
        tex.sprint(tex.box[0].width)
      end
    end)
  end
end)

```

Say that we were to run the bad snippet:

```

for i=1,10000 do
  tex.sprint("\\setbox0\\hpack{x}")
  tex.sprint(tex.box[0].width)
  for i=1,3 do
    tex.sprint("\\setbox0\\hpack{xx}")
    tex.sprint(tex.box[0].width)
  end
end

```

This executes in only 0.12 seconds in both engines. So what if we run this:

```

\dorecurse{10000}{%
  \setbox0\hpack{x}
  \number\wd0
  \dorecurse{3}{%
    \setbox0\hpack{xx}
    \number\wd0
  }}

```

contextgroup > context meeting 2018

Pure T_EX needs 0.30 seconds for both engines but there we lose 0.13 seconds on the loop code. In the Lua example where we yield, the loop code takes hardly any time. As we need only 0.05 seconds more it demonstrates that when we use the power of Lua, the performance hit of the switch is quite small: we yield 40.000 times! In general, such differences are far exceeded by the overhead: the time needed to typeset the content (which `\hpack` doesn't do), breaking paragraphs into lines, constructing pages and other overhead involved in the run. In ConT_EXt we use a slightly different variant which has 0.30 seconds more overhead, but that is probably true for all Lua usage in ConT_EXt, but again, it disappears in other runtime.

Here is another example:

```
\def\TestWord#1%
  {\directlua{
    tex.routine(function()
      tex.sprint("\setbox0\hbox{\tttf #1}")
      tex.yield()
      tex.sprint(math.round(100 * tex.box[0].width/tex.hsize))
      tex.sprint(" percent of the hsize: ")
      tex.sprint("\box0")
    end)
  }}
```

```
The width of next word is \TestWord {inline}!
```

The width of next word is 9 percent of the hsize: inline!

Now, in order to stay realistic, this macro can also be defined as:

```
\def\TestWord#1%
  {\setbox0\hbox{\tttf #1}%
   \directlua{
     tex.sprint(math.round(100 * tex.box[0].width/tex.hsize))
   } %
   percent of the hsize: \box0\relax}
```

We get the same result: “The width of next word is 9 percent of the hsize: inline!”.

We have been using a Lua-T_EX mix for over a decade now in ConT_EXt and have never really needed this mixed model. There are a few places where we could (have) benefited from it and now we might use it in a few places, but so far we have done fine without it. In fact, in most cases typesetting can be done fine at the T_EX end. It's all a matter of imagination.