Reversing GO binaries like a pro

Published: 2016-09-21 · Archived: 2026-04-05 22:08:17 UTC

GO binaries are weird, or at least, that is where this all started out. While delving into some Linux malware named Rex, I came to the realization that I might need to understand more than I wanted to. Just the prior week I had been reversing Linux Lady which was also written in GO, however it was not a stripped binary so it was pretty easy. Clearly the binary was rather large, many extra methods I didn't care about - though I really just didn't understand why. To be honest - I still haven't fully dug into the Golang code and have yet to really write much code in Go, so take this information at face value as some of it might be incorrect; this is just my experience while reversing some ELF Go binaries! If you don't want to read the whole page, or scroll to the bottom to get a link to the full repo, just go here.

To illistrate some of my examples I'm going to use an extremely simple 'Hello, World!' example and also reference the Rex malware. The code and a Make file are extremely simple;

Hello.go

```
1  package main
2  import "fmt"
3  func main() {
4      fmt.Println("Hello, World!")
5  }
```

Makefile

```
1  all:
2          GOOS=linux GOARCH=386 go build -o hello-stripped -ldflags "-s" hello.go
3          GOOS=linux GOARCH=386 go build -o hello-normal hello.go
```

Since I'm working on an OSX machine, the above `GOOS` and `GOARCH` variables are explicitly needed to cross-compile this correctly. The first line also added the `ldflags` option to strip the binary. This way we can analyze the same executable both stripped and without being stripped. Copy these files, run `make` and then open up the files in your disassembler of choice, for this blog I'm going to use IDA Pro. If we open up the unstripped binary in IDA Pro we can notice a few quick things;

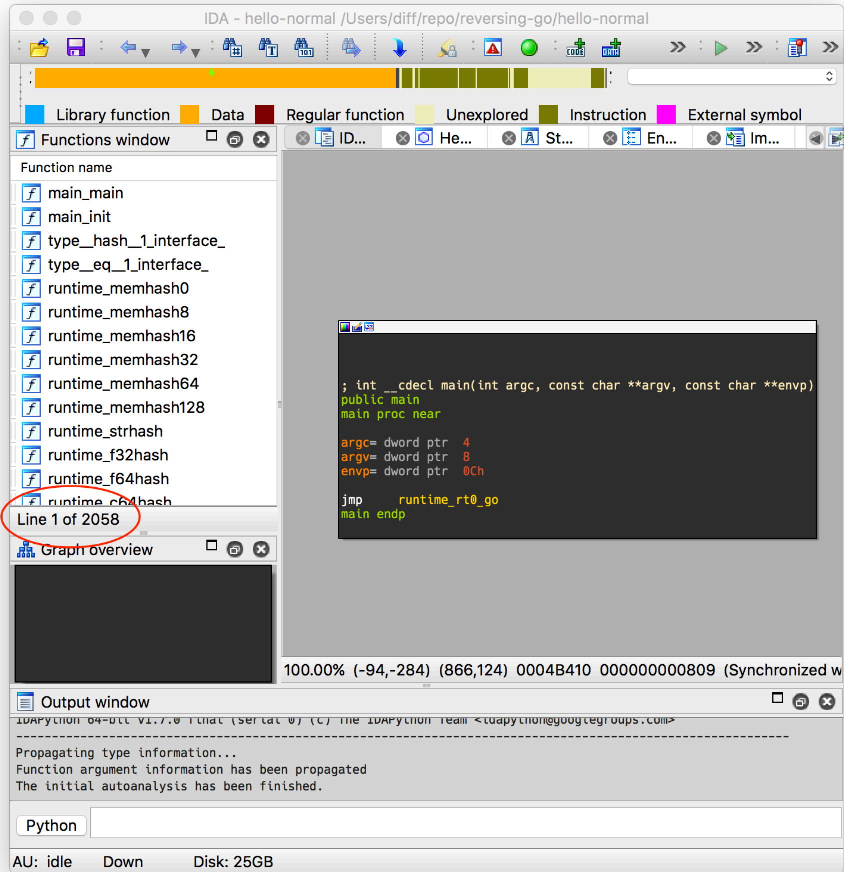

Well then - our 5 lines of code has turned into over 2058 functions. With all that overhead of what appears to be a runtime, we also have nothing interesting in the `main()` function. If we dig in a bit further we can see that the actual code we're interested in is inside of `main_main` ;

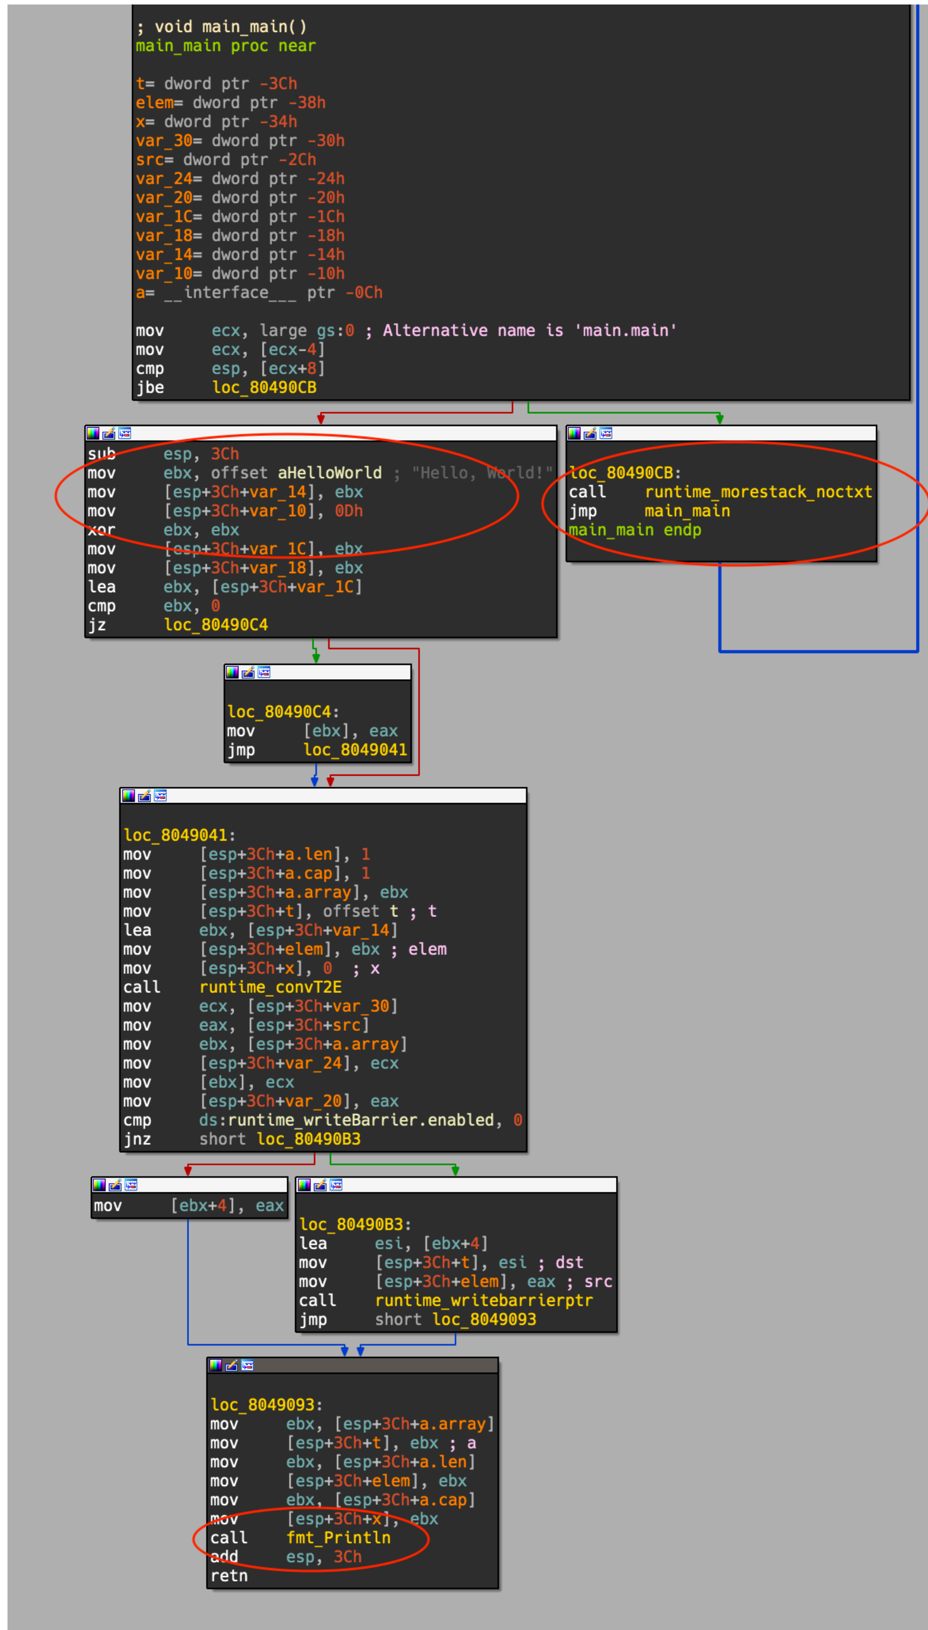

This is, well, lots of code that I honestly don't want to look at. The string loading also looks a bit weird - though IDA seems to have done a good job identifying the necessary bits. We can easily see that the string load is actually a set of three `mov` s;

String load

```
1    mov     ebx, offset aHelloWorld ; "Hello, World!"
2    mov     [esp+3Ch+var_14], ebx ; Shove string into location
```

```
3    mov     [esp+3Ch+var_10], 0Dh ; length of string
```

This isn't exactly revolutionary, though I can't off the top of my head say that I've seen something like this before. We're also taking note of it as this will come in handle later on. The other tidbit of code which caught my eye was the `runtime_morestack_context` call;

morestack_context

```
1    loc_80490CB:
2    call    runtime_morestack_noctxt
3    jmp     main_main
```

This style block of code appears to always be at the end of functions and it also seems to always loop back up to the top of the same function. This is verified by looking at the cross-references to this function. Ok, now that we know IDA Pro can handle unstripped binaries, lets load the same code but the stripped version this time.

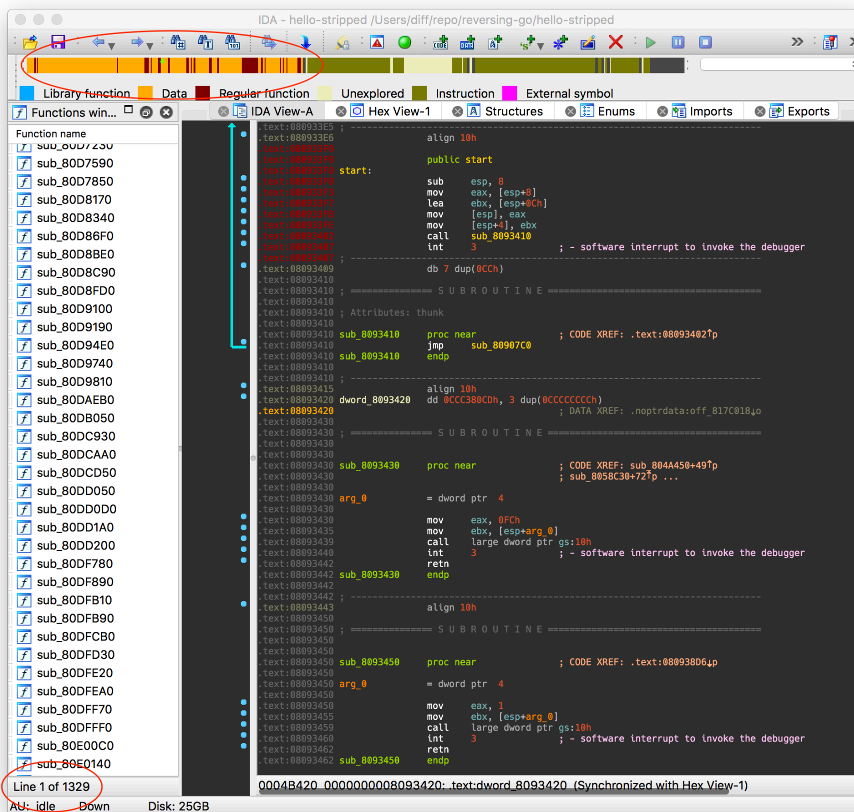

Immediately we see some, well, lets just call them "differences". We have 1329 functions defined and now see some undefined code by looking at the navigator toolbar. Luckily IDA has still been able to find the string load we are looking for, however this function now seems much less friendly to deal with.

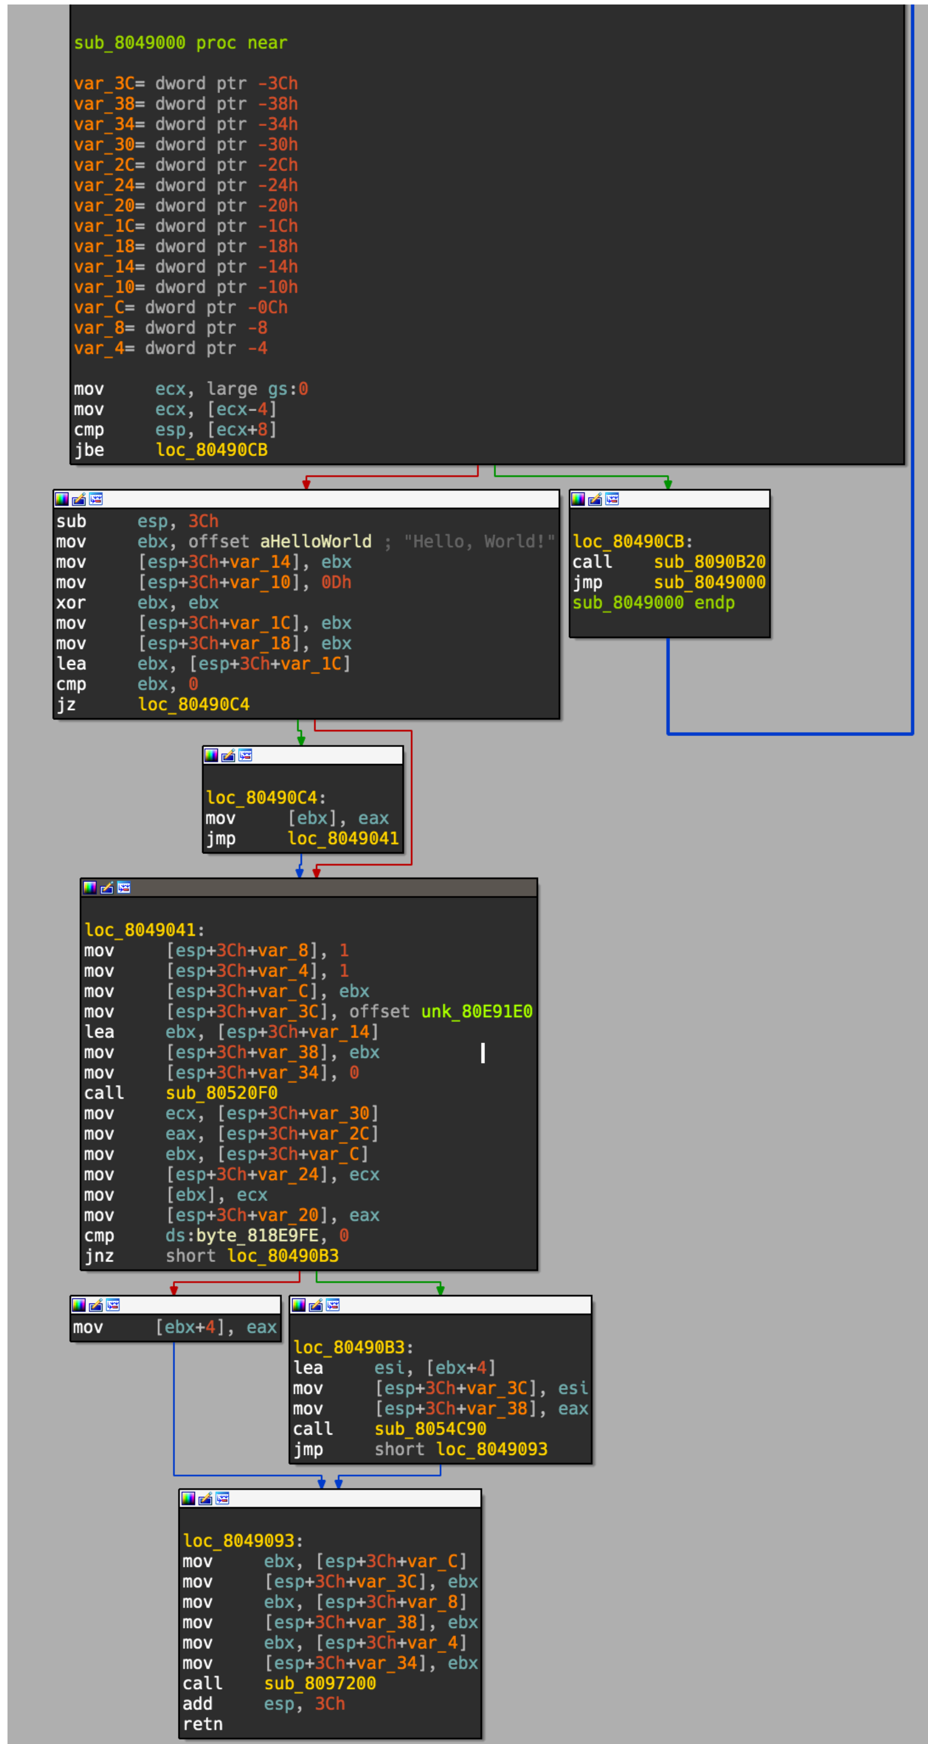

We now have no more function names, however - the function names appear to be retained in a specific section of the binary if we do a string search for `main.main` (which would be repesented at `main_main` in the previous screen shots due to how a `.` is interpreted by IDA);

.gopclntab

```
1  .gopclntab:0813E174                db  6Dh ; m
2  .gopclntab:0813E175                db  61h ; a
3  .gopclntab:0813E176                db  69h ; i
4  .gopclntab:0813E177                db  6Eh ; n
5  .gopclntab:0813E178                db  2Eh ; .
6  .gopclntab:0813E179                db  6Dh ; m
7  .gopclntab:0813E17A                db  61h ; a
8  .gopclntab:0813E17B                db  69h ; i
9  .gopclntab:0813E17C                db  6Eh ; n
```

Alright, so it would appear that there is something left over here. After digging into some of the Google results into `gopclntab` and tweet about this - a friendly reverser [George (Egor?) Zaytsev] showed me his IDA Pro scripts for [renaming function and adding type information]. After skimming these it was pretty easy to figure out the format of this section so I threw together some functionally to replicate his script. The essential code is shown below, very simply put, we look into the segment `.gopclntab` and skip the first 8 bytes. We then create a pointer (`Qword` or `Dword` dependant on whether the binary is 64bit or not). The first set of data actually gives us the size of the `.gopclntab` table, so we know how far to go into this structure. Now we can start processing the rest of the data which appears to be the `function_offset` followed by the (function) `name_offset`). As we create pointers to these offsets and also tell IDA to create the strings, we just need to ensure we don't pass `MakeString` any bad characters so we use the `clean_function_name` function to strip out any badness.

renamer.py

```
1   def create_pointer(addr, force_size=None):
2   if force_size is not 4 and (idaapi.get_inf_structure().is_64bit() or force_size is 8):
3           MakeQword(addr)
4   return Qword(addr), 8
5   else:
6           MakeDword(addr)
7   return Dword(addr), 4
8   STRIP_CHARS = [ '(', ')', '[', ']', '{', '}', ' ', '"' ]
9   REPLACE_CHARS = ['.', '*', '-', ',', ';', ':', '/', '\xb7' ]
10  def clean_function_name(str):
11      str = filter(lambda x: x in string.printable, str)
12  for c in STRIP_CHARS:
13          str = str.replace(c, '')
14  for c in REPLACE_CHARS:
15          str = str.replace(c, '_')
16  return str
17  def renamer_init():
18      renamed = 0
19      gopclntab = ida_segment.get_segm_by_name('.gopclntab')
20  if gopclntab is not None:
21          addr = gopclntab.startEA + 8
22          size, addr_size = create_pointer(addr)
```

```
        addr += addr_size
        early_end = addr + (size * addr_size * 2)
while addr < early_end:
            func_offset, addr_size = create_pointer(addr)
            name_offset, addr_size = create_pointer(addr + addr_size)
            addr += addr_size * 2
            func_name_addr = Dword(name_offset + gopclntab.startEA + addr_size) + gopclntab.startEA
            func_name = GetString(func_name_addr)
            MakeStr(func_name_addr, func_name_addr + len(func_name))
            appended = clean_func_name = clean_function_name(func_name)
            debug('Going to remap function at 0x%x with %s - cleaned up as %s' % (func_offset, func_name, clean_func_n
if ida_funcs.get_func_name(func_offset) is not None:
if MakeName(func_offset, clean_func_name):
                    renamed += 1
else:
                    error('clean_func_name error %s' % clean_func_name)
return renamed
def main():
    renamed = renamer_init()
    info('Found and successfully renamed %d functions!' % renamed)
```

The above code won't actually run yet (don't worry full code available in [this repo](#)) but it is hopefully simple enough to read through and understand the process. However, this still doesn't solve the problem that IDA Pro doesn't know *all* the functions. So this is going to create pointers which aren't being referenced anywhere. We do know the beginning of functions now, however I ended up seeing (what I think is) an easier way to define all the functions in the application. We can define all the functions by utilizing `runtime_morestack_noctxt` function. Since every function utilizes this (basically, there is an edgecase it turns out), if we find this function and traverse backwards to the cross references to this function, then we will know where every function exists. So what, right? We already know where every function started from the segment we just parsed above, right? Ah, well - now we know the end of the function *and* the next instruction after the call to `runtime_morestack_noctxt` gives us a jump to the top of the function. This means we should quickly be able to give the bounds of the start and stop of a function, which is required by IDA, while seperating this from the parsing of the function names. If we open up the window for cross references to the function `runtime_morestack_noctxt` we see there are many

more undefined sections calling into this. 1774 in total things reference this function, which is up from the 1329 functions IDA has already defined for us, this is highlighted by the image below;

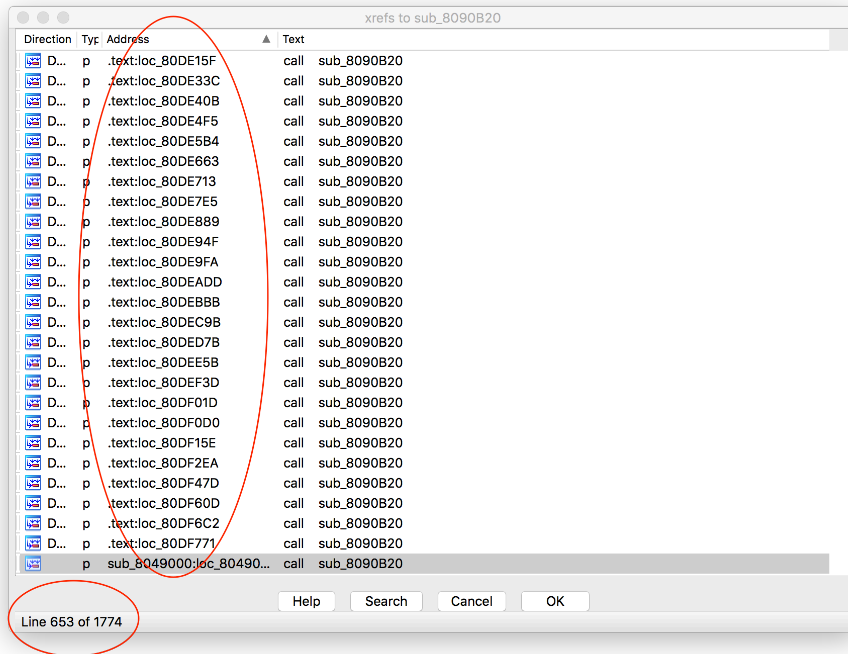

After digging into mutliple binaries we can see the `runtime_morestack_noctext` will always call into `runtime_morestack` (with context). This is the edgecase I was referencing before, so between these two functions we should be able to see cross refereneces to ever other function used in the binary. Looking at the larger of the two functions, `runtime_more_stack`, of multiple binaries tends to have an interesting layout;

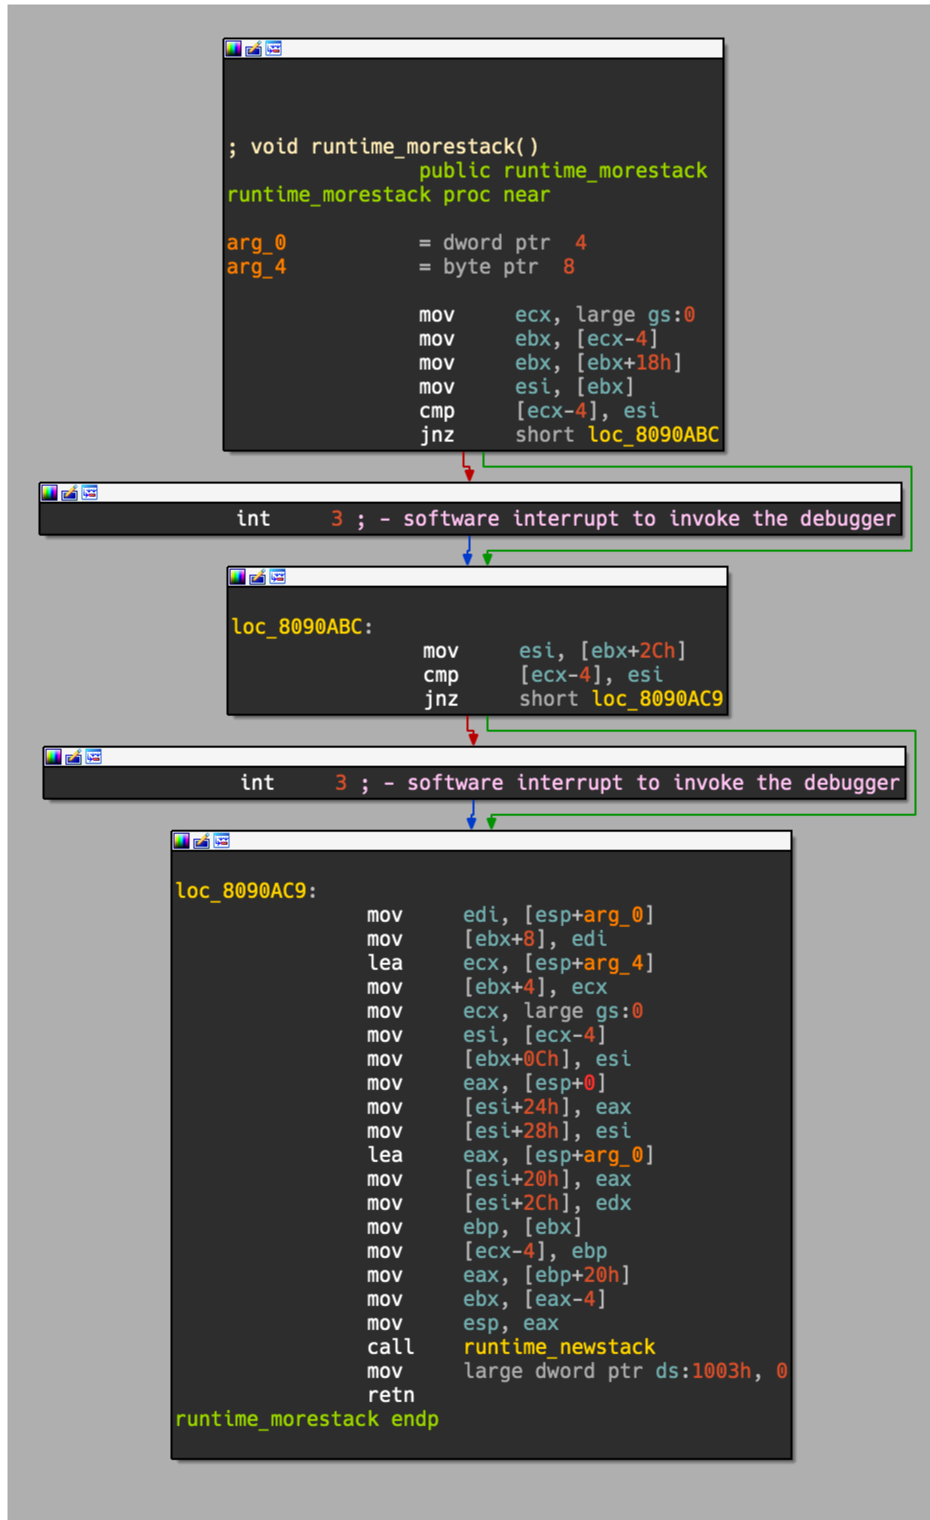

The part which stuck out to me was `mov large dword ptr ds:1003h, 0` - this appeared to be rather constant in all 64bit binaries I saw. So after cross compiling a few more I noticed that 32bit binaries used `mov qword ptr ds:1003h, 0`, so we will be hunting for this pattern to create a "hook" for traversing backwards on. Lucky for us, I haven't seen an instance where IDA Pro fails to define this specific function, we don't really need to spend much brain power mapping it out or defining it outselves. So, enough talk, lets write some code to find this function;

find_runtime_morestack.py

```
1   def create_runtime_ms():
2       debug('Attempting to find runtime_morestack function for hooking on...')
```

```
    text_seg = ida_segment.get_segm_by_name('.text')
    runtime_ms_end = ida_search.find_text(text_seg.startEA, 0, 0, "word ptr ds:1003h, 0", SEARCH_DOWN)
    runtime_ms = ida_funcs.get_func(runtime_ms_end)
if idc.MakeNameEx(runtime_ms.startEA, "runtime_morecontext", SN_PUBLIC):
        debug('Successfully found runtime_morecontext')
else:
        debug('Failed to rename function @ 0x%x to runtime_morestack' % runtime_ms.startEA)
return runtime_ms
```

After finding the function, we can recursively traverse backwards through all the function calls, anything which is not inside an already defined function we can now define. This is because the structure always appears to be;

golang_undefined_function_example

```
.text:08089910                                          ; Function start - however undefined currently according to IDA
.text:08089910 loc_8089910:                             ; CODE XREF: .text:0808994B
.text:08089910                                          ; DATA XREF: sub_804B250+1A1
.text:08089910                mov     ecx, large gs:0
.text:08089917                mov     ecx, [ecx-4]
.text:0808991D                cmp     esp, [ecx+8]
.text:08089920                jbe     short loc_8089946
.text:08089922                sub     esp, 4
.text:08089925                mov     ebx, [edx+4]
.text:08089928                mov     [esp], ebx
.text:0808992B                cmp     dword ptr [esp], 0
.text:0808992F                jz      short loc_808993E
.text:08089931
.text:08089931 loc_8089931:                             ; CODE XREF: .text:08089944
.text:08089931                add     dword ptr [esp], 30h
.text:08089935                call    sub_8052CB0
.text:0808993A                add     esp, 4
.text:0808993D                retn
.text:0808993E ; ---------------------------------------------------------------------------
.text:0808993E
.text:0808993E loc_808993E:                             ; CODE XREF: .text:0808992F
.text:0808993E                mov     large ds:0, eax
.text:08089944                jmp     short loc_8089931
.text:08089946 ; ---------------------------------------------------------------------------
.text:08089946
.text:08089946 loc_8089946:                             ; CODE XREF: .text:08089920
.text:08089946                call    runtime_morestack ; "Bottom" of function, calls out to runtime_morestack
```

```
28  .text:0808994B                 jmp     short loc_8089910 ; Jump back to the "top" of the function
```

The above snippet is a random undefined function I pulled from the stripped example application we compiled already. Essentially by traversing backwards into every undefined function, we will land at something like line `0x0808994B` which is the `call runtime_morestack`. From here we will skip to the next instruction and ensure it is a jump above where we currently are, if this is true, we can likely assume this is the start of a function. In this example (and almost every test case I've run) this is true. Jumping to `0x08089910` is the start of the function, so now we have the two parameters required by `MakeFunction` function;

traverse_functions.py

```
1   def is_simple_wrapper(addr):
2   if GetMnem(addr) == 'xor' and GetOpnd(addr, 0) == 'edx' and  GetOpnd(addr, 1) == 'edx':
3           addr = FindCode(addr, SEARCH_DOWN)
4   if GetMnem(addr) == 'jmp' and GetOpnd(addr, 0) == 'runtime_morestack':
5   return True
6   return False
7   def create_runtime_ms():
8       debug('Attempting to find runtime_morestack function for hooking on...')
9       text_seg = ida_segment.get_segm_by_name('.text')
10      runtime_ms_end = ida_search.find_text(text_seg.startEA, 0, 0, "word ptr ds:1003h, 0", SEARCH_DOWN)
11      runtime_ms = ida_funcs.get_func(runtime_ms_end)
12  if idc.MakeNameEx(runtime_ms.startEA, "runtime_morestack", SN_PUBLIC):
13          debug('Successfully found runtime_morestack')
14  else:
15          debug('Failed to rename function @ 0x%x to runtime_morestack' % runtime_ms.startEA)
16  return runtime_ms
17  def traverse_xrefs(func):
18      func_created = 0
19  if func is None:
20  return func_created
21      func_xref = ida_xref.get_first_cref_to(func.startEA)
22  while func_xref != 0xffffffffffffffff:
23  if ida_funcs.get_func(func_xref) is None:
24              func_end = FindCode(func_xref, SEARCH_DOWN)
25  if GetMnem(func_end) == "jmp":
26                  func_start = GetOperandValue(func_end, 0)
27  if func_start < func_xref:
28  if idc.MakeFunction(func_start, func_end):
29                          func_created += 1
30  else:
31                          error('Error trying to create a function @ 0x%x - 0x%x' %(func_start, func_end))
32  else:
33              xref_func = ida_funcs.get_func(func_xref)
34  if is_simple_wrapper(xref_func.startEA):
35                  debug('Stepping into a simple wrapper')
```

```
36                  func_created += traverse_xrefs(xref_func)
37  if ida_funcs.get_func_name(xref_func.startEA) is not None and 'sub_' not in ida_funcs.get_func_name(xref_func.startEA)
38                  debug('Function @0x%x already has a name of %s; skipping...' % (func_xref, ida_funcs.get_func_name(xre
39  else:
40                  debug('Function @ 0x%x already has a name %s' % (xref_func.startEA, ida_funcs.get_func_name(xref_func.
41          func_xref = ida_xref.get_next_cref_to(func.startEA, func_xref)
42  return func_created
43  def find_func_by_name(name):
44      text_seg = ida_segment.get_segm_by_name('.text')
45  for addr in Functions(text_seg.startEA, text_seg.endEA):
46  if name == ida_funcs.get_func_name(addr):
47  return ida_funcs.get_func(addr)
48  return None
49  def runtime_init():
50      func_created = 0
51  if find_func_by_name('runtime_morestack') is not None:
52          func_created += traverse_xrefs(find_func_by_name('runtime_morestack'))
53          func_created += traverse_xrefs(find_func_by_name('runtime_morestack_noctxt'))
54  else:
55          runtime_ms = create_runtime_ms()
56          func_created = traverse_xrefs(runtime_ms)
57  return func_created
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
```

```
80
81
82
```

That code bit is a bit lengthy, though hopefully the comments and concept is clear enough. It likely isn't necessary to explicitly traverse backwards recursively, however I wrote this prior to understanding that `runtime_morestack_noctxt` (the edgecase) is the only edgecase that I would encounter. This was being handled by the `is_simple_wrapper` function originally. Regardless, running this style of code ended up finding all the extra functions IDA Pro was missing. We can see below, that this creates a much cleaner and easier experience to reverse;

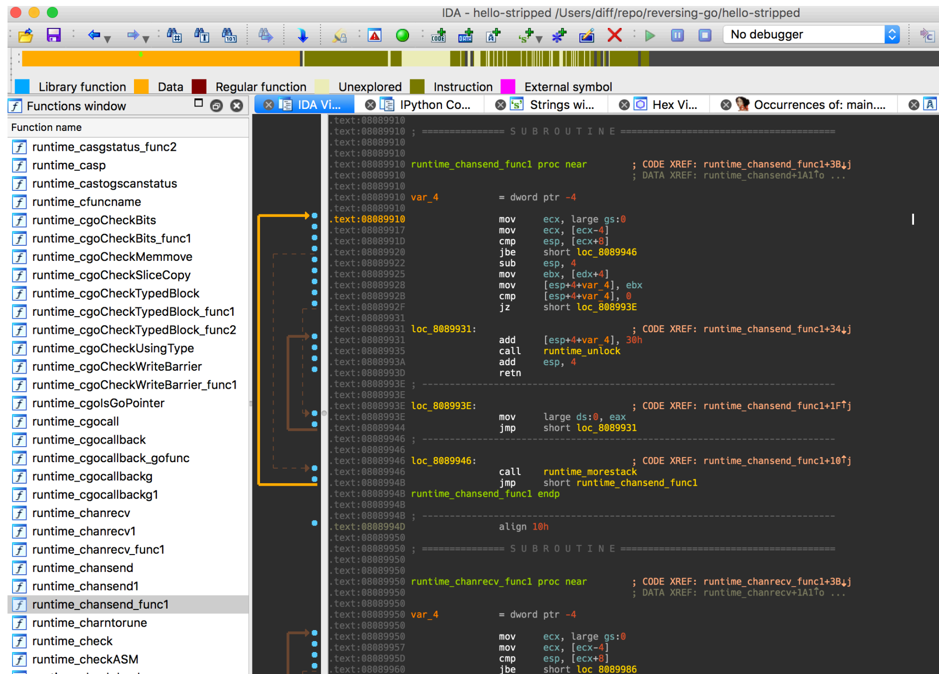

This can allow us to use something like [Diaphora] as well since we can specifically target functions with the same names, if we care too. I've personally found this is extremely useful for malware or other targets where you *really* don't care about any of the framework/runtime functions. You can quiet easily differentiate between custom code written for the binary, for example in the Linux malware "Rex" everything because with that name space! Now onto the last challenge that I wanted to solve while reversing the malware, string loading! I'm honestly not 100% sure how IDA detects most string loads, potentially through idioms of some sort? Or maybe because it can detect strings based on the `\00` character at the end of it? Regardless, Go seems to use a string table of some sort, without requiring null character. The appear to be in alpha-numeric order, group by string length size as well. This means we see them all there, but often don't come across them correctly asserted as strings, or we see them asserted as extremely large blobs of strings. The hello world example isn't good at illistrating this, so I'll pull open the `main.main` function of the Rex malware to show this;

```
loc_80494D8:            ;
mov     ebx, offset unk_8600920 ; pointer to a string (undefined currently)
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 5 ; string length
mov     byte ptr [esp+0F0h+var_E8], 0
mov     ebx, 860AB34h    ; constant... though this is actually pointing to a string as well
mov     dword ptr [esp+0F0h+var_E8+4], ebx
mov     [esp+0F0h+var_E0], 10h ; string length
call    flag_Bool
mov     ebx, [esp+0F0h+var_DC]
mov     [esp+0F0h+var_90], ebx
mov     ebx, offset unk_86001AD
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 4
mov     dword ptr [esp+0F0h+var_E8], 0
mov     ebx, 861DC4Ch
mov     dword ptr [esp+0F0h+var_E8+4], ebx
mov     [esp+0F0h+var_E0], 31h
call    flag_Int
mov     ebx, [esp+0F0h+var_DC]
mov     [esp+0F0h+var_B8], ebx
mov     ebx, 8602175h
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 6
mov     ebx, offset unk_8604841
mov     dword ptr [esp+0F0h+var_E8], ebx
mov     dword ptr [esp+0F0h+var_E8+4], 9
mov     ebx, offset unk_860551F
mov     [esp+0F0h+var_E0], ebx
mov     [esp+0F0h+var_DC], 9
call    flag_String
mov     ebx, [esp+0F0h+var_D8]
mov     [esp+0F0h+var_B4], ebx
mov     ebx, offset unk_860456A
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 8
mov     ebx, 8601F23h
mov     dword ptr [esp+0F0h+var_E8], ebx
mov     dword ptr [esp+0F0h+var_E8+4], 6
mov     ebx, 8617547h
mov     [esp+0F0h+var_E0], ebx
mov     [esp+0F0h+var_DC], 22h
```

I didn't want to add comments to everything, so I only commented the first few lines then pointed arrows to where there should be pointers to a proper string. We can see a few different use cases and sometimes the destination registers seem to change. However there is definitely a pattern which forms that we can look for. Moving of a pointer into a register, that register is then used to push into a (d)word pointer, followed by a load of a lenght of the string. Cobbling together some python to hunt for the pattern we end with something like the pseudo code below;

string_hunting.py

```
1   VALID_REGS = ['ebx', 'ebp']
2   VALID_DEST = ['esp', 'eax', 'ecx', 'edx']
3   def is_string_load(addr):
4       patterns = []
5   if GetMnem(addr) == 'mov':
6   if GetOpnd(addr, 0) in VALID_REGS and not ('[' in GetOpnd(addr, 1) or 'loc_' in GetOpnd(addr, 1)) and('offset ' in Get
7               from_reg = GetOpnd(addr, 0)
8               addr_2 = FindCode(addr, SEARCH_DOWN)
9   try:
10                  dest_reg = GetOpnd(addr_2, 0)[GetOpnd(addr_2, 0).index('[') + 1:GetOpnd(addr_2, 0).index('[') + 4]
11  except ValueError:
12  return False
13  if GetMnem(addr_2) == 'mov' and dest_reg in VALID_DEST and ('[%s' % dest_reg) in GetOpnd(addr_2, 0) and GetOpnd(addr_2
14                  addr_3 = FindCode(addr_2, SEARCH_DOWN)
15  if GetMnem(addr_3) == 'mov' and (('[%s+' % dest_reg) in GetOpnd(addr_3, 0) or GetOpnd(addr_3, 0) in VALID_DEST) and 'o
16  try:
17                          dumb_int_test = GetOperandValue(addr_3, 1)
```

```
18  if dumb_int_test > 0 and dumb_int_test < sys.maxsize:
19  return True
20  except ValueError:
21  return False
22  def create_string(addr, string_len):
23      debug('Found string load @ 0x%x with length of %d' % (addr, string_len))
24  if GetStringType(addr) is not None and GetString(addr) is not None and len(GetString(addr)) != string_len:
25          debug('It appears that there is already a string present @ 0x%x' % addr)
26          MakeUnknown(addr, string_len, DOUNK_SIMPLE)
27  if GetString(addr) is None and MakeStr(addr, addr + string_len):
28  return True
29  else:
30          MakeUnknown(addr, string_len, DOUNK_SIMPLE)
31  if MakeStr(addr, addr + string_len):
32  return True
33          debug('Unable to make a string @ 0x%x with length of %d' % (addr, string_len))
34  return False
35
36
37
38
39
40
41
42
43
44
45
46
47
```

The above code could likely be optimized, however it was working for me on the samples I needed. All that would be left is to create another function which hunts through all the defined code segments to look for string loads. Then we can use the pointer to the string and the string length to define a new string using the `MakeStr`. In the code I ended up using, you need to ensure that IDA Pro hasn’t mistakenly already create the string, as it sometimes tries to, incorrectly. This seems to happen sometimes when a string in the table contains a null character. However, after using code above, this is what we are left with;

```
loc_80494D8:              ; "debug"
mov     ebx, offset aDebug
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 5
mov     byte ptr [esp+0F0h+var_E8], 0
mov     ebx, offset aEnableDebuggin ; "enable debugging"
mov     dword ptr [esp+0F0h+var_E8+4], ebx
mov     [esp+0F0h+var_E0], 10h
call    flag_Bool
mov     ebx, [esp+0F0h+var_DC]
mov     [esp+0F0h+var_90], ebx
mov     ebx, offset aWait ; "wait"
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 4
mov     dword ptr [esp+0F0h+var_E8], 0
mov     ebx, offset aWaitForPidToEx ; "wait for PID to exit before starting (0"...
mov     dword ptr [esp+0F0h+var_E8+4], ebx
mov     [esp+0F0h+var_E0], 31h
call    flag_Int
mov     ebx, [esp+0F0h+var_DC]
mov     [esp+0F0h+var_B8], ebx
mov     ebx, offset aTarget ; "target"
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 6
mov     ebx, offset a0_0_0_00 ; "0.0.0.0/0"
mov     dword ptr [esp+0F0h+var_E8], ebx
mov     dword ptr [esp+0F0h+var_E8+4], 9
mov     ebx, offset aTargetS ; "target(s)"
mov     [esp+0F0h+var_E0], ebx
mov     [esp+0F0h+var_DC], 9
call    flag_String
mov     ebx, [esp+0F0h+var_D8]
mov     [esp+0F0h+var_B4], ebx
mov     ebx, offset aStrategy ; "strategy"
mov     [esp+0F0h+var_F0], ebx
mov     [esp+0F0h+var_EC], 8
mov     ebx, offset aRandom ; "random"
mov     dword ptr [esp+0F0h+var_E8], ebx
mov     dword ptr [esp+0F0h+var_E8+4], 6
mov     ebx, offset aScanStrategyRa ; "scan strategy [random, sequential]"
mov     [esp+0F0h+var_E0], ebx
mov     [esp+0F0h+var_DC], 22h
```

This is a much better piece of code to work with. After we throw together all these functions, we now have the golang_loader_assist.py module for IDA Pro. A word of warning though, I have only had time to test this on a few versions of IDA Pro for OSX, the majority of testing on 6.95. There is also very likely optimizations which should be made or at a bare minimum some reworking of the code. With all that said, I wanted to open source this so others could use this and hopefully contribute back. Also be aware that this script can be painfully slow depending on how large the `idb` file is, working on a OSX El Capitan (10.11.6) using a 2.2 GHz Intel Core i7 on IDA Pro 6.95 - the string discovery aspect itself can take a while. I've often found that running the different methods seperately can prevent IDA from locking up. Hopefully this blog and the code proves useful to someone though, enjoy!

Source: https://rednaga.io/2016/09/21/reversing_go_binaries_like_a_pro/