

# About GCC printf optimization

Peter Seiderer

November 21, 2005

## Abstract

As you may have noticed, your GCC compiled source code is not always calling the functions you wrote in the source code. This can easily be observed with `printf` in the *hello world* example. Starting from this some more GCC optimizations will be examined and analyzed in detail.

## 1 Introduction

For a start we examine what the *GCC* compiler [GCCa] will do with the *hello world* example when different optimization level are used.

```
1 #include <stdio.h>
2 int main(int argc, char *argv[]) {
3     printf("hello world\n");
4     return 0;
5 }
```

Figure 1: example1.c

Take a look at the `nm`<sup>1</sup> output when compiled without optimization.

```
> gcc-3.3.6 -O0 example1.c
> nm --undefined-only a.out
     w __gmon_start__
     w _Jv_RegisterClasses
     U __libc_start_main@GLIBC_2.0
     U printf@GLIBC_2.0
```

As expected the dynamically linked executable *a.out* will call the function *printf*. Now we compile the same source code with optimization on.

```
> gcc-3.3.6 -O1 example1.c
> nm --undefined-only a.out
     w __gmon_start__
```

---

<sup>1</sup>`nm` - list symbols from object files

```
w _Jv_RegisterClasses
U __libc_start_main@GLIBC_2.0
U puts@GLIBC_2.0
```

The *nm* output indicates that *puts* is called instead of the original *printf* from the source code. To verify this we take a look at the intermediate assembler listing. First with optimization off.

```
> gcc-3.3.6 -O0 -S example1.c
```

Here the resulting *example1.s* file:

```
1      .file   "example1.c"
2      .section .rodata
3      .LC0:
4      .string "hello world\n"
5      .text
6      .globl main
7      .type   main, @function
8      main:
9      pushl  %ebp
10     movl   %esp, %ebp
11     subl   $8, %esp
12     andl   $-16, %esp
13     movl   $0, %eax
14     subl   %eax, %esp
15     movl   $.LC0, (%esp)
16     call  printf
17     movl   $0, %eax
18     leave
19     ret
20     .size  main, .-main
21     .ident "GCC: (GNU) 3.3.6"
```

Second with optimization on.

```
> gcc-3.3.6 -O1 -S example1.c
```

And the resulting *example1.s* file:

```
1      .file   "example1.c"
2      .section .rodata
3      .LC0:
4      .string "hello world"
5      .text
6      .globl main
7      .type   main, @function
8      main:
9      pushl  %ebp
```

```

10      movl    %esp, %ebp
11      subl   $8, %esp
12      andl   $-16, %esp
13      movl   $.LC0, (%esp)
14      call   puts
15      movl   $0, %eax
16      movl   %ebp, %esp
17      popl   %ebp
18      ret
19      .size   main, .-main
20      .ident  "GCC: (GNU) 3.3.6"

```

The string constant in line 4 changed from `"hello world\n"` to `"hello world"`. The `call printf` on line 16 changed to `call puts` on line 14.

## 2 Optimization according to the source code

Where is the optimization implemented in the GCC source code:

- GCC-3.3.6: `gcc/c-common.c` line 1324
- GCC-3.4.4: `gcc/builtins.c` line 4651
- GCC-4.0.2: `gcc/builtins.c:4654`.

Step by step guide along the source code lines (for GCC-3.4.4). In the following examples we will omit the first two lines

```

w __gmon_start__
w _Jv_RegisterClasses

```

of the `nm` output.

### 2.1 No optimization if return code is used

No optimization if the return code of `printf` is evaluated [line 4662-4664].

```

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      return printf("hello world\n");
4  }

```

Figure 2: `example2.c`

```

> gcc-3.4.4 -O1 example2.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0

```

## 2.2 Format string must be a literal constant

The format string must be a literal string constant [line 4666-4677].

```
1 #include <stdio.h>
2 static const char *f() {
3     return "hello world\n";
4 }
5 int main(int argc, char *argv[]) {
6     printf(f());
7     return 0;
8 }
```

Figure 3: example3.c

```
> gcc-3.4.4 -O1 example3.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0
```

Note that the GCC-4.0.2 is smart enough to optimize this example if an optimization level greater or even  $-O2$  is given.

```
> gcc-4.0.2 -Wall -O2 example3.c
nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U puts@@GLIBC_2.0
```

## 2.3 "%s\n"

A *printf* call with the format string "%s\n" [line 4679-4687] is converted to a *puts*() call.

```
printf("%s\n", "hello world"); // converted to puts("hello world");
```

```
1 #include <stdio.h>
2 int main(int argc, char *argv[]) {
3     printf("%s\n", "hello world");
4     return 0;
5 }
```

Figure 4: example4.c

```
> gcc-3.4.4 -O1 example4.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U puts@@GLIBC_2.0
```

## 2.4 "%c"

A `printf` call with the format string "%c" [line 4688-4696] is converted to a `putchar()` call.

```
printf("%c", 'A'); // converted to putchar('A');

1 #include <stdio.h>
2 int main(int argc, char *argv[]) {
3     printf("%c", 'A');
4     return 0;
5 }
```

Figure 5: example5.c

```
> gcc-3.4.4 -O1 example5.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U putchar@@GLIBC_2.0
```

## 2.5 No '%' in format string

No optimization if one or more '%' are detected in the format string [line 4699-4701].

```
1 #include <stdio.h>
2 int main(int argc, char *argv[]) {
3     printf("hello%%world\n");
4     return 0;
5 }
```

Figure 6: example6.c

```
> gcc-3.4.4 -O1 example6.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0
```

## 2.6 The empty format string

Omit `printf` call if format string is empty [line 4706-4708].

```
printf(""); // converted to empty statement

> gcc-3.4.4 -O1 example7.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
```

```

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      printf("");
4      return 0;
5  }

```

Figure 7: example7.c

## 2.7 The one character string

A `printf` call with a format string of length one [line 4709-4718] is converted to `putchar()` call.

```
printf("A"); // converted to putchar('A');
```

```

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      printf("A");
4      return 0;
5  }

```

Figure 8: example8.c

```

> gcc-3.4.4 -O1 example8.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U putchar@@GLIBC_2.0

```

## 2.8 The string not ending with '\n'

No optimization if the string is not ending with '\n' [line 4721-4739].

```

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      printf("hello world");
4      return 0;
5  }

```

Figure 9: example9.c

```

> gcc-3.4.4 -O1 example9.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0

```

## 2.9 The string ending with '\n'

A `printf` call with a simple format string ending with '\n' [line 4721-4739] is converted to a `puts()` call.

```
printf("hello world\n"); // converted to puts("hello world");

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      printf("hello world\n");
4      return 0;
5  }
```

Figure 10: example10.c

```
> gcc-3.4.4 -O1 example10.c
> nm --undefined-only a.out
     U __libc_start_main@@GLIBC_2.0
     U puts@@GLIBC_2.0
```

## 3 Fine differences

### 3.1 %s and NULL argument

Example taken from [BUG]. The glibc [GLIa] `printf` implementation has the nice feature to print '(null)' in case of the format string '%s' and a pointer to NULL is given as argument.

```
1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      char *p = NULL;
4      printf("%s\n", p);
5      return 0;
6  }
```

Figure 11: bug\_15685.c

```
> gcc-3.3.6 -Wall -O0 bug_15685.c
./a.out
(null)

> gcc-3.3.6 -Wall -O1 bug_15685.c
./a.out
Segmentation fault
```

As stated in the bug report by Andrew Pinski this is not a bug, because both behaviors are allowed by the standard. If the standard states *undefined behavior* the compiler and/or library is free in what to implement. You should not depend on the one or the other behavior in a portable program.

From the glibc manual: “If you accidentally pass a null pointer as the argument for a %s conversion, the GNU library prints it as (null). We think this is more useful than crashing. But it’s not good practice to pass a null argument intentionally.” [GLib].

## 4 Comparison of GCC versions

In figure 12 note that the GCC-4.0.2 is doing this optimization even if *-O0* (no optimization) is given on the command line.

compiler	-O0	-O1	-O2	-O3
GCC-3.3.6	printf	puts	puts	puts
GCC-3.4.4	printf	puts	puts	puts
GCC-4.0.2	puts	puts	puts	puts

Figure 12: Results for *example1.c*

To get the original *printf* call with GCC-4.0.2 you must provide *-fno-builtin* or *-fno-builtin-printf* as command line argument.

```
> gcc-4.0.2 -fno-builtin -O0 example1.c
> nm --undefined-only a.out
     w __gmon_start__
     w _Jv_RegisterClasses
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0

> gcc-4.0.2 -fno-builtin-printf -O0 example1.c
> nm --undefined-only a.out
     w __gmon_start__
     w _Jv_RegisterClasses
     U __libc_start_main@@GLIBC_2.0
     U printf@@GLIBC_2.0
```

## 5 Performance advantage

Some simple measurement<sup>2</sup> for the hello world example.

<sup>2</sup>On AMD Athlon(TM) XP 1800+ system



```

1  #include <stdio.h>
2  int main(int argc, char *argv[]) {
3      int i;
4      for(i = 0; i < 10000000; i++) {
5          printf("hello world\n");
6      }
7      return 0;
8  }

```

Figure 13: loop\_printf.c

```

> gcc-3.3.6 -Wall -O0 loop_printf.c
> time ./a.out > /dev/null
real    0m2.197s
user    0m2.186s
sys     0m0.010s
> gcc-3.3.6 -Wall -O1 loop_printf.c
> time ./a.out > /dev/null
real    0m1.007s
user    0m0.997s
sys     0m0.009s

```

In figure 14 we see that the *puts* loop consumes roughly half the time of the *printf* loop.

	GCC-3.3.6	GCC-3.4.4	GCC-4.0.2
-O0	real 0m2.197s user 0m2.186s sys 0m0.010s	real 0m2.164s user 0m2.147s sys 0m0.015s	real 0m1.026s user 0m1.012s sys 0m0.014s
-O1	real 0m1.007s user 0m0.997s sys 0m0.009s	real 0m1.006s user 0m0.996s sys 0m0.011s	real 0m1.007s user 0m0.998s sys 0m0.010s

Figure 14: Times for *loop\_printf.c*

## 6 Future work

Document the remaining optimizations mentioned in *gcc/builtins.c*. Explore advanced optimization like suggested in [GCCb].

## References

[BUG] [http://gcc.gnu.org/bugzilla/show\\_bug.cgi?id=15685](http://gcc.gnu.org/bugzilla/show_bug.cgi?id=15685).

[GCCa] <http://gcc.gnu.org>.

[GCCb] <http://gcc.gnu.org/ml/gcc-patches/1999-01/msg00187.html>.

[GLIa] <http://www.gnu.org/software/libc/libc.html>.

[GLIb] [http://www.gnu.org/software/libc/manual/html\\_mono/libc.html.gz#Other%20Output%20Conve](http://www.gnu.org/software/libc/manual/html_mono/libc.html.gz#Other%20Output%20Conve)

## 7 Appendix

### 7.1 Source code *gcc-3.4.4/gcc/builtins.c*

```
4645  /* Expand a call to printf or printf_unlocked with argument list ARGLIST.
4646     Return 0 if a normal call should be emitted rather than transforming
4647     the function inline.  If convenient, the result should be placed in
4648     TARGET with mode MODE.  UNLOCKED indicates this is a printf_unlocked
4649     call.  */
4650  static rtx
4651  expand_builtin_printf (tree arglist, rtx target, enum machine_mode mode,
4652                       bool unlocked)
4653  {
4654     tree fn_putchar = unlocked
4655                 ? implicit_builtin_decls[BUILT_IN_PUTCHAR_UNLOCKED]
4656                 : implicit_builtin_decls[BUILT_IN_PUTCHAR];
4657     tree fn_puts = unlocked ? implicit_builtin_decls[BUILT_IN_PUTS_UNLOCKED]
4658                           : implicit_builtin_decls[BUILT_IN_PUTS];
4659     const char *fmt_str;
4660     tree fn, fmt, arg;
4661
4662     /* If the return value is used, don't do the transformation.  */
4663     if (target != const0_rtx)
4664         return 0;
4665
4666     /* Verify the required arguments in the original call.  */
4667     if (! arglist)
4668         return 0;
4669     fmt = TREE_VALUE (arglist);
4670     if (TREE_CODE (TREE_TYPE (fmt)) != POINTER_TYPE)
4671         return 0;
4672     arglist = TREE_CHAIN (arglist);
4673
4674     /* Check whether the format is a literal string constant.  */
4675     fmt_str = c_getstr (fmt);
4676     if (fmt_str == NULL)
4677         return 0;
4678
4679     /* If the format specifier was "%s\n", call __builtin_puts(arg).  */
4680     if (strcmp (fmt_str, "%s\n") == 0)
4681     {
4682         if (! arglist
4683             || TREE_CODE (TREE_TYPE (TREE_VALUE (arglist))) != POINTER_TYPE
4684             || TREE_CHAIN (arglist))
4685             return 0;
4686         fn = fn_puts;
```

```

4687     }
4688     /* If the format specifier was "%c", call __builtin_putchar(arg). */
4689     else if (strcmp (fmt_str, "%c") == 0)
4690     {
4691         if (! arglist
4692             || TREE_CODE (TREE_TYPE (TREE_VALUE (arglist))) != INTEGER_TYPE
4693             || TREE_CHAIN (arglist))
4694             return 0;
4695         fn = fn_putchar;
4696     }
4697     else
4698     {
4699         /* We can't handle anything else with % args or %% ... yet. */
4700         if (strchr (fmt_str, '%'))
4701             return 0;
4702
4703         if (arglist)
4704             return 0;
4705
4706         /* If the format specifier was "", printf does nothing. */
4707         if (fmt_str[0] == '\0')
4708             return const0_rtx;
4709         /* If the format specifier has length of 1, call putchar. */
4710         if (fmt_str[1] == '\0')
4711         {
4712             /* Given printf("c"), (where c is any one character,)
4713              convert "c"[0] to an int and pass that to the replacement
4714              function. */
4715             arg = build_int_2 (fmt_str[0], 0);
4716             arglist = build_tree_list (NULL_TREE, arg);
4717             fn = fn_putchar;
4718         }
4719     else
4720     {
4721         /* If the format specifier was "string\n", call puts("string"). */
4722         size_t len = strlen (fmt_str);
4723         if (fmt_str[len - 1] == '\n')
4724         {
4725             /* Create a NUL-terminated string that's one char shorter
4726              than the original, stripping off the trailing '\n'. */
4727             char *newstr = (char *) alloca (len);
4728             memcpy (newstr, fmt_str, len - 1);
4729             newstr[len - 1] = 0;
4730
4731             arg = build_string_literal (len, newstr);
4732             arglist = build_tree_list (NULL_TREE, arg);

```

```
4733         fn = fn_puts;
4734     }
4735     else
4736         /* We'd like to arrange to call fputs(string,stdout) here,
4737            but we need stdout and don't have a way to get it yet. */
4738         return 0;
4739     }
4740 }
4741
4742 if (!fn)
4743     return 0;
4744 return expand_expr (build_function_call_expr (fn, arglist),
4745                  target, mode, EXPAND_NORMAL);
4746 }
```